IBM Virtualization Engine TS7700 Release 1.4a:

Tape Virtualization for System z Servers

Consolidate and protect your data with a Three-Cluster Grid configuration

Prepare for Disaster Recovery and implement Copy Export

Protect your data using TS1120 Tape Encryption

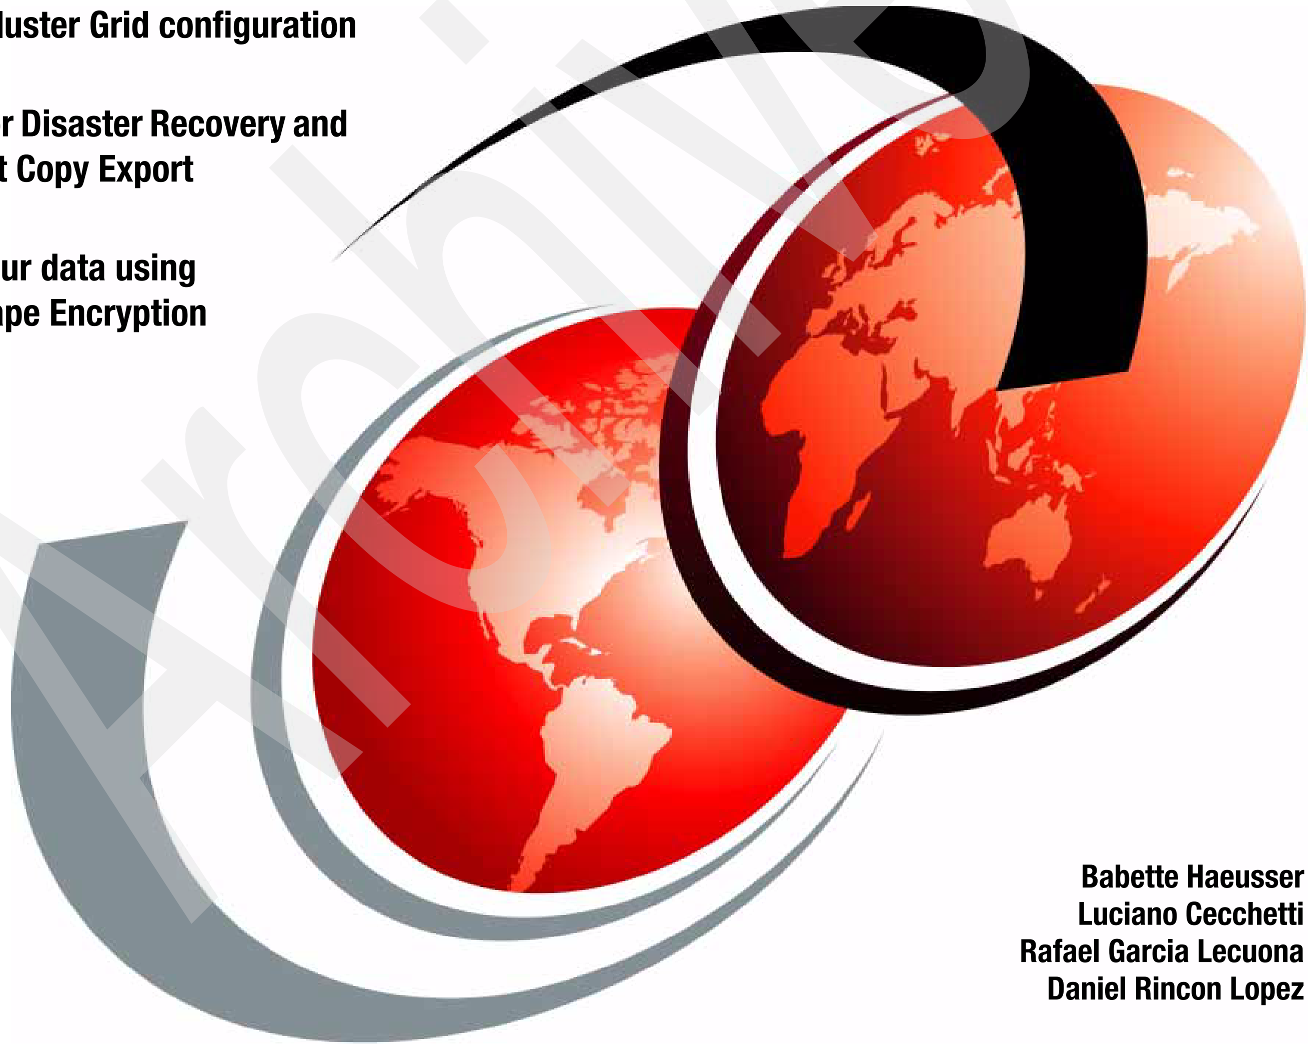

Babette Haeusser
Luciano Cecchetti
Rafael Garcia Lecuona
Daniel Rincon Lopez

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM Virtualization Engine TS7700 Release 1.4a: Tape Virtualization for System z Servers

February 2009

SG24-7312-02

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

Third Edition (February 2009)

This edition applies to the IBM Virtualization Engine TS7700 Release 1.4a and IBM Tape Drives and Tape Libraries current at the time of general availability of TS7700 R1.4a.

Contents

Notices

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

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

IBM might have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

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

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You might copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks might also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

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

AIX®
AS/400®
CICS®
DB2®
DS8000®
ESCON®
FICON®
GDPS®
Geographically Dispersed Parallel Sysplex™
IBM®
iSeries®
Netfinity®
OS/400®
Parallel Sysplex®
PowerPC®
pSeries®
RACF®
Redbooks®
Redbooks (logo) ®
RETAIN®
RS/6000®
S/390®
System i®
System p5®
System p®
System Storage™
System x®
System z9®
System z®
Tivoli®
TotalStorage®
VM/ESA®
WebSphere®
xSeries®
z/OS®
z/VM®
z/VSE™
z9®
zSeries®

The following terms are trademarks of other companies:

Snapshot, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

ELC, Java, JavaScript, S24, Sun, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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

Preface

The IBM® Virtualization Engine TS7700 provides tape virtualization for the System z® environment. As the follow-on product to the highly successful IBM TotalStorage® Virtual Tape Server, the TS7700 Virtualization Engine is designed to provide improved performance and capacity to help lower the total cost of ownership for tape processing. It introduces a new modular, scalable, high-performing architecture for mainframe tape virtualization.

The TS7700 Virtualization Engine integrates the advanced performance, capacity, and data integrity design of the 3592 Tape Drives, industry-leading tape technology from IBM, with high-performance disk and a new advanced System p® server to form a storage hierarchy managed by robust storage management firmware with extensive self management capability. It includes functions such as advanced policy management to control physical volume pooling, cache management, dual copy, dual copy across a grid network, and copy mode control. The TS7700 offers a new standards-based management interface and enhanced statistical reporting, compared to the VTS.

This IBM Redbooks® publication provides a detailed description of the TS7700 Virtualization Engine. In addition, it discusses the tasks involved with planning, implementing, and operating the TS7700 Virtualization Engine and describes migration considerations.

This book covers the TS7700 Virtualization Engine releases up to and including TS7700 R1.4a and includes information about supported tape drives and libraries current and supported at the time of writing. For detailed information about the IBM Virtualization Engine TS7740 R1.5 and the disk-only TS7720 Virtualization, refer to *IBM Virtualization Engine TS7740 R1.5 and TS7720: New Virtualization Options for Mainframe Servers,* SG24-7712, which will be available in early 2009.

The team that wrote this book

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

Babette Haeusser is an IBM Certified IT Specialist at the International Technical Support Organization, San Jose Center. She writes extensively and teaches IBM classes worldwide on all areas of Enterprise and Open Systems Tape, Tape Libraries, Tape Encryption, and Tape Virtualization. Babette joined IBM in 1973 as an application programmer. In 1987, she became an MVS Systems Engineer and specialized on IBM Storage hardware and software, which she has supported in various job roles since then. Before joining the ITSO in early 2005, Babette worked in the Advanced Technical Sales Support EMEA. She led a team of specialists for Enterprise Storage. Babette focused on Enterprise Tape, including tape libraries and Virtual Tape Servers.

Luciano Cecchetti is an IBM Senior Accredited Product Service, working since 1985 in technical support roles in Rome, Italy. In 1988, he provided second-level support for ES/9000 during an international assignment at the IBM European Product Support Group in Montpellier, France. As a specialist member of the RMSS SWAT Team, he supported and coordinated the first 3494 VTS/Library installations in Portugal and Greece. Since 2002 he holds a Specialist (CS) degree in High End Tape Solutions from the IBM Professional Certification Program. Since 2007 he holds a nomination for Senior membership of

World-Wide Product Services Profession. His areas of expertise include CPUs, disks, and tapes in complex mainframe environments.

Rafael Garcia Lecuona is a Senior Accredited IT Specialist in IBM Spain. Rafael joined IBM in 1984 in HW Customer Engineering in Spain. In 1987, he moved to the MVS Software Support Center. He has more than 20 years of experience in Storage Solutions and in implementing Tape, DASD, and Storage Software IBM solutions. He currently works in Service Delivery in the ITS organization.

Daniel Rincon Lopez is a Field Technical Sales Specialist in Mexico. He joined IBM in 1998 after finishing Electronic Engineering studies. Daniel also has a master degree in Business Administration and 10 years of experience in Systems, with the last 8 years in storage. When starting in IBM, he did SW Support and Services for MVS for 2 years. From 1999 to 2003 he performed Storage Implementation Services, including ESS and 3494/VTS implementations. Since 2004, Daniel has worked as a Field Technical Sales Specialist with primary focus on tape solutions including VTS and TS7700. As a Regional Designated Specialist for tapes in Mexico, Daniel is in charge of the necessary training for technical sales specialists in Latin America.

The team: Daniel, Rafael, Babette, and Luciano

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

Ann Lund, Alex Osuna, Sangam Racherla
International Technical Support Organization, San Jose Center

Wayne Carlson, Jon Peake, Tom Bish, Jesse Thrall, Greg Kishi, Stefan Lehmann, Jennifer Mason, Jim Fisher, Lawrence Fuss, Erika Dawson, Corie Neri,
IBM Tucson

Thanks to the authors of the previous editions of this book.

- Authors of the first edition, *IBM Virtualization Engine TS7700 Release 1.4a: Tape Virtualization for System z Servers*, published in April 2007, were:

  Luciano Cecchetti, Babette Häußer, Axel Melber, Stefan Neff

- Authors of the second edition, *IBM Virtualization Engine TS7700 Release 1.4a: Tape Virtualization for System z Servers*, published in February 2008, were:

  Soren Aakjar, Jim Fisher, Babette Häußer, Mauro Silvestri

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Summary of changes

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

Summary of Changes
for SG24-7312-02
for *IBM Virtualization Engine TS7700 Release 1.4a: Tape Virtualization for System z Servers*
as created or updated on February 17, 2009.

February 2009, Third Edition

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

New information

- Enhancements for TS7700 R1.4a, including the removal of a cluster from a grid

February 2008, Second Edition

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

New information

- TS1120 Tape Encryption
- Three-Cluster Grid configuration
- Copy Export function
- IBM 3494 Tape Library support
- 1,000,000 logical volumes
- Secure Data Erase
- Host Console Request Facility

Changed information

- Migration and Upgrade paths
- Statistics and Bulk Volume Information Retrieval

1

Introducing the TS7700 Virtualization Engine

The IBM Virtualization Engine TS7700 is the newest member of the IBM TS7000 Virtualization Family. It represents the fourth generation of IBM Tape Virtualization for mainframe systems and replaces the highly successful IBM TotalStorage Virtual Tape Server (VTS).

The TS7700 Virtualization Engine is designed to provide improved performance and capacity to help lower the total cost of ownership for tape processing. It introduces a new modular, scalable, high-performing architecture for mainframe tape virtualization. It integrates the advanced performance, capacity, and data integrity design of the IBM 3592 Family of Tape Drives, with high-performance disks and a new advanced IBM System p server to form a storage hierarchy managed by robust storage management firmware with extensive self-management capabilities.

The TS7700 Virtualization Engine uses outboard policy management to manage physical volume pools, cache management to control selective dual copy, dual copy across a grid network, copy mode control, encryption, and copy export.

The TS7700 offers a new standards-based management interface and enhanced statistical reporting, compared to the VTS.

The TS7700 Virtualization Engine integrates the following components into the virtual tape solution:

- One *IBM Virtualization Engine TS7740 Server* Model V06 (3957 Model V06)
- One *IBM Virtualization Engine TS7740 Cache Controller* Model CC6 (3956 Model CC6)
- Zero, one, or three *IBM Virtualization Engine TS7740 Cache Drawers* Model CX6 (3956 Model CX6)

Important characteristics of these components are summarized here:

- The TS7740 Server provides host connections of up to four FICON® channels, and connections to the tape library and tape drives for backend tape processing.
- A TS7700 with Grid Enablement features can be interconnected with one or two other TS7700s to provide peer-to-peer copy capability between Virtualization Engines for tape

using IP network connections. It also has the removal of a cluster from a grid functionality for specific requirements such as data center movement.

- The TS7740 Cache, comprised of the TS7740 Model CC6 and the TS7740 Model CX6, provides from 1 TB to 6 TB of tape volume cache capacity before compression.
- Each TS7700 supports up to a maximum of 256 3490E virtual tape drives.
- A TS7700 Single or Multi Cluster Grid supports up to 1,000,000 logical volumes, each logical volume having a maximum capacity of 1.2 GB to 12 GB (assuming 3:1 compression and using the 400 to 4000 MB volume sizes).

Figure 1-1 shows the main components of the TS7700 Virtualization Engine.

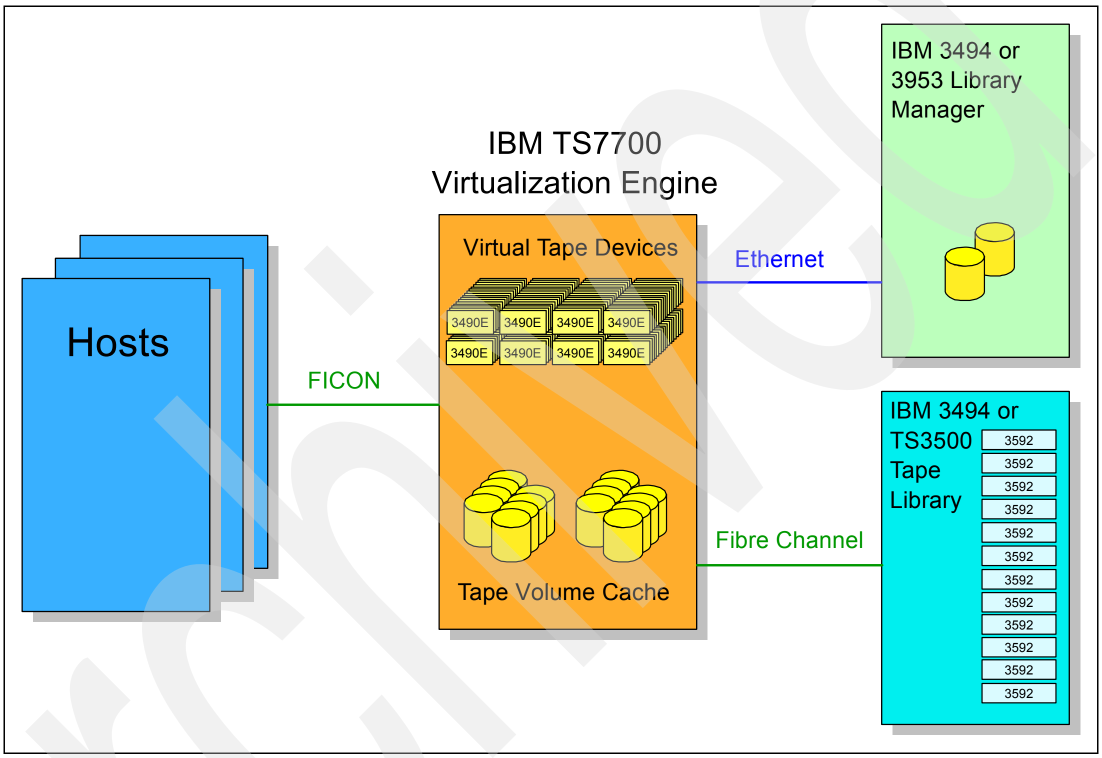

Figure 1-1 Main components of the TS7700 Virtualization Engine

1.1 Concepts of storage virtualization

A virtual tape subsystem presents emulated tape drives to the host and stores tape data on emulated tape volumes in a disk-based cache rather than on physical tape media. The TS7740 Node emulates the function and operation of IBM 3490 Enhanced Capacity (3490E) tape drives and uses a RAID5 disk subsystem to store volumes written by the host. The disk space provided is called *Tape Volume Cache (TVC)*.

Emulated tape drives are also called *virtual drives*. To the host, virtual 3490E tape drives look exactly the same as physical 3490E tape drives. Emulation is transparent to the host and to applications. The host always writes to and reads from virtual tape drives, it never accesses the physical tape drives in the backend. In fact, it does not need to know that these tape drives exist. As a consequence of this, even an application that only supports 3490E tape technology can use the TS7700 Virtualization Engine without any changes to the application and still benefit from high capacity/high performance tape drives in the backend. Tape virtualization provides you with a large number of virtual devices, up to 256 per TS7700

Virtualization Engine cluster. When jobs are contending for native attached physical drives, tape virtualization can solve this issue by providing a large number of virtual devices for concurrent use.

Because the host exclusively accesses the virtual tape drives, all data must be written to or read from emulated volumes in the disk-based Tape Volume Cache (TVC). We call these emulated tape volumes in the Tape Volume Cache *virtual volumes*.

When the host requests a volume that is still in cache, the volume will be virtually mounted. No physical mount is required. When the virtual mount is complete the host can access the data at disk speed. Mounting of scratch tapes also is virtual and does not require a physical mount.

Although you define maximum sizes for your volumes, a virtual volume takes up just the space in cache that the data on the volume actually requires. Later, when a virtual volume is copied from disk to tape, it will also need only the amount of tape capacity occupied by the data. Tape virtualization makes very efficient use of disk and tape capacity.

Another benefit you get from tape virtualization is the large number of drives available to applications. Each TS7700 Virtualization Engine provides you with a maximum of 256 virtual tape devices. If you have a situation where applications are contending for tape drives and jobs have to wait because no physical tape drive is available, tape virtualization can efficiently address these issues by providing a large number of virtual tape drives.

The TS7740 Node manages the physical tape drives and physical volumes in the tape library and controls the movement of data between physical and virtual volumes.

Data that is written from the host into the tape volume cache is scheduled for being copied to tape at a later point in time. The process of copying data to tape that still exists in cache is called *premigration*. When a volume has been copied from cache to tape, the volume on the tape is called a *logical volume*. A physical volume can contain a large number of logical volumes. The process of putting several logical volumes on one physical tape is called stacking, and a physical tape containing logical volumes is, therefore, referred to as a *stacked volume*.

Because many applications are not able to fill the high capacity media of today's tape technology, you can end up with a large number of under-utilized cartridges, wasting a lot of space on your physical media and requiring an excessive number of cartridge slots in your tape automation system. Tape virtualization reduces the space required by volumes and fully utilizes the capacity of current tape technology. Tape virtualization enables you to exploit the full potential of modern tape drive and tape media technology without making any changes to your applications or JCL.

When space is required in the TVC for new data, volumes that already have been copied to tape will be removed from the cache. By default, removal is based on a *least recently used* (LRU) algorithm to ensure that no new data or recently accessed data is removed from cache. The process of copying volumes from cache to tape and afterwards deleting them is called *migration*. Accordingly, volumes that have been deleted in the cache, and now only exist on tape, are called *migrated volumes*.

When the host needs to access a volume that was migrated previously, the volume has to be copied back from tape into the TVC, because the host has no direct access to the physical tapes. When the complete volume has been copied back into the cache, the host is able to access the data on this volume. The process of copying data back from tape to the Tape Volume Cache is called *recall*.

Figure 1-2 summarizes these processes described.

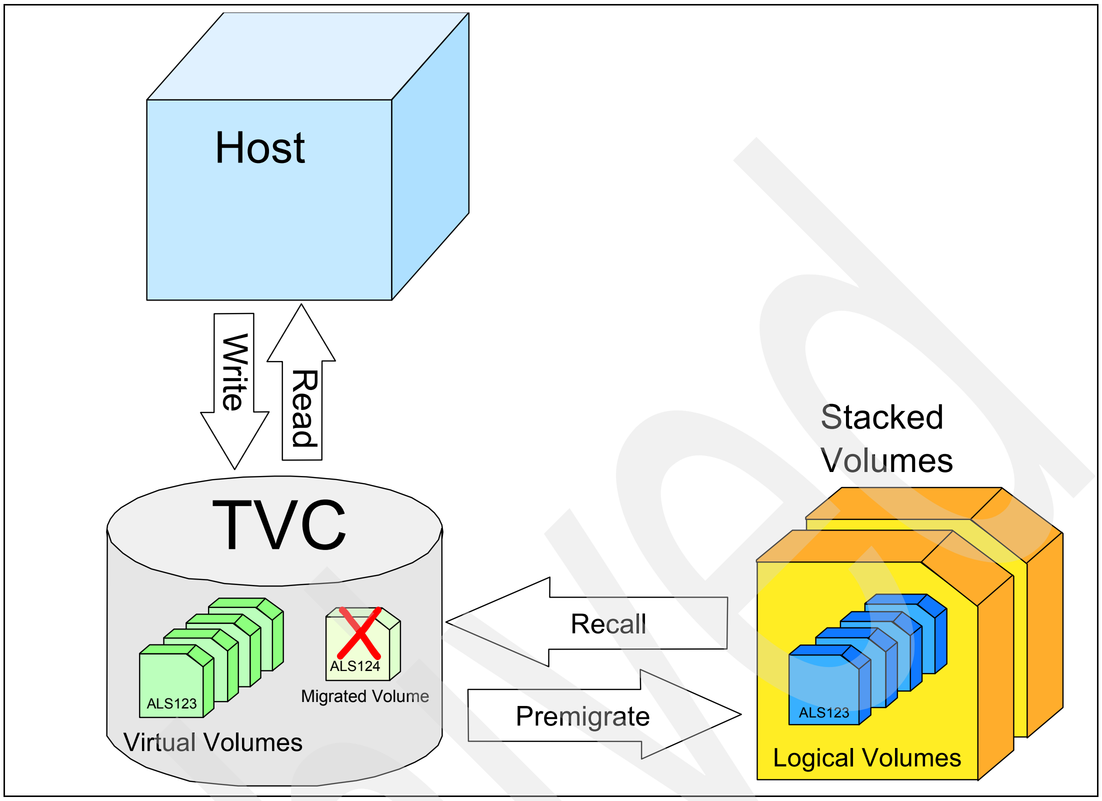

Figure 1-2 Tape Volume Cache processing

Two or three TS7700 Virtualization Engines can be interconnected with two 1 Gb Ethernet links to form a *Multi Cluster Grid configuration*. Logical volume attributes and data are replicated across the clusters in a grid. Any data replicated between the clusters is accessible through any other cluster in a grid configuration. Through remote volume access, you can reach any virtual volume through any virtual device, even if a replication has not been made.

By setting up policies on the TS7700s you define where and when you want to have a secondary or tertiary copy of your data. You can also specify, for certain kinds of data, for example test data, that you do not need a secondary or tertiary copy.

Depending on the configuration, multiple TS7700 Virtualization Engines forming a grid provide you with a High Availability, Disaster Recovery, or High Availability/Disaster Recovery solution.

A Multi Cluster Grid configuration presents itself to the attached hosts as one large library with a maximum of 512 virtual devices for a Two-Cluster Grid or 768 virtual tape devices for a Three-Cluster Grid. The copying of the volumes in a grid configuration is handled by the clusters and is completely transparent to the host. Each TS7700 Virtualization Engine in a grid manages its own set of physical volumes and maintains the relationship between logical volumes and the physical volumes on which they reside.

In a Multi Cluster Grid configuration, you can permanently have one cluster removed from the grid in case you no longer require one of the TS7700s to participate in the grid. After removal, the removed cluster is disabled and cannot be returned to active use until a Cluster Cleanup process has been performed.

1.2 Benefits of tape virtualization

You can expect the following main benefits from tape virtualization:

- High Availability and Disaster Recovery configurations
- Fast access to data through caching on disk
- Utilization of current tape drive, tape media, and tape automation technology
- Capability of filling high capacity media 100%
- Large number of tape drives available for concurrent use
- No additional software required
- Reduced Total Cost of Ownership (TCO)

2

Architecture, components, and functional characteristics

In this chapter, we provide a detailed description of the components, the processes, the underlying concepts and the architecture specific to the IBM System Storage™ TS7700 Virtualization Engine.

We cover the following topics:

- Terms and expressions that are used throughout this book describing the IBM System Storage TS7700 Virtualization Engine
- Hardware components that make up a TS7700 Virtualization Engine
- Attachment of the TS7700 Virtualization Engine to an IBM System Storage Tape Library TS3500 using the IBM System Storage 3953 Tape System
- Attachment of the TS7700 Virtualization Engine to an IBM 3494 Tape Library
- Underlying concepts of tape virtualization with the TS7700
- Architecture of the TS7700 Virtualization Engine and the architecture's potential for future development
- Functional characteristics of the TS7700 Virtualization Engine

For more information about TS3500 Tape Library attachment to the TS7700, refer to *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

2.1 Terminology and definitions

In this section we offer short descriptions and definitions of terms that are used throughout this book.

If you know the IBM Virtual Tape Server, the predecessor of the TS7700 Virtualization Engine, many of the terms will look familiar to you. Some of the terms have not changed, some have a new meaning, and others were introduced with the TS7700 Virtualization Engine. In order to avoid confusion we point out similarities and differences to the previous virtualization solution wherever it seems appropriate.

2.1.1 TS7700-specific terminology

In this section we discuss terms that are introduced with the TS7700 Virtualization Engine.

Nodes

The TS7700 Virtualization Engine is built on a distributed node architecture. The nodes perform either virtualization (vNode) or hierarchical data storage management (hNode). Based on the node architecture, a vNode or hNode can run on separate virtualization hardware or can be combined to run on the same virtualization hardware. When a vNode and hNode are combined and run on the virtualization hardware, this combination of nodes is referred to as a gNode (general node). The current version of the TS7700 Virtualization Engine runs a gNode. For detailed information about node architecture, refer to 2.5, “Architectural capabilities” on page 73.

vNode

The vNode (virtualization node) presents the image of virtual drives to a host system. It receives tape drive and library requests from the host and processes them as real devices would. It also translates the tape requests through a virtual drive and uses a file in a file system to represent the virtual tape image.

hNode

The hNode performs all management of a logical volume residing in disk cache or physical tape after it has been created or altered by the host system through a vNode. The hNode is the only node that is aware of physical tape resources and the relationships between the logical volumes and physical volumes. It is also responsible for any replication of the logical volumes and their attributes across site boundaries.

gNode

The gNode can be considered a vNode and hNode sharing the same physical controller. It can also be considered the equivalent of the IBM Virtual Tape Server in terms of a single controller having both the virtualization and hierarchical storage management capabilities.

Cluster

The TS7700 Virtualization Engine Cluster combines the Virtualization Engine with a disk subsystem TS7740 cache, as shown in Figure 2-1. The architecture allows additional disks or nodes to be added in the future to expand the capabilities of the system. A cluster provides FICON host attachment and 256 virtual tape device addresses. The cluster also includes that part of the 3494 or TS3500 Tape Library which is assigned to the TS7700 Virtualization Engine.

Figure 2-1 shows the components of a TS7700 cluster.

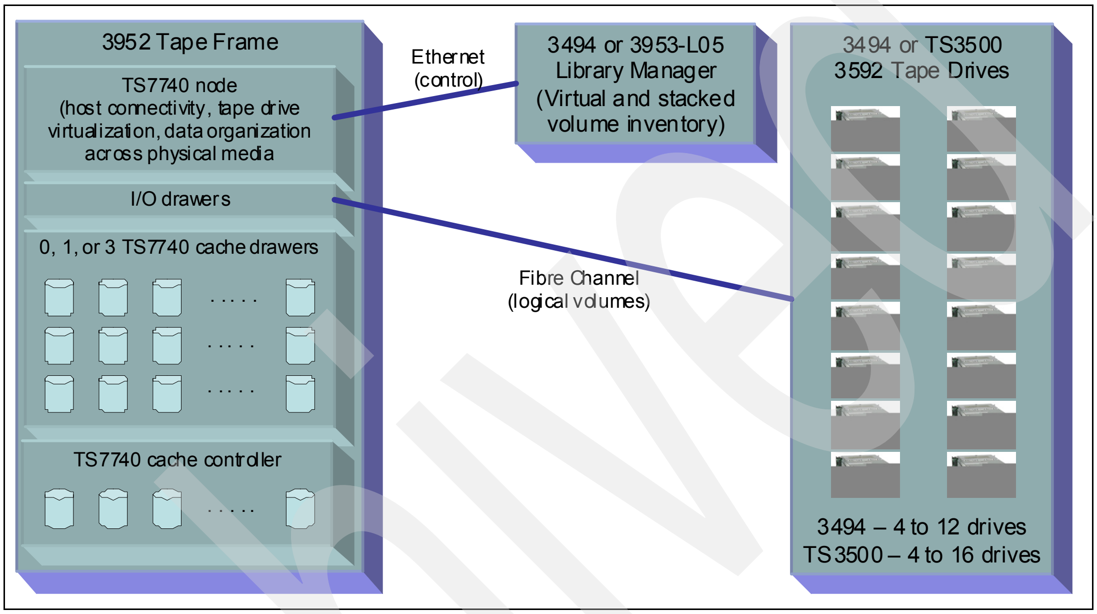

Figure 2-1 TS7700 cluster

Grid

Two or three TS7700 clusters can be interconnected to provide a disaster recovery/high availability solution. The grid enablement feature must be installed on all TS7700 Virtualization Engines in the grid, and an Ethernet connection must be established between the clusters. Each cluster has two 1 Gbps Ethernet adapters and uses TCP/IP for communication between them. The 1 Gbps Ethernet adapters are available in RJ45 copper and short wave optical fiber versions.

Logical volume attributes and data are replicated across the clusters in a grid configuration to ensure the continuation of production work, should a single cluster become unavailable. Any data replicated between the clusters is accessible through any of the other clusters in a grid configuration. Through remote file access, you can get to a remote volume's copy of the data through any virtual device, assuming the entire Grid WAN infrastructure is not failed.

A Multi Cluster Grid configuration presents a single Composite Library image to the host, such that the entire subsystem appears as a single tape subsystem to the attached hosts. The host does not see the underlying distributed libraries just as with the prior generation's Peer-to-Peer Virtual Tape Server.

Multiple TS7700 Virtualization Engine grid configurations can be attached to host systems and operate independently of one another. TS7700 Virtualization Engine provides the following grid configurations:

Single Cluster Grid A single TS7700 Virtualization Engine. The TS7700 architecture has the grid functionality built in. Although the Single Cluster Grid is not actually a grid, it is prepared to participate in a Multi Cluster Grid.

Two-Cluster Grid Consists of two single clusters attached across a LAN or WAN through the 1 Gbps Ethernet links, also known as a *Dual Cluster Grid*.

Three-Cluster Grid Consists of three single clusters attached to each other across a LAN or WAN through the 1 Gbps Ethernet links. Each cluster can access the other two clusters.

Multi Cluster Grid Consists of a Two- or Three-Cluster Grid. The architecture is designed to allow more than three clusters in a Multi Cluster Grid configuration in the future.

Figure 2-2 shows a Two-Cluster Grid configuration. Each cluster contains the TS7700 Virtualization engine and a tape library with tape drives. The tape library in each cluster can be either the IBM 3494 Tape Library or the IBM TS3500 Tape Library.

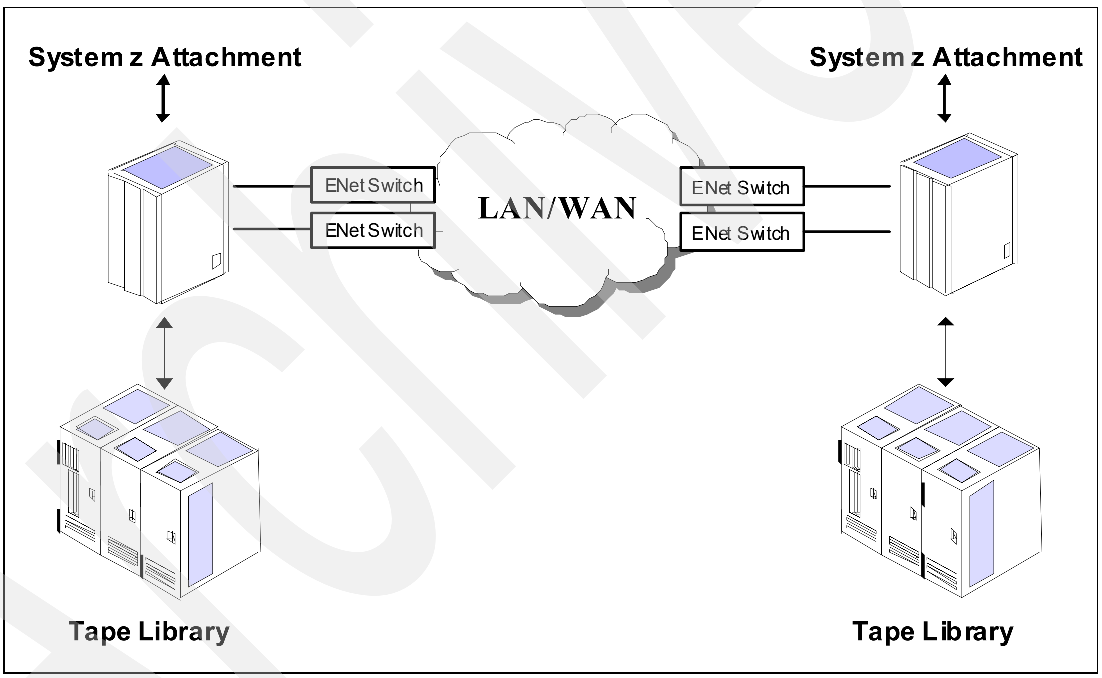

Figure 2-2 TS7700 Two-Cluster Grid

Figure 2-3 shows a Three-Cluster Grid. Just like the Two-Cluster Grid, each cluster contains a TS7700 Virtualization Engine and a tape library with tape drives. The tape library in each cluster can be either the IBM 3494 Tape Library or the IBM TS3500 Tape Library.

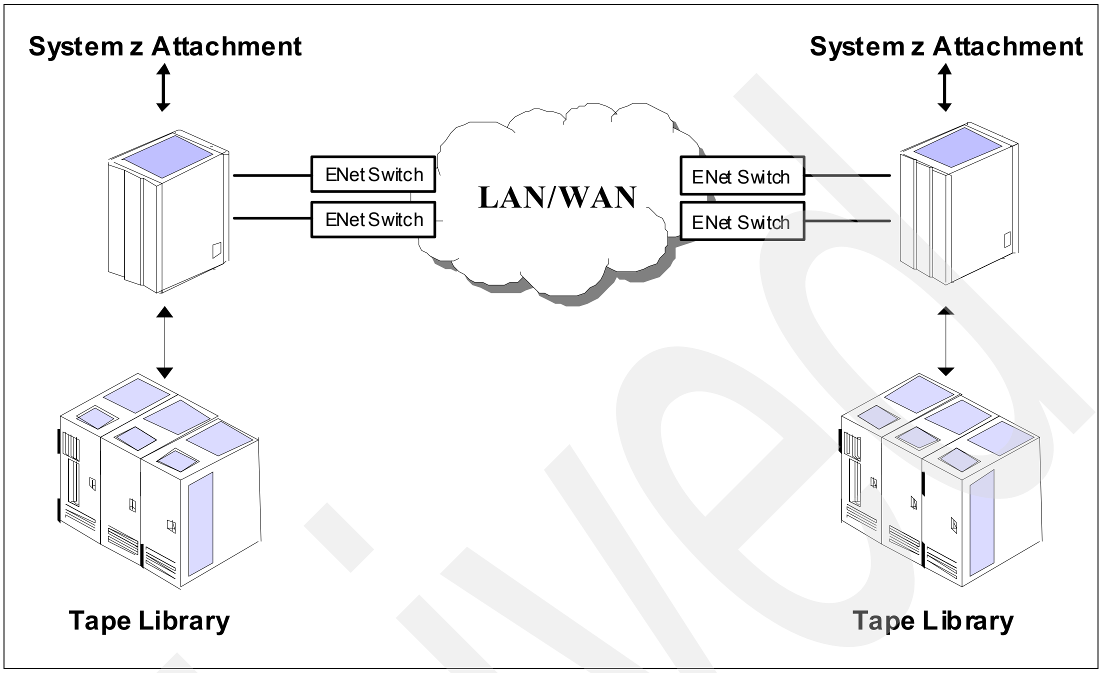

Figure 2-3 TS7700 Three-Cluster Grid

Ownership

The concept of ownership has been introduced with TS7700 to ensure data integrity. At any point in time a logical volume is owned by one cluster. The owning cluster controls access to the volume and the attributes of a volume. Ownership can change dynamically. If a cluster needs to mount a logical volume on one of its virtual devices and it is not the owner of that volume, it has to obtain ownership for this specific volume first. We describe these processes in detail in 2.4.1, “Data integrity by volume ownership” on page 51.

I/O Tape Volume Cache

All vNodes in a grid have direct access to all logical volumes in the grid, either from the local cluster’s Tape Volume Cache using Fibre Channel or from a remote cluster’s TVC using a WAN connection. During mount processing a TVC is selected as I/O TVC. All I/O operations associated with the virtual tape drive the mount was issued on are routed from its vNode to the TVC on the cluster selected. For details, see “I/O TVC selection” on page 57.

Management interface

The Management Interface is a Web-based interface, which was introduced with the TS7700 Virtualization Engine. It is used to monitor and configure the system, to manage access, and to manage logical volumes. Many of the tasks that were previously done using the Library Manager Console have been consolidated into this interface. We cover the Management Interface in more detail in Chapter 4, “Hardware implementation” on page 131 as well as in Chapter 7, “Operation” on page 303.

Note: Insertion and deletion of logical volumes is managed through the Management Interface, not through the Library Manager Console as for the Virtual Tape Server.

User roles

Users of the TS7700 Virtualization Engine can be assigned one or more roles associated with performance of certain system functions.

User roles are levels of access, assigned by the administrator, that allow users to perform certain functions. User roles are created using the TS7700 Management Interface. When, as an administrator, you create a new user account, you must specify an initial password for the account. You cannot assign multiple roles to one user.

Administrators can assign users to the following roles:

- **Administrator**

  The administrator has the highest level of authority, including the authority to add or remove user accounts. The administrator has access to all service functions and TS7700 Virtualization Engine resources.

- **Lead Operator**

  The lead operator has permissions nearly identical to those of the administrator, but *cannot* execute the following actions:

  – Perform Standalone Mount/Demount Operations
  – Configure Cluster Nodes to Modify Cluster Node Nicknames
  – Configure Set Write Protect Mode
  – Configure Modify Copy Policy Override
  – Configure Modify Cluster Network Settings
  – Configure Modify Cluster Identification Properties
  – Configure Modify Grid Identification Properties
  – Configure Feature License Modify
  – Add a user
  – Modify a user
  – Remove a user
  – Modify Custom Roles
  – Modify Security Policy
  – Force the system online
  – Set Service Mode
  – Modify Ownership Takeover Mode Manager

- **Manager**

  The manager has access to health and monitoring information, jobs and processes information, and performance data and functions but is restricted from changing most settings, including those for logical volume management, network configuration, and feature licenses as listed below:

  – Insert Logical Volumes
  – Delete Logical Volumes
  – Perform Standalone Mount/Demount operations
  – Configure Cluster Nodes to Modify Node Nicknames
  – Configure Set Write Protect Mode
  – Configure Modify Copy Policy Override
  – Configure Modify Cluster Network Settings
  – Configure Modify Cluster Identification Properties
  – Configure Modify Grid Identification Properties
  – Configure Feature License Modify
  – Force the system online
  – Set Service Mode
  – Modify Ownership Takeover Mode

- **Operator**

  The operator has access to monitoring information, jobs and processes information, and performance and statistics information, and can perform some network configuration tasks, but is restricted from changing settings for performance, feature licenses, user accounts or custom roles. The operator is also restricted from inserting or deleting logical volumes. The list of restrictions is as follows:

  – Insert Logical Volumes
  – Delete Logical Volumes
  – Perform Standalone Mount/Demount operations
  – Configure Cluster Nodes to Modify Node Nicknames
  – Configure Set Write Protect Mode
  – Configure Modify Cluster Network Settings
  – Configure Modify Cluster Identification Properties
  – Configure Modify Grid Identification Properties
  – Configure Feature License Modify
  – Add a user
  – Modify a user
  – Remove a user
  – Modify Custom Roles
  – Modify Security Policy
  – Force the system online
  – Set Service Mode
  – Modify Ownership Takeover Mode

- **Custom roles**

  The administrator can name and define two custom roles by selecting the individual tasks permitted to each custom role. All available tasks are selectable for a custom role, with the exception of creating, modifying, and deleting a user. Figure 2-4 shows the Roles and Permissions panel, including the custom roles.

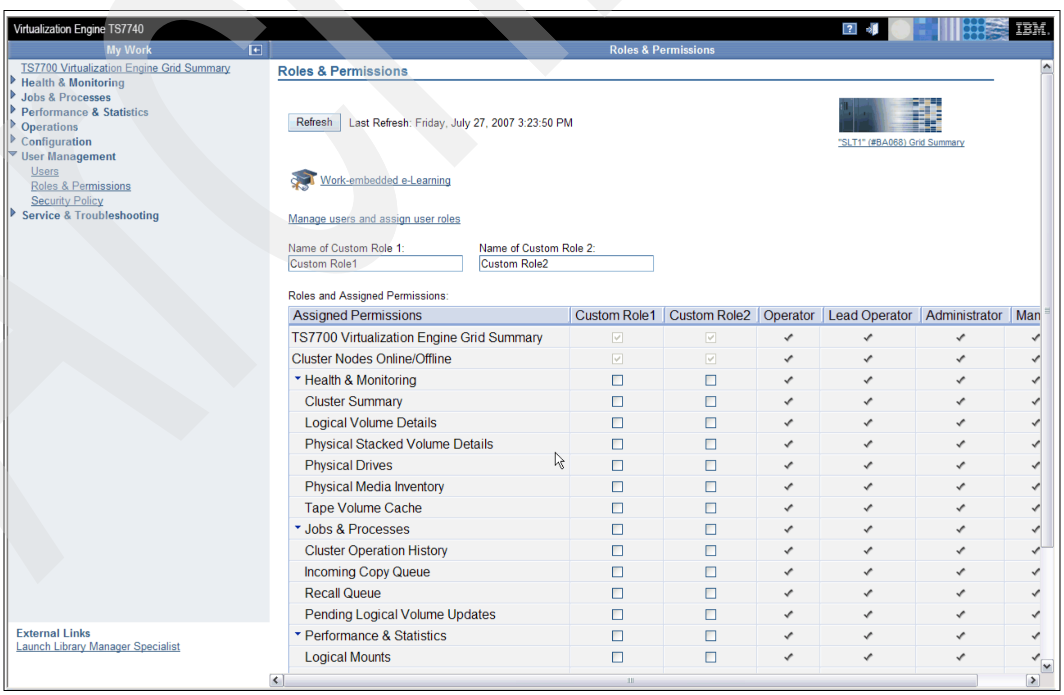

Figure 2-4 TS7700 MI Roles and Permissions panel

2.1.2 Multi Cluster Grid-related terms

In this section we explain terms that are specific to a Multi Cluster Grid configuration.

Composite Library

The Composite Library is the logical image of the grid, which is presented to the host. As opposed to the IBM Virtual Tape Server, a *Single Cluster* TS7700 Virtualization Engine also has a Composite LIBRARY-ID defined. From an architectural perspective, a standalone TS7700 is considered a grid consisting of just one cluster. We refer to such a cluster as a *Single Cluster Grid*.

Both with a TS7700 Single Cluster Grid configuration and a TS7700 Multi Cluster Grid configuration, the Composite Library is presented to the host. In case of a standalone TS77000, the host sees a logical tape library with sixteen 3490E tape control units, each with sixteen IBM 3490E tape drives, attached through two or four FICON channel attachments.

Figure 2-5 illustrates the host view of a Single Cluster Grid configuration.

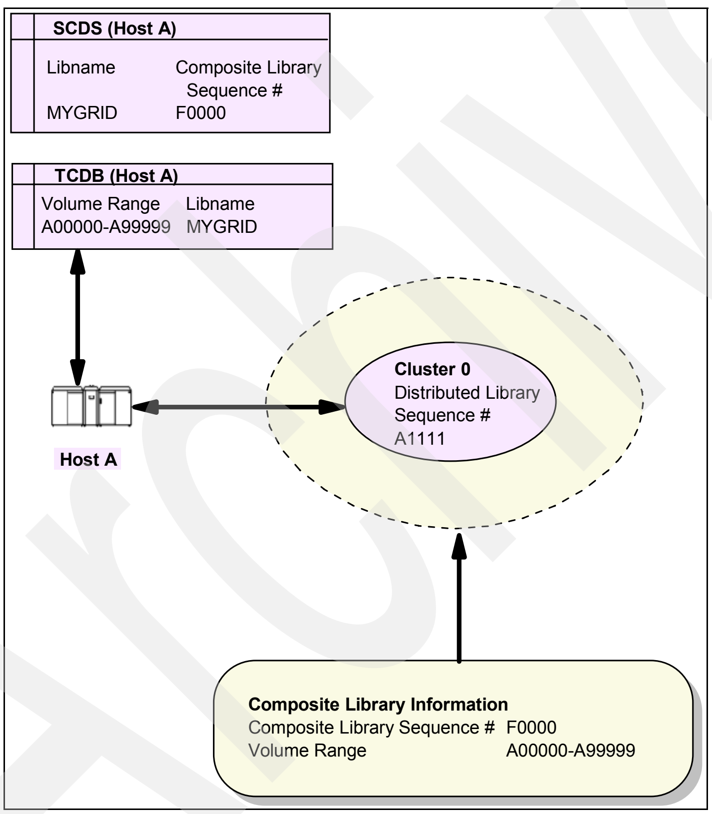

Figure 2-5 TS7700 Single Cluster Grid configuration

In the case of a Multi Cluster Grid, the host sees a logical tape library with 16 3490E tape control units per cluster, each with 16 IBM 3490E tape drives, attached through four, six, eight, 10, or 12 FICON channel attachments.

Figure 2-6 illustrates the host view of a Two-Cluster Grid configuration.

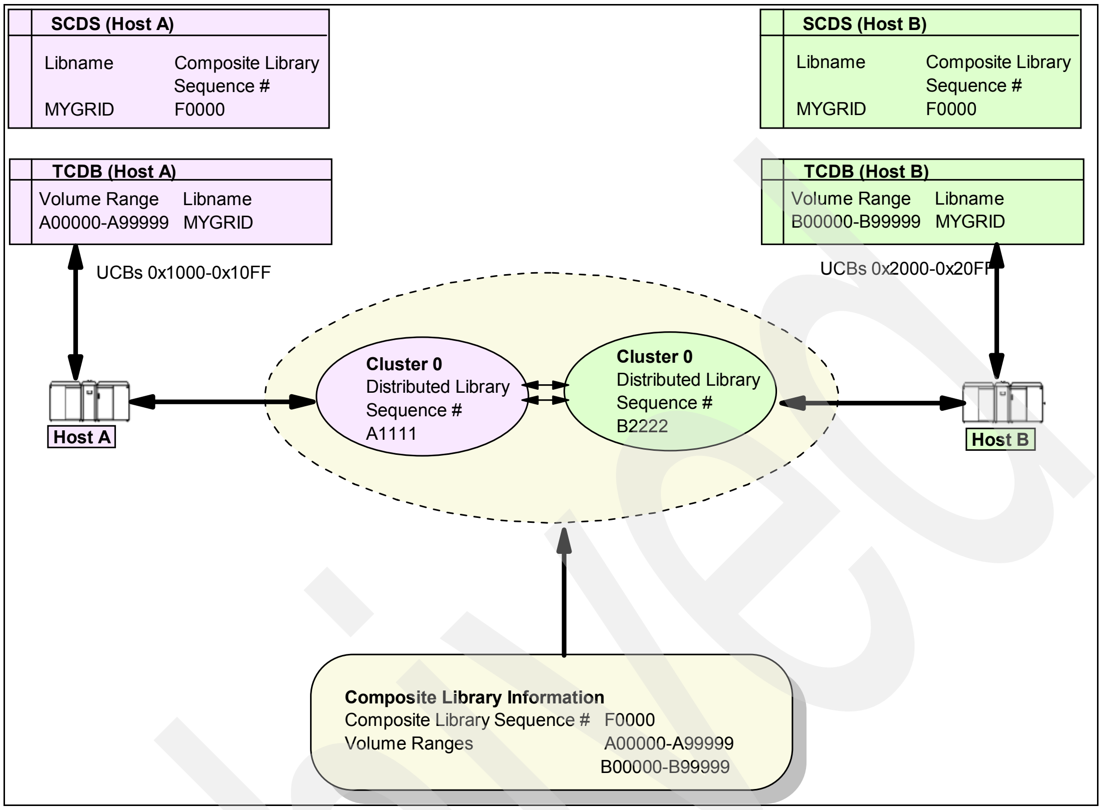

Figure 2-6 TS7700 Two-Cluster Grid configuration

Figure 2-7 illustrates the host view of a Three-Cluster Grid configuration.

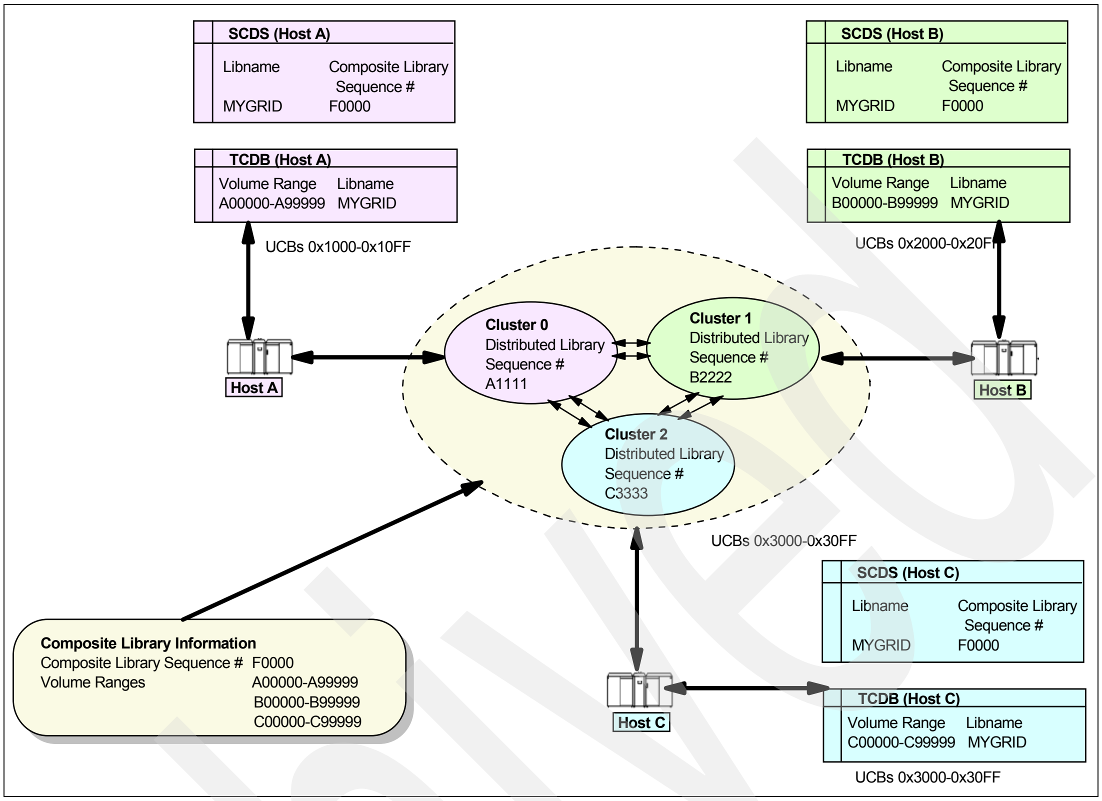

Figure 2-7 TS7700 Three-Cluster Grid configuration

Important: A Composite Library-ID needs to be defined for both a Multi Cluster Grid and also for a Single Cluster Grid. For a Single Cluster Grid, the Composite LIBRARY-ID must not be the same as the Distributed LIBRARY-ID. For a grid configuration, the Composite LIBRARY-ID needs to be different from any of the Distributed LIBRARY-IDs.

Distributed Library

Each cluster in a grid is a Distributed Library, which consists of a TS7700 Virtualization Engine and its attached physical tape library. The Library Manager, either in a 3494 or a TS3500/3953 tape library, interfaces with the TS7700. Within the 3494 or 3953 Library Manager, the TS7700 occupies a Library Manager partition in the same manner as a VTS occupies a Library Manager partition.

A Distributed Library comprises the hardware components of a cluster and consists of:

- Virtualization Engine
- Disk Cache controller
- Disk expansion drawers
- Frame
- Attachment to a physical library through a controller
- Number of physical drives

In a grid configuration that includes multiple TS7700 Virtualization Engines, each of which has the Grid Enablement Feature installed, a Distributed Library must be defined at the host for each cluster. The host has sufficient knowledge about the distributed libraries to allow appropriate console message handling of messages from the Library Manager of a

Distributed Library. On the host, the Distributed Library is only defined to SMS. It is defined using the existing Interactive Storage Management Facility (ISMF) panels and has no tape devices defined. The virtual tape devices are defined for the Composite Library only.

Copy Consistency Point

In a grid environment, you can specify where and when you want to have a copy of your data. Currently there are three different settings, two Copy Consistency Points and an option of having no copy at all. These are the settings:

RUN The copy will occur as part of the Rewind-UNload (RUN) operation and before the rewind-unload operation at the host completes. This mode is comparable to the immediate copy mode of the Peer-to-Peer VTS.

Deferred The copy will occur some time after the rewind-unload operation at the host. This mode is comparable to the deferred copy mode of the Peer-to-Peer VTS.

No Copy No copy will be made.

On each cluster in a Multi Cluster Grid you specify a setting for this cluster itself and one for each of the other clusters. The settings need not be the same on all clusters. When a volume is mounted on a virtual tape device, the Copy Consistency Point policy of the cluster to which the virtual device belongs will be honored.

For a detailed description, see 2.4.2, “Copy policy management” on page 54.

Dynamic Grid Network Balancing

Dynamic Grid Network Balancing functionality controls Grid performance in adverse conditions. This function helps to compensate unbalanced network performance in a Grid configuration due path length that can be different or differences in the equipment that provides the path connection.

Figure 2-8 shows a Two-Cluster Grid with two links connected with two paths with different length. The Primary path in the example is 600 km long, and the secondary path is 800 km long. Dynamic Grid Network Balancing helps to compensate unbalanced network performance due to this condition.

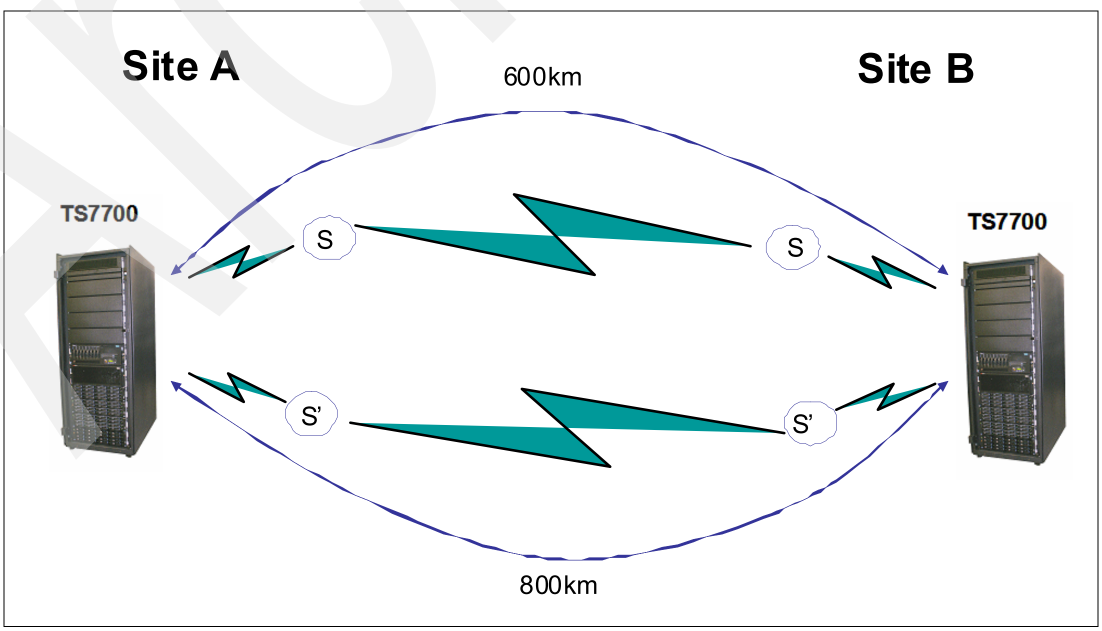

Figure 2-8 Dynamic Grid Network Balancing in a Two-Cluster Grid

With Dynamic Grid Network Balancing, link capabilities are evaluated every 5 minutes, and a host warning message is sent to the console if grid links are degraded severely (this is less than 75% of the link capability).

2.1.3 Tape Virtualization general terms

We now explain the general terms related to Tape Virtualization. If you know the predecessor of the TS7700, the VTS, you will be familiar with some of these terms.

Tape Volume Cache

The Tape Volume Cache (TVC) is disk space where logical volumes are stored. The TVC is under complete and exclusive control of the TS7700.

When a virtual volume in the TVC is closed and demounted, it is scheduled to be copied to a stacked volume. Volumes that have previously been copied to physical tape can be removed from the cache to provide space for new virtual volume data or for the recall of existing logical volumes. Candidates for removal from cache are by default selected by using a least recently used (LRU) algorithm. The only external interface with which to influence this is tape volume cache management. For further details, see 2.3.2, “Tape Volume Cache Management” on page 36.

Virtual drives

From a host perspective, a TS7700 looks like sixteen logical IBM 3490E tape control units, each with sixteen drives attached through FICON channels. Virtual tape drives are defined just like physical IBM 3490 controller addresses, through hardware configuration definition (HCD). Defining a preferred path for the virtual drives gives you no benefit because the IBM 3490 control unit functions inside the TS7700 are emulated to the host.

Each virtual drive has the following characteristics (just like real tape drives):

- It has a host device address.
- It is included in the I/O generation for the system.
- It is varied online or offline to the host.
- It signals ready when a virtual volume is loaded.
- It responds and processes all IBM 3490E I/O commands.
- It becomes not ready when a virtual volume is rewound and unloaded.

For software transparency reasons, the functionality of the 3490E integrated cartridge loader (ICL) is also included in the virtual drive's capability. All virtual drives indicate that they have an ICL. For scratch mounts, there is no benefit in using emulated ICL in the TS7700 to preload virtual cartridges in the ICL, because the virtual volume is created directly in the tape volume cache without the need to copy the data from a stacked volume when the fast-ready attribute is set. No mechanical operation is required to mount a scratch volume.

Physical drives

The physical TS1120 Model E05 and IBM 3952-J1A tape drives used by a TS7700 are installed in the IBM 3494 or IBM TS3500 Tape Library. The physical tape drives are not addressable by any attached host system. They are completely under the control of the TS7700 Virtualization Engine.

If tape drives of both 3592 models (J1A and E05) are attached to the same TS7700 cluster, the TS1120 Model E05 tape drives will operate in J1A Emulation mode. In J1A Emulation mode, the TS1120 drives provide capacity characteristics the same as the J1A tape drive. The data rate drops from 100 MB/s to 50 MB/s. The higher access performance characteristics (load/unload) of the TS1120 are still realized, however.

The TS1120 Tape Drive supports Tape Encryption in native E05 mode.

The Advanced Library Management System (ALMS) of the TS3500 Tape Library allows for dynamic reconfiguration of tape drives, which can be useful in an Open Systems environment, where you might want to assign drives to different hosts temporarily or share them between hosts on a time schedule base.

Important: Do not use ALMS to reassign physical tape drives attached to a TS7700 in an IBM TS3500 Tape Library. This function is used for Open Systems hosts and can cause unpredictable complications when used to reconfigure TS7700 attached tape drives without the assistance of trained IBM service personnel.

Virtual volumes

A virtual volume is created in the Tape Volume Cache when the host writes data to the TS7700 subsystem. All host interaction with tape data in a TS7700 Virtual Engine is through virtual volumes and virtual tape drives.

Each virtual volume, like a real volume, has the following characteristics:

- Has a unique volume serial (VOLSER) number known to the host.
- Is loaded and unloaded on a virtual device.
- Supports all tape write modes, including Tape Write Immediate.
- Contains all standard tape marks and data blocks.
- Supports an IBM standard label (and non-labeled tape as well).
- Can be appended to after it is initially written from the beginning of tape.
- The application is notified that the write operation is complete when the data has been written in the TVC.
- Each host-written record has a logical block ID.
- End of volume is signaled when the total number of bytes written into the TVC after compression has reached 400 MB for an emulated cartridge system tape (CST), 800 MB for an emulated enhanced capacity cartridge system tape (ECCST) volume or 1000, 2000, or 4000 MB using the larger volume size options.

Note: With the use of large logical volume sizes of 2000 or 4000 MB, recalls can take considerably longer, if a volume is not resident in the Tape Volume Cache and has to be copied from physical tape to the cache.

Copy time between clusters is also affected by larger logical volumes because more data must be copied when the amount of data is larger.

Also consider the consequences of large logical volumes and the size of your cache. The TS7700 does not restrict the logical volume size based on the cache configuration. With a 1 TB configuration and 4000 MB logical volumes, where the logical volumes are filled, opening 256 virtual devices with this size of volume would fill the cache and cause throttling. With smaller cache configurations, do not use filled larger logical volume sizes for all devices at the same time.

With data compression based on the IBMLZ1 algorithm by the FICON channel card in a TS7740 Node, the actual host data stored on a virtual CST or ECCST volume can vary from 1.2 GB up to 12 GB (assuming a 3:1 compression ratio). The default logical volume sizes of

400 MB or 800 MB, still used at insert time, can be overwritten at every individual scratch mount by the use of a Data Class construct.

Virtual volumes can exist only in a TS7700. You can direct data to a virtual tape drive by directing it to a system-managed storage (SMS) tape Storage Group inside the TS7700, using the automatic class selection (ACS) routines in a system-managed tape environment. You can also pass Data Class name, Management Class name, Storage Class name, and Storage Group name. The TS7700 then uses these names to determine how to manage a virtual volume. You can define up to 1 000 000 virtual volumes per Single or Multi Cluster Grid. The VOLSERs for the logical volumes are defined through the Management Interface.

You should also associate a "Fast Ready" attribute with a scratch category of VTS virtual volumes, to speed the scratch mount process. Although you might want to use logical volume sizes larger than 400 or 800 MB, you still define CST or ECCST emulated cartridges to the TS7700, simulating MEDIA1 with 400 MB capacity, or MEDIA2 with 800 MB capacity. Virtual volumes go through the same cartridge entry processing as native cartridges inserted in a library.

Logical volumes

When a virtual volume is copied from the TVC to a physical tape cartridge, it becomes a logical volume. When a logical volume is moved from a physical cartridge to the TVC, the process is called recall and the volume becomes a virtual volume again.

Although the TS7700 emulates a 3490E tape of a specific size, 400, 800, 1000, 2000, or 4000 MB, the space used in the TVC is no more than that needed for the number of bytes of data written to the virtual volume. When the virtual volume is written to the physical tape, it uses only the space occupied by the data. In other words, neither the TVC nor the physical tapes are partitioned in any way into 400 MB, 800 MB, 1000 MB, 2000 MB, or 4000 MB segments.

As virtual volumes are copied from the TVC to a physical cartridge, they are stacked on the cartridge end to end, taking up only the space written by the host application. This arrangement maximizes utilization of a cartridge's storage capacity. The storage management software within the TS7740 Node manages the location of the logical volumes on the physical cartridges. The only control the user has over the location of the data is where volume pooling is used. For details, see "Physical Volume Pooling" on page 40.

Through the use of the Data Class constructs, you have the choice of additional logical volume sizes. The default logical volume sizes of 400 and 800 MB are extended to 1000, 2000, and 4000 MB. The default logical volume sizes which are still used at insert time, can be overwritten at every individual scratch mount according to the size specified for the Data Class. All definitions are outboard and can be defined at the Library Manager console or through the Web specialist.

Stacked volume

The physical cartridges used by the TS7700 to store logical volumes are completely under the control of the TS7740 node and are not known to the hosts. The physical volumes are called stacked volumes. The stacked volumes must have unique machine-readable VOLSERs and external labels like any other cartridges in a tape library.

Through the Library Manager, you define which physical cartridges are to be used by the TS7700. Logical volumes stored on those cartridges are mapped by the internal storage management software. When you use pooling, your stacked volumes can be assigned to individual pools. Logical volumes can then be assigned to the stacked volume pools. In the case of out-of-scratch scenarios, pools can be set up to enable "borrowing" from other pools.

The methodology and configuration of volume pooling is covered in “Physical Volume Pooling” on page 40.

Tokens

Tokens are used to track changes to the ownership, data, or properties of a logical volume. The tokens are independent at each site and represent the state of the logical volume at that site:

- Every logical volume has a corresponding token.
- The grid component manages updates to the tokens.
- Tokens are maintained in a DB/2 database coordinated by the local hNodes.

Tokens are an internal data structure which are not directly visible to you.

Hot Tokens

Hot Tokens are special tokens that track changes against sites that are not operational or that are in service mode. In a Two-Cluster Grid, there is only one non-operational cluster and one remaining operational cluster. In a Three-Cluster Grid up to two clusters can be in service mode at one time. We recommend that only one cluster be in service mode at one time to maintain no single points of failure.

The remaining sites track logical volume updates in their own database. When the down site returns back into the grid, all sites share their Hot Tokens. The returning site builds a Hot Token list of outstanding updates it needs to make, and processes them in the background.

A site can come online while Hot Tokens are being processed. Unlike a Peer to Peer VTS, any unresolved tokens are now dynamically resolved if a host issues a command and specifies a volume that has Hot Tokens associated with it.

Service Prep

The transition of a cluster into service is called Service Prep, which allows for a cluster to be gracefully removed from the grid. The remaining sites are given permission to acquire ownership of the volumes owned by the site going into service while it is away from the grid. Composite Library operation ability is surrendered by the site going into service, before transition to service is complete. The host device addresses associated with the site in service send up Device State Change attentions, to have the host remove them from device selection for new mounts.

Nodes at the site in service are still online, but are removed from communication with other sites, allowing service personnel to perform service on the site’s nodes, verify hardware, and so forth, without impacting other sites.

Only one Service Prep can occur within a Composite Library at any one time. If a second Service Prep is attempted at the same time, it is failed. When Service Prep for one cluster has completed and the cluster is in service, then another cluster can be placed in Service Prep.

A site in Service Prep automatically cancels if its peer in the grid experiences an unexpected outage while the Service Prep process is still active. Service Prep can be cancelled by an IBM SSR.

When a cluster has completed Service Prep and is in service, it will stay in service until explicitly taken out of service again by the IBM SSR.

We recommend that only one cluster be placed into service mode at a time. This eliminates the creation of a single point of failure within the grid. Each cluster in a grid can be upgraded in a serial manner. Two of three clusters can be placed in service mode at the same time but we recommend that this be done only during emergencies.

Time coordination

All nodes in the entire subsystem *must* coordinate their time with one another. All nodes in the system keep track of time in relation to Coordinated Universal Time (UTC), also known as Greenwich Mean Time (GMT). Statistics are also reported in relation to UTC.

The TS7700 logs its UTC time in the Library Manager logs every 10 minutes. This provides a means to line up the Library Manager and TS7700 logs.

The preferred method to keep nodes in sync is with a Network Time Protocol (NTP) Server. The NTP server can be a part of the Grid/WAN infrastructure, it can be a part of a customer intranet, or it can be a public server in the Internet (see Figure 2-9).

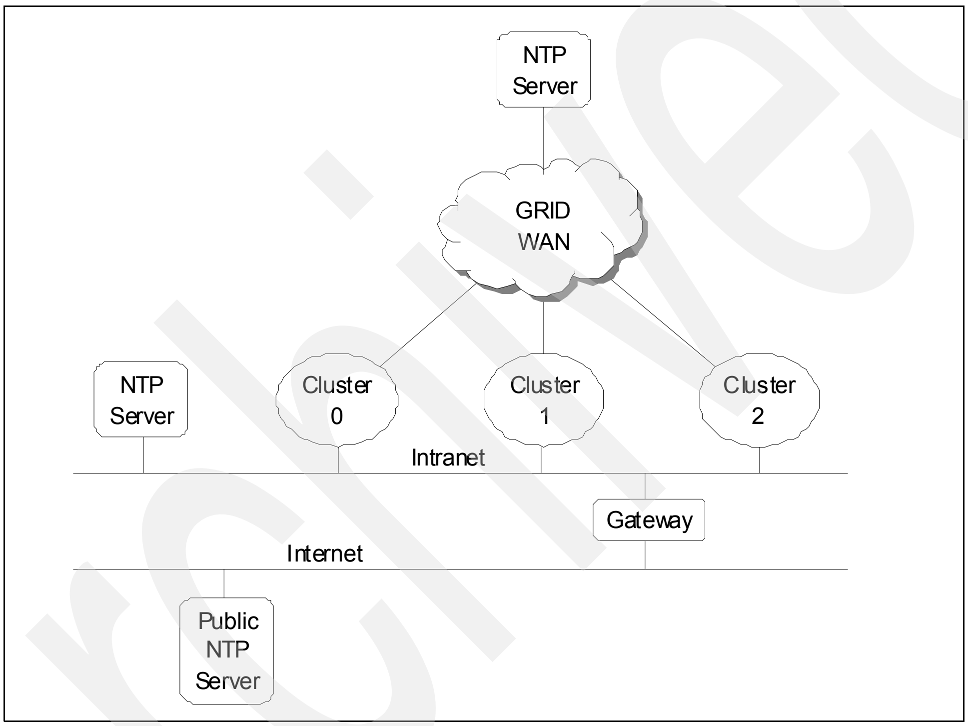

Figure 2-9 Time coordination with NTP servers

The NTP server address is configured into system VPD on a system-wide scope, such that all nodes will access the same NTP server. All of the clusters in a grid need to be able to communicate with the same NTP server defined in VPD. In the absence of an NTP server, all nodes will coordinate time with Node 0 of Cluster 0.

Note: Library Managers are not in sync with the grid, because they can be shared between TS7700s, Virtual Tape Servers, and ESCON/FICON Controllers. Library Managers should be set to local time.

2.2 Hardware components

In this section we discuss the hardware that makes up a TS7700 cluster. Figure 2-1 on page 9 shows you the components of a TS7700 Virtualization Engine.

These components include:

- An IBM 3952 Tape Frame which houses:
  - The TS7740 nodes
  - The TS7740 Cache Drawers
  - The TS7740 Cache Controller
  - Two Ethernet routers
- When attached to a TS3500 Tape Library, the following components are part of the cluster:
  - The IBM 3953-L05 Library Manger, which is installed inside an IBM 3953-F05 Tape Frame
  - The TS3500 Tape Library, which houses the TS1120 or 3592-J1A tape drives attached to the TS7700
- When attached to a 3494 Tape Library, the Library Manager is an integral part of the tape library.

2.2.1 IBM 3952 Tape Frame

The IBM 3952 Tape Frame (3952-F05) is a multipurpose frame providing up to 36 EIA units of usable space to contain components of various tape offerings. In a TS7700 configuration it is used for the installation of the TS7740 Node, the TS7740 Cache Controller, the TS7740 Cache Modules, and two Ethernet routers (see Figure 2-10 on page 23). These components are discussed in detail in the following sections.

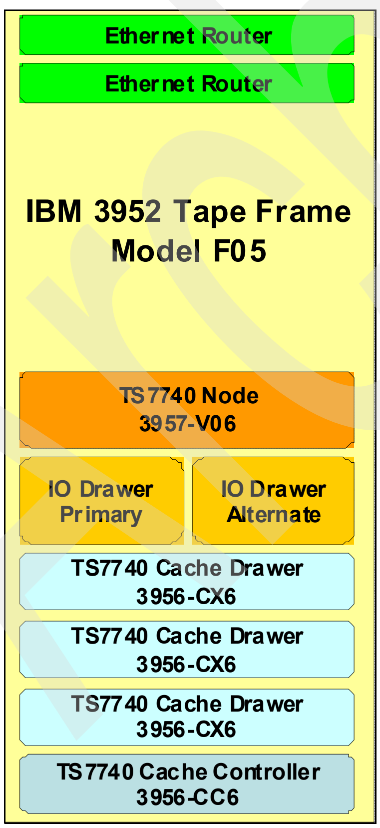

Figure 2-10 IBM 3952 Tape Frame layout

Note: The IBM 3952 Tape Frame is a multipurpose frame, which is also used for installation of the IBM System Storage TS7500 Series Virtualization Engine and the IBM System Storage TS1120 Tape Controller. In a TS7700 configuration you cannot install any non-TS7700 components in the IBM 3952 Tape Frame.

2.2.2 TS7740 Node

The TS7740 Node (3957-V06) is based on an IBM System p5® 520 server with two IBM 7311-D11 I/O drawers. The TS7740 Node offers the following features:

- Two dual-core, 64-bit, 1.9 GHz processors with 36 MB of L3 cache memory
- Eight GB of 266 MHz DDR1 memory
- Two hot-pluggable AC/DC power supplies with imbedded enclosure cooling units
- The following integrated features:
  - Dual SCSI adapter (not used)
  - Two USB ports (not used)
  - Two serial ports (used for service and maintenance only)
  - Six hot-pluggable PCI adapter slots used to house SCSI adapter for mirrored internal disks
  - Eight hot-pluggable DASD (2 × 4-pack)
  - One hot-pluggable DVD-ROM drive
  - One hot-pluggable DVD-RAM drive

One of the IO drawers acts as primary drawer and the other as secondary. The primary drawer contains the same cards in the same slots as the secondary. The cards are installed in Blind Swap Modules, which allow for replacement of the cards without moving the drawer in the frame. The following cards are installed in each of the IO drawers:

- One or two FICON adapters for host attachment
- One Dual Port Fibre Channel Adapter for the attachment of tape drives
- One Dual Port Fibre Channel Adapter for the attachment of the cache
- One Ethernet adapter for communication with the Library Manager (or the Library Managers in a configuration with redundant Library Managers), the cache controller, the TS3000 Service Console, and the Management Interface

Figure 2-11 shows the IBM 3952 frame, including the TS7740 Node.

Figure 2-11 TS7740 Node

2.2.3 TS7740 Cache

The cache controller and cache drawers are self-contained units that mount in the IBM 3952 Tape Frame. The tape frame holds one cache controller and zero, one, or three cache drawers. The cache controller and cache drawers make up the cache. The cache enablement feature determines the size of the usable cache, which ranges from 1 to 6 TB uncompressed in 1 TB increments. The cache enablement features can only enable an amount of cache that is equal to or less than the amount of physical cache installed. Table 2-1describes the relationship between the amount of physical cache and the number of cache enablement features that can be applied.

Table 2-1 Physical cache versus number of cache enablement features

Number of 3956-CC6 Cache Controllers	Number of 3956-CX6 Cache Drawers	Amount of physical cache	Number of 1TB Cache Enablement Features that can be applied
1	0	1.5 TB	1
1	1	3.0 TB	1, 2, 3
1	3	6.0 TB	1, 2, 3, 4, 5, 6

Cache capacity upgrades are available for TS7700 configurations with zero and one cache drawers. See 3.1, “Hardware configurations” on page 78 for more information.

The cache holds data temporarily before writing it to tape, caches the data afterwards to provide fast subsequent access, and provides source data when replicating to a peer.

TS7740 Cache Controller

The TS7740 Cache Controller (3956-CC6) is a self-contained 3U enclosure that mounts in the 3952 Tape Frame. It offers the following features:

- Redundant FC processing cards and processors on a single card (two of those cards)
- PowerPC® 750GX 1 GHz processor
- Two battery backup units (one for each processor card)
- Two AC/DC power supplies with imbedded enclosure cooling units
- Sixteen disk drive modules, each with a storage capacity of 146 GB, for a total raw storage capacity of 2.34 TB
- Attachment to up to three TS7740 cache drawers, each with 16 disk drive modules

Figure 2-12 TS7740 Cache Controller (back view)

Figure 2-13 TS7740 Cache Controller (front view)

TS7740 Cache Drawer

The TS7740 Cache Drawer (3956-CX6) is a self-contained 3U enclosure that mounts in the 3952 Tape Frame. The frame contains zero, one, or three cache drawers, each offering the following features:

- Two Fibre Channel processor cards
- Two AC/DC power supplies with imbedded enclosure cooling units
- Sixteen disk drive modules, each with a storage capacity of 146 GB, for a total raw storage capacity of 2.34 TB per drawer
- Attachment to the TS7740 cache controller

Figure 2-14 TS7740 Cache Drawer (back view)

Figure 2-15 TS7740 Cache Drawer (front view)

2.2.4 Ethernet routers

Two Ethernet routers are installed in the IBM 3952 Tape Frame. Each has eight internal LAN ports used for communication with the Library Manager in the IBM 3494 Tape Library or in the IBM 3953-F05 frame, and two WAN ports used for attachment to the customer LAN and to the TS3000 System Console.

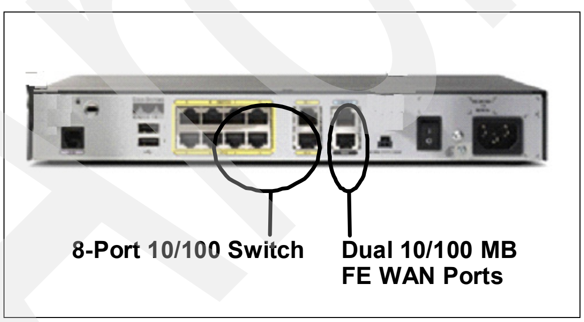

Figure 2-16 Ethernet router

2.2.5 TS7700 Grid communications

The connection path between two TS7700 Virtualization Engines in a grid configuration is a pair of 1 Gbps copper connections or, two shortwave optical fiber connections.

2.2.6 IBM 3953 Tape System

When attached to a TS3500 Tape Library, an external IBM 3953 Model L05 Library Manager is required to control a library partition dedicated to System z hosts. The Library Manager is installed in a separate frame, the IBM 3953 Model F05 Frame.

IBM 3953 Model L05 Library Manager

The Library Manager controls all or part of the IBM 3584 Tape Library according to how the logical libraries are set up. For example, the physical library can be divided between Open Systems hosts as well as System z. One Library Manager controls one logical library, but there can be up to four Library Managers that control one logical library each in one IBM 3584 Tape Library. Because each Library Manager supports the attachment of two TS7700 Virtualization Engines, a maximum of eight TS7700s can attach to a single TS3500 Tape Library.

The 3953 Library Manager is supported with the following configurations:

- IBM TS7700
- IBM 3494 VTS Model B10 or B20, and PtP VTS
- IBM System Storage TS1120 Tape Controller (3592-C06) or 3592-J70 Controller

> **Note:** IBM System Storage 3592-J70 Controller was withdrawn from marketing, effective 16 June 2006.

IBM 3953 Model F05 Frame

The IBM 3953 Tape Frame comes in two different versions:

- IBM 3953-F05 Base Frame

  This is the first 3953 frame in a configuration. It houses one Library Manager or, optionally, two Library Managers for redundancy. It also houses one Model C06 Controller (optional) and the TS3000 System Console, if installed, as well as Ethernet switches, up to six Fibre Switches, one KVM Switch, and monitor and keyboard.

  The following are the components in a 3953 Base Frame (Figure 2-17 on page 29).

  – IBM 3953 Library Manager A Model L05
  – IBM 3953 Library Manager B Model L05 (optional)
  – TS3000 System Console (if no external console is installed)
  – KVM Switch for Library Managers and TS3000
  – Two Ethernet switches
  – TS1120 Tape Controller Model C06 or IBM 3592 Model J70 controller
  – Up to six fiber switches, depending on the number of TS7700s and other tape controllers attached

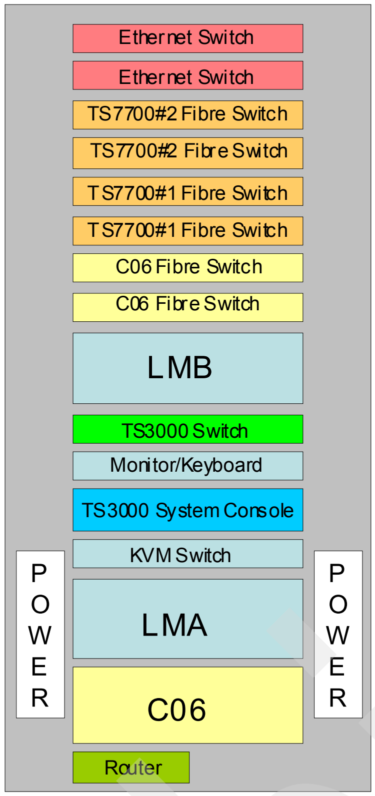

Figure 2-17 Contents housed in the 3953 Model F05 Base Frame

- IBM 3953-F05 Expansion Frame

  The IBM 3953 Model F05 Expansion Frame can be used to install additional IBM TS1120 Tape Controllers Model C06 or IBM 3592 Model J70 Controllers. A single expansion frame can house up to three 3592-J70 or TS1120 Controllers Model C06, up to two Ethernet switches, and up to six fibre switches. Both the Base Frame and the Expansion Frame have the same model type F05 but are differentiated by a descriptive feature code. There can be up to five Expansion Frames associated with one IBM 3953 Base Frame and one IBM 3584 logical library. These frames do not need to be installed next to each other, nor do they need to be alongside the IBM TS3500 Tape Library. The first frame must be a Base Frame with at least one Library Manager.

2.2.7 TS3500 Tape Library

The IBM System Storage TS3500 Tape Library is part of a family of tape libraries designed for large automated tape storage and backup solutions.

Originally delivered in 2000 at the same time as Linear Tape Open (LTO) Ultrium technology, the IBM TS3500 offered a robust enterprise library solution for mid-range and high-end open systems. Since its introduction, the library has been enhanced to accommodate different drive types and operating platforms, including the attachment of System z hosts and tape drive controllers.

Now the IBM TS3500 Tape Library is also available to connect drives to host systems with FICON or ESCON® attachments, as well as any combination of Fibre Channel and Ultra2/Wide Low Voltage Differential (LVD) SCSI. The combination of proven reliable tape handling with technology and functional enhancements has resulted in an optimal design for a robust enterprise solution, and outstanding retrieval performance; a typical cartridge move time is less than three seconds. Using IBM TS1120, TS1130 or LTO Ultrium high density cartridge technology, the IBM TS3500 provides a powerful and robust tape storage solution for the whole enterprise, yet it is contained in a minimal footprint.

Refer to *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789 for more information regarding the IBM System Storage TS1130 Tape Drive.

Figure 2-18 IBM System Storage TS3500 Tape Library

In summary, IBM TS3500 provides:

- Modular, scalable, automated Tape Library, combining IBM tape and automation for open systems and mainframe hosts, using a variety of IBM drive types
- Attachment to IBM System z, IBM System i®, IBM System p, RS/6000®, IBM System x®, Netfinity®, Sun™, Hewlett-Packard, and other non-IBM servers
- Connectivity using FICON, ESCON, Fibre Channel, Low Voltage Differential (LVD) SCSI, and High Voltage Differential (HVD) SCSI
- IBM Multi-Path Architecture designed to support redundant control paths, mixed drive configurations, and library sharing between multiple applications.

TS3500 Frame models

The TS3500 frames come in different models, depending on whether or not it is the first frame in the library, which contains the accessor as well as tape drives and storage cells, and on which type of tape drive it is installed in the frame. All currently available frame models can be intermixed in the same TS3500 as previously installed frame models. Previous frame models include the L22, L32, L52, D22, D32, and D52.

> **Note:** Models L32 and D32 were withdrawn from marketing, effective October 1, 2004. Models L22/D22 and L52/D52 were withdrawn from marketing effective September 29, 2006.

The following list describes the currently available frame models:

- TS3500 Tape Library Model L23: A base frame that contains up to 12 IBM 3592 tape drives and up to 260 IBM TotalStorage 3592 tape cartridges.
- TS3500 Tape Library Model D23: An expansion frame that can house up to 12 IBM 3592 Tape Drives and up to 400 IBM TotalStorage 3592 tape cartridges. Up to 15 expansion frames can be installed with a base frame.
- TS3500 Tape Library Model L53: A base frame that contains up to 12 Ultrium tape drives and up to 287 LTO Ultrium tape cartridges.
- TS3500 Tape Library Model D53: An expansion frame that contains up to 12 Ultrium tape drives and up to 440 LTO Ultrium tape cartridges. Up to 15 expansion frames can be installed with a base frame.
- TS3500 Tape Library Model D24: A driveless storage-only high density expansion frame for up to 1000 3592 tape cartridges. Up to 15 expansion frames can be installed with a base frame. Advanced Library Management System (ALMS) is required for any library with a Model S24™ frame. This frame can optionally be configured as service bay B.
- TS3500 Tape Library Model D54: A driveless storage-only high density expansion frame for up to 1320 Ultrium (LTO) tape cartridges. Up to 15 expansion frames can be installed with a base frame. Advanced Library Management System (ALMS) is required for any library with a Model S54 frame. This frame can optionally be configured as service bay B.
- IBM 3584 Model HA1: An optional high availability frame. In conjunction with service bay features on the TS3500 Tape Library Models D23 and D53, this supports the installation of a second accessor in the TS3500.

The Lxx frames also contain an I/O station for 16 cartridges. If both LTO and TS1120 or IBM 3592 tape drives are installed inside the TS3500 Tape Library, the optional second I/O station is required for the second media format. The second I/O station is installed below the first I/O station. The drive type in the Lxx frame determines which I/O station is in the top position. In an L53 frame, this would be the I/O station for LTO cartridges, and in an L23 frame, the I/O station for 3592 cartridges would be in the top position.

TS3500 attached to TS7700

The TS3500 Tape Library houses the physical tape drives for a TS7700 Virtualization Engine. The tape drive types currently supported for System z attachment in the TS3500 Tape Library are the TS1120 Tape Drive and the 3592 Model J1A Tape Drive. These tape drives also support Open Systems hosts attachment, but there are some differences between System z and Open Systems attachment:

- System z hosts are attached to controllers, which in turn have a set of drives attached, rather than each drive having a separate attachment to a host channel.
- System z-attached drives must be in a separate partition (logical library) in the IBM TS3500 Tape Library.
- The interface with the library robotics is managed by a Library Manager that is physically external to the IBM TS3500 Tape Library.

A TS7700 connects to a TS3500 Tape Library using the IBM 3953 Library Manager, which provides all necessary elements to manage automated cartridge and physical tape drive handling.

All the components needed for accessor control are integrated in the TS3500 Tape Library and all physical motions, inventory processing, and error recovery are handled by the library itself. The Library Manager, needed for the control of the logical library of the TS3500 Tape

Library and for the partitioning and sharing of System z-attached logical library partitions, resides in an IBM 3953 Tape Frame Model F05.

The Library Manager controls a logical library of the TS3500. Up to two TS7700 Virtualization Engines (or 3494-B10/B20 VTS) and up to sixteen controllers can attach to one logical library controlled by a Library Manager. The total number of attached subsystems (TS7700 + VTS + controllers) per logical library cannot exceed 16.

Note: Advanced Library Management System (ALMS) must be enabled in a TS3500 library to support attachment of System z hosts.

The TS7700 requires a minimum of four dedicated physical tape drives and supports a maximum of 16 drives. These drives can reside in the TS3500 Model L23/L22 Base Frame and the Model D23/D22 Expansion Frames. Up to twelve drives can be installed in one frame. TS7700-attached drives need not be installed in contiguous positions. Theoretically, ALMS allows for installation of 16 TS7700-attached drives in 16 different TS3500 frames. Although this is certainly not a favored configuration, we recommend to spread the drives across at least two frames for availability reasons. TS7700-attached drives cannot be shared with other systems, but they can share a frame with tape drives attached to other TS7700s, to tape controllers, or to Open Systems hosts.

The TS7700 Virtualization Engine supports TS1120 Tape Drives and IBM 3592 Model J1A Tape Drives. IBM 3592-J1A drives are withdrawn from marketing, but existing drives can be used for TS7700 attachment. TS1120 drives assigned to a TS7700 run in J1A Emulation mode, when intermixed with J1A tape drives on the same TS7740, or in native E05 mode, when only E05 drives are attached to a TS7740.

Logical libraries in the TS3500/3953

The patented IBM Multi-Path Architecture enables simple SCSI medium changer libraries to be shared by heterogeneous host platforms, without middleware or a dedicated server acting as a Library Manager. This facility is implemented on a number of IBM libraries, including the IBM TS3500.

The Multi-Path Architecture makes sharing possible by letting you partition the library's storage slots and tape drives into logical libraries (or partitions). Servers can then run separate applications for each logical library. Heterogeneous hosts share the library robotics, but have the physical library drives and slots divided between them. Each host sees the same accessor, but has a different view of the physical attributes of the library. This existing IBM TS3500 concept of logical libraries is used to segregate the physical devices associated with TS7700 Virtualization Engines and the 3592-J70 and TS1120 Tape Controllers from the Open Systems drives. The System z host is allocated its own logical library partition in the TS3500 Tape Library; there can be other logical libraries in the same TS3500, connected to another IBM 3953 Library Manager or dedicated to Open System hosts.

The minimum requirement for a logical library definition is that it contain at least one drive and one cartridge slot. Because the maximum number of drives that can be installed in the IBM TS3500 is 192, 12 drives in each of 16 frames, it follows that up to 192 logical libraries can be defined, overall. However, the IBM TS3500 supports only a maximum of four System z logical library partitions.

The term logical library is used in two different contexts in the IBM 3953 environment. It refers to both of the following:

- Splitting a single IBM TS3500 Tape Library by putting logical walls between tape drives and cartridge slots within a single physical library enclosure. This is managed within the IBM TS3500 hardware itself.
- The library name or ID seen by a System z host when it is attached to a tape library. For each library, up to three library IDs, called logical libraries, are supported. This is managed through the Library Manager and the System z host configuration software.

In the IBM TS3500 you must first identify a hardware-based logical library for the IBM 3953 to manage. Note that there can be only one logical library defined, which occupies the entire physical library and is dedicated to System z.

The IBM 3953 Library Manager and the System z host are aware of one library only. When one or two TS7700s, as well as native 3592 tape drives are attached, the Library Manager defines its own logical libraries or partitions: one for the native drives, and one for each attached TS7700 Virtualization Engine or B10/B20 VTS.

In both cases the terms logical library and partition are used interchangeably. It is very important when referring to logical libraries or partitions to understand in what context the term is being used. From a host operating system point of view there is no difference whether it connects to TS7700 and tape controller partitions through a Library Manager in a 3953 frame, or to the three different library partitions in a 3494 library through the internal Library Manager of that library.

Each 3953 Tape System can support up to three logical library partitions consisting of two TS7700s (or Virtual Tape Server B10/B20) and one native logical library with 3592-J70 or TS1120 Tape Controllers. These three logical library partitions themselves make up one logical library within a 3584 Tape Library. A total of four 3953 subsystems can be connected to one TS3500 for a maximum of 12 different 3953-managed logical library partitions. With two TS7700s per Library Manager a maximum of eight TS7700s can attach to one TS3500 library. TS3500 library partitions for System z attachment can coexist with partitions for Open Systems attachment. For a detailed description of this connectivity, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

2.2.8 IBM TotalStorage 3494 Tape Library

You can attach the TS7700 Virtualization Engine to the 3494 Tape Library in a manner similar to attaching a VTS. Here is a list of characteristics related to the TS7700 in a 3494 Tape Library:

- The TS7700 occupies a partition in the library just like a VTS, either VTS partition 1 or 2.
- The TS7700 can coexist with a VTS of any model type (B16, B18, B10, or B20) in the 3494.
- The TS7700 attaches to 3592-J1A or TS1120 drives.
- The TS7700 attaches to a minimum of four and a maximum of twelve 3592 drives.
- All of the 3592 drives must reside in the same frame. Drives cannot be split across frames.

Refer to *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632 for more information about implementing a virtualization solution in a 3494.

2.2.9 IBM System Storage TS1120 Tape Drive

The IBM System Storage TS1120 Tape Drive (3592-E05) is the second generation in the IBM Tape Drive 3592 family with, again, dramatically enhanced performance and capacity characteristics when compared to the Model J1A Tape Drive. When compared to previous tape drive families and generations, the TS1120 is unique in the following ways:

- The E05 models can read cartridges written by J1A drives and can also write in J1A Emulation mode. The drive can change its mode of operation dynamically per physical cartridge.
- The 3592-E05 tape drives can be intermixed with 3592-J1A tape drives:
  - In the same tape library frame
  - Behind the same TS7700 Virtualization Engine
  - Behind the same Model TS1120 or 3592-J70 Controller
  - Behind the same VTS
- The TS1120 supports encryption of the physical cartridges when attached to the TS7740 Virtualization Engine when the following conditions exist:
  - E05 drives are the only drives attached to the TS7740.
  - The drives have encryption enabled.
  - The TS7740 (V06) has the encryption feature enabled.

> **Note:** When intermixed with 3592-J1A tape drives behind the same TS7700 Virtualization Engine or other controller, or in the same IBM 3494 tape library frame, the TS1120 Model E05 Tape Drive must always operate in J1A Emulation mode.
>
> To use the full capacity and functionality, such as Encryption, of the TS1120 Model E05, do not intermix it with J1A tape drives behind the same TS7740.

The TS1120 Tape Drive, when attached to a TS7700 Virtualization Engine and operating in E05 mode, supports these cartridge types:

- JAMEDIA5, with an uncompressed capacity of 500 GB
- JBMEDIA9, with an uncompressed capacity of 700 GB
- JJMEDIA7, with an uncompressed capacity of 100 GB

> **Note:** When operating in J1A Emulation mode, the TS1120 Tape Drive can read or write JA and JJ cartridges only. JB media is not supported. In this case, the JA cartridge capacity is 300 GB of uncompressed data, and the JJ cartridge capacity is 60 GB.
>
> If the drives in a TS7700 are upgraded from J1A to all TS1120s, the TS1120s reads the J1A formatted tapes and, when they are re-used as physical scratch, they are formatted at the native TS1120 capacity.

2.2.10 IBM System Storage TS1130 Tape Drive

The IBM System Storage TS1130 Tape Drive is the third generation in the IBM 3592 Tape Drive family that provides higher capacity and performance compared to previous generations. It provides native data rate up to 160 MB/s and uncompressed capacity of 1 TB with the Extended Data Cartridge (JB) versus 100 MB/s data rate and 700 GB uncompressed capacity of the TS1120 Tape Drive.

> **Note:** The IBM Virtualization Engine TS7740 Release 1.4a does not support the TS1130 Tape Drive. TS7700 R1.5 provides support for TS1130.
>
> Virtual Tape Server (3494-B10 and 3494-B20) does not support the TS1130 Tape Drive.
>
> For more information regarding the TS1130 Tape Drive, refer to the *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

The TS1130 is available in two models:

- 3592-E06 for new Tape Drives
- 3592-EU6 as a MES from TS1120 (3592-E05) Tape Drives

The TS1130 Tape Drive also supports data encryption and key management. It is downward read compatible (n-2) to the 3592 Model J1A and is downward write compatible (n-1) to the 3592 Model E05 formats. The TS1130 Tape Drive uses the same IBM 3592 cartridges that are available for TS1120 and 3592-J1A. Attachments to System z and Open Systems platforms are maintained.

TS1130 drives cannot be intermixed with E05 or J1A drives behind thee TS7740 R1.5, because the TS1130 drives do not have an emulation mode. When a TS1130 drive writes to an empty tape, it always writes in E06 format.

Other new capabilities that are included in the TS1130 Tape Drive include:

- Giant Magneto-Resistive (GMR) Read/Write head, which provides fewer data errors and improve robustness
- Head Overcoat technology, which provides improved corrosion protection and wear characteristics
- Improved Standby power management, which delivers energy savings through lower power dissipation

2.2.11 IBM System Storage TS3000 System Console

The IBM System Storage TS3000 System Console (TSSC), formerly known as the TotalStorage Master Console, enables remote monitoring by IBM Engineering of attached units to expedite microcode updates, reduce service times, and enhance local service. It is required for the TS7700 and can be ordered as an integrated feature in the IBM 3953-F05 Base Frame or as a desktop console. You can choose to attach an external TSSC that is already installed for other IBM equipment in the same physical location.

A folding keyboard, video monitor, and mouse are provided in the Base Frame for the operator console of the Library Manager. If a TSSC or a second Library Manager is installed in the Base Frame, then a KVM switch is required, because they all share the same keyboard, mouse, and monitor.

2.3 TS7700 management concepts

In this section, we explain the main management concepts of the TS7700 Virtualization Engine.

2.3.1 Buffering data in the tape volume cache

The tape volume cache of the TS7700 Virtualization Engine is the key element that enables the utilization of high-capacity tape technology. Buffering host-created volumes on disk and subsequently stacking them on high-capacity tape allows for the utilization of the full cartridge capacity provided by current tape technology.

The tape volume cache is a disk buffer where emulated tape volumes in 3490E format are written before they are copied to physical tape cartridges. The host operating system sees IBM 3490E tape drives, but actually these tape drives are virtual, and the 3490 tape volumes are represented by storage space in a fault-tolerant disk subsystem. All host interaction is through the virtual control unit. The host never writes directly to the physical tape drives attached to a TS7700.

While resident in the tape volume cache, the user data is protected by RAID-5. The cache is configured into two RAID groups. This RAID configuration gives continuous data availability to users. If one data disk in a RAID group becomes unavailable, the user data can be recreated dynamically from the data contents and the parity information about the remaining disks. In addition, the RAID groups contain already mounted hot spare disks to replace any damaged disk, so the RAID controller can rebuild the data of the damaged disk on the hot spare. In this way, the TS7700 can continue working while the IBM service representative replaces the damaged hard disk in the TS7740 cache controller or cache module.

In addition to making full use of high-capacity tape cartridges, there are these benefits:

- Emulated 3490E volumes are accessed at disk speed. Tape commands such as space, locate, rewind, and unload are mapped into disk commands that are completed in tens of milliseconds rather than the tens of seconds required for traditional tape commands.
- Multiple, different, emulated 3490E volumes can be accessed in parallel because they physically reside in the tape volume cache. (A single virtual volume cannot be shared by different jobs or systems at the same time.)

2.3.2 Tape Volume Cache Management

Cache Management allows management of the contents of the tape volume cache to improve average mount response time. You can use DFSMS policy management to assign virtual volumes to preference groups. These groups can control the priority of when volumes are deleted from cache after they are copied to tape when space is needed in cache. Cache management can also be used to improve recovery time by keeping required volumes in cache.

TVC management processes

Four processes manage the Tape Volume Cache of the TS7700 Virtualization Engine.

- Premigration Management

  This comes into effect when the amount of TVC data not copied to tape reaches a predefined threshold. It is intended to make sure that the TVC does not become completely full of data that has not been backed up to physical tape. It is the mechanism that takes the TS7700 from peak mode to sustained mode.

- Free-space Management

  This comes into effect when the amount of unused (free) TVC space reaches another predetermined threshold. It is intended to make sure that the TVC does not become completely full of data, copied to physical tape or not. It is the mechanism that keeps the input to the TVC from over-running the available free space. It results in the second of the “small volume” throughput limitations.

- TS7700 Copy Management

  This applies only to a grid configuration and comes into effect when the amount of un-copied data in the TVC reaches a predefined threshold. It applies in particular to deferred copy mode, and when invoked will reduce the incoming host data rate independently of premigration or free-space management. Its purpose is to prevent logical volumes from being migrated to physical tape prior to being copied to the other TS7700, which could result in a recall operation prior to the copy to the remote cluster in the grid. This process is also called *Copy Throttling*.

- TS7700 Copy Time Management

This also applies only to a grid configuration, and in particular to Rewind/Unload (RUN) copy mode. It comes into effect only if specifically invoked using a service panel. When invoked, it limits the host input rate when a copy has not completed within one of two selectable time periods. It is intended to prevent a copy from exceeding the Missing Interrupt Handler (MIH) time-out value for the host job. The job completes before the MIH, and an alert is posted to the host console that RUN copy has been moved to deferred.

Preference level 0

Preference level 0 is assigned to volumes that are unlikely to be accessed after being created, for example volumes holding DASD image copies. There is no need to keep them in cache any longer than necessary to copy them to physical tape. Informal studies suggest that the proportion of data that is unlikely to be accessed can be as high as 80%.

When a volume is assigned preference level 0, the TS7700 gives it preference to be copied to physical tape. Then, when space is needed in the cache, the TS7700 will first select a preference level 0 volume that has been copied to a physical volume. Preference level 0 volumes are selected by largest size first, independent of how long they have resided in cache. If there are no preference level 0 volumes that have been copied to physical volumes to remove, then the TS7700 will select preference level 1 volumes.

In addition to removing preference level 0 volumes from the cache when space is needed, the TS7700 will also remove them if the subsystem is relatively idle. There is a small amount of internal processing overhead to remove a volume from cache, so there is some benefit in removing them when extra processing capacity is available. In the case where the TS7700 removes PG 0 volumes during idle times, it selects them by smallest size first.

Preference level 1

Preference level 1 is assigned to volumes that are likely to be accessed after being created, for example volumes that contain master files created as part of the nightly batch runs. Because they are likely to be used as input for the next night’s batch runs, it would be beneficial for them to stay in the TVC for as long as possible.

When a volume is assigned preference level 1, the TS7700 adds it to the queue of volumes to copy to physical tape after a four minute time delay and after any volumes assigned to preference level 0. The four minute time delay is to prevent unnecessary copies to be performed when a volume is created, then quickly mounted and appended to again. When space is needed in the cache, the TS7700 will first see if there are any preference level 0 volumes that can be removed. If this is not the case, the TS7700 selects preference level 1 volumes to remove based on a “least recently used” algorithm, so volumes that have been copied to physical tape and have been in cache the longest without access are removed first.

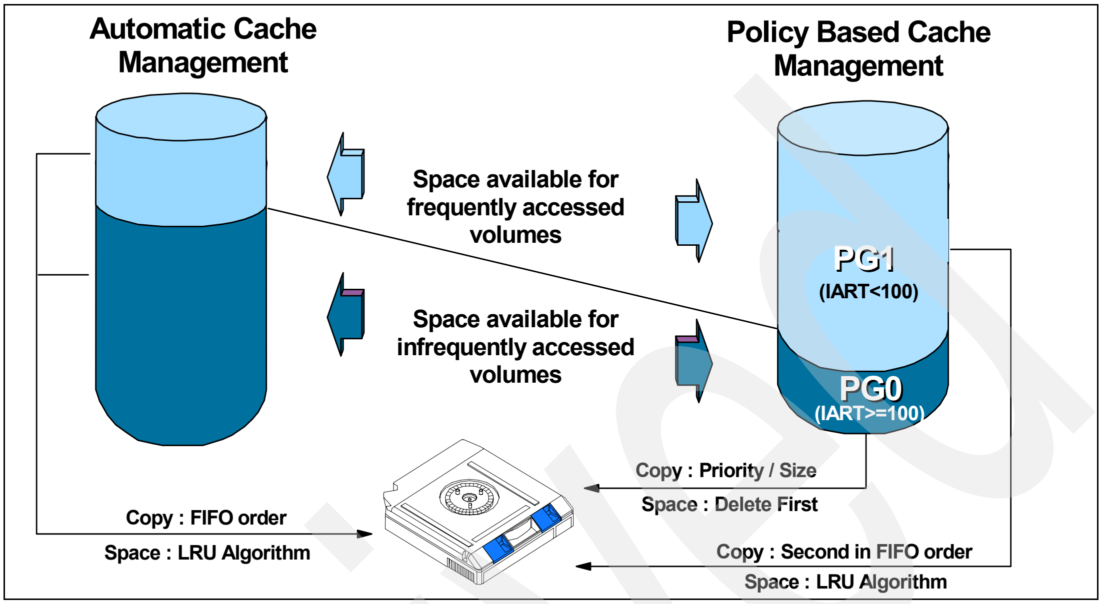

Figure 2-19 Cache utilization with policy-based cache management

When a preference level has been assigned to a volume, that assignment is persistent until the volume is re-used for scratch and a new preference selection is made. For example, this means that a volume assigned a preference level 0 will maintain that preference level when it is subsequently recalled into cache.

The Storage Class name assigned to a volume in the ACS routine is directly passed to the Library Manager. Figure 2-20 shows this process. For non-DFSMS environments, you can assign a Storage Class to a range of logical volumes, by inserting them using the Management Interface and afterwards changing the attributes using the Library Manager Console.

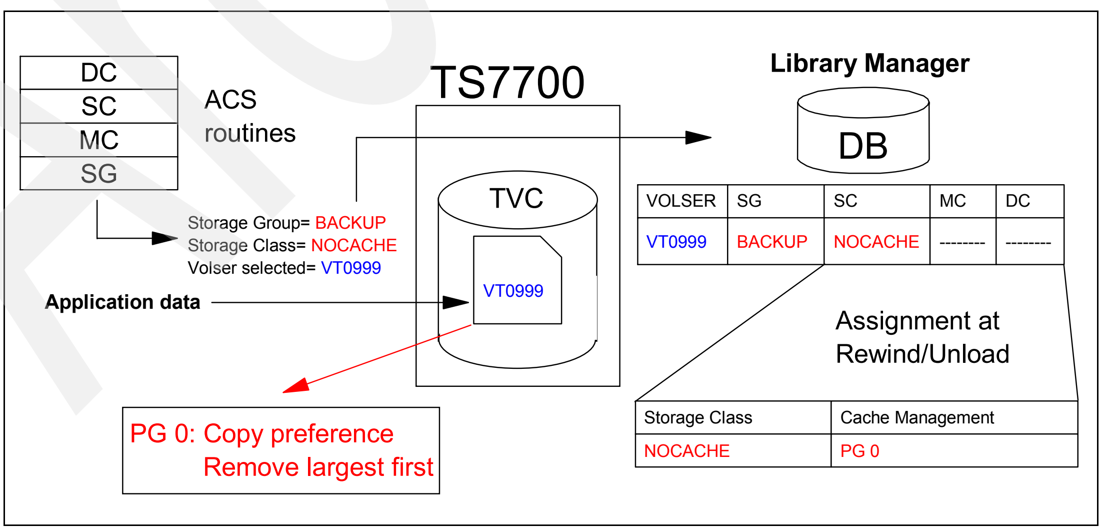

Figure 2-20 TVC Management through Storage Class

In the ETL Specialist, you define one or more Storage Class names and pick preference level 0 or 1 for them. To be compatible with the IART method to set the preference level, the Storage Class definition also allows a *Use IART* selection to be made.

Even before outboard policy management was made available for the VTS, you had the ability to assign a preference level to virtual volumes by using the Initial Access Response Time Seconds (IART) attribute of the Storage Class. The IART is a Storage Class attribute, which was originally added to specify the desired response time (in seconds) for an object using the Object Access Method (OAM). If you wanted a virtual volume to remain in cache, you would assign a Storage Class to the volume whose IART value is 99 seconds or less. Conversely, if you wanted a virtual volume preferably out of cache, you would assign a Storage Class to the volume whose IART value is 100 seconds or more.

Assuming that the *Use IART* selection has not been specified, the TS7700 sets the preference level for the volume based on the preference level 0 or 1 definition for the Storage Class assigned to the volume.

If the Storage Class name passed to the Library Manager for a volume has not previously been defined, the Library Manager will create it using the definitions for the default Storage Class.

Business continuance settings

For customers who have installed or are considering the installation of a Multi Cluster Grid for business continuance, the standard management algorithms of the subsystem might not meet all of their business or recovery needs. To better meet those needs, the following aspects of the grid management algorithms can be tailored by the IBM SSR.

Copy files preferenced to reside in cache

Normally, the caches in the clusters in a Multi Cluster Grid are managed as one to increase the likelihood that a needed volume will be in cache. By default, the volume on the TVC selected for I/O operations is preferred to stay in the cache on that cluster, whereas the copy made on the other cluster is preferred to be removed from cache. Preferred to stay in cache means that when space is needed for new volumes, the oldest volumes are removed first (Least Recently Used algorithm). Preferred to remove from cache means that when space is needed for new volumes, the largest volumes are removed first, regardless of when they were written to the cache.

For example, in a Two-Cluster Grid, when you have set up a Copy Consistency Point Policy of RUN/RUN and the host has access to all virtual devices in the grid, the selection of virtual devices combined with I/O TVC selection criteria will automatically balance the distribution of original volumes and copied volumes across the TVCs. The original volumes (newly created or modified) will be preferred to stay in cache, while the copies will be preferred to be removed from cache. The result is that each TVC is filled with unique newly created or modified volumes, thereby roughly doubling the amount of cache the host sees.

For a Multi Cluster Grid used for remote business continuation, particularly when the local clusters are used for all I/O (remote virtual devices varied offline), this default management method might not be desired. In the case where the remote cluster of the grid is used for recovery, the recovery time is minimized by having most of the needed volumes already in cache. What is really needed is to have the most recent copy volumes remain in cache, not be preferred out of cache.

Based on customer requirements, the IBM SSR can set or modify this control through the service menu of the remote cluster.

When off, which is the default, copy files are managed as preference group 0 volumes (prefer out of cache first by largest size). When on, copy files are managed based on the Storage Class construct associated with the volume defined at the Library Manager associated with the TS7700 receiving the copy. If you define the Storage Class construct names the same on all Library Managers in the grid configuration, the copied volumes will then be managed in the cache the same as the originally created volumes. For example, the storage group constructs could be SCBACKUP=Pref Group 1, SCARCHIV=Pref Group 0. All logical volumes written that specify SCARCHIV would be treated as PG0 in both the local and remote (copy) caches. All logical volumes written that specify SCBACKUP would be treated as PG1 in both the local and remote caches.

Recalls preferenced for cache removal

Normally, a volume recalled into cache is managed as though it were newly created or modified, because it resides in the TVC selected for I/O operations on the volume. A recalled volume displaces other volumes in cache. In the case where the remote cluster of a grid is used for recovery, the recovery time is minimized by having most of the needed volumes in cache.

However, it is likely that not all of the volumes to restore will be resident in the cache and that some number of recalls will be required. Unless the customer can explicitly control the sequence of volumes to be restored, it is likely that recalled volumes will displace cached volumes that have not yet been restored from, resulting in further recalls at a later time in the recovery process. After a restore completes from a recalled volume, that volume is no longer needed. What is needed is to remove the recalled volumes from the cache after they are accessed so that they minimally displace other volumes in the cache.

Based on customer requirements, the IBM SSR can set or modify this control through the service menu on the remote cluster:

- When off (default), recalls are managed as preference group 1 volumes (LRU).
- When on, recalls are managed as preference group 0 volumes (prefer out of cache first by largest size).

This control is independent of and not affected by cache management controlled through the Storage Class construct. Storage Class cache management affects only how the volume in the I/O TVC is managed.

2.3.3 Logical and stacked volume management

Logical volumes and physical cartridges are managed inside the TS7700 and the Library Manager. For example, you set up management policies at the Library Manager, and you insert logical volumes through the Management Interface of the TS7700 Virtualization Engine.

Physical Volume Pooling

Logical volumes can be assigned to selected Storage Groups. These Storage Groups point to primary storage pools. The pool assignments are stored in the Library Manager database. When a logical volume is copied to tape, it is written to a stacked volume that is assigned to a Storage Pool as defined by the Storage Group constructs at the Management Interface.

Physical pooling of stacked volumes is identified by defining a pool number. Some of this definition can be set through the LM or the tape library specialist. Pool encryption settings are set using the TS7700 Management Interface (MI). Each TS7700 attached to a 3494 or TS3500 has its own set of pools. There is a Common Scratch Pool (pool 00), which is a reserved pool containing only scratch stacked volumes. There are also 32 general purpose pools (pools 1-32).

By default there is one pool, Pool1, and the TS7700 stores virtual volumes on any stacked volume available to it. This creates an intermix of logical volumes from differing sources, for example, LPAR, customer, and applications. The user cannot influence the physical location of the logical volume. Physical Volume Pooling addresses the following concerns:

- Data from different customers on the same physical volume can contravene some outsourcing contracts.
- Customers like to be able to "see, feel, and touch" their data by having only their data on their media.
- Charging for tape is complicated. Traditionally users are charged by the number of volumes they have in the tape library. With Physical Volume Pooling users can create and consolidate multiple logical volumes on a smaller number of stacked volumes and so reduce their media charges.
- The TS7700 will use any media type available to it. Small logical volumes on the Enterprise Tape Cartridges (JA) or the Enterprise Extended Tape Cartridges (JB) will take a longer time to recall than volumes on the Economy Cartridge (JJ).
- Protecting data through encryption can be set on a per pool basis. This enables you to encrypt all or some of your data when it is written to the back-end tapes.

Volume pooling allows the administrator to define pools of stacked volumes within the TS7700, and the user can use these pools through use of SMS constructs. There can be up to 32 general purpose pools (1-32) and one common pool (0). Pool 0 is a source of scratch stacked volumes for the other pools, which can be configured to borrow from pool 0. The pool can then return the "borrowed volume" when it becomes scratch or keep the "borrowed" volume. Each pool can be defined to borrow single media (JA, JB, or JJ), mixed media, or have a "first choice" and a "second choice".

VOLSER ranges can be defined with a home pool at insert time. Changing the home pool of a range has no effect on existing volumes in the library.

Common scratch pool (Pool 00)

The Common Scratch Pool (CSP) is a pool which contains only scratch stacked volumes. It also serves as a reserve pool from which scratch stacked carts can be borrowed either on a temporary or permanent basis for the primary pools as they run out of scratch stacked volumes. The borrowing options can be set at the Library Manager Console or the ETL Specialist when defining Stacked Volume Pool Properties.

General purpose pools (Pools 01-32)

There are 32 general purpose pools available per TS7700 cluster. These pools can contain both physical scratch and private stacked volumes. All volumes, private and scratch, are assigned category FF04. When initially creating these pools, it is important to ensure that the correct borrowing properties are defined to the pool. See 4.4.1, "Defining stacked volume pool properties" on page 164.

Using this facility, you can also:

- Move stacked volumes to different pools
- Set reclamation threshold at the pool level
- Force reclamation policies for stacked volumes
- Eject stacked volumes from specific pools
- Intermix or segregate media types
- Map different Storage Groups to the same primary pools

Note: Primary Pool 01 is the default private pool for TS7700 stacked volumes.

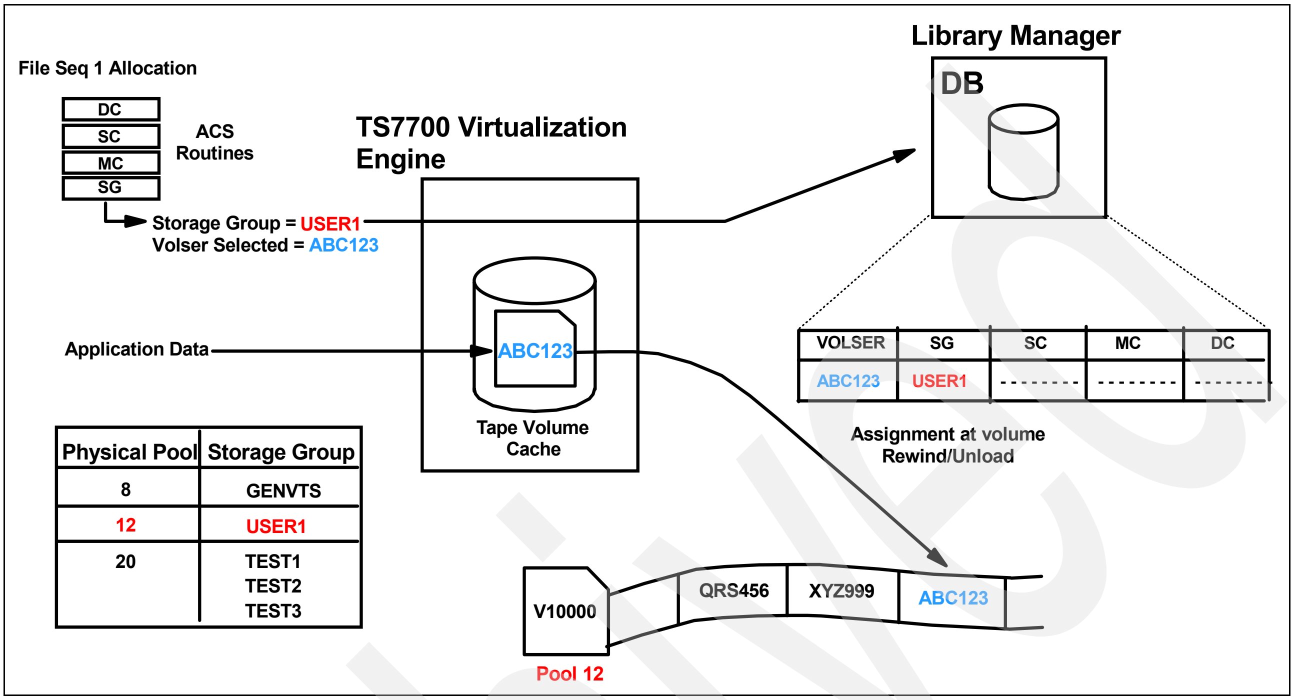

Figure 2-21 Logical volume allocation to specific physical volume pool flow

Borrowing and returning

With the concept of borrowing and returning, out of scratch scenarios can be automatically managed. With borrowing, stacked volumes can move from pool to pool and back again to the original pool. In this way the TS7700 can manage out of scratch and low scratch scenarios which can occur within any TS7700 from time to time.

Stacked volume pool properties

Logical volume pooling supports cartridge type selection. This can be used to create separate pools of 3592 tape cartridges with a variety of capacities from 60 GB up to 700 GB, depending upon the type of media and tape drive technology used.

You might want to use the lower capacity JJ cartridge to provide fast access to applications such as HSM, Content Manager, and the higher capacity JA or JB cartridges to address archival requirements such as full volume dumps.

Selective Dual Copy

In a Single Cluster Grid, a logical volume usually exists as a single entity on a single stacked volume. If the stacked volume is damaged you can lose access to the data on the logical volume. Without the Selective Dual Copy Function the only way to ensure data availability is to use host software to duplex the logical volume, or to set up a grid environment.

With the Selective Dual Copy function, storage administrators have the option to selectively create dual copies of logical volumes within a TS7700. In a grid environment logical volumes can be duplexed across multiple clusters.

The selective dual copy function can also be used along with the Copy Export function to provide a copy of data for disaster recovery purposes. Refer to 2.3.4, “Copy Export” on page 46 for more details concerning Copy Export.

The copy of the logical volume is created in a second physical pool ensuring physical cartridge separation. Control of Dual Copy is through the Management class construct (see 4.4.3, “Creating Management Classes” on page 168). The second copy is created when the original volume is pre-migrated.

The second copy is only available when the primary volume cannot be recalled, it cannot be addressed separately and cannot be used if the primary volume is being used by another application, or it has been corrupted by an application.

A copy of the logical volume is made to a secondary pool as defined to the TS7700:

- The copy feature is enabled by the Management Class setting at the ETL specialist, where you define the secondary pool.
- Secondary and primary pools can be intermixed:
  - A primary pool for one logical volume can be the secondary pool for another logical volume unless the secondary pool is used as a Copy Export pool.
  - Multiple primary pools can use the same secondary pool.
- At rewind or unload time, the secondary pool assignment is determined and the copy of the logical volume is scheduled. The scheduling of the backup is determined by the premigration activity occurring in the TS7700.
- The copy is created prior to the primary volume being migrated.

Mounting a scratch virtual volume

The Library Manager uses categories to group volumes. After virtual volumes are defined through the Management Interface, they are placed in the insert category and handled exactly as native cartridges. When the IBM TS7700 is varied online or during insert processing, the host operating system assigns scratch volumes to a particular category. When a request for a scratch is issued to the TS7700, the request specifies a mount from category. The Library Manager selects a virtual VOLSER from the candidate list of scratch volumes in the category.

If the volume is mounted for the first time (that is, the VOLSER has never been used before as a virtual volume), the TS7700 Virtualization Engine generates a set of records for the volume in the TVC as though the volume was initialized using the EDGINERS or IEHINITT program, with an IBM standard label (that is, the volume will contain a VOL1 record, an HDR1 record and a tape mark).

Copying the virtual volume to tape

After the host closes and unloads a virtual volume, the storage management software inside the TS7700 schedules the virtual volume to be copied (also known as pre-migration) onto a physical tape cartridge. The TS7700 attempts to maintain a mounted stacked volume to which virtual volumes are copied. Therefore, mount activity is reduced because only one physical cartridge is mounted to service several scratch mount requests from the host.

Mounting a specific virtual volume

The process followed in the TS7700 to mount a specific tape is similar to the scratch mount scenario. The TS7700 storage management software first checks the TVC for the tape VOLSER. If the virtual volume exists in the TVC, a physical tape mount is not required; the mount is signaled as complete and the host can access the data immediately.

If the volume has already been removed from the TVC, a recall of the logical volume to the TVC is required to create a virtual volume that the host can then access directly. Recall typically requires a physical mount unless the stacked volume is already mounted following another request. Mount completion is signaled to the host system only after the complete volume is available in the TVC. The virtual volume will remain in the TVC until it becomes the Least Recently Used (LRU) volume, unless the volume was assigned a preference group of 0. Then its TVC space can be made available for new data. If modification of the virtual volume did not occur when it was mounted, the VTS does not schedule another copy operation and the current copy of the logical volume on the original stacked volume remains active.

In a z/OS® environment, in order to mount a specific volume in the TS7700, that volume must reside in a private category within the library. DFSMS OAM prevents a scratch volume from being mounted in response to a specific mount request.

The library itself, although it does not consider what the category is, looks to see if the category has the fast-ready attribute set and will treat the specific volume mount as a fast-ready mount if it is set.

Expired virtual volumes

When a virtual volume is returned to scratch status, it is placed on an expire list. Prior to the expired volume management function, the data associated with an expired logical volume remained available on the stacked volume until the logical VOLSER was reused as a scratch volume and new data was written on it, which occurred only if unused virtual volumes were not available.

The disadvantage of this method of managing the data is that these “user expired” logical volumes needlessly consume physical stacked volume resources therefore requiring more physical stacked volumes in a TS7700. Also, because the data is still considered active, the time until a physical volume falls below the reclamation threshold is increased and, potentially, customer expired data will be moved during a reclaim.

With expired volume management you can set a “grace period” for expired volumes, ranging from 24 hours to approximately 144 weeks. After that period has elapsed, expired volumes become candidates for deletion. The deletion of expired logical volume data eliminates the need for the TS7700 to manage logical volume data that has already been expired at the host. For details on expired volume management, see 4.3.6, “Define logical volume Expiration Time” on page 156.

Secure Data Erase

As described above, expired data on a physical volume remains readable until the volume has been completely overwritten with new data. Some customers are concerned that a court order could expose them to liability and cost to be able to try to *find* an old version of a data volume. Another concern is security of old data. The TS7700 provides physical volume erasure on a physical volume pool basis controlled by an additional reclamation policy. With Advanced Policy Management, which is standard on a TS7700, all reclaimed physical volumes in that pool are erased with a random pattern prior to being reused. A physical cartridge is not available as a scratch cartridge as long as its data is not erased.

The Secure Data Erase function supports the erasure of a physical volume as part of the reclaim process. The erasure is performed by writing a random data pattern on the physical volume being reclaimed. A random data pattern is written on the physical volume being reclaimed so that the logical volumes that had been written to the physical volume prior to the volume being reclaimed are no longer readable. As part of this data erase function, an additional reclaim policy is added. The policy specifies the number of days a physical volume can contain invalid logical volume data before the physical volume becomes eligible to be reclaimed.

The data associated with a logical volume is considered invalidated as follows:

- A host has assigned the logical volume to a scratch category. The volume is subsequently selected for a scratch mount and data is written to the volume. The older version of the volume is now invalid.
- A host has assigned the logical volume to a scratch category that has the fast-ready attribute set, the category has a nonzero delete expired data parameter value, the parameter value has been exceeded, and the TS7700 has deleted the logical volume.
- A host has modified the contents of the volume. This could be a complete rewrite of the volume, or appending to it. The new version of the logical volume will be migrated to a different physical location, and the older version is now invalid.

When a physical volume contains encrypted data, the TS7700 is able to perform a fast erase of the data, under certain circumstances, by "shredding" the encryption keys on the cartridge. This shredding of the encryption keys method is used when a stacked volume was encrypted on its previous use. This reduces the erasure time down to a couple of minutes instead of several hours. The first erasure of an encrypted volume causes the keys to be shredded and a pattern to be written to the tape.

The TS7700 keeps track of the amount of active data on a physical volume. It starts at 100% when a volume becomes full. Although the granularity of the percentage of full TS7700 tracks is 0.1%, it rounds down, so even *one byte* of inactive data drops the percentage to 99.9%. TS7700 keeps track of the time that the physical volume went from 100% full to less than 100% full by:

- Checking on an hourly basis for volumes in a pool with a non-zero setting
- Comparing this time against current time to determine if the volume is eligible for reclaim

This data erase function is enabled on a pool basis. It is enabled when a nonzero value is specified for the data erase reclaim policy. When enabled, all physical volumes in the pool are erased as part of the reclaim process, independent of under which reclaim policy the volume became eligible for reclaim.

Any physical volume that has a status of read-only is not subject to this function and is not designated for erasure as part of read-only recovery.

If a customer uses the eject stacked volume function, no attempt is made to erase the data on the volume prior to ejecting the cartridge. The control of expired data on an ejected volume becomes a customer responsibility.

Volumes tagged for erasure cannot be moved to another pool until erased, but they can be ejected from the library, because such a volume is usually removed for recovery actions.

The usage of the *Move function of the LM* will also cause a physical volume to be erased, even though the number of days specified has not yet elapsed. This includes returning borrowed volumes.

Virtual volume reconciliation and reclamation

Every time a logical volume is modified, the data from the previous use of this logical volume, which is on a stacked volume, becomes obsolete. When you modify a logical volume, you have a new virtual volume in the TVC and the copy on the stacked volume is invalidated but still exists in its current state on the physical volume.

The reconciliation process checks for invalidated volumes. A reconciliation is that period of activity by the TS7700 when the most recent instance of a logical volume is determined as the active one and all other instances of that volume are deleted from the active volume list. This process automatically adjusts the active data amount for any stacked volumes that hold invalidated logical volumes.

A reconciliation normally occurs sometime during a 24-hour period, before every reclamation process. This makes the reclamation process more effective because it deletes more invalidated logical volumes from the active volume list and therefore reduces the percentage of active data on more stacked cartridges.

Reclamation, which consolidates active data and frees stacked volumes for return to scratch use, is part of the internal management functions of a TS7700.

2.3.4 Copy Export

One of the key reasons to use tape is for recovery of critical operations in the event of a disaster. The TS7700, in a grid configuration, provides for automatic, remote replication of data that supports recovery time and recovery point objectives measured in seconds. For customers that do not require the recovery times that can be obtained in a grid configuration, a new function called Copy Export is being introduced for the TS7700. With Copy Export, logical volumes written to a TS7700 can be removed from the TS7700 and taken to an offsite location to be used for disaster recovery.

The Copy Export function is supported on all configurations of the TS7700, including grid configurations. In a grid configuration, each TS7700 is considered a separate source TS7700. This means that only the physical volume exported from a source TS7700 can be used for recovery of a source TS7700. Physical volumes from more than one source TS7700 in a grid configuration cannot be combined to use in recovery. Recovery is only to a single cluster configuration. After recovery, the grid MES offering can be applied to recreate a grid configuration.

To perform a Copy Export operation, the TS7700 must have a minimum of four available physical tape drives.

> **Note:** Support for the Copy Export function in a standalone configuration is available with TS7700 Release 1.3. Support for a grid configuration is available with TS7700 Release 1.4.

The Copy Export function allows a copy of selected logical volumes written to the TS7700 to be removed and taken offsite for disaster recovery purposes. The benefits of volume stacking, which places many logical volumes on a physical volume, are retained with this function. In addition, because the data being exported is a copy of the logical volumes, the logical volume data remains accessible by the production host systems.

During a Copy Export operation, all of the physical volumes with active data on them in a specified secondary pool are removed from the library associated with a specific TS7700 performing the operation. Only the logical volumes that are valid on that TS7700 are considered during the execution of the operation (if they are in the cache, but have not yet been migrated to the secondary pool, copies are performed as part of the Copy Export operation). If the TS7700 is in a grid configuration, copies that have not been completed to the TS7700 performing the Copy Export operation are not considered during the execution of the operation. It is expected that Copy Export operations will be run on a periodic basis resulting in multiple groups of physical volumes that contain the copies of the logical volume from the TS7700.

During the Copy Export operation, a copy of the current TS7700's database is written to the exported physical volumes. To restore access to the data on the physical volumes removed, all exported physical volumes for a source TS7700 are placed into a library that is attached to an empty TS7700. A disaster recovery procedure is then performed that restores access using the latest copy of the database.

The physical volumes exported during a Copy Export operation continue to be managed by the source TS7700 with regard to space management. As logical volumes that are resident on the exported physical volumes are expired, rewritten, or otherwise invalidated, the amount of valid data on a physical volume will decrease until the physical volume becomes eligible for reclamation based on the customer's provided criteria. The exported physical volumes to be reclaimed are not brought back to the source TS7700 for processing. Instead, a new secondary copy of the remaining valid logical volumes is made using the primary logical volume copy as a source.

The next time the Copy Export operation is performed, the physical volumes with the new copies are also exported. The physical volumes that were reclaimed (that are offsite) no longer are considered to have valid data and can be returned to the source TS7700 to be used as new scratch volumes.

The host that initiates the Copy Export operation first creates an export list logical volume on the TS7700 that will perform the operation. The export list volume contains instructions regarding the execution of the operation as well as a reserved file that the TS7700 will use to provide completion status and export operation information. As part of the Copy Export operation, the TS7700 creates response records in the reserved file that list the logical volumes exported and the physical volume that they reside on. This information can be used by the customer as a record for what data is offsite. The TS7700 also writes records in the reserved file on the export list volume that provide the current status for all physical volumes with a state of copy exported.

The Bulk Volume Information Retrieval function can also be used to obtain a current list of exported physical volumes for a secondary pool. For each exported physical volume, information is provided regarding the amount of active data that each contains.

We now describe the basic flow of the Copy Export function. Figure 2-22 shows the flow of the Copy Export volumes and how they are used in the event of a disaster recovery. We discuss a disaster recovery scenario where the production site is lost and the latest set of Copy Export volumes is used to restore the TS7700 at a remote site.

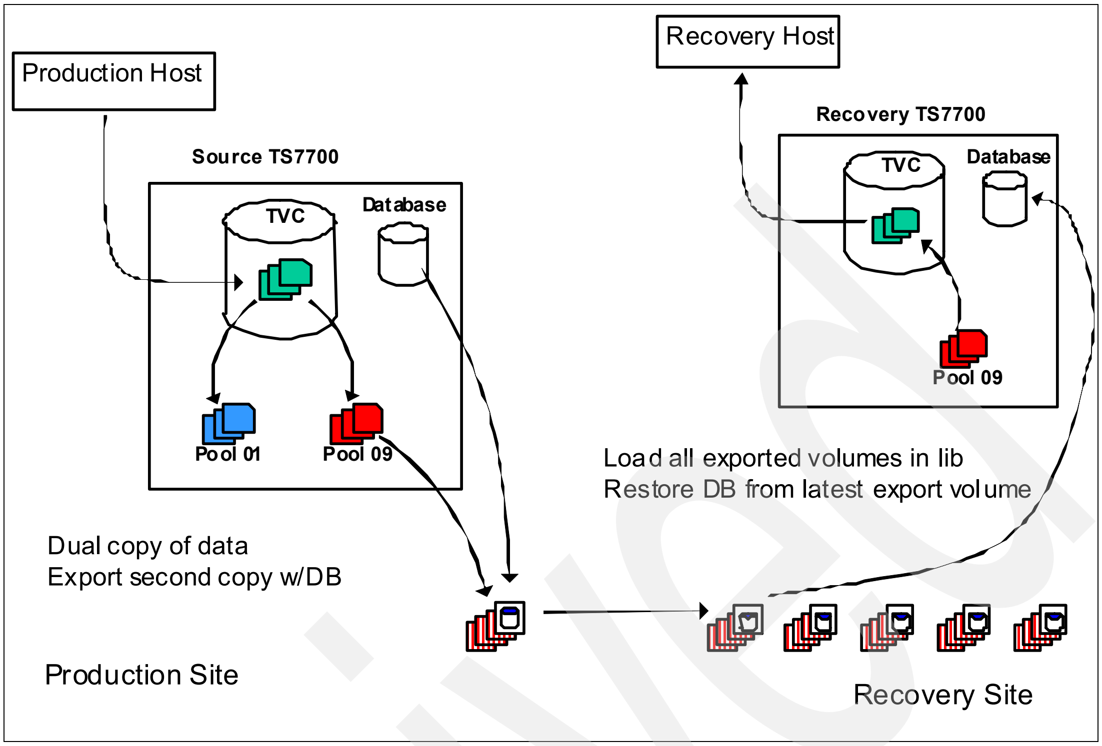

Figure 2-22 Copy export overview

You will set up the Storage Group and Management Class constructs to use separate pools for the primary (pool 01) and secondary (pool 09) copies of the logical volume. The existing Management Class construct, which is part of Advanced Policy Management (APM), is used to create a second copy of the data to be copy exported. You configure the management class actions at the Library Manager, using the operator panel or the Web specialist. A new control on the panel allows you to designate a secondary pool as a Copy Export pool. As logical volumes are written, the secondary copy of the data is pre-migrated to stacked volumes in the Copy Export pool. Our example uses pool 09 as the Copy Export pool.

When the time comes to initiate a Copy Export, you will run a Copy Export job from the production host. The TS7700 will pre-migrate any logical volumes in the Copy Export pool that have not been pre-migrated. Any new logical volumes written after the Copy Export operation is initiated will not be included in the Copy Export set of physical volumes. The TS7700 then writes a complete TS7700 database to each of the physical volumes in the Copy Export set.

Your Copy Export job can specify whether the stacked volumes in the Copy Export set should be ejected immediately or placed into the export-hold category. When the Copy Export set is in the export-hold category, you will need the operator to request that the export-hold volumes be ejected. The choice to eject as part of the Copy Export job or to eject them later from the export-hold category will be based on your operational procedures. You will transport the ejected Copy Export set to the disaster recovery site or vault. Your Recovery Point Objective (RPO) will determine the frequency of the Copy Export operation.

You will set up a rotation schedule for the Copy Export sets at the disaster recovery site. You will periodically return sets of stacked volumes to the production system TS7700 where they will be reintegrated. The TS7700 will reconcile the logical volumes on the returning stacked volumes based on activity while the stacked volumes were outside the library.

In the event of a disaster, you will perform a recovery process through the TS7700 Management Interface on an empty TS7700. As part of that recovery process, you will insert all of the copy exported stacked volumes for the source TS7700 into the library. If there are multiple pools that have been exported, you will recover one pool at a time. You will also identify the stacked volume that contains the latest copy of the source TS7700's database. The TS7700 then restores the TS7700 database from the stacked volume containing the TS7700 database. The TS7700 then restores the Library Manager database. At this point the Disaster Recovery host can start operations.

2.3.5 Encryption

The importance of data protection has become increasingly apparent with news reports of security breaches, loss and theft of personal and financial information, and government regulation. Encryption of backstore tapes helps control the risks of unauthorized data access without excessive security management burdens or subsystem performance issues.

The encryption solution for tape virtualization consists of several components.

Tape encryption solutions from IBM use an Encryption Key Manager (EKM) as a central point from which all encryption key information is managed and served to the various subsystems. The EKM communicates with the TS7700 Virtualization Engine as well as tape libraries, control units, and open-systems device drivers.

The TS1120 Model E05 encryption-enabled tape drive provides hardware that performs the cryptography function without reducing the data-transfer rate.

The TS7700 Virtualization Engine provides the means to manage the use of encryption and what keys are used on a storage pool basis. It also acts as a proxy between the tape drives and the EKMs, using Ethernet to communicate with the EKMs and Fibre Channel connections with the drives. Encryption support is enabled with Feature Code 9900 on the TS7740.

Encryption on the TS7700 Virtualization Engine is controlled on a storage pool basis. “Storage Group” and “Management Class” DFSMS constructs specified for logical tape volumes determine, through mapping in the Library Manager, which storage pools are used for the primary and secondary (if used) copies of the logical volumes. The storage pools, originally created for management of physical media, have been enhanced to include encryption characteristics.

For encryption support, all drives attached to the TS7700 must be encryption capable and encryption-enabled. The TS1120 drives must also be enabled to run in their native format mode. The management of encryption is performed on a physical volume pool basis. Through the management interface, one or more of the 32 pools can be enabled for encryption. As part of enabling a pool for encryption, the customer enters one or two key labels for the pool and a key mode. A key label can be up to 64 characters. Key labels do not have to be unique per pool and the management interface provides the capability to assign the same key label to multiple pools.

For each key, a key mode can be specified. The supported key modes are Label and Hash. As part of the configuration set up for encryption, the customer provides I/P addresses for a primary and optional secondary Key Manager through the TS7700 Management Interface.

For logical volumes that contain data that is to be encrypted, host applications direct them to a specific pool that has been enabled for encryption using the Storage Group storage construct name. All data directed to a pool enabled for encryption will be encrypted when it is pre-migrated to the physical stacked volumes. The storage group construct name is assigned to a logical volume when it is mounted as a scratch volume. Through the library manager, the storage group name is associated with a specific pool number. When the data for a logical volume is to be copied from the TVC to a physical volume in an encryption enabled pool, the TS7740 determines whether a new physical volume needs to be mounted. If one is required, as part of the mount process, the TS7740 directs the drive to use encryption and provides it with the key labels specified by the customer for the pool. When the first write data is received by the drive, a connection is made to a key manager to obtain the key needed to perform the encryption. Physical volumes are encrypted with the keys that are in effect when they are (re)written from beginning of the tape. Any partially-filled physical volumes continue to use the encryption setting at the time the tape was written from the beginning of the tape until they are reclaimed and rewritten again from the beginning of the tape.

Figure 2-23 on page 50 illustrates that the method for communicating with a key manager is through the Ethernet interface that is used to connect the TS7700 to a customer's network for access to the management interface. The request for an encryption key is directed to the I/P address of the primary key manager. Responses are passed through the TS7700 to the drive. In the event the primary Key Manager did not respond to the key management request, the optional secondary key manager I/P address will be used. After the TS1120 drive has completed the key management communication with the key manager, it will accept data from the TVC.

When a logical volume needs to be read from a physical volume in a pool enabled for encryption, either as part of a recall or reclamation operation, the TS7700 uses the key management interface to obtain the necessary information to decrypt the data.

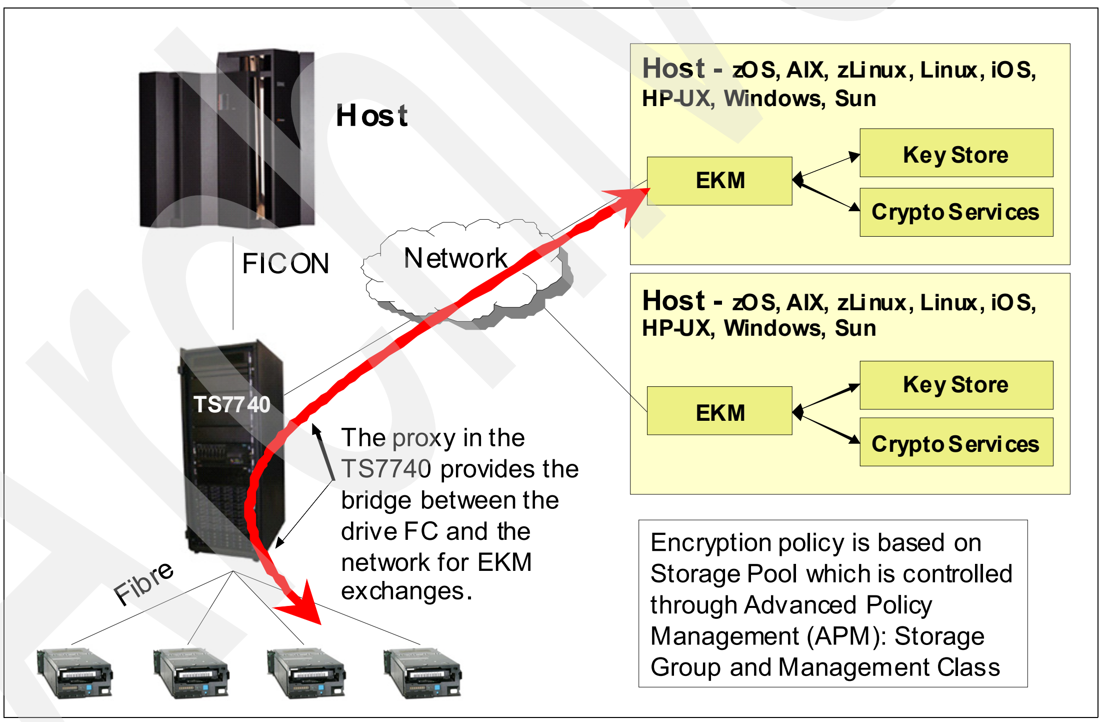

Figure 2-23 TS7700 encryption

The TS1130 Tape Drive, supported with TS7700 R1.5 and higher, provides the same Encryption support as the TS1120 when operating in E05 mode.

2.3.6 Host Console Request

This function provides a host console command to request information about the current operational status of the TS7700, its logical and physical volumes, and its physical resources. The intent of the information allows an MVS console operator to perform basic problem determination without having to access a Web-based specialist. It is also intended to allow you to automate monitoring of the physical resources of the TS7700 using MVS facilities.

Note: With the TS7700 machine code Release 1.4a, the Host Copy Control function has been introduced. This function is an extension of the Host Control Request command and allows host control of logical volume copying on a cluster basis.

Refer to 7.3.3, "Host Console Request" on page 375 for more details concerning the Host Console Request command.

2.4 TS7700 Multi Cluster Grid

In this section, we explain logical volume management and Multi Cluster Grid configuration options. When multiple clusters are configured for copies in a grid, logical volumes and their copies reside on some or all of the other clusters.

2.4.1 Data integrity by volume ownership

Any logical volume or any copies of it among the clusters in a grid, can be accessed from any virtual device in the system regardless whether the cluster from which the host is accessing the volume has a copy. This is subject to *volume ownership*. At any point in time a logical volume is "owned" by one cluster. The owning cluster has control over access to the volume and over changes to the attributes associated with the volume (such as volume category or storage constructs).

The cluster that has ownership of a logical volume can change dynamically. The TS7700 node passes the ownership of the logical volume as part of mount processing, as required to ensure that the cluster with the virtual device associated with the mount has ownership. When a mount request is received on a virtual device address, the TS7700 cluster for that virtual device must have ownership of the volume to be mounted or must obtain the ownership from the cluster that currently owns it.

If the TS7700 clusters in a grid configuration and the communication paths between them are operational, the change of ownership and the processing of logical volume-related commands are transparent with regard to the operation of the TS7700. If a TS7700 cluster has a host request for a logical volume which it does not own, and it cannot communicate with the owning cluster, the operation against that volume will fail, unless some additional direction is given. In other words, clusters will not automatically assume or take over ownership of a logical volume without being directed.

The volume ownership relationship protects the volume from being accessed or modified by different clusters simultaneously. If more than one cluster has ownership of a volume, it could result in the volume's data or attributes being changed differently on each cluster resulting in a data integrity issue with the volume. If a TS7700 cluster has failed or is known to be unavailable (for example, it is being serviced), its ownership of logical volumes is transferred to another TS7700 cluster, as follows:

- Read-only Ownership Takeover

  When Read-only Ownership Takeover (ROT) is enabled for a failed cluster, ownership of a volume is allowed to be taken from a TS7700 cluster that has failed. Only read access to the volume is allowed through the other TS7700 clusters in the grid. Scratch mounts are failed to a cluster in this mode. When ownership for a volume has been taken in this mode, any operation attempting to modify data on that volume or change its attributes is failed. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster has been restored.

- Write Ownership Takeover

  When Write Ownership Takeover (WOT) is enabled for a failed cluster, ownership of a volume is allowed to be taken from a cluster that has been marked as failed. Full access is allowed through the requesting TS7700 cluster in the grid. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster has been restored.

- Service preparation or service mode

  When a TS7700 cluster is placed in service preparation mode or is in service mode, ownership of its volumes is allowed to be taken by the other TS7700 clusters. Full access is allowed. The mode for the cluster in service remains in place until it has been taken out of service mode.

You can set the level of ownership takeover, Read-only or Write Ownership, through the TS7700 Management Interface. Note that you cannot set a cluster in service preparation after it has failed already.

Autonomic Ownership Takeover

Autonomic Ownership Takeover provides a means of allowing a subsystem to automatically take ownership of the volumes owned by a failed site. For an Autonomic Ownership Takeover a TSSC is required at each cluster. The path between the TSSCs is used to determine whether the remote node is down, or whether the communication path between the sites is down. The TSSCs at each site can be configured to automatically perform Volume Ownership transfer.

TSSC is used as another path to the remote site. To be able to determine whether the TSSC is able to reach the TSSC at the remote site, the remote TSSC determines the state/health of the remote node, and passes information back to the local site.

If a site failure is determined, then the logical volume acquiring ownership can be automatically enabled by the system, without requiring an intervention by an operator or by service personnel. Automatic Volume Ownership will be enabled in either R/O or R/W capability, depending on the customer-configured policy (see Figure 2-24).

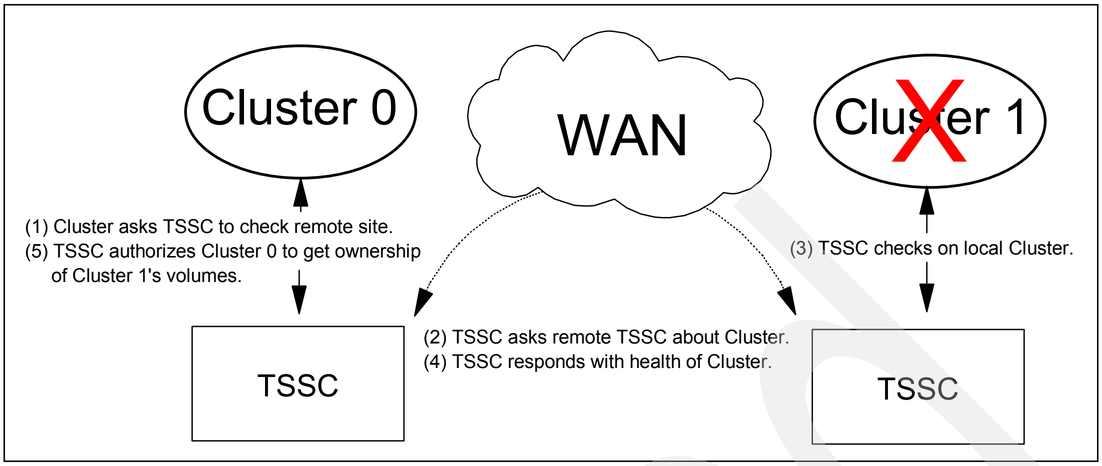

Figure 2-24 Successful Autonomic Ownership Takeover

If the TSSC is unable to ascertain the health of the remote site, then Autonomic Volume Ownership Takeover is not performed (see Figure 2-25).

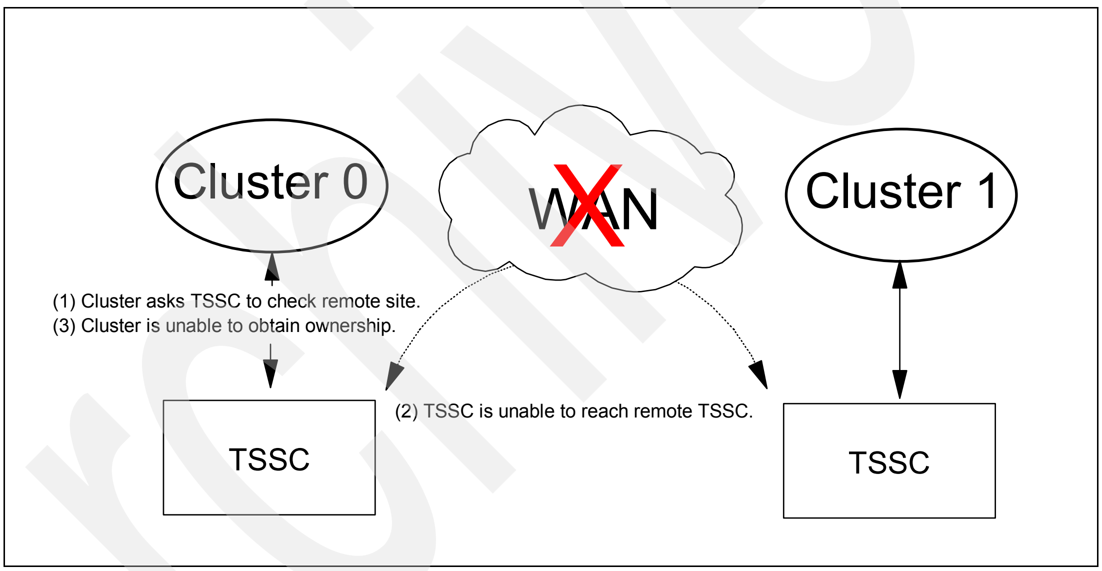

Figure 2-25 Unsuccessful Autonomic Ownership Takeover

Through options provided with SMIT, the SSR can enable or disable this function and select which ownership takeover mode is to be entered if the other TS7700 is determined to not be operational.

With a Three-Cluster Grid, an additional vote is available for peer outage detection (see Figure 2-26 on page 54). For example, if cluster 0 thinks that cluster 2 is unavailable, but cluster 1 disagrees, then a network issue must be present. In this case the ownership takeover will be failed. If the TSSCs of the available clusters agree that the third cluster is unavailable, then ownership takeover can occur.

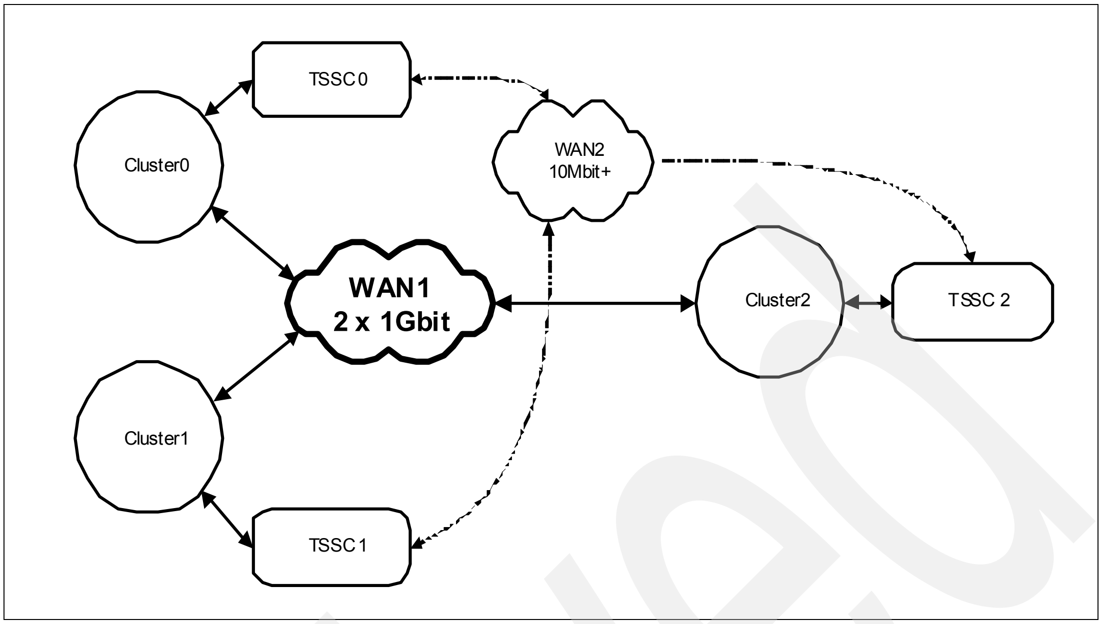

Figure 2-26 Autonomic Ownership Takeover for a Three-Cluster Grid

2.4.2 Copy policy management

When a TS7700 Virtualization Engine is part of a Multi Cluster Grid configuration, there are several customer policies and settings that can be used to influence the location of data copies, when the copies are to be made, and the performance of access to the data. This section describes the customer controls available. The implementation of multiple copies and their management is different than that of the prior generation PtP VTS. Differences will be noted throughout the section.

To put the concepts described below in context, a comparison of some specific aspects of the architecture of the prior generation PtP VTS and the TS7700 is useful to do first. A configuration of the prior generation PtP VTS included two sets of storage server hardware, Virtual Tape Controllers (VTCs) and Virtual Tape Servers (VTSs). The VTCs served two purposes: to provide an access point for hosts to write and read data resident in the VTSs and to provide for replication of data between the VTSs.

With the TS7700 architecture, the hardware for the VTCs has been eliminated and its function embedded into the Virtualization Engine Controller. The access point function is now part of the vNodes and the replication function is part of the hNodes. Associated with each cluster of vNodes and hNodes is the Tape Volume Cache (TVC) and, like the prior generation VTS, this is where host data resides when being accessed.

A significant difference with the TS7700 architecture is that any vNode in the grid can access the data within the TVC of any cluster directly. A copy of the data does not have to be in the TVC in the same cluster as the vNode the host is using for data access. Because of this difference, the copy policy management and where the data is written and read are tied together differently.

With the prior generation PtP VTS, a VTC selected one of the two VTSs to act as the I/O VTS for a mounted volume based upon explicit direction in the Management Class, internal algorithms, and IBM service representative settings on each VTC. With the TS7700, the determination of which cluster's TVC will act as I/O TVC is based upon the copy requirements of the Management Class, internal algorithms, and customer settings. The differences are somewhat subtle and are described in more detail below.

Copy management

With the TS7700's architectural capability of having more than two clusters peered together, copy policy management needs are different than those of the prior generation PtP VTS. In a TS7700 Grid, you might want to have multiple copies of a virtual volume on different clusters within the grid. You might also want to specify when the copies are performed relative to the job that has written to a virtual volume and have that be different for each cluster.

Like the prior generation PtP VTS, copy management is controlled through the Management Class storage construct. Using the Manage Management Classes panel through the Library Manager user interface, you can create Management Classes that define where copies reside and when they will be synchronized to the host job that created them. Depending on your business needs for multiple copies of your data, multiple Management Classes, each with a different set of definitions, can be created. There are two key questions that determine the copy management of the TS7700:

- Where do I want my copies to reside?
- When do I want my copies to become consistent with the originating data?

Where copies will reside

With the prior generation PtP VTS, you had the ability using a Management Class to specify the VTS that is to be the I/O VTS and to indicate whether a copy should reside on the other VTS. With the TS7700, this concept is expanded to include multiple sites. What is now specified for a Management Class is where copies are to reside after all replication has been completed.

When a TS7700 is included in a grid configuration, the Management Class definition panel lists each cluster by its distributed library name and allows a copy policy selection for each. For example, assume there are three clusters in the grid: LIBRARY1, LIBRARY2, and LIBRARY3. A portion of the Management Class definition panel will include the cluster name and allow a copy consistency point to be specified for each cluster. If a copy is to reside on a cluster's TVC, then you indicate a copy consistency point. If you do not want a cluster to have a copy of the data, then you specify the No Copy option.

> **Note:** The architecture of the TS7700 allows for more than two clusters in a grid. The current version of the TS7700 supports three clusters in a Multi Cluster Grid configuration.

When copies become consistent with the originating data

As part of the volume mount process within a grid, the TVC of one of the clusters is selected as the I/O TVC (for details on the selection process and the factors considered as part of the process, see "I/O TVC selection" on page 57).

All host read and write operations are routed to that TVC, which can be in a different cluster than the mount vNode (the vNode providing the virtual tape drive address the I/O operations are being issued against). To the application writing data to a volume, the data on the I/O TVC is always consistent with the application. Based on the policies defined for the Management Class assigned to the volume, the data from the I/O TVC is replicated to other TVCs associated with the other clusters in the grid. Two copy consistency points are being provided, at volume Rewind/Unload time and deferred after Rewind/Unload time.

There is also the option to not have a copy of the data reside on a cluster. If a copy consistency point of Rewind/Unload is defined for a cluster in the Management Class assigned to the volume, a consistent copy of the data must reside in that cluster's TVC before command completion is indicated for the Rewind/Unload command.

If multiple clusters have a copy consistency point of Rewind/Unload, then, optionally, all of their associated TVCs must have a copy of the data before command completion is indicated for the Rewind/Unload command. Refer to “Override settings” on page 58 for the description of the option regarding indicating command completion of the Rewind/Unload command when multiple clusters specify the Rewind/Unload copy consistency point.

If a copy consistency point of *deferred* is defined, the copy to that cluster's TVC can occur any time after the Rewind/Unload command has been processed for the I/O TVC. A mixture of copy consistency points can be defined for a Management Class.

Consistency point example

Figure 2-27 shows the panel of the ETL specialist where you can set the consistency points for a Management Class. In our example we set the consistency points for Management Class MCPOOL4D on cluster 1 of the grid. Every time a logical volume is mounted on cluster 1, which has this Management Class assigned, we want a consistent copy of the data to reside in the TVC of clusters 0 and 1 before command completion of Rewind/Unload is indicated to the host and a deferred copy on cluster 2. Therefore, we specify RUN for both the local cluster 1 and the remote cluster 0 and specify Deferred for cluster 2. Keep in mind that the consistency point settings for cluster 0 and cluster 2 can be different for the same Management Class.

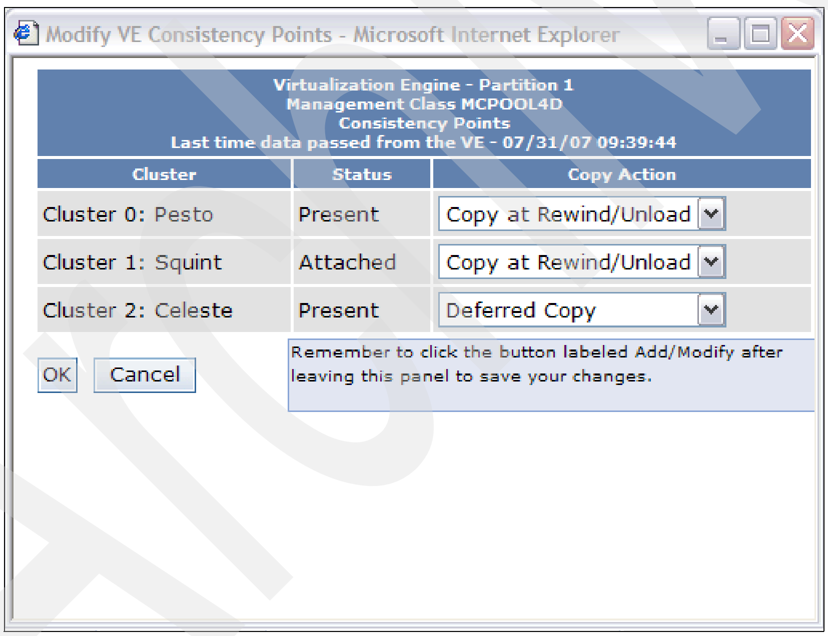

Figure 2-27 Setting cluster copy data consistency points

You can specify the following consistency points:

No Copy If a data consistency point of No Copy is specified, the data created on one TS7700 is not copied to the other TS7700. If the TS7700 cluster the data was written to fails, the data for that logical volume is inaccessible until that TS7700 cluster's operation is restored.

Copy at Rewind Unload If a data consistency point of RUN is specified, the data created on one TS7700 is copied to the other TS7700 as part of successful rewind unload command processing, meaning that for completed

	jobs, a copy of the volume will exist on both TS7700s. Access to data written by completed jobs (successful Rewind/Unload) prior to the failure is maintained through the other TS7700 cluster. Access to data of incomplete jobs that were in process at the time of the failure is not provided.
Deferred Copy	If a data consistency point of Deferred is specified, the data created on one TS7700 is copied to the other TS7700 after successful Rewind/Unload command processing. Access to the data through the other TS7700 cluster is available before the copy completes, as long as the source TS7700 is still available. Because there will be some delay in performing the copy, access might or might not be available when a failure occurs.

The default Management Class action is to have a Rewind/Unload consistent copy reside only on the local cluster. The other clusters would have a No Copy option specified. The default actions are applied whenever a new construct is defined through the Library Manager or specified in a **`mount`** command and it was not previously defined.

Management Class locality to a cluster

Similar to the prior generation PTP VTS, Management Classes for the TS7700 are created at each Library Manager associated with a cluster. The same Management Class can be defined differently at each cluster, and there are valid reasons for doing so. For example, one of the functions controlled through Management Class is to have a logical volume copied to two different physical volume pools.

You might want to have two separate physical copies of your logical volumes on one of the clusters and not on the others. Through the Library Manager associated with the cluster where you want the second copy, you would specify a secondary pool when defining the Management Class. For the Management Class definition on the other clusters, you would not specify a secondary pool. For example, you might want to use the Copy Export function to remove a copy of data from the cluster to be taken to a disaster recovery site.

It is important to note that during mount processing, the copy consistency point information that is used for a volume is taken from the Management Class definition for the cluster the mount vNode is associated with. We recommend that you define the copy consistency point definitions of a Management Class the same on each cluster so as to avoid confusion where copies will reside. You can come up with a scenario in which you would want to define different copy consistency points for the same Management Class on each of the clusters, just be aware that the location of copies and when the copies are consistent with the host that created the data are going to be different depending upon which cluster a mount is issued to.

I/O TVC selection

The TVC associated with one of the clusters in the grid is selected as the I/O TVC for a given tape mount request. All I/O operations associated with the virtual tape drive the mount was issued on are routed from its vNode to the TVC on the cluster selected. The vNode is referred to as the mount vNode. The I/O TVC is selected as part of the processing for the mount request. Similar to the prior generation PTP VTS's I/O VTS selection process, the TS7700 takes into consideration such factors as data validity, cluster availability, type of mount request (Fast Ready (scratch) or non-Fast Ready), performance, local preferencing, and override settings. The many factors are evaluated and an I/O TVC is selected. The TS7700 also takes into consideration the copy consistency points defined in the Management Class associated with the volume being mounted. A cluster with a Rewind/Unload copy consistency point requirement will be weighted heavier than one with a deferred copy consistency point.

For example, Management Class MCPROD01 has been defined with the following copy data consistency points:

- LIBRARY1 Rewind/Unload
- LIBRARY2 Deferred

Assume a scratch (category) mount request is received specifying a Management Class of MCPROD01 on a virtual drive address associated with cluster LIBRARY2 and no customer overrides have been established for that cluster. The selection process for picking the I/O TVC will select the TVC associated with cluster LIBRARY1 as long as it is available because it has the highest copy consistency point requirement. The reason for weighting the copy consistency point in this way is so that when the Rewind/Unload is received, the data will already be consistent and no delay will be introduced in performing a copy to another cluster's TVC prior to indicating that the Rewind/Unload command has completed.

If a Management Class had been defined that specified more than one Rewind/Unload cluster copy consistency point, those clusters' TVCs would be given equal weight and other factors such as performance between the mount vNode and the TVCs would steer the decision. For example, a Management Class MCPROD02 could be defined as follows:

- LIBRARY1 Rewind/Unload
- LIBRARY2 Rewind/Unload

If a scratch mount specifying MCPROD02 is received for a vNode in cluster LIBRARY2, the TVC associated with LIBRARY2 would be selected because it meets the cluster consistency point requirement and is expected to have the best performance because the TVC is local. The TS7700 estimates the "best" performing TVC based on many factors including whether a TVC is full and throttling (because this is a scratch mount, other factors like consistency and cache residency are not considered). It could be possible that the TVC in LIBRARY2 is full and throttling. In this case the TVC in LIBRARY1 would likely be selected to avoid job throttling.

With the prior generation PTP VTS, you could force a specific VTS to be the I/O VTS by specifying that option in the Management Class definition. With the TS7700, you can also direct that the TVC for a specific cluster be used for the I/O TVC. The method of defining this capability is different than for the prior generation's method and allows for more flexibility. For example, say that you want the data to only reside on the cluster associated with LIBRARY2—you would define a Management Class that only has a copy consistency point for LIBRARY2.

The other cluster would be specified as No Copy. Regardless of which vNode a mount using that Management Class is received on (assuming that you have defined the Management Class the same on all clusters), the TVC associated with cluster LIBRARY2 would be selected. If the LIBRARY2 cluster is not available, the mount would fail. It might be that failing the mount is not desired. If that is the case, you could define the copy point for LIBRARY3 as Rewind/Unload and another cluster's copy consistency point as Deferred. With that Management Class definition, if the LIBRARY2 cluster is not available, the TVC associated with the other cluster will be chosen and the mount is not failed. A deferred copy will still be made to LIBRARY2 when it is again available.

Note: The copy consistency point is considered for both scratch and specific mounts. This is different from the PTP VTS, which only looked at the copy policy for scratch mounts.

Override settings

With the prior generation PTP VTS, there were several optional override settings that influenced how an individual VTC selected a VTS to perform the I/O operations for a mounted tape volume. The override settings were determined by the customer but set by a service representative. With the TS7700, the customer defines and sets the optional override settings that influence the selection of the I/O TVC and replication responses.

TS7700 overrides I/O TVC selection and replication response

The settings are specific to a cluster, meaning that each cluster can have different settings if desired. The settings take effect for any mount requests received after the settings were saved. Mounts already in progress are not affected by a change in the settings. The following override settings are supported:

- Prefer Local Cache for Fast Ready Mount Requests

  This override will select the TVC local to the mount vNode cluster as the I/O TVC as long as it is available and a copy consistency point other than No Copy is specified for that cluster in the Management Class specified with the mount. The cluster does not have to have a valid copy of the data for it to be selected for the I/O TVC.

  The default behavior of the TS7700 is to select an I/O TVC based on the highest copy consistency point defined for the Management Class associated with the mount request as well as availability considerations. If the cluster of the mount vNode has a lower cluster copy consistency point, a TVC remote from the cluster local to the mount VNode would be selected and that might have undesired performance impacts on the application.

  > **Note:** The copy consistency rules are not overridden with this override. For example, if the cluster associated with the I/O TVC has a copy consistency point of Deferred and another cluster has a copy consistency point of Rewind/Unload, completion of the Rewind/Unload command is not indicated until the copy has been completed to the other cluster.

- Prefer Local Cache for non-Fast Ready Mount Requests

  This override selects the TVC local to the mount VNode cluster as the I/O TVC as long as it is available and the cluster has a valid copy of the data, even if the data is only resident on a physical tape. Having an available, valid copy of the data overrides all other selection criteria. If the local cluster does not have a valid copy of the data, then the default selection criteria applies.

  The default behavior of the TS7700 is to select an I/O TVC based upon the following criteria:

  – Availability of a valid copy
  – Cache residency
  – The copy consistency point defined for the volume (highest is best)
  – Performance characteristics between clusters
  – Cluster availability

  > **Note:** If the volume is modified, the copy consistency rules are not overridden with this override.

- Force Local TVC to have a copy of the data

  The default behavior of the TS7700 is to only make a copy of the data based on the definitions of the Management Class associated with the volume mounted and to select an I/O TVC that was defined to have a copy. If the mount vNode is associated with a cluster for which the specified Management Class defined a copy consistency point of No Copy, a copy is not made locally and all data access is to a remote TVC.

  This override has two effects, depending on the type of mount requested. For a non-Fast Ready mount, a copy is performed to the local TVC as part of the mount processing. For a

Fast Ready mount, it has the effect of "ORing" the specified Management Class with a copy consistency point of Rewind/Unload for the cluster. The override does not change the definition of the Management Class, it serves only to influence the selection of the I/O TVC or force a local copy.

> **Note:** If the volume is modified, the copy consistency rules are not overridden with this override.

- Copy Count Override

  This override limits the number of RUN consistency points in a two or three cluster grid. Initial status or device end is surfaced the copy count is satisfied. Only copy consistency points of RUN are counted. For example, in a Three-Cluster Grid, if the Management Class specifies copy consistency points of RUN, RUN, RUN, and the override is set to two, then initial status or device end is presented after two copies have been made. The third RUN is downgraded to a deferred copy after the first two copies are complete. The third site that has its copy consistency point downgraded to deferred is called the *floating deferred site*. The floating deferred site or sites is the site or sites that have not completed their copy when the Copy Count has been reached.

  This override has limited use in a Two-Cluster Grid. The only limitation available is to limit the number of RUN consistency points to one. The override will only come into play for a Management Class with copy consistency points of RUN, RUN.

Overrides for Geographically Dispersed Sysplex (GDPS)

The default behavior of the TS7700 is to follow the Management Class definitions as well as considerations to provide the best overall job performance. In a Geographically Dispersed Parallel Sysplex™ (GDPS®), all I/O must be local to the mount vNode. There can be other customer requirements, such as disaster recovery testing, where all I/O must only go to the local TVC to ensure that the correct copy policies have been implemented and data is available where required.

In a GDPS environment, you have to set all three of the following overrides to ensure that the local TVC is selected for all I/O:

- Prefer Local for Fast Ready Mounts
- Prefer Local for non-Fast Ready Mounts
- Force Local TVC to have a Copy of the Data

Consistency Configuration example 1

In the scenario shown in Figure 2-28, we have two sites, with a System z host and a TS7700 Virtualization Engine installed in each. The two clusters are part of a Multi Cluster Grid configuration. The hosts attach to both TS7700s. In our example we assume that the FICON connections between the two sites only provide limited bandwidth. For this reason we vary off the virtual devices on the remote cluster for each host during normal operation. Only in case of disaster recovery we vary these devices online. This means that during normal operation host A always accesses virtual devices on cluster 0, and host B accesses virtual devices on cluster 1.

We want to have two consistent copies of the data to provide for continued data access, when a volume with a certain Management Class assigned to it is closed by the application. As we want to have the same behavior, no matter in which site a volume is created or modified, we specify for each cluster copy data consistency points RUN/RUN (indicated by RR in Figure 2-28).

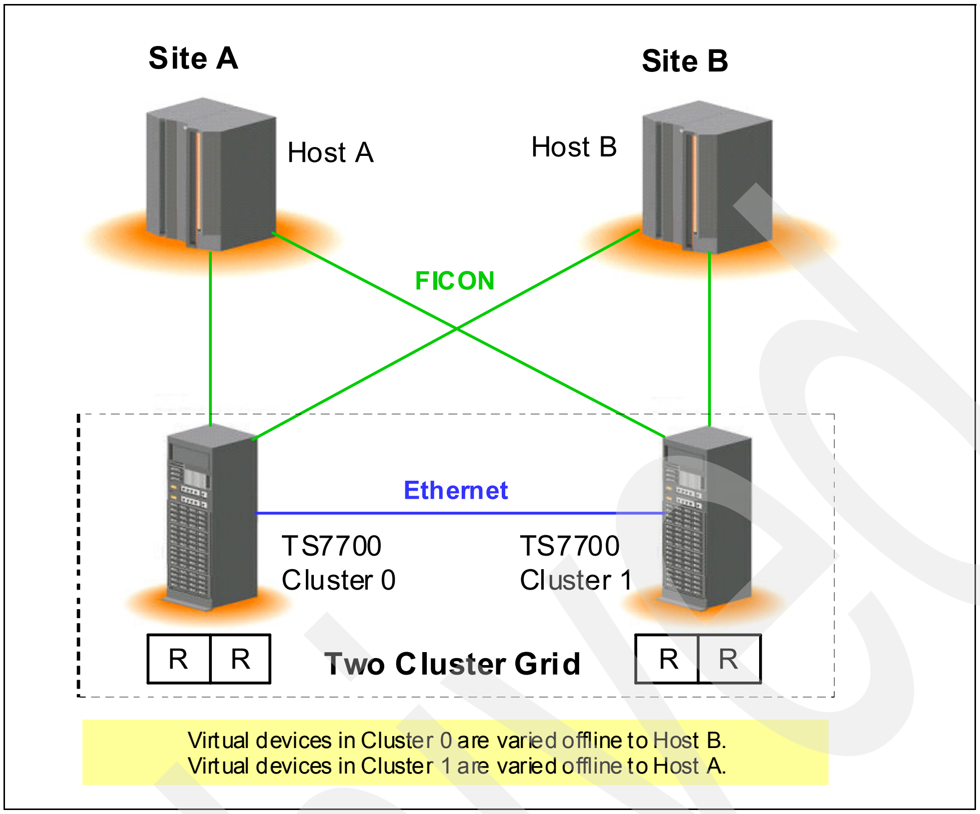

Figure 2-28 Example 1: Copy consistency settings RUN/RUN

Let us have a look at what happens when Host A mounts a logical volume on a virtual device in cluster 0. The mount vNode of cluster 0 has access to both its local TVC and to the TVC of remote cluster 1, so the local TVC will not necessarily be selected as I/O TVC. The cluster copy consistency point settings are the same for both clusters, so there is no preference as far as copy consistency is concerned. In this case locality will probably tip the scales in favor of the local TVC, as the Fibre attached drives of the local TVC give you a better performance than the Ethernet connected remote TVC under normal circumstances. You should be aware that selection of the remote TVC as I/O TVC is still possible, because numerous other factors such as cache residency or validity of data are also taken into account and can overrule the locality factor. The same considerations apply to host B and cluster 1.

If you want to ensure that the local TVC is always selected as the I/O TVC, you have to fine-tune the system by using the override settings described in this section above.

Although there is no one-to-one correspondence of TS7700 copy consistency points to the Peer-to-Peer Copy Modes of the previous generation of IBM Tape Virtualization, this TS7700 settings (RUN/RUN on both clusters) roughly compare to Immediate Copy/Balanced Mode of the Peer-to-Peer VTS.

Consistency configuration example 2

In this scenario (see Figure 2-29) we have a System z host and a TS7700 Virtualization Engine in Site A, and a remote TS7700 in Site B. The two Virtualization Engines form a Multi Cluster Grid. The host has access to all virtual devices of cluster 0 and of cluster 1.

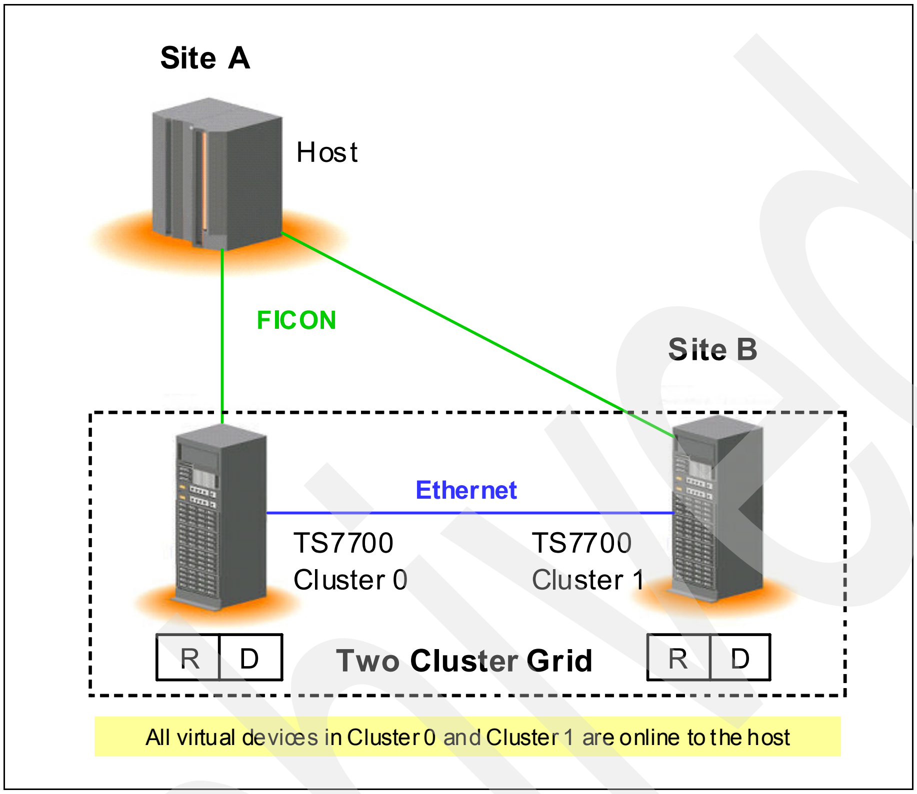

Figure 2-29 Example 2: Copy consistency point settings RUN/Deferred

We assume that a copy of the data is required on both clusters. Contrary to the previous example, we need a valid copy of the data at Rewind/Unload only on cluster 0 in the local site. The data can be copied to cluster 1 at a later time. Therefore, we define on both clusters the copy consistency points as RUN for cluster 0 and Deferred for cluster 1. When we look at a Fast Ready Mount, for example, we now have to differentiate between two cases:

- The host allocates a virtual drive on cluster 0.

  During the mount process the vNode of cluster 0 selects the I/O TVC. One of the factors it takes into consideration for this decision is the copy consistency points defined in the Management Class associated with the volume being mounted. The RUN setting for cluster 0 weighs more than the Deferred setting for cluster 1. The vNode of cluster 0 will therefore select its local TVC as I/O TVC. Locality also speaks for the TVC of cluster 0, but the decisive factors here are the copy consistency points.

- The host allocates a virtual drive on cluster 1.

  During the mount process the mount vNode in cluster 1 selects the I/O TVC. It evaluates the copy consistency points for the decision, and again the TVC of cluster 0 is selected as I/O VTC. The difference is, that now the remote TVC is selected (remote from the perspective of the selecting vNode). Locality would speak for the local TVC as I/O TVC, but copy consistency point settings weigh heavier than locality. So data is written to the TVC of cluster 0 first, using the Ethernet connection between the two clusters. The copy to the TVC of cluster 1 will happen at a later time.

You can modify the behavior of the clusters by using the Override Settings explained in "Override settings" on page 58.

Although there is no one-to-one correspondence of TS7700 copy consistency points and the Peer-to-Peer Copy Modes of the previous generation of IBM Tape Virtualization, these TS7700 settings (RUN/Deferred on both clusters) roughly compare to Deferred Copy/Preferred Mode of the Peer-to-Peer VTS.

Consistency configuration example 3

As we mentioned before, the copy consistency point settings can be different for the two clusters in a Multi Cluster Grid configuration. Though we generally recommend to have the same settings on both clusters, there might be scenarios with special requirements, where different settings are appropriate.

The scenario shown in Figure 2-30 is very much the same as the scenario in Example 1 (Figure 2-28 on page 61). The scenarios differ only in the copy consistency point settings. Here we have different copy consistency point settings defined on cluster 0 and cluster 1 of the grid configuration.

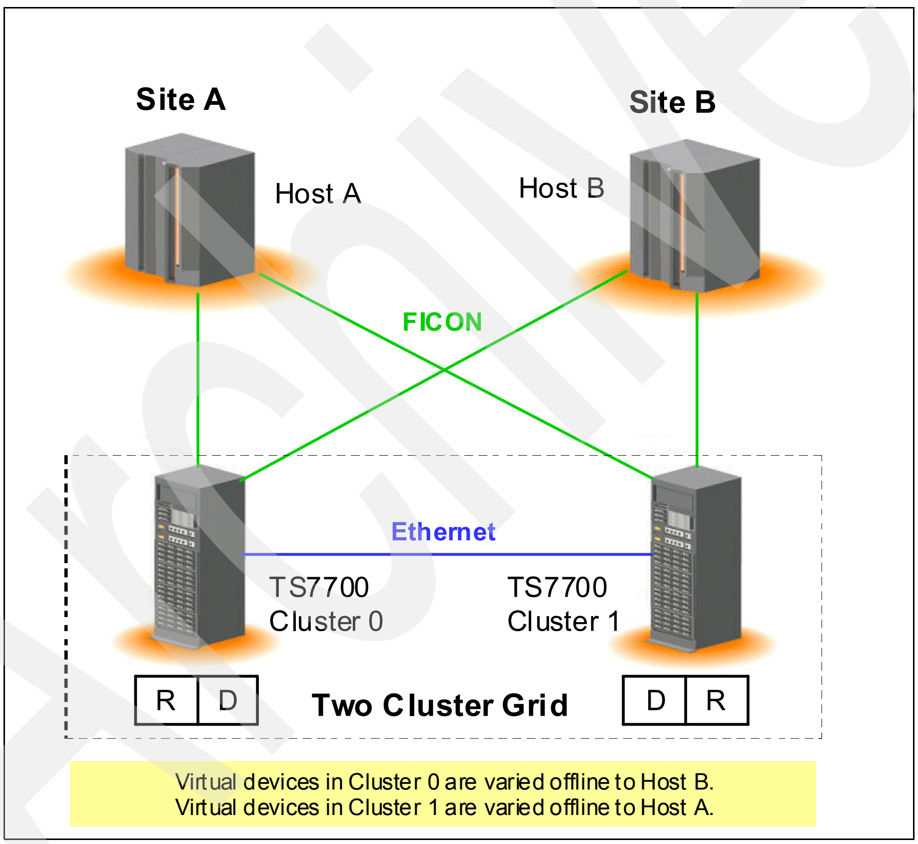

Figure 2-30 Example 3: Different consistency point settings in a grid configuration

The hosts in this scenario have access to the virtual devices of their local cluster only, as the remote virtual devices are varied offline to the hosts. When a host mounts a logical volume on one of these virtual devices, we want like to see that the respective local TVC is selected as the I/O TVC, that a copy exists in the local TVC before Rewind/Unload is indicated to the host, and that a deferred copy will be made to the TVC of the remote cluster at a later time.

To achieve this goal, we set the copy consistency points on cluster 0 to RUN for cluster 0 and to Deferred for cluster 1(RD in Figure 2-30 on page 63). On cluster 1 we reverse these settings and define Deferred for cluster 0 and RUN for cluster 1 (DR in Figure 2-30 on page 63). Regardless of the site where a virtual volume is mounted, the local TVC is preferred to the remote TVC because I/O TVC and a copy must exist in the local TVC prior to Rewind/Unload, because of the copy consistency point setting RUN for the local cluster. The copy consistency point setting Deferred for the remote cluster causes a deferred copy to be made to the remote TVC.

In the same scenario, if you wanted to always have two valid copies at Rewind/Unload for mounts on cluster 0, and just one valid copy in the local TVC for mounts on cluster 1, you would set the copy consistency points as follows:

- Cluster 0: RUN for cluster 0, RUN for cluster 1
- Cluster 1: No Copy for cluster 0, RUN for cluster 1

2.4.3 High availability and disaster recovery: Two-Cluster Grid

With a Two-Cluster Grid, you can configure the grid for disaster recovery, high availability, or both. In the following sections, we describe configuration considerations for Two-Cluster Grids. The scenarios that we present are typical configurations. However, other configurations are possible and might be better suited for your environment.

Disaster Recovery configuration

This section provides information needed to plan for a TS7700 Two-Cluster Grid configuration to be used specifically for disaster recovery purposes.

We assume that some natural or human-caused event has made the local site’s TS7700 cluster unavailable. The two TS7700 clusters reside in different locations, separated by a distance dictated by your company's requirements for disaster recovery. The only connection between the local site and the disaster recovery site are the grid interconnections. There is no host connectivity between the local hosts and the FICON channels on the disaster recovery site’s TS7700 Virtualization Engine. Figure 2-31 summarizes this configuration.

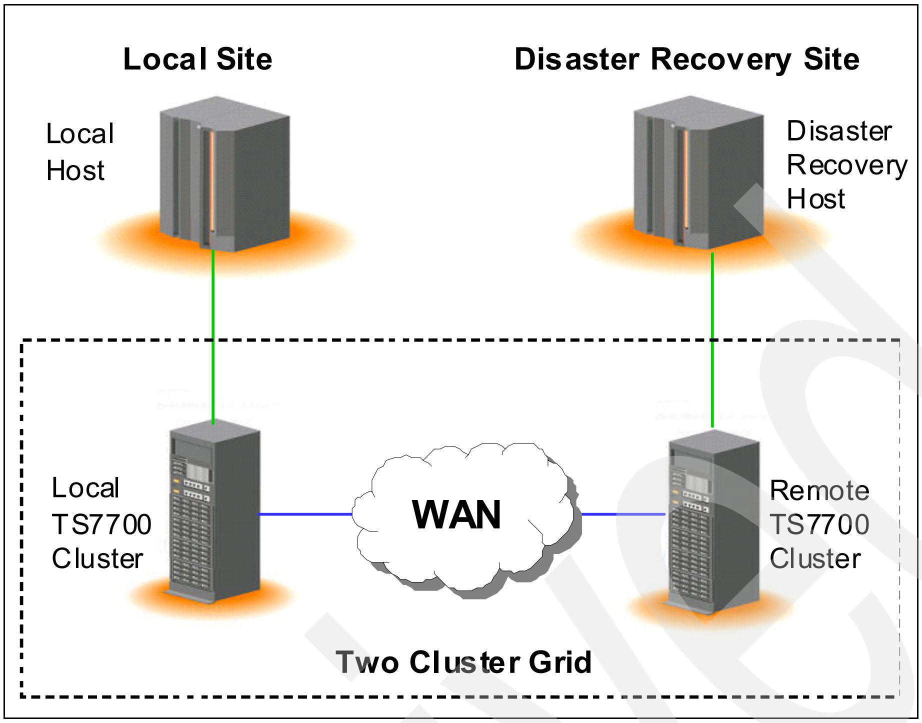

Figure 2-31 Disaster Recovery configuration

As part of planning a TS7700 Grid configuration to address this solution, you need to consider the following:

- Plan for the necessary wide area network infrastructure and bandwidth to meet the copy requirements that you need. You generally will need more bandwidth if you are primarily using a copy consistency point of RUN because any delays in copy time caused by bandwidth limitations will result in an elongation of job run times. If you have limited bandwidth available between sites, have data that is critical copied with a consistency point of RUN, with the rest of the data using the deferred copy consistency point.
- Plan for host connectivity at your disaster recovery site with sufficient resources to perform your critical workloads.
- Design and code the DFSMS Automatic Class Selection routines to control what data gets copied and by which copy consistency point.
- Prepare procedures that your operators would execute in the event the local site becomes unusable. The procedures would include such tasks as bringing up the disaster recovery host, varying the virtual drives online, as well as placing the disaster recovery TS7700 cluster in one of the ownership takeover modes.
- If the local TS7700 cluster becomes unavailable, there is no local host access to the data in the disaster recovery site's TS7700 cluster.

Configuring for high availability

This section provides information needed to plan for a TS7700 Two-Cluster Grid configuration to be used specifically for high availability. The assumption is that continued access to data is critical and no single point of failure, repair, or upgrade can impact the availability of data.

In a high availability configuration, both TS7700 clusters are located in close proximity to one another at the same site. These clusters are connected through a local area network. If one of them becomes unavailable because it has failed or is undergoing service or being updated, data can be accessed through the other TS7700 cluster until the unavailable cluster is again available. As part of planning a TS7700 Grid configuration to address this solution, you will need to consider the following:

- Plan for the virtual device addresses in both clusters to be configured to the local hosts. In this way, a total of 512 virtual tape devices are available for use (256 from each TS7700 cluster).
- Set up a copy consistency point of RUN for all data to be made highly available. With this copy consistency point, as each logical volume is closed, it is copied to the other TS7700 cluster.
- Design and code the DFSMS Automatic Class Selection routines to set the necessary copy consistency point.
- Prepare procedures that your operators would execute in the event one of the TS7700 clusters becomes unavailable. The primary task will be to place the remaining cluster in the mode so it can take over ownership of logical volumes that were currently owned by the unavailable TS7700 cluster.

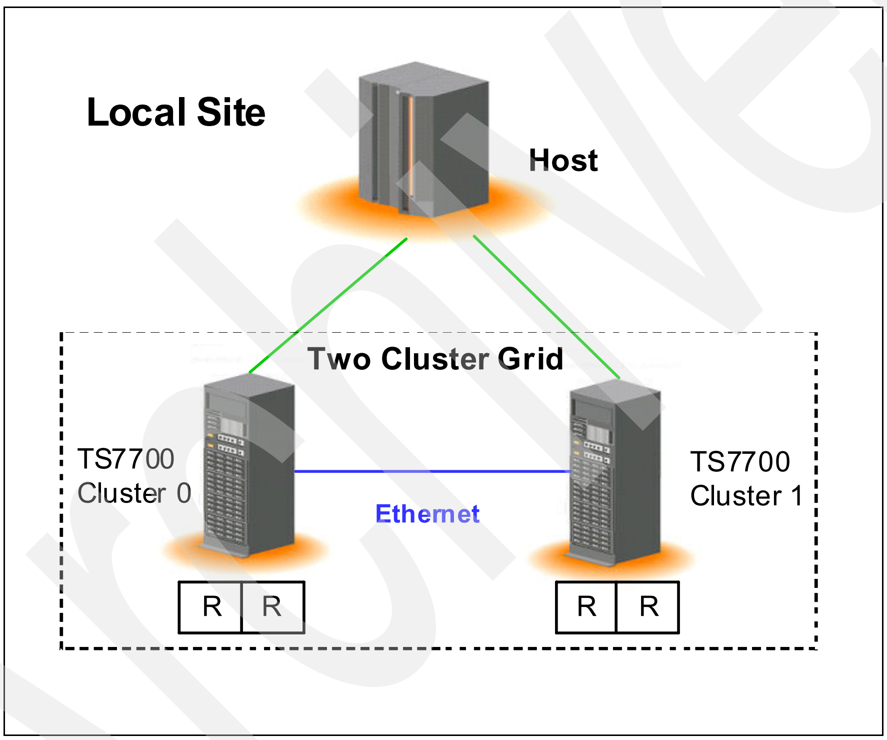

Figure 2-32 Availability configuration

Configuring for disaster recovery and high availability

It is possible to configure a TS7700 Two-Cluster Grid configuration to provide for both disaster recovery and high availability solutions.

The assumption is that the two TS7700 clusters will reside in different locations, separated by a distance dictated by your company's requirements for disaster recovery. In addition to the considerations for configuring for disaster recovery, you will need to plan for the following:

- Access to the FICON channels on the TS7700 cluster located at the disaster recovery site from your local site's hosts. This can involve connections using DWDM or channel extension equipment, depending on the distance separating the two sites. If the local TS7700 cluster becomes unavailable, you would use this remote access to continue your operations using the remote TS7700 cluster.

- Because the virtual devices on the remote TS7700 cluster are connected to the host through a channel extension, there can be a difference in read or write performance as compared to the virtual devices on the local TS7700 cluster. If performance differences are a concern, you should consider only using the virtual device addresses in the remote TS7700 cluster when the local TS7700 is unavailable. If that is an important consideration, in addition to the ownership takeover procedure, you would need to provide for operator procedures to vary online and offline the virtual devices in the remote TS7700.
- You might want to have separate copy consistency policies for your disaster recovery data versus your data that requires high availability.

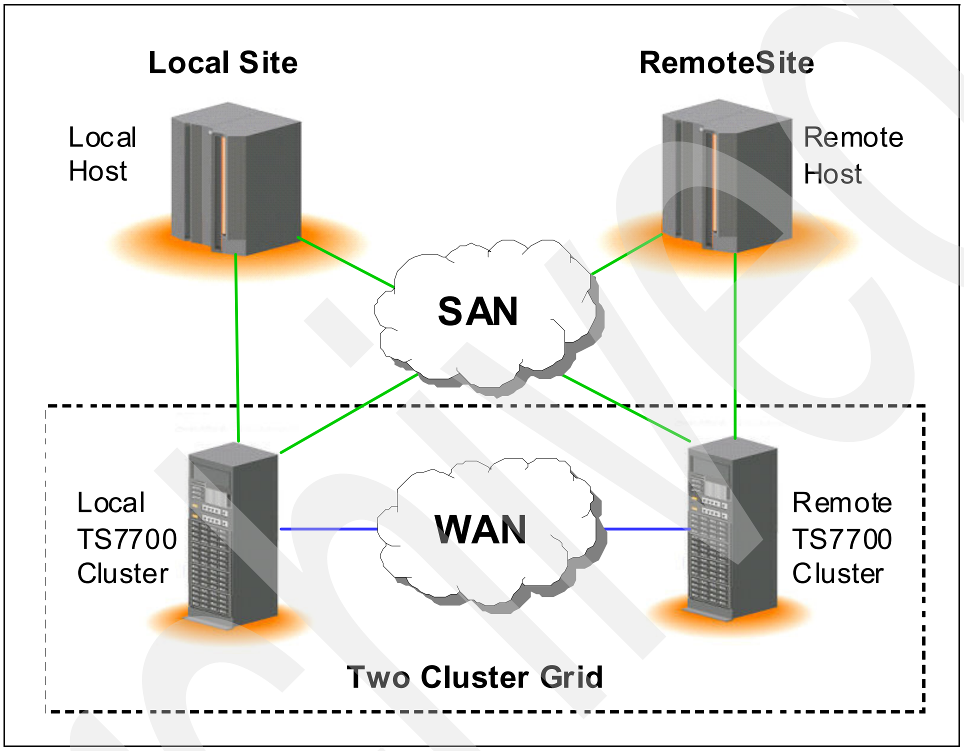

Figure 2-33 Availability and disaster recovery configuration

2.4.4 High availability and disaster recovery: Three-Cluster Grid

With a Three-Cluster Grid, you can configure the grid for disaster recovery and high availability as well as dual production sites that share a common disaster recovery site. In the following sections, we describe configuration considerations for Three-Cluster Grids. The scenarios that we present are typical configurations. However, other configurations are possible and might be better suited for your environment.

The planning considerations for a Two-Cluster Grid also apply to a Three-Cluster Grid.

High availability and disaster recovery

Figure 2-34 on page 68 illustrates a combined high availability and disaster recovery solution for a Three-Cluster Grid. In this example cluster 0 and 1 are the high availability clusters and are local to each other, typically less than 50 kilometers apart. Cluster 2 is at a remote site at a distance away from the production site. The virtual devices in clusters 0 and 1 are online to the host and the virtual devices in cluster 2 are offline to the host. The host accesses the 512 virtual devices provided by cluster 0 and 1. Host data written to cluster 0 is copied to cluster

1at RUN time. Host data written to cluster 1 is written to cluster 0 at RUN time. Host data written to cluster 0 or 1 is copied to cluster 2 on a deferred basis.

The copy consistency points at the disaster recovery site (NNR) are set to only create a copy of host data at cluster 2. Copies of data are not made to clusters 0 and 1. This allows for disaster recovery testing at cluster 2 without touching the production site clusters.

Figure 2-34 also shows an optional host connection that can be established to remote cluster 2 using DWDM or channel extenders. With this configuration, you need to define an additional 256 virtual devices at the host for a total of 768 devices.

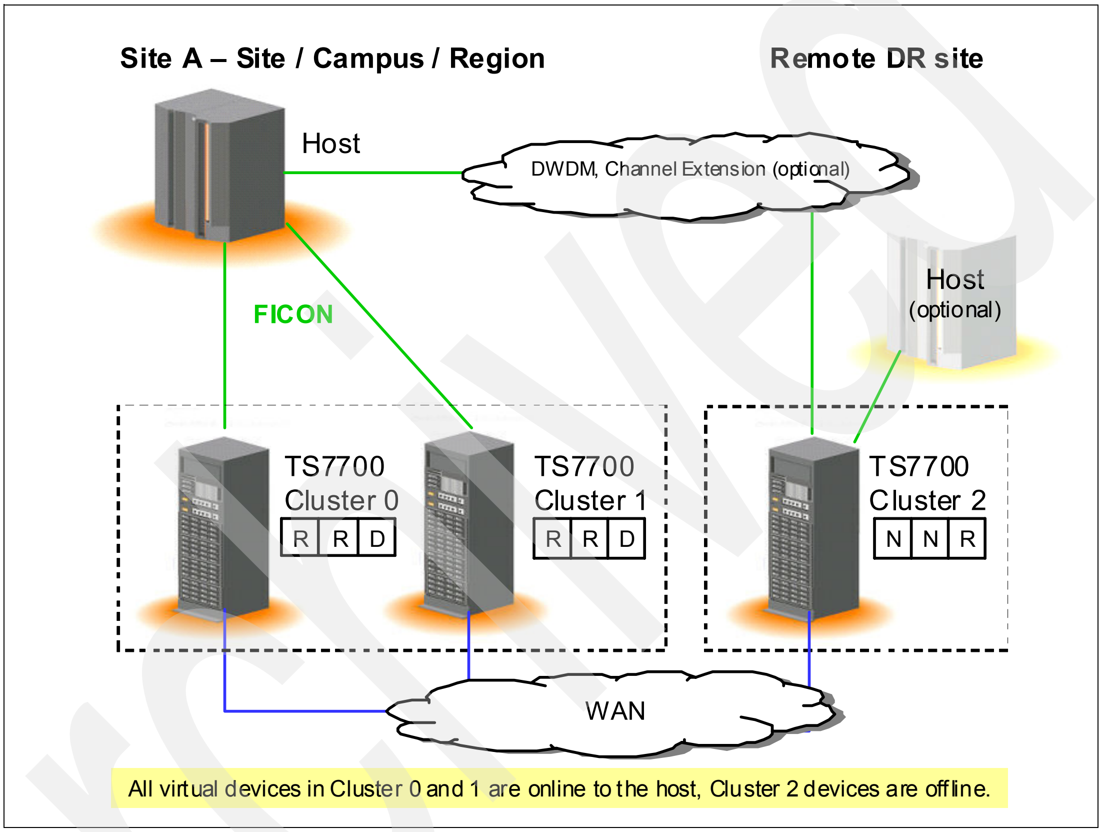

Figure 2-34 High availability and disaster recovery configuration

Dual production site and disaster recovery

Figure 2-35 illustrates dual production sites that are sharing a disaster recovery site in a Three-Cluster Grid. In this example cluster 0 and 1 are separate production systems that can be local to each other or can be a large distance from each other. The disaster recovery cluster, cluster 2, is at a remote site at a distance away from the production sites. The virtual devices in cluster 0 are online to host A and the virtual devices in cluster 1 are online to host B. The virtual devices in cluster 2 are offline to both hosts. Host A and host B access their own set of 256 virtual devices provided by their respective clusters. Host data written to cluster 0 is not copied to cluster 1. Host data written to cluster 1 is not written to cluster 0. Host data written to cluster 0 or 1 is copied to cluster 2 on a deferred basis.

The copy consistency points at the disaster recovery site (NNR) are set to only create a copy of host data at cluster 2. Copies of data are not made to clusters 0 and 1. This allows for disaster recovery testing at cluster 2 without touching the production site clusters.

Figure 2-35 also shows an optional host connection that can be established to remote cluster 2 using DWDM or channel extenders.

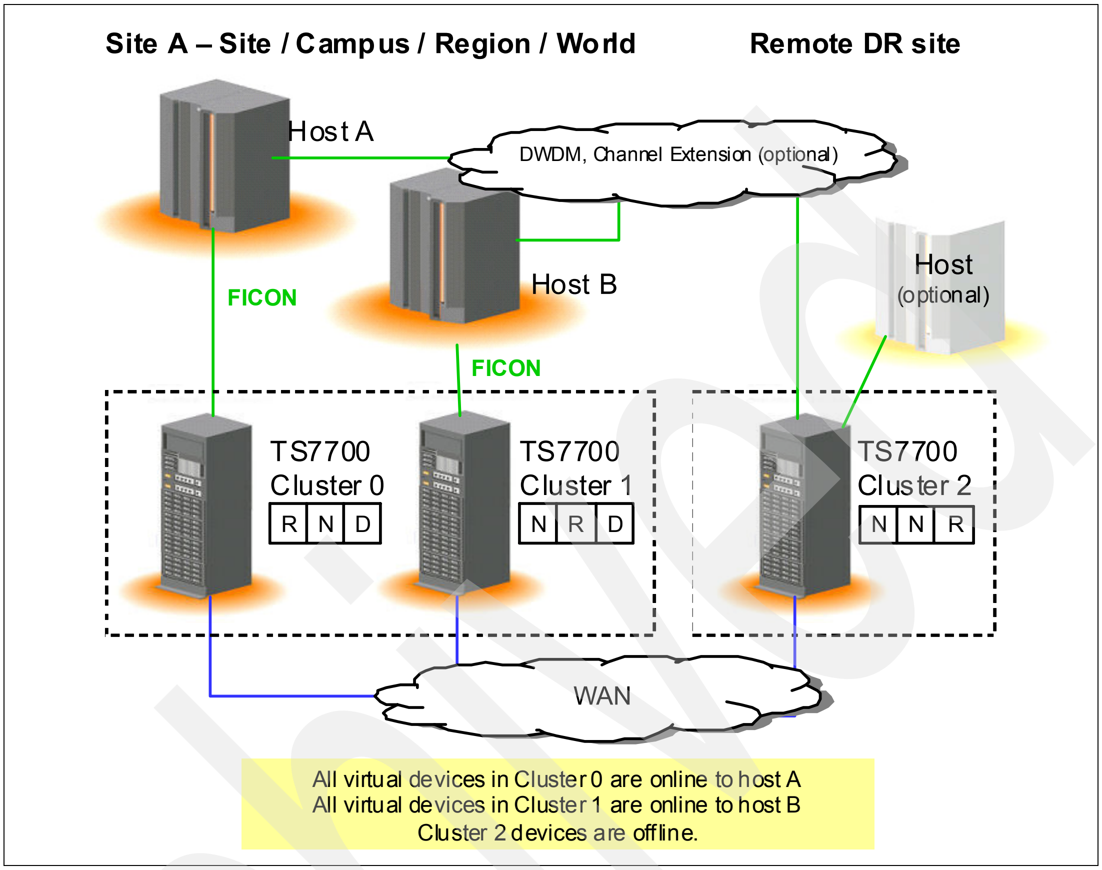

Figure 2-35 Dual production site with disaster recovery

2.4.5 Removal of a Cluster from a Grid and cluster cleanup

It is possible to remove one cluster from a Grid if you have a Multi Cluster Grid configuration and one of the TS7700 clusters is no longer required to participate in the grid.

Removal of a Cluster from a Grid is a feature that delivers instructions for a one-time process to remove or unjoin a cluster from a Grid configuration. This function can be used to remove one cluster from a two or three cluster domain. Subsequent invocations can be executed to remove two clusters from a Three-Cluster Grid.

Removal of a Cluster from a Grid requires that all data in the cluster that is going to be removed must be copied, removed or expired before the cluster removal. After the removal, the removed cluster is disabled and cannot be returned to active until a Cluster Cleanup is performed on this cluster.

Cluster cleanup cleans the database, logically deletes volumes from the Tape Volume Cache, and removes the configuration data for host connections from a TS7700 cluster. However, it does not physically delete user data from the Tape Volume Cache and from the back-end stacked tape cartridges. Cluster cleanup can also be performed on a standalone cluster that never has been a member of a grid.

The cluster cleanup feature is required in order to reuse the removed cluster. This feature is a one-use feature and returns the removed cluster to a state similar to one received from manufacturing.

Removal of a Cluster from a Grid and Cluster Cleanup can be used if you are planning a data center consolidation or if you need to ungrid a TS7700 for a specific requirement or need.

Data center consolidation scenario

We use this scenario to consolidate data centers by collecting the data from remote data centers and using the TS7700 Grid to move the data to the centralized data center. In this scenario, Cluster 2 is removed, but any cluster within the grid cannot be removed.

We assume that there are two clusters at the primary data center for high availability. The third cluster is located at a remote data center, from which you want to move data. Figure 2-36 shows this initial status.

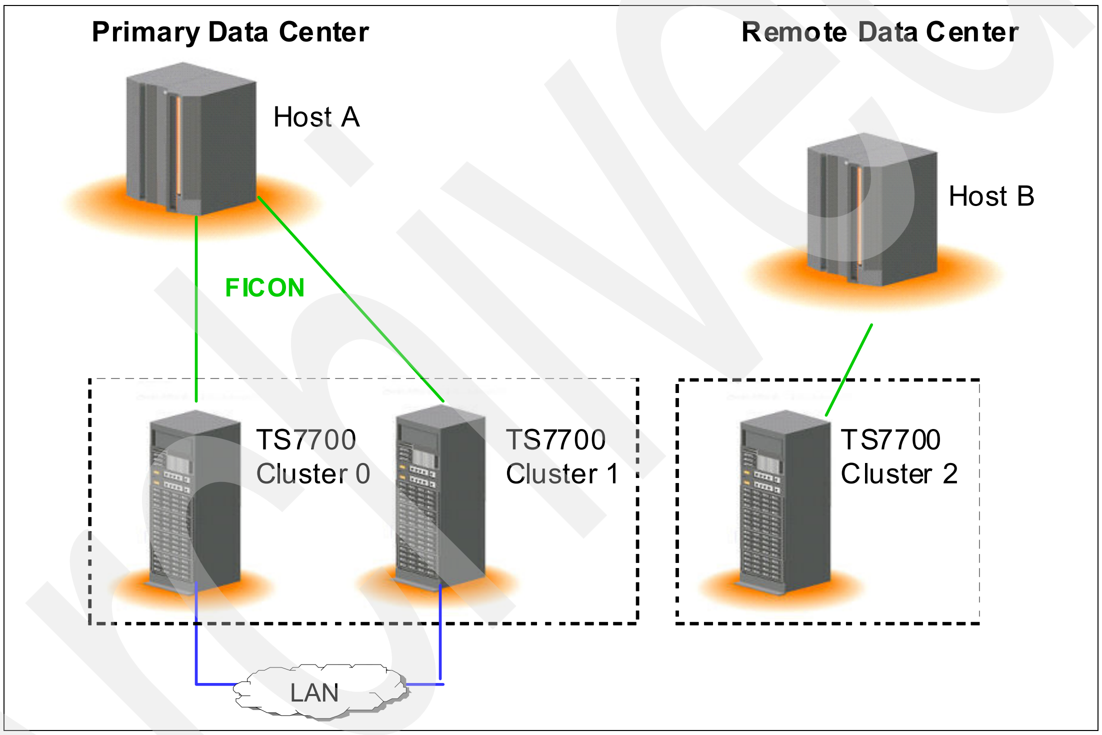

Figure 2-36 Data center consolidation initial status

The third cluster is joined with the two existing clusters, and you use the remote hosts to copy data from their existing infrastructure into the third cluster. The data is then moved to their primary data center using Grid replication. Figure 2-37shows this phase.

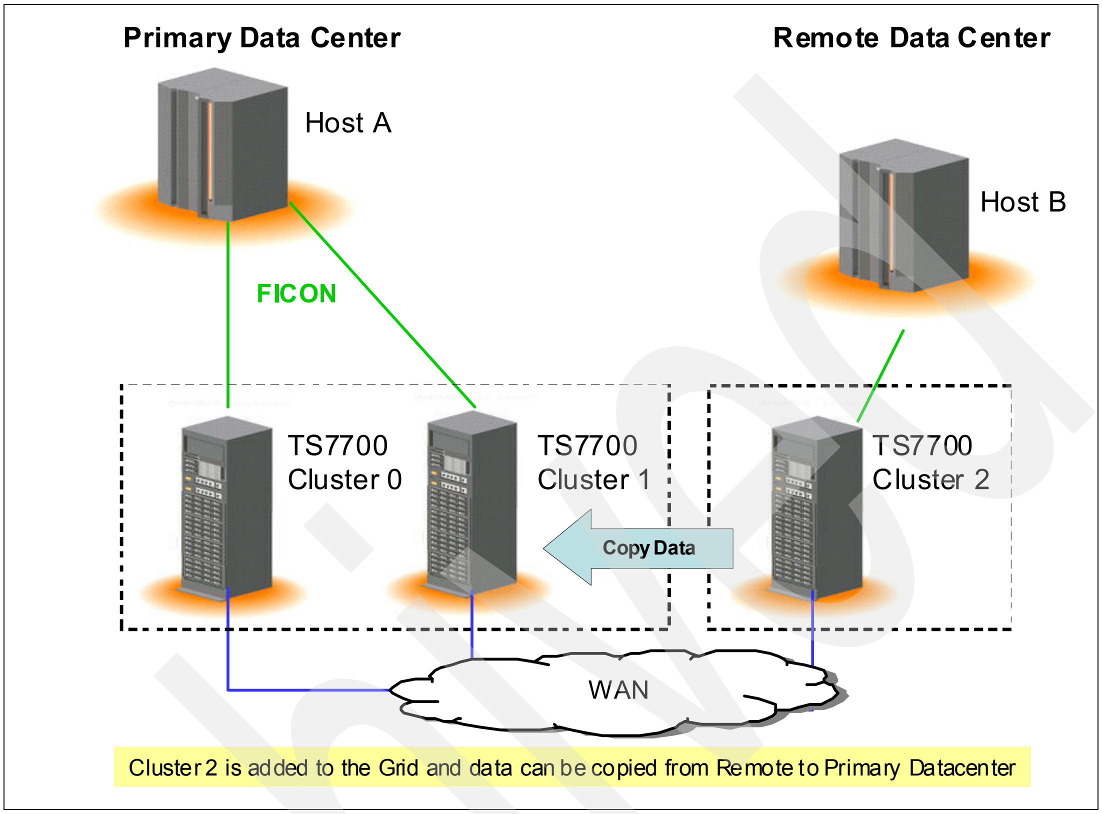

Figure 2-37 Data center Consolidation join to Grid

After all of your data is copied to the primary data center's TS7700s, you can remove the third cluster from the remote data center and clean up the data from it. You can now relocate this TS7700 and repeat the process. Figure 2-38 shows this final status.

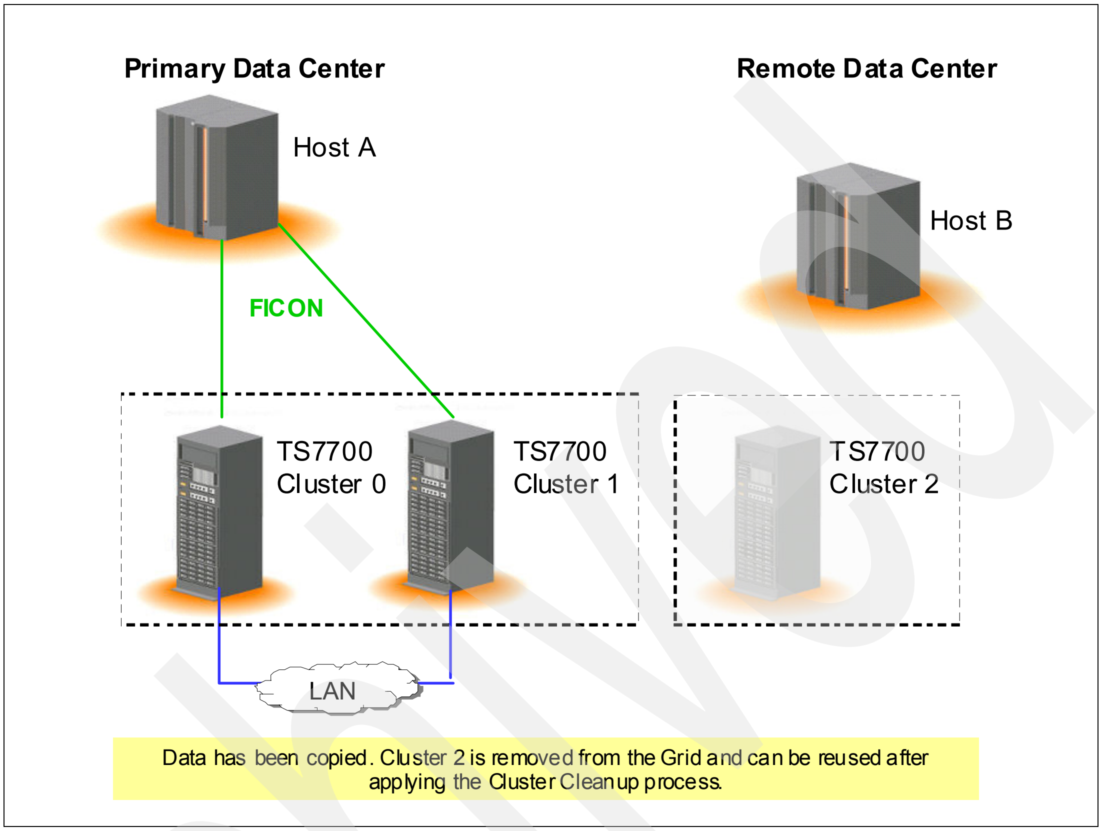

Figure 2-38 Data Center Consolidation Final

TS7700 Reuse

For a TS7700 Reuse scenario, we assume that you have a Multi Cluster Grid configuration with multiple sites and that one TS7700 Cluster is no longer required at one site. You can remove this TS7700 Cluster (after copying, removing, or expiring all the required data) and use this resource in another role. Before you can use the cluster, you need to remove it from the grid domain and perform cleanup.

It is important to mention that the cleanup process does not include cluster reinstallation. If you are planning to reuse a TS7700 after you remove it from a Grid, you must consider a service contract for a subsequent reinstallation of the cluster.

It is the customer's responsibility to determine how to handle the volumes that have a only copy consistency point at the cluster that is being removed (that is, you need to eject them, move them to the scratch category, or activate a management class change on a mount or demount to get a copy on another cluster). This needs to be done prior to starting the remove process. A new Bulk Volume Information Retrieval (BVIR) option (Copy Audit) is provided to generate the list of inconsistent volumes. See 8.6, "Bulk Volume Information Retrieval (BVIR)" on page 442 for details.

The removal of the cluster from the grid is concurrent with customer operations on the remaining clusters, except for certain steps where inserts, ejects, and exports are inhibited. We recommend that you complete the removal of a cluster from the grid with minimal or very low activity, so that, for example, the removal occurs while few or no jobs are running to the remaining clusters.

No data, on cache or on tapes, on the removed cluster will be available after the cluster has been removed with the completion of FC 4016. The cluster cannot be rejoined with the existing data.

No secure erase or low level format will be done on the tapes or the cache as part of Cluster Removal or Cluster Cleanup.

2.5 Architectural capabilities

When you compare a Peer-to-Peer VTS with a TS7700 Multi Cluster Grid, the differences are quite obvious. With the TS7700 you do not need any virtual tape controllers and the peers in a grid are connected through Ethernet rather than through ESCON or FICON as before. When you look at a standalone VTS and a Single Cluster Grid the differences are not as apparent at first glance. They look very much the same from a functional perspective, though under the cover the architectural design of a TS7700 is completely different from that of a VTS. In this section we discuss the architectural design of the TS770 and its capabilities. To help you understand the changes in architecture we start with a short description of the architecture of a VTS.

> **Note:** Here we discuss the potential capabilities a modular design offers for the future. The current version of the TS7700 Virtualization Engine is limited to one gNode per cluster and to three clusters in a Multi Cluster Grid configuration.

Monolithic design of a VTS

A VTS performs all functionality within a single pSeries® server. The VTS also serves as the RAID disk controller. The RAID system is tightly integrated into the system. When you want to expand the capabilities and functionality, these expansions must be fit within the capabilities of the pSeries server. There are a number of drawbacks to that architectural concept:

- There are a limited number of expansion slots available.
- The amount of RAM supported by the platform is limited.
- The CPU capabilities are limited.
- Increasing the VTS capabilities requires more powerful pSeries servers.

The *law of diminishing returns* comes into play for upgrades. After you reach a certain point, a small incremental increase in capabilities will require a large investment in a more powerful server. As code components are tightly integrated with one another, implementing new functions affects large amounts of code. Finally, the system must always be upgraded or extended as a whole, because of the tight integration of its components.

All these concerns have been addressed in the architectural design of the TS7700, which we now describe.

Modular design of the TS7700 Virtualization Engine

The modular design of the TS7700 separates the functionality of the system into smaller components with well defined, open interfaces. It provides a platform where the smaller components can be tailored from a small solution into an extremely large one, so it can eventually give you the capability to grow the solution with your needs.

In the future, you might be able to plug in different components of the same functionality to provide a solution for specific environments.

vNode

The *vNode* refers to a code stack that performs all of the actions needed to present a library image and drive images to a host. The vNode code was designed to run alongside the hNode code in the same controller, or in a separate controller. In the current version of the TS7700 both the vNode and the hNode run on the same System p server. As the architectural design of the TS7700 allows for vNode and hNode to run on separate hardware for horizontal scalability, it uses standardized interfaces to talk with outside components (TCP/IP).

hNode

The *hNode* refers to a code stack that performs all of the actions needed to coordinate the contents of the disk cache with the data on backend tape. It also includes the logic for managing changes and replication of the data across different sites. The hNode code was designed to run alongside the vNode code in the same controller, or in a separate controller. An hNode uses standardized interfaces (TCP/IP) to talk with outside components.

gNode

A *gNode* can be considered a vNode and an hNode sharing the same physical controller. A gNode can be considered the equivalent of a VTS, such that a single controller has both the vNode and hNode capabilities.

Standards-based interface

The standards-based interface between the elements is FCP for data movement and TCP/IP for command and control (including the management interface). The current functions of the VTCs in a PTP VTS have been integrated into the hNode in something we call the *grid layer*. The grid layer interconnects all of the nodes whether they are in the same cluster or in another cluster.

A comparison of the modular design of a TS7700 and the monolithic design of a VTS is shown in Figure 2-39.

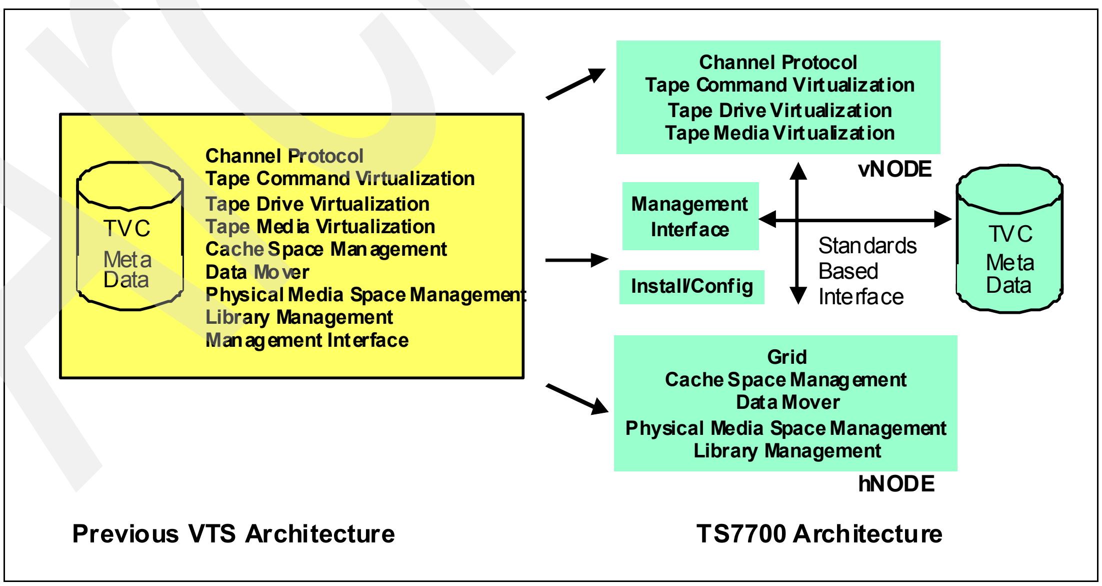

Figure 2-39 Monolithic design versus modular design

With the modular design of the TS7700 it will be possible in the future to add gNodes, or to add vNodes and hNodes separately to an existing cluster configuration. You might want to add vNodes to increase the number of virtual device addresses, or to add vNodes with some special functionality to perform additional tasks. More hNodes would give you better performance and increase availability. Instead of replacing one monolithic system with another, bigger one, this would allow you to grow just according to your needs by horizontally scaling the components of the system.

Serviceability would also profit from having more than one node in a cluster. A nondisruptive code update for a Single Cluster Grid would require two gNodes (or a pair of vNodes and a pair of hNodes). Redundant nodes would eliminate or minimize the risk of planned and unplanned outages.

As the cache is a modular component of the system, you can upgrade the cache with minimal impact to the whole system.

Multi Cluster Grid

When we talked about a Multi Cluster Grid so far, we always referred to a Two- and Three-Cluster Grid, because these are the only grid configurations we support today. From an architectural perspective there is no need to limit the number to three clusters in a grid.

In a Multi Cluster Grid with more than two clusters, Copy Consistency Point policies, especially in combination with Override settings, are even more important. You will have to plan carefully where and when you want a copy of your data to reside.

3

Pre-installation planning and sizing

This chapter provides the information that is required for planning the installation and implementation of the IBM System Storage TS7700 Virtualization Engine.

We cover the following topics:

- Hardware configurations
- Hardware installation and infrastructure planning
- Remote installations and switch support
- Planning for software implementation
- Planning for logical and physical volumes
- Tape analysis and sizing the TS7700
- Education and training

To assist you with installation, we include pre-installation checklists in Appendix D, "TS3500 checklists" on page 557 and in Appendix E, "3494 Checklists" on page 575.

> **Note:** For this chapter, we use the term *Tape Library* when the subject applies to both the 3494 Tape Library and the TS3500/3953 Tape Library. We use the specific library name when a subject is specific to a particular tape library.
>
> For this chapter, we use the term *Library Manager* when the subject applies to both the 3494 and 3953 Library Manager. We use the specific Library Manager name when the subject is specific to a particular library manager.

3.1 Hardware configurations

An IBM System Storage TS7700 Virtualization Engine is composed of:

- pSeries Server as Node (3957 Model V06)
- Cache Controller (3956 Model CC6)
- Cache Drawer (3956 Model CX6)
- Frame (3952-F05)
- I/O Expansion drawers
- Redundant network routers

All components are installed in a unique IBM 3952 frame. The Virtualization Engine is connected to the host through FICON channels and to the Library Manager through Ethernet links in a Single Cluster Grid configuration. In a Multi Cluster Grid configuration, we have two or three TS7700 Virtualization Engines interconnected with two 1 Gb Ethernet links per cluster.

We have summarized the hardware components and configuration requirements for a Single Cluster Grid and for a Multi Cluster Grid configuration below.

3.1.1 TS7700 configuration requirements

In this section we describe the specific features of the TS7700 in a standalone and grid configuration.

Single Cluster TS7700

For a standalone TS7700, the feature codes listed below per machine type and model are mandatory:

- One *3952 Tape Frame Model F05* with the following required features:
  - TS7700 Base Frame (FC7312)
  - Plant Install 3957 Model V06 (FC5628)
  - Plant Install 3956 Model CC6 (FC5638)
  - Zero, one or three of plant install 3956 Model CX6 (FC5648)

    Note: With Release 1.4a, the following field cache upgrade options are available:

    - **Zero to One TS7740 Cache Drawers** (3956-CX6)
      by ordering one FC 5649 to an existing 3956-CC6 cache controller.
    - **One to Three TS7740 Cache Drawers** (3956-CX6)
      by ordering two FC 5649 to an existing two-drawer TS7740 Cache subsystem (one 3956-CC6 and one 3956-CX6).

  - Integrated Control Path (FC5759)

    Provides two routers, two switches, and cables to create the control path connection between the TS7700 and the Library Manager.
  - Dual AC Power (FC1903)

    Provides additional Power Distribution Unit and cables to allow connection to independent branch power circuits.
  - One power cord appropriate for the country of installation must be selected from feature FC9954 through FC995.8

- One *TS7740 Server (3957 Model V06)* with the following required features:
  - One to six of 1 TB Cache Enablement features (FC5267)

    > **Note:** With release 1.4a, you can order from one to six 1 TB cache enablement features. This feature can be plant or field installed.
    >
    > If you are performing a field installation of this feature, you enter the feature license from the TS7700 Management Interface (MI). See 4.5.2, "Cache Enablement and Performance Increment license key entry" on page 175 for details about how to enter the feature license.

  - One to six of 100 MB/s Increment (FC5268)

    > **Note:** If you are performing a field installation of this feature, you enter the feature license from the TS7700 Management Interface (MI). See 4.5.2, "Cache Enablement and Performance Increment license key entry" on page 175 for details about how to enter the feature license.
    >
    > When all six Cache Enablement features are installed, the performance of the TS7700 is no longer restricted. The performance can exceed 600 MB/s.

  - Attach 3592 Tape Drives (FC5240)
  - Mainframe Attach (FC9000)
  - Attach to 3953 Library Manager (FC9217) or Attach to 3494 Library Manager (FC9218). One of these two features must be ordered.
  - Plant Install V06 in 3952 F05 (FC9350)
  - Either two or four Host to V06 FICON cables (FC0201 through FC0206), or one No Factory cable (FC9700)
  - Two or four of either FICON Shortwave Attachment (FC3441) or FICON Longwave Attachment (FC3442), or FICON 10 km Longwave Attachment (FC3443)
  - Either TS3000 System Console (FC2720 and FC2719), Console Expansion (FC2714), or Console Attachment (FC2715)

  > **Notes:**
  >
  > 1. 3957 Model V06 features FC2714 and FC2715 cannot be attached to a TotalStorage Master Console for service (FC2713).
  > 2. The TS3000 System Console (FC2721) can be installed in the 3953 Model F05 Tape Frame, connecting to the TS7740 Server using Console Expansion (FC2714) or Console Attachment (FC2715).
  > 3. When features FC2714 or FC2715 are installed on the 3957 Model V06, the Console Upgrade (FC2719) is required on the machine type model where feature FC2718, FC2720, or FC2721 is installed.

- One *TS7700 Cache Controller (3956 Model CC6)* with the following required features:
  - 1.7 TB Fibre Storage (FC7120)
  - Plant Install CC6 in 3952 F05 (FC9352)
  - Attach to 3957 V06 (FC9230)
  - Four of Intraframe Fibre Cable (FC6003)

- Zero, one or three *TS7700 Cache Drawers (3956 Model CX6)* with the following required features:
  - 1.7 TB Fibre Storage (FC7120)
  - Plant Install CX6 in 3952 F05 (FC9354)
  - Four of Intraframe Fibre Cable (FC6000)

> **Note:** With zero cache drawers (1.5 TB of physical cache on 3956-CC6 cache controller only), you can order one Cache Enablement feature (FC5267) to use 1 TB of the cache or two FC5267 to use all the available cache of 1.5 TB.
>
> With one TS7700 Cache Drawer (3 TB of physical cache on 3956-CC6 and 3956-CX6), you can order one, two, or three Cache Enablement features (FC5267).
>
> With three cache drawers (6 TB of physical cache on 3956-CC6, three 3956-CX6), you can order from one to six Cache Enablement features.
>
> Refer to Table 3-1 on page 89 for more information.

- One *3953 Tape Frame Model F05* with the following required features:
  - TS7700 Attach (FC9013)
  - Two 4 Gb Fibre Channel switches (FC3488 or FC4897)
- For the TS3500 Tape Library, from four to sixteen 3592- J1A or 3592-E05 Tape Drives must be installed in the TS3500 and connected to the TS7700 through the 4 Gbps FC Switches in the 3953-F05 Tape Frame for TS3500 attach.
- For the 3494 Tape Library, from four to twelve 3592-J1A or 3592-E05 Tape Drives must be installed in a single frame of the 3494 and connected to the TS7700 through the 4 Gbps FC Switches in the 3494 drive frame containing the associated tape drives.

Multi Cluster TS7700 configuration

For a Multi Cluster TS7700, the following feature codes are mandatory:

- TS7740 Server (3957 Model V06)
  - Two 1 Gbps Grid Copper Connections (FC1030) or two 1 Gbps Grid Optical SW Connections (FC1031) provide 1000 BaseT adapters for Grid communications between TS7740 Server. 1 Gbps Grid Copper or Optical SW Connections must be installed on each TS7740 Server. Each TS7740 Server in a Multi Cluster Grid can have copper or optical connections. The TS7740 Servers in a Multi Cluster Grid do not have to have the same type of connection. The customer must provide Cat5e or Cat6 cables to attach the 1 Gbps Grid Copper or Optical Connection adapters to the communication grid. Grid Enablement (FC4015) is required when FC1030 or FC1031 is installed. Valid quantities of FC1030 or FC1031 are zero and two. These distances are supported:

Up to 100 m	with copper connections
Up to 260 m	with 62.5 micron multimode optical fiber cable
Up to 500 m	with 50.0 micron multimode optical fiber cable

  - Grid Enablement (FC4015) allows this TS7740 Server to communicate to other TS7740 Servers through the grid. Grid Enablement must be installed on each TS7740 Server that participates in the communication grid. 1 Gbps Grid Copper Connection (FC1030) or 1 Gbps Optical SW connection (FC1031) is required when FC4015 is installed.

Note: More than one situation has occurred where a TS7700 was ordered with feature code 4015 (Grid enablement) and the intention is to initially use the TS7700 in a standalone test environment. This TS7700 should be “install complete” (service code 20 CIA 1) before being used in this test configuration. This use of a grid enabled TS7700 in a standalone environment is an unintended use and requires that the TS7700 be “Reset to Factory Settings,” also sometimes referred to as “manufacturing cleanup,” before using the TS7700 in a grid configuration.

The “Reset to Factory Settings” function is not currently part of the TS7700 product. This function might be delivered in a future release. The activity to perform a reset (to wipe the data or to wipe out both the data and the configuration) is not a valid warranty item nor expense. This reset must be done by BDA, PFE, or development and is billable to the IBM Account Team. An ICA must be created to cover the cost of BDA/PFE/development involvement.

With release 1.3 code installed, a valid alternative is to order the TS7700 without feature code 4015. When the time comes to use the TS7700 in a grid configuration, it must be connected to an IP network and feature code 4015 must be ordered as an MES, plus a new TS7700 (with feature code 4015) must be ordered as its grid partner. The MES feature code's cost would be used to offset the additional expense for the SSR to upgrade the unit. Release 1.4 code includes support to merge an existing standalone cluster with a Two-Cluster Grid to form a three-way grid.

- Valid total quantities of FICON Shortwave Attachment (FC3441), FICON Longwave Attachment (FC3442), and FICON 10 km Longwave Attachment (FC3443) are two or four per cluster.
- The number of Cache Enablement features (FC5267) might be different on the clusters in a grid configuration.

► Cables

A TS7740 Server with the FICON Attachment features (FC3441, FC3442 or FC3443) can attach to FICON channels of host System z9®, zSeries®, or S/390® systems using Fibre Cable features ordered on the TS7740 Server. A maximum of four Fibre Cables, each 31 meters in length, can be ordered with the following Host System Attachment Cable features:

- 4 Gbps FICON Long-Wavelength Attachment feature (FC3442 and FC3443): The FICON long-wavelength adapter shipped with feature number 3442 (4 Gbps FICON Long-Wavelength Attachment) or feature number 3443 (4 Gbps FICON 10 km Long-Wavelength Attachment) has an LC Duplex connector, and can connect to FICON long-wavelength channels of System z9, zSeries, or S/390 servers using a 9 micron single-mode fiber cable. If host attachment cables are desired they can be specified with one of the following feature numbers:
  - FC0201: 9 Micron LC/LC 31 meter Fibre Cable
  - FC0202: 9 Micron LC/SC 31 meter Fibre Cable

– 4 Gbps FICON Short-Wavelength Attachment feature (FC3441): The FICON short-wavelength adapter shipped with feature number 3441 has an LC Duplex connector, and can connect to FICON short wave length channels of System z9, zSeries, or S/390 servers utilizing a 50 micron or 62.5 micron multimode Fibre Cable. If host attachment cables are desired they can be specified with one of the following feature numbers:

  - FC0203: 50 Micron LC/LC 31 meter Fibre Cable
  - FC0204: 50 Micron LC/SC 31 meter Fibre Cable
  - FC0205: 62.5 Micron LC/LC 31 meter Fibre Cable
  - FC0206: 62.5 Micron LC/SC 31 meter Fibre Cable

Additional cables, fabric components, and cabling solutions: Conversion cables from SC Duplex to LFCC Duplex are available as features on the System z servers for customers currently using cables with SC Duplex connectors that now require attachment to fiber components with LC Duplex connections. Refer to the IBM 2064 Sales Manual for more information about these conversion kits. In many situations, the configuration and facility specifications might require different cable lengths or the installation of cables. Product Support Services offered by IBM Global Services can provide additional fiber optic components and fiber optic cabling solutions.

Grid WAN requirements

The following requirements apply to the configuration of the WAN in a Multi Cluster Grid:

- Every node in the system requires two network connections to a customer WAN for site-site operations.
- Cross-site network infrastructure is 1 Gb Copper Ethernet (RJ45) or shortwave optical fiber connection as the first connection to the customer's WAN. The customer’s WAN can be a combination of copper and fiber connections.
- The grid Network must be configured such that the nodes of each site can address the nodes of all other sites.

3.1.2 TS7700 configuration options

The TS7700 Virtualization Engine is currently available with the following configuration options:

Host attachment Two or four 4 Gb FICON channels

Cache sizes 1 to 6 TB of cache (one to six FC5267: 1TB Cache Enablement) There are three physical sizes, 1.5 TB, 3 TB, and 6 TB, corresponding to the three allowed configurations. There are then six cache enablement configurations based on how many features FC5267 are installed.

> **Note:** You can order from one to six 1 TB cache enablement features. This feature can be plant or field installed.

Bandwidth 100 MB/s to 600 MB/s (One to six FC5268: 100 MB/s Increments)

Virtual drives 256 virtual drives for a Single Cluster Grid configuration
512 virtual drives for a Two-Cluster Grid configuration (256 per site).

A Three-Cluster Grid configuration can have up to 768 virtual devices (256 per site). However, a typical configuration will have 512 online virtual devices at the production sites and 256 offline devices at the disaster recovery cluster.

Physical drives Minimum four, maximum of 16 for the TS3500 Tape Library and a maximum of 12 for 3494 Tape Library.

For Tape Encryption, you must also order FC9900 on the TS77040

Single Cluster configuration

The TS7700 can be installed as a single site configuration, a so-called cluster, where a single TS7700 is connected to a single 3494 or TS3500 Tape Library. The single site configuration is comparable to the standalone Virtual Tape Server.

Every TS7700 configuration has multi-site capabilities and logic built into it, even a single site configuration. Therefore, even a single site configuration requires the definition of a composite and a distributed LIBRARY-ID. See Chapter 4, “Hardware implementation” on page 131 for more details on the definition.

Multi Cluster Grid configuration

For Multi Cluster Grid configurations, each TS7700 Virtualization Engine must be installed in physically separate Tape Libraries. These clusters can be connected to one another using a TCP/IP WAN, to form a grid. The interconnections between the sites use the standard TCP/IP networking infrastructure.

Any logical volume can be mounted and accessed from any virtual device in the subsystem. A grid configuration provides capabilities comparable to a Peer-to-Peer Virtual Tape Server configuration.

A grid TS7700 can be shared by multiple System z systems in the same way that a single site TS7700 or a physical tape library can be shared.

FICON channels

The TS7700 FICON attachment supports 4 Gbps link speed attached to a System z server, or an appropriate FICON/FC switch, with 4 Gbps FICON features and the appropriate levels of software to utilize the 4 Gbps capability. Extended distance can be provided with channel extenders certified for use by providers. If the customer infrastructure does not support 4 Gbps link speeds, the adapters automatically match the link’s slower capability.

Two or four of the following FICON attachments are supported and can be ordered with the following feature codes:

- FICON Shortwave Attachment (FC3441)
- FICON Longwave Attachment (FC3442)
- FICON 10 km Longwave Attachment (FC3443)

In case of two FICON adapter cards, they will be installed in slot 6 of each I/O Drawer. Both must be the same type (Longwave or Shortwave).

In case of four FICON adapter cards, they will be installed in slot 3 and 6 of each I/O Drawer. The same pair of slot numbers must have the same type of card (Longwave or Shortwave).

A field upgrade from two to four FICON adapters is available with using the same feature codes. The different FICON adapters can be mixed in a TS7700, but they must be installed in pairs. This means you can have two of one type and two of another.

TS1120 and IBM 3592 Tape Drives

The TS7700 Virtualization Engine operates with 3592-J1A and TS1120 Model E05 Tape Drives. Currently, the TS7700 Virtualization Engine can attach to the following models of tape drives installed in a Tape Library:

- IBM TotalStorage 3592 Tape Drive Model J1A
- IBM System Storage TS1120 Tape Drive

> **Note:** The TS1130 Tape Drive is supported starting with TS7700 Release 1.5. For more information, see *IBM Virtualization Engine TS7740 R1.5 and TS7720: New Virtualization Options for Mainframe Servers,* SG24-7712, which will be available in early 2009.

When both tape drive models are intermixed on the same TS7700 Virtualization Engine, the TS1120 must operate in J1A emulation mode, providing similar capacity and data rate as 3592-J1A Tape Drives. However, the access time performance benefits of the TS1120 drives are realized, even when the TS1120 is emulating a 3592-J1A drive.

If only TS1120 drives are attached to the TS7700, the SSR will set them to *native mode* so that they operate in E05 mode providing significantly higher performance and capacity than in J1A emulation mode. To use the full performance and capacity capabilities of the TS1120 Tape Drive, you should not intermix TS1120 drives with 3592-J1A drives on the same TS7700.

If only encryption capable TS1120 drives are attached to the TS7700, encryption is supported by the TS7700. If any of the drives are not TS1120 drives or are TS1120 drives that are not encryption capable, the TS7700 does not support encryption. After all of the TS1120 drives are encryption-capable and encryption-enabled, FC9900 must be installed on the TS7740 to allow the TS7740 to encrypt data.

In a Multi Cluster Grid, the tape drives in each of the grids do not have to be the same type. One cluster can contain only encryption capable TS1120 drives and can support encryption, while another cluster in the grid might not have all encryption capable TS1120 drives and will not support encryption.

Supported media

The TS1120 and the 3592-J1A use the following rewritable media:

- JA cartridge (MEDIA5) with a native capacity of 300 GB in J1A Emulation mode and 500 GB in E05 native mode
- JJ cartridge (MEDIA7) with a native capacity of 60 GB in J1A Emulation mode and 100 GB in E05 native mode
- JB cartridge (MEDIA9) with a native capacity of 700 GB (TS1120 only)

> **Restrictions:**
>
> - JB Media is only supported on TS1120 drives running in E05 native mode.
> - Write Once Read Many (WORM) cartridges are not supported on TS7700 attached TS1120 or 3592 tape drives.
> - Performance Scaling cannot be used on stacked volumes in a TS7700.

IBM System Storage TS3000 System Console

The IBM System Storage TS3000 System Console (TSSC) is *mandatory* in an IBM Virtualization Engine TS7700 environment. The TS7700 has a function to call the IBM support center with summary information when a problem is detected. In addition, the TS7700 can be configured to automatically execute a health check and send the results to the TSSC for transmission to IBM for review and preventive actions.

The TSSC integrates service monitoring of up to 43 controllers. It centralizes maintenance and service terminals, and enhances remote support capability. It is the only method for electronic call home, problem determination, and repair of the 3956 models CC6 and CX6 Cache.

The TSSC provides a service interface to the TS7700 through:

- Ethernet using Console Launcher
- Serial connection to the TS7700
- GUI Interface for servicing the 3956-CC6/CX6

You can have two types of system consoles:

- Tower server (xSeries® 206m)

  This is the new, standalone TSSC that can be attached to IBM 3494 models Lxx, Bxx, CX1, 3590 Model A60, 3592 Model J70, TS1120 Model C06, TS3500, 3957 Model V06, and 3956 Model CC6.

  The tower server is shown on the left side of Figure 3-1.

- Rack-optimized server

  This is the TSSC that is installed inside an IBM 3953-F05 frame. It can be a Series x 306 or the new Series x 306m.

  The rack-optimized server is shown on the right side of Figure 3-1.

Figure 3-1 Tower server and rack-optimized TSSC

To better understand the feature codes that apply to the IBM TS3000 System Console, see the following feature code explanations and refer to Figure 3-1 for a tower system console and for a rack-optimized internal console.

If FC2718 or FC2721 (rack-optimized console) is already present in an existing 3953-F05 Tape Frame, or if FC2720 is used on other tape products, then the following feature codes must be ordered:

- FC2714 Console expansion: This feature provides an Ethernet hub and product attachment cables for expanding the number of components that can attach to the TSSC that is provided by FC2713, System console for service.

  Or:

- FC2715 Console Attachment: This feature provides a cable for attaching a TS7700 Node to the Ethernet hub provided by FC2718, FC2720, FC2721, or by FC2714, Console expansion. A single TSSC facility can support a maximum of forty instances of this feature.

  And:

- FC2719 Console Upgrade - FC2719, Console upgrade, provides a second Ethernet network interface card for the TSSC to allow redundant attachment to the service network. This feature provides a memory upgrade to 2 GB total RAM and a second Ethernet card for the service console to allow redundant connections into the service network. This feature only applies to consoles shipped with features FC2718, FC2720, and FC2721. It is required on any console used by TS7700 features FC2714, FC2715, or FC2720. This feature is only available on TSSCs shipped before 26 might 2006; TSSCs shipped after this date will include the second Ethernet network interface card.

Note: If you have FC2713 (system console for service) installed, you must order FC2720 against the TS7740, or FC2721 against the 3953 Model F05 Frame.

If no TS3000 System Console is present at your account, or it is a new installation, then you can order:

- FC2720, TS3000 System Console (standalone tower server), which provides the IBM TS3000 System Console, an Ethernet hub, and a cable and connectors to enable remote enhanced service connection of a TS7700 to an IBM-supplied modem. The Ethernet hub provides 14 additional connections for cables supplied with FC2714, Console expansion, or FC2715, Console attachment. You should specify FC2720, TS3000 System Console on the first unit in an installation connected to a master console facility.

  Or:

- FC2721 (rack-optimized), which will be mounted on a 3953 F05 Frame.

Upgrade considerations

Currently, the upgrade options for the TS7700 Virtualization Engine include changing the number of tape drives, the grid connection, the number of virtual tape devices, the amount of cache, and the performance increments.

Number of tape drives

For a TS3500, if less than 16 tape drives are included in the initial TS7700 configuration, you can add tape drives up to a maximum of 16. For a 3494, if less than 12 drives are included in the initial TS7700 configuration, you can add tape drives up to a maximum of 12 drives. An intermix of 3592 models J1A and E05 attached to the same TS7700 Virtualization Engine is supported. See the restrictions listed in “TS1120 and IBM 3592 Tape Drives” on page 84 and “Supported media” on page 84.

Grid connection

In a grid configuration, you can change your grid connection from two 1 Gbps Copper connections (two FC1030s) to two 1 Gbps optical SW fiber connections by ordering feature code 1031. Feature code 1031 provides two shortwave optical adapters and instructions to replace the copper adapters. You can also change your grid connection from two 1 Gbps Optical SW connections (two FC1031s) to two 1 Gbps Copper connections by ordering feature code 1030.

You cannot intermix copper and optical fiber adapters in a TS7700 grid configuration.

> **Note:** Upgrading the grid connection requires the TS7700 to be offline.

Number of virtual device addresses

TS7700 R1.0, the initial release of the TS7700, supported 128 virtual device addresses per cluster. With R1.1 and higher, the microcode level delivered after January 26, 2007, The TS7700 supports 256 virtual drives per cluster.

When R1.1 or higher is installed and activated on a TS7700, the number of virtual device addresses that are configured is increased from 128 to 256 automatically. These steps are required to upgrade to 256 virtual devices per cluster:

1. Performing a Library Teach operation to define 256 devices to the Library Manager
2. Defining the TS7740 for 256 devices
3. Defining the host IODEF to be able to use the 256 devices

The IBM SSR performs Steps 1 and 2 during the installation of the R1.1 or higher microcode. Step 3 enables the use of 256 virtual devices when you redefine your IODF for 256 devices. If you choose to continue with 128 virtual drives after the upgrade to R1.1 or higher, you can activate the remaining 128 virtual drives for a total of 256 virtual drives at any time later. Figure 3-2 on page 88 illustrates the upgrade options.

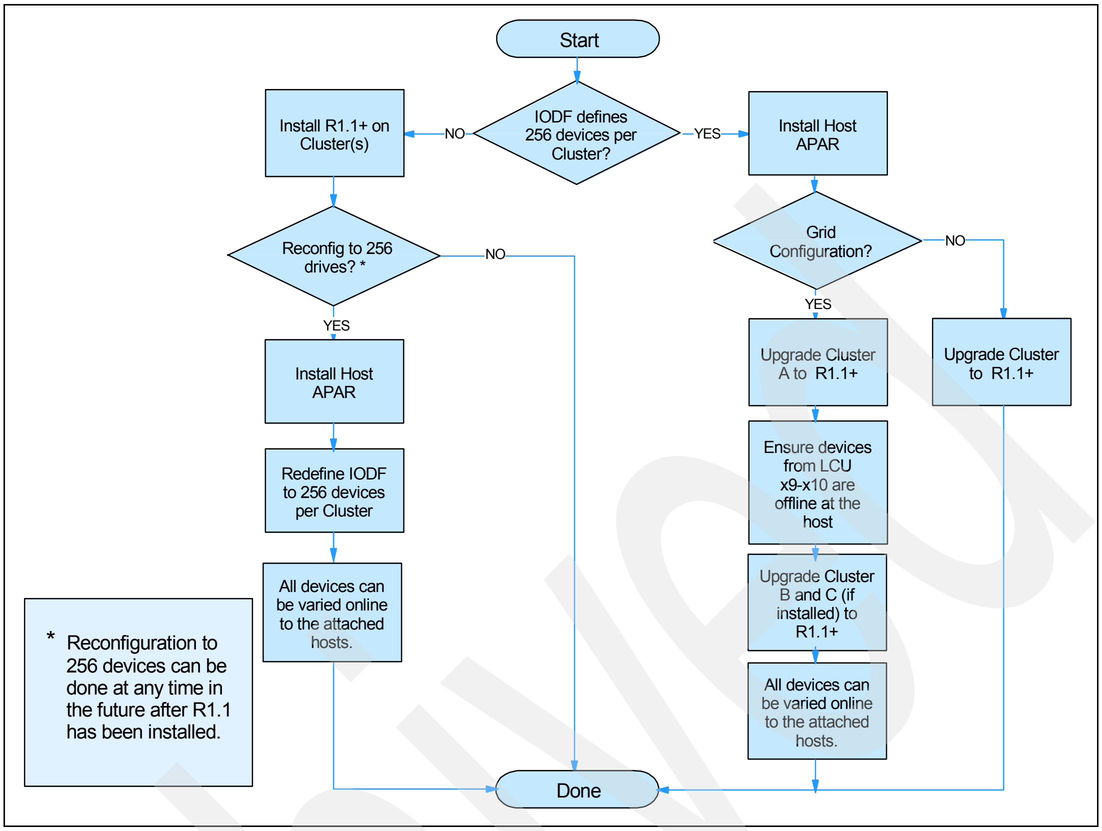

Figure 3-2 Upgrade to R1.1 and 256 virtual devices

APAR OA19061 is required for support of more than 256 virtual devices per Composite Library. See 3.4.2, "Software requirements" on page 104 for more details. You must apply this APAR before you redefine the IODF to 256 devices. If you already have 256 devices defined before starting the upgrade to R1.1 or higher, you must apply the APAR as well before starting the upgrade.

During a grid upgrade, if you already have 256 virtual devices defined, we recommend that you do not use the upper 128 devices (LCU x9 - x10) of cluster A which has been upgraded until cluster B is upgraded as well.

Cache increments

The TS7700 Virtualization Engine supports from 1 TB to 6 TB of tape volume cache in 1 TB increments using feature code 5267. A TS7700 supports physical cache sizes of 1 TB, 3 TB and 6 TB. The 1 TB configuration consists of the 3956-CC6 Cache Controller and no 3956-CX6 cache drawers. The 3 TB configuration consists of one 3956-CC6 cache controller and one 3956-CX6 cache drawer. The 6 TB configuration consists of one 3956-CC6 cache controller and three 3956-CX6 cache drawers.

Note: You can order from one to six 1 TB cache enablement features. This feature can be plant- or field-installed.

Table 3-1 lists the physical cache options, the possible cache enablement features, and whether the cache can be upgraded.

Table 3-1 Physical cache size and cache enablement options

Starting Physical Cache Size	Cache Enablement options FC 5267	Can be Field cache upgrade?	Final Physical Cache Size
1 TB (3956-CC6)	1 TB, 2 TB	Yes - One 3956-CX6 can be added by FC 5649	3 TB
3 TB (3956-CC6, 1x 3956-CX6)	1 TB, 2 TB, 3 TB	Yes - Two 3956-CX6 can be added by 2x FC 5649	6 TB
6 TB (3956-CC6, 3x 3956-CX6)	1 TB, 2 TB, 3 TB, 4 TB, 5 TB, 6 TB	No	

Cache increments are activated using feature licenses. One license key is required for each 1 TB of enablement. Refer to 4.5.2, “Cache Enablement and Performance Increment license key entry” on page 175 for details concerning how the feature licenses are installed. Cache increments can be either plant- or field-installed.

Performance increments

The TS7700 Virtualization Engine supports host data throughput performance from 100 MB/s to 600 MB/s in 100 MB/s increments. The performance increments are ordered using feature code 5268. The TS7700 limits host throughput for the 100 MB/s through 500 MB/s performance increments. The TS7700 turns off throughput limiting when all six performance increments are installed.

Performance increments are activated using feature licenses. One license key is required for each 100 MB/s of enablement. Refer to 4.5.2, “Cache Enablement and Performance Increment license key entry” on page 175 for details concerning how the feature licenses are installed. Performance increments can be either plant- or field-installed.

3.2 Hardware installation and infrastructure planning

The topics in this section provide planning information related to your IBM System Storage TS7700 Virtualization Engine. Topics covered include system requirements, infrastructure requirements, host compatibility, and available feature codes.

3.2.1 System requirements

You must ensure that your facility meets the system requirements for the TS7700 Virtualization Engine when planning for installation. System requirements for installation include such things as requirements for power; cooling; floor leveling, loading, and distribution; clearance; environmental conditions; and acoustics.

IBM 3952 Tape Frame specifications

The 3952 Tape Frame houses the components of the TS7700 Virtualization Engine. Table 3-2 shows the specifications for the 3952 Tape Frame, including its standard components.

Table 3-2 IBM 3952 Tape Frame specifications

	Side in mm (in.)	Front in mm (in.)	Height in mm (in.)
Dimensions	1102 (43.4)	644 (25.4)	1804 (71.0)
Operator/Service clearance	0 (0)	690 (27.2)	660 (26.0)
Weight	Source Power		Heat Output
565.6 kg (1247 lb.)	240 V AC	15 amp	2.5 kVA

Environmental operating requirements

Your facility should meet specified temperature and humidity requirements before installing the TS7700 Virtualization Engine. Table 3-3 shows recommended environmental conditions for the IBM TS7700.

Table 3-3 Environmental specifications

Condition	Air temperature	Altitude	Relative humidity	Wet bulb temperature
Operating	10-32°C (50-80°F)	Up to 1.523m (5.000ft.) amsl	20%-80%	23°C (73°F)
Operating	10-28°C (50-82.4°F)	1.524-2.133m (5001-700ft.) amsl	20%-80%	23°C (73°F)
Power Off		n/a	8%-80%	27°C (80°F)
Storage	1-60°C (33.8°F-140°C)	n/a	5%-80%	29°C (84°F)

Note: For a complete list of system requirements, refer to *Introduction and Planning Guide,* GA32-0567.

Power considerations

Your facility must ensure an available power supply to meet the input voltage requirements for the TS7700 Virtualization Engine.

The 3952 Tape Frame houses the components of the TS7700 Virtualization Engine. The standard 3952 Tape Frame ships with one internal power distribution unit. However, FC1903, Dual AC power, is required to provide two power distribution units to support the availability characteristics of the TS7700 Virtualization Engine.

Table 3-4 lists the input power requirements for the 3952 Tape Frame, which houses the components of the TS7700 Virtualization Engine.

Table 3-4 3952 Tape Frame input power requirements

Power requirement	**Value**
Nominal input voltage	200-240 V AC, single phase
Input frequencies	50 ± 3.0 Hz 60 ± 3.0 Hz

3.2.2 TS7700 and LM microcode levels

To support the TS7700 in a TS3500 Tape Library with the IBM 3953 Library Manager, minimum microcode levels are required on all components. Refer to Table 3-5 for the correct levels for TS7700 R1.0.

Table 3-5 Minimum microcode level for R1.0

Component	**Microcode level**
TS7700 Virtualization Engine	8.0.0.203
IBM 3953-L05 Library Manager	534.01, EC H27461
TS3500 Tape Library	6130
IBM 3952 Model J1A Tape Drive	85A
TS1120 Model E05 Tape Drive	942

The minimum microcode levels for TS7700 R1.1, including JB media and TS1120 native support, are listed in Table 3-6.

Table 3-6 Minimum microcode levels for R1.1

Component	**Microcode level**
TS7700 Virtualization Engine	8.0.1.24
IBM 3953-L05 Library Manager	534.23
TS3500 Tape Library	6830
IBM 3952 Model J1A Tape Drive	D3I0_904
TS1120 Model E05 Tape Drive	D3I1_A33

The minimum microcode levels for TS7700 R1.2, including encryption and 3494 support, are listed in Table 3-7.

Table 3-7 Minimum microcode levels for R1.2

Component	**Microcode level**
TS7700 Virtualization Engine	8.2.0.33
IBM 3953-L05 Library Manager or IBM 3494 Library Manager	534.33
TS3500 Tape Library	7260
IBM 3952 Model J1A Tape Drive	D3I0_A0D
TS1120 Model E05 Tape Drive	D3I1_B25

The minimum microcode levels for TS7700 R1.3, including Three-Cluster Grid, 1,000,000 logical volumes and copy export for Single Cluster Grid, are listed in Table 3-8.

Table 3-8 Minimum microcode levels for R1.3

Component	Microcode level
TS7700 Virtualization Engine	8.3.0.106
IBM 3953-L05 Library Manager or IBM 3494 Library Manager	535.02
TS3500 Tape Library	7270
IBM 3952 Model J1A Tape Drive	D3I0_B07
TS1120 Model E05 Tape Drive	D3I1_C13

The minimum microcode levels for TS7700 R1.4, including copy export for Multi Cluster Grid and single cache drawer configuration, are listed in Table 3-9.

Table 3-9 Minimum microcode levels for R1.4

Component	Microcode level
TS7700 Virtualization Engine	8.4.0.xxx
IBM 3953-L05 Library Manager or IBM 3494 Library Manager	535.xx
TS3500 Tape Library	7xxx
IBM 3952 Model J1A Tape Drive	D3I0_xxx
TS1120 Model E05 Tape Drive	D3I1_xxx

The minimum microcode levels for TS7700 R1.4a, including Dynamic Grid Network Balancing, Host Copy Control, removal of a cluster from a grid, Cluster Cleanup, and upgrading a single cache unit configuration to a two cache unit configuration are listed in Table 3-10.

Table 3-10 Minimum microcode levels for R1.4a

Component	Microcode level
TS7700 Virtualization Engine	8.4.1.37
IBM 3953-L05 Library Manager or IBM 3494 Library Manager	535.23
TS3500 Tape Library	7368
IBM 3952 Model J1A Tape Drive	D3I0_C0D
TS1120 Model E05 Tape Drive	D3I1_D12

3.2.3 TCP/IP configuration considerations

In this section, we discuss configuration considerations and LAN/WAN requirements for the TS7700 Virtualization Engine. We cover Single and Multi Cluster Grid configurations.

Single Node configuration

The NAT routers specified are embedded into the TS7700 for its internal network, for connection to the Management Console, and for users to access the Management Interface from their intranet.

You need to supply three TCP/IP addresses and two network cables for the physical connection from your intranet for assignment to the TS7700 for access to the management interface. Two of those IP addresses are static, assigned to each of the two physical connections within the two NAT routers. The third IP is used as a virtual IP address, which will automatically move between the two static ones in case an outage should occur to one of the routers. Therefore, we recommend that you only use this third, so-called *virtual IP address,* for your access to the TS7700 Management Interface, although you could technically use any of the three IP addresses.

> **Important:** All three provided IP addresses that will be assigned to one TS7700 cluster for the Management Interface access must be in the same subnet, for example:
>
> - 192.168.1.1
> - 192.168.1.2
> - 192.168.1.3
>
> with subnet mask 255.255.255.0.

Grid configuration

In a grid environment all network components are duplicated from the gNode point of view, depending also on how many hosts the grid is composed of (see Figure 3-3).

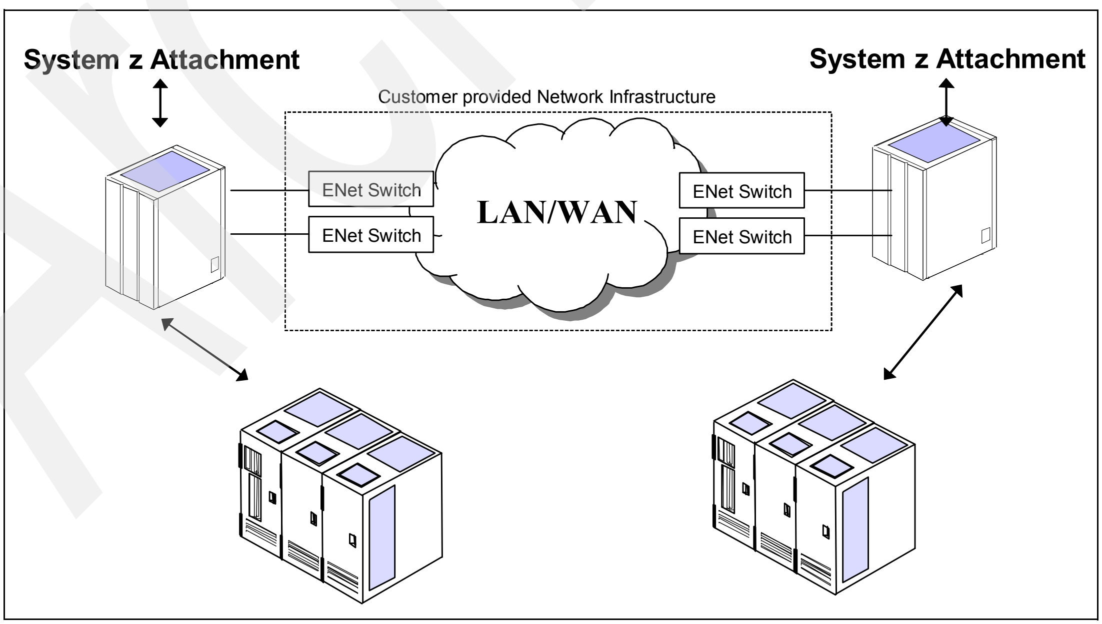

Figure 3-3 TS7700 Two-Cluster Grid

The connections between the two TS7700 peers shown in Figure 3-3 are two separate 1 Gbit Ethernet connections. The TS7700 is connected to the customer-provided switches through Cat6 (recommended) or Cat5E (not recommended) cables. Standard TCP/IP protocol is used.

Note: For performance reasons and utilization of the full 1 Gbps, we strongly recommend that you use Cat6 cables.

Then also we highly recommend to use *two* links of 1 Gbit Ethernet between Clusters, in a Grid, for performance reasons and to apply with 1.4a enhancements of Dynamic Grid Network Balancing.

Figure 3-4 shows the customer and TSSC network connections and the IP addresses in a grid environment.

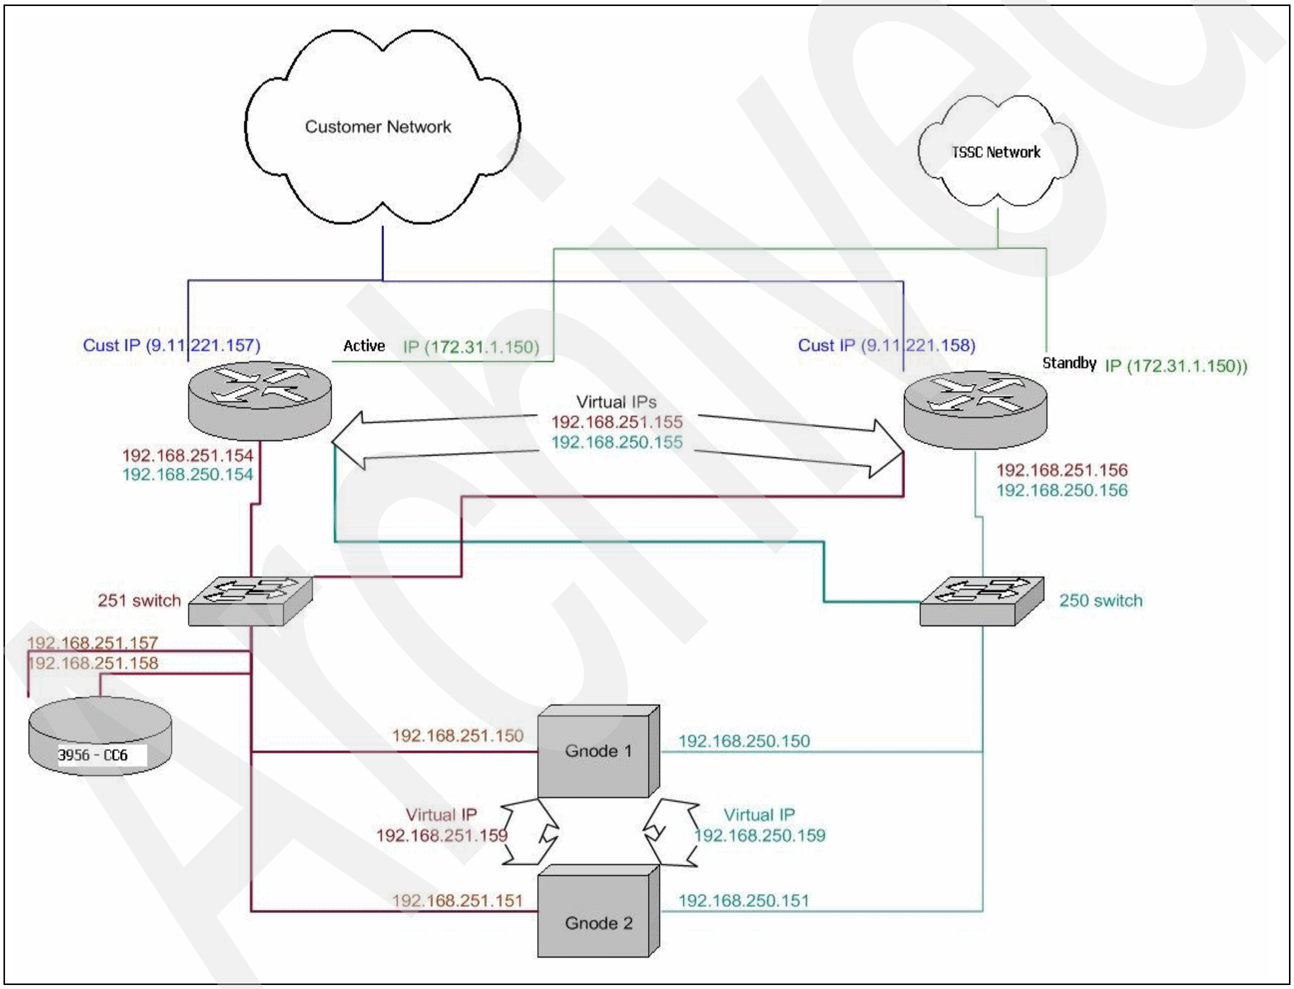

Figure 3-4 Grid network

IP addresses for the Management Interfaces in a Multi Cluster Grid

In a TS7700 Multi Cluster Grid environment, you need to supply three TCP/IP addresses and two network cables per cluster for the physical connection from your intranet for assignment to the TS7700s for access to the Management Interfaces.

Similar to a Single Cluster configuration, three of the IP addresses are used for each cluster as well as two physical cable connections per cluster. Again, two of those IP addresses per cluster are static, assigned to each of the physical connections within the two TS7700 NAT routers. The third IP is used as a virtual IP address that will automatically move between the two static ones in case an outage occurs on one of the routers.

Therefore, we recommend that you only use this third, so-called *virtual IP address* per cluster for your access to each TS7700 Management Interface. Technically you can use any of the three IP addresses per cluster.

> **Attention:** All three provided IP addresses that will be assigned to one TS7700 cluster for the Management Interface access must be in the same subnet.
> In the example 192.168.1.1, 192.168.1.2, 192.168.1.3 with subnet mask 255.255.255.0.

In a TS7700 Multi Cluster Grid, for access to each cluster's Management Interface, it is possible to use a separate subnet for each cluster for redundancy reasons. The remote TS7700 cluster's Management Interface can also be accessed through the internal Gbit Ethernet Link, should the outside network to one of the clusters be unavailable.

This function is handled automatically by each cluster's Management Interface under the cover. Figure 3-5 shows a sample setup for a Two-Cluster Grid. For a Three-Cluster Grid simply use addresses 192.168.3.1, 192.168.3.2, and 192.168.3.3 for the third cluster.

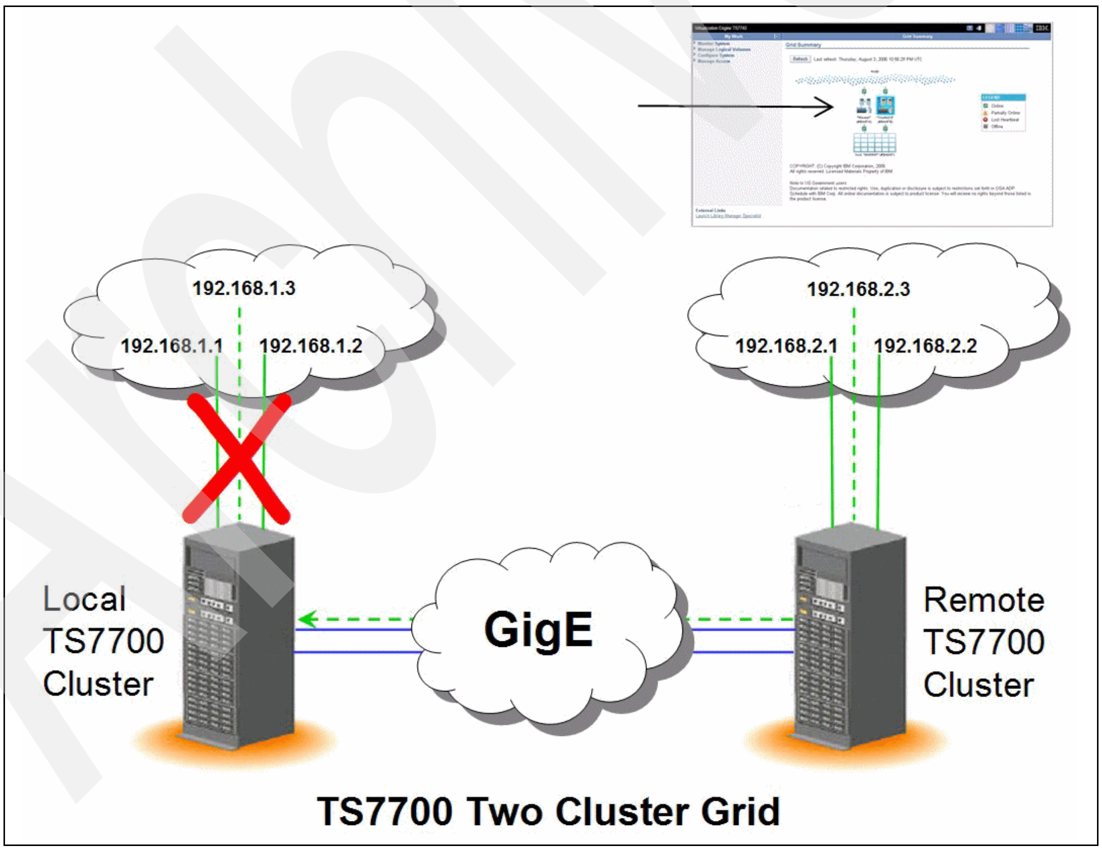

Figure 3-5 TS7700 MI access from a remote cluster

IP addresses for Cross Site Replication within a Multi Cluster Grid

In a TS7700 Multi Cluster Grid environment, you need to supply two additional TCP/IP addresses and two additional Gbit Ethernet network connections per cluster for the physical connection for the TS7700 Cross Site Replication communication.

The physical Gbit Ethernet links should be configured with TCP/IP addresses acting in different and independent logical subnets for proper failure and error detection of the Gbit Ethernet links. In other words, each Gbit Ethernet link operates with two TCP/IP addresses and a subnet mask different from the second Gbit Ethernet link.

Important: We recommend that you configure each Gbit Ethernet link between the TS7700 clusters as a separate subnet. For example, use:

- 172.1.1.1 and 172.1.1.2 and 172.1.1.3 for the first Gb Ethernet link
- 172.1.2.1 and 172.1.2.2 and 172.1.2.3 for the second Gb Ethernet link
- 255.255.255.0 as subnet mask for both

See Figure 3-6 for a sample configuration for a Two-Cluster Grid.

Note: A Three-Cluster Grid simply adds another cluster to Figure 3-6 that uses, for example, IP addresses of 172.1.1.3 and 172.1.2.3.

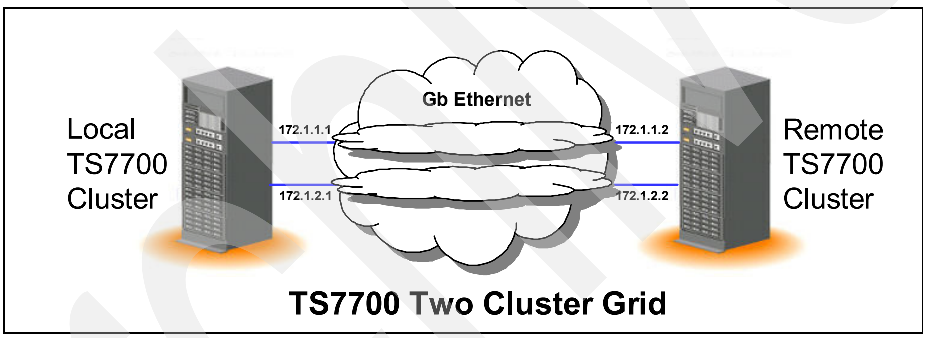

Figure 3-6 TS7700 Gbit Ethernet link subnets

IP addresses for the TSSC

The TS7700 reserves a block of ten IP addresses on the TSSC network. Seven of these IP addresses are currently used. Three are reserved for future use.

The seven IP addresses are:

- Two IP addresses for the routers
- Two IP addresses for the IBM 3956-CC6
- One IP address for the IBM 3957-V06
- One ASMI IP address
- One IP address for the Management Interface

Note: The IP addresses used for the TSSC network must be entered as an increment of 10 between 10 and 240. The router configuration will use this address and the next nine sequential addresses for TSSC configuration. To prevent network problems, do not configure these addresses on this TSSC for another system.

3.2.4 LAN/WAN requirements

The TS7700 Virtualization Engine uses your LAN/WAN to replicate logical volumes, access logical volumes remotely, and perform cross-site messaging. The LAN/WAN bandwidth should have adequate bandwidth to deliver the throughput necessary for customer requirements. The cross-site grid network can be Shortwave optical fiber or 1 Gbit copper Ethernet, which can be Cat 5E or 6 cabling. Cat 6 cabling is recommended to achieve the highest throughput when using copper connections. The TS7700 Virtualization Engine does not encrypt the data it sends over the LAN/WAN.

Cross-site dedication

Your LAN/WAN must be dedicated to cross-site grid communication and not shared with other Internet traffic or high bandwidth applications. The TS7700 connection to the cross-site grid network can be Shortwave optical fiber or 1 Gbit copper Ethernet. With copper connections, Cat6 cabling is highly recommended.

The system will use all available bandwidth up to 1 Gbit per link for its use. Sharing the WAN infrastructure with other high bandwidth applications might cause contentions and adversely affect performance.

Network protection

The security of the data exchanged between the sites should be considered as part of the WAN topology and implementation. For example, consider to:

- Isolate your LAN/WAN from the general employee network
- Encrypt or otherwise secure data in flight

Bandwidth requirements

Your LAN/WAN must be able to support the bandwidth requirements of your grid. These requirements are a factor of the amount of data requiring replication and the expected time in which replication should be accomplished. Recommended bandwidth capacity is 1 Gbps per redundant paths. Work with your sales team to determine a recommended minimum bandwidth based on your requirements.

Network redundancy

The TS7700 Virtualization Engine provides two independent 1 Gb Ethernet connections. We recommend that you connect each through an independent WAN interconnection to be protected from a single point of failure that would disrupt service to both WAN paths from a node.

3.2.5 ETL Specialist prerequisites

The Library Manager must be connected to your system's LAN using FC5220 (Ethernet Adapter). During the installation process, the IBM System Service Representative (SSR) will set up TCP/IP on the Library Manager to use your assigned TCP/IP host name and TCP/IP address (and router information, if necessary). You can help the installation process if you obtain the following information before the installation starts:

- TCP/IP host name
- TCP/IP address
- Subnet mask (or network mask)
- Router address (or Gateway address)*
- Domain name*
- Nameserver address*

> **Note:** Items with an asterisk (*) are optional. Their use depends on your system’s LAN configuration.

You need to have a commonly-used browser to view the information provided by the Specialist. For example, Microsoft® Internet Explorer® Version 5.0 or Netscape Navigator Version 4.7 or higher with JavaScript™ and Java™ enabled provide compatible capability. A text-based Web browser is not supported.

The Specialist resides in the Library Manager and is part of the microcode. When attached to the customer network, it is available to the customer through a Web browser. There is no special "setup" required on the browser; however, the initialization phase will take a little longer when you start the connection for the first time on any client. This is expected due to the Java code download, but it should only occur once.

Some of the panels can be used for online queries about the status of the TS7700 Virtualization Engine, its components, and the distributed libraries. Others, which require Administration access, can be used to change your existing environment.

3.3 Remote installations and switch support

The TS7700 Virtualization Engine attaches to the System z host through FICON channel attachments. There are three basic types of switch connections that can be used between the host and TS7700:

- Direct connect
- Single switch
- Cascaded switches

You can also use Dense Wave Division Multiplexers (DWDMs) or FICON channel extenders between the System z host and the TS7700. See Figure 3-7 on page 102 for more details about the distances supported.

3.3.1 Factors that affect performance at a distance

Fibre Channel distances depend on many factors, which include:

- Type of laser used: Longwave or Shortwave
- Type of fiber optic cable: Multi-mode or Single-mode
- Quality of the cabling infrastructure in terms of dB signal loss:
  - Connectors
  - Cables
  - Bends and loops in the cable

Native shortwave FC transmitters have a maximum distance of 500 m with 50 micron diameter, multi-mode, optical fiber. Although 62.5 micron, multi-mode fiber can be used, the larger core diameter has a greater dB loss and maximum distances are shortened to 300 m. Native longwave FC transmitters have a maximum distance of 10 km when used with 9 micron diameter single-mode optical fiber.

Link extenders provide a signal boost that can potentially extend distances to up to about 100 km. These link extenders simply act as a very big, fast pipe. Data transfer speeds over link extenders depend on the number of buffer credits and efficiency of buffer credit management in the FC nodes at either end of this fast pipe. Buffer credits are designed into the hardware for each FC port. FC provides flow control that protects against collisions.

This is extremely important for storage devices, which do not handle dropped or out-of-sequence records. When two FC ports begin a conversation they exchange information about their buffer capacities. An FC port will send only the number of buffer frames for which the receiving port has given credit. This not only avoids overruns, but also provides a way to maintain performance over distance by filling the “pipe” with in-flight frames or buffers. The maximum distance that can be achieved at full performance is dependent on the capabilities of the FC node that is attached at either end of the link extenders.

This relationship is very vendor-specific. There should be a match between the buffer credit capability of the nodes at either end of the extenders. A host bus adapter (HBA) with a buffer credit of 64 communicating with a switch port with only eight buffer credits would be able to read at full performance over a greater distance than it would be able to write. This is because, on the writes, the HBA can send a maximum of only eight buffers to the switch port, while on the reads, the switch can send up to 64 buffers to the HBA. Until recently, a rule of thumb has been to allot one buffer credit for every 2 km in order to maintain full performance.

Buffer credits within the switches and directors have a large part to play in the distance equation. The buffer credits in the sending and receiving nodes heavily influence the throughput that is attained in the Fibre Channel. Fibre Channel architecture is based on a flow control that ensures a constant stream of data to fill the available pipe. A rule-of-thumb says that to maintain acceptable performance, one buffer credit is required for every 2 km distance covered. Refer to *Introduction to SAN Distance Solutions*, SG24-6408.

> **Note:** For performance reasons and utilization of the full 1 Gbps, we strongly recommend that you use Cat6 cables.
>
> Then also we highly recommend to use *two* links of 1 Gb Ethernet between Clusters, in a Grid, for performance reasons and to apply with 1.4a enhancements of Dynamic Grid Network Balancing.

3.3.2 FICON Director support

All FICON Directors are supported for Single and Multi Cluster Grid configurations with 1 Gbps, 2 Gbps, or 4 Gbps links. The components will auto-negotiate to the highest speed allowed.

You cannot mix different vendors (McData (CNT & Inrange), CISCO, and Brocade) but you can mix models of one vendor. See the switch Web pages for specific intermix combinations supported.

Table 3-11 lists the FICON Director vendors and models supported with the TS7700.

Table 3-11 FICON Directors support

Director Vendor and Model	Director level of microcode supported	TS7740 level of microcode supported
Brocade 14000 IBM 2109-M14	5.1.0c or 5.2.0b	8.0.0.x - 8.4.1.x
Brocade Silkworm 48000 IBM 2109-M48	5.1.0c or 5.1.0d or 5.2.0b or 5.1.0c_rc2 or 5.3.0b	8.0.0.x - 8.4.1.x
Brocade 3900 IBM 2109-F32	5.2.0b	8.3.0.x - 8.4.1.x
Brocade 4100 IBM 2005-B32	5.2.0b FOS 6.1.0c (FTA not supported)	8.3.0.x - 8.4.1.x

Director Vendor and Model	Director level of microcode supported	TS7740 level of microcode supported
Brocade DCX IBM 2499-384 (SAN768B) Nortel 10 GB/s MOTR card supported	FOS 6.1.0c (FTA not supported)	8.4.1.x
Brocade 4900 IBM 2005-B64	FOS 6.1.0c (FTA not supported)	8.4.1.x
Brocade 5000 IBM 2005-B5K (SAN32B-3)	FOS 6.1.0c (FTA not supported)	8.4.1.x
Brocade 5100 IBM 2498-B40 (SAN40B-4)	FOS 6.1.0c (FTA not supported)	8.4.1.x
Brocade 5300 IBM 2498-B80 (SAN80B-4)	FOS 6.1.0c (FTA not supported)	8.4.1.x
Brocade 7500 IBM 2005-R18 (SAN18B-R)	FOS 6.1.0c (FTA not supported)	8.4.1.x
Brocade ED-6140 (McData) IBM 2027-140	EOS 9.0 EOS 9.0.1 EOS 9.02.00 EOS 9.02.01 EOSc 9.7.2	8.0.0.x - 8.4.1.x
Brocade i10K (McData) IBM 2027-256	EOS 6.03.02 EOS 9.0.1 EOS 9.01.00 9.02.00 EOSn 9.7.2	8.0.0.x - 8.4.1.x
Brocade 4700 (McData) IBM 2026-432	EOS 9.02.00 EOSc 9.7.2	8.3.0.x - 8.4.1.x
Brocade ED-6064 (McData) IBM 2032-064	EOS 9.02.00 EOSc 9.6.2	8.3.0.x - 8.4.1.x
Cisco MDS9506 IBM 2062-D04 IBM 2054-E04	3.0.2 3.0.3 3.1(2a) 3.2(2c) 12, 24 & 48 port blades only NOTE: FICON Tape Acceleration feature is not supported.	8.0.0.x - 8.4.1.x
Cisco MDS9509 IBM 2062-D07 IBM 2054-E07	3.0.2 3.0.3 3.2(2c) 12, 24 & 48 port blades only NOTE: FICON Tape Acceleration feature is not supported.	8.0.0.x - 8.4.1.x

Director Vendor and Model	Director level of microcode supported	TS7740 level of microcode supported
Cisco MDS9513 IBM 2062-E11 IBM 2054-E11	3.0.2 3.0.3 3.0(3b) 3.1(2a) 3.2(2c) 12, 24 & 48 port blades only NOTE: FICON Tape Acceleration feature is not supported.	8.0.0.x - 8.4.1.x
Cisco MDS9134	3.2(2c)	8.4.1.x

For the most recent updates of the matrix in Table 3-11, refer to:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ116133

See the Technote *Performance Considerations for a Cascaded FICON Director Environment,* which is available at:

http://www-1.ibm.com/servers/eserver/zseries/library/techpapers/pdf/gm130237.pdf

3.3.3 FICON channel extenders

FICON channel extenders are available working in one of the following modes:

- Frame shuttle or tunnel mode
- Emulation mode

Using the *shuttle* or *tunnel* mode, the extender receives and forwards FICON frames without doing any special channel or control unit processing. The performance is limited to the distance between the sites and the normal round trip delays in FICON channel programs.

Emulation mode can go unlimited distances, and it monitors the I/O activity to devices. The channel extender interfaces emulate a control unit by presenting command responses and CE/DE status ahead of the controller and emulating the channel when running the pre-acknowledged write operations to the real remote tape device. Thus, data is accepted early and forwarded to the remote device to maintain a full pipe throughout the write channel program.

The channel extenders listed in Table 3-12 are supported between the System z host and the TS7700 Virtualization Engine.

Table 3-12 FICON channel extender support

Channel Extender Vendor and Model	Channel Extender level of code supported	TS7740 level of microcode supported
Ciena CN2000	V4.1.2 Note: Flow control needs to be enabled to get optimal performance.	8.0.0.x - 8.4.1.x
Brocade (McDATA) Edge	3.1.5 or V3.2.0.4.3 or V3.2.0.5	8.0.0.x - 8.4.1.x
Brocade (McDATA) USD and USD-x	3.4.2.0.6 or V4.1.0.7.3 or V4.1.0.11	8.0.0.x - 8.4.1.x

3.3.4 Supported distances

When directly attaching to the host, the TS7700 can be installed at a distance of up to 10 km from the host. With FICON Switches, also called FICON Directors, or Dense Wave Division Multiplexers (DWDMs), the TS7700 can be installed at extended distances from the host.

Supported FICON extended distances are summarized in Figure 3-7.

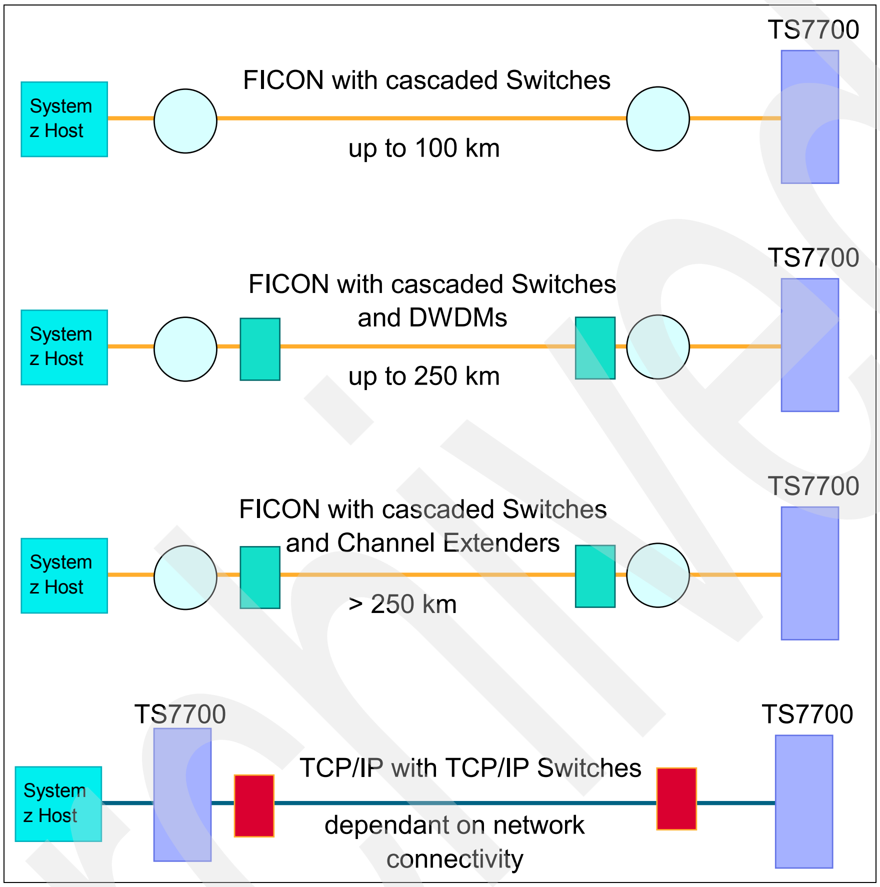

Figure 3-7 TS7700 Extended Distance Support

In a Multi Cluster Grid configuration, the TS7700 Virtualization Engines are connected through TCP/IP connections. These connections are uncritical as long as sufficient bandwidth is available.

3.3.5 Implementing cascaded switches

The following list summarizes the general configuration rules for configurations with cascaded switches:

- Director Switch ID

  This is defined in the Director GUI setup.

  The inboard Director Switch ID is used on the SWITCH= parameter in the CHPID definition. The Director Switch ID does not have to be the same as the Director Address. We recommend that you keep them the same to reduce configuration confusion and simplify problem determination work, although the example uses a different ID and address for clarity.

The following allowable Director Switch ID ranges have been established by the manufacturer:

– McDATA must be in the range x'61' to x'7F'
– CNT/Inrange must be in the range x'01' to x'EF'
– Brocade must be in the range x'01' to x'EF'

- Director Address

  This is defined in the Director GUI setup.

  The Director Domain ID is the same as the Director Address that is used on the LINK parameter in the CNTLUNIT definition. The Director Address does not have to be the same as the Director ID, but again, we recommend that you keep them the same to reduce configuration confusion and simplify PD work.

  The following allowable Director Address ranges are established by the manufacturer:

  – McDATA must be in the range x'61' to x'7F'.
  – CNT/Inrange must be in the range x'01' to x'EF'.
  – Brocade must be in the range x'01' to x'EF'.

- Director Ports

  The Port Address might not be the same as the Port Number. The Port Number identifies the physical location of the port, and the Port Address is used to route packets.

  The Inboard Director Port is the port to which the CPU is connected. The Outboard Director Port is the port to which the control unit is connected. It is combined with the Director Address on the LINK parameter of the CNTLUNIT definition:

  – Director Address (hex) combined with Port Address (hex) (two bytes)
  – Example: LINK=6106 would indicate a Director Address of x'61' and a Port Address of x'06'

- External Director connections

  – Inter-Switch Links (ISLs) connect to E Ports.
  – FICON Channels connect to F Ports.

- Internal Director connections

Port type and Port-to-Port connections are defined using the Director's GUI setup.

3.4 Planning for software implementation

In this section, we provide information for planning tasks that are related to host configuration and software requirements for use with the TS7700 Virtualization Engine.

3.4.1 Host configuration definition

You must define the hardware to the host using the Hardware Configuration Definition (HCD) dialog boxes. We recommend that you use LIBRARY-ID and LIBPORT-ID.

LIBRARY-ID

In a grid configuration used with the IBM Virtualization Engine TS7700, each virtual device that is attached to a System z host reports back the same library sequence number, known as the *Composite Library-ID*. The Composite Library-ID allows the host to consider the grid configuration as a single library.

Each cluster in a grid configuration possesses a unique library sequence number, known as the *distributed library ID*, which identifies it among the clusters in its grid. This distributed library ID is reported to the System z host upon request and is used to distinguish one cluster from another in a given grid.

Subsystem identification

Each logical control unit, or 16-device group, must present a unique subsystem identification to the System z host. This ID is a 1 byte field that uniquely identifies each device associated with the logical control unit within the larger storage controller and is called LIBPORT-ID. The value of this ID cannot be 0. Table 3-13 shows the definitions of the LIBPORT-IDs in a Multi Cluster Grid with 768 virtual devices.

For more details, refer to 5.1.4, "Library names, library IDs, and port IDs" on page 196.

Table 3-13 Subsystem identification definitions

Cluster	Logical CU (hex)	LIBPORT-ID (hex)
0	0-F	01-10
1	0-F	41-50
2	0-F	81-90

Virtual tape drives

The IBM Virtualization Engine TS7700 presents a tape drive image of a 3490 C2A that is identical to that of current VTS and Peer-to-Peer (PtP) subsystems. The command set, responses to inquiries, and accepted parameters match the defined functional specifications of a 3490 drive.

Each TS7700 cluster provides 128 or 256 virtual devices. When two clusters are connected to a grid configuration, the grid provides up to 512 virtual devices. When three clusters are connected to a grid configuration, the grid provides up to 768 virtual devices.

3.4.2 Software requirements

The TS7700 Virtualization Engine is supported with the following operating system levels:

- z/OS V1R7 and higher

  With z/OS, the TS7700 server is, for the most part, transparent to the host software. Refer to the existing VTS and PtP PSP buckets for the latest information about recommended software maintenance. The recommended software maintenance for the existing virtualization products also applies to the TS7700.

  In addition to the PSP buckets previously mentioned, also refer to the new 3957 PSP bucket.

> **Notes**:
>
> - APAR OA19061 is required if a host system defines more than 32 subsystem pools to a Composite Library. More than 16 subsystems were introduced with TS7700 R1.1.
>
>   APAR OA20065 applies to a Multi Cluster Grid configuration where up to 768 devices or 48 subsystems can be associated with a Composite Library.
>
>   APAR OA20065 and device services APARs OA20066, OA20067 and OA20313 are required to support the Host Console Request function.
>
>   APAR OA24966 applies to new host warning message for significant degradation in grid links.
> - z/OS R1V4, R1V5, and R1V6 support the TS7700, but these levels are out of service. An extended service contract is recommended for support.

- z/VM® 5.1.0 and higher

  With z/VM, the TS7700 is transparent to host software. z/VM V5R1, or later, is required for both guest and native VM support. DFSMS/VM Function Level 221 is required for native VM tape library support and EREP V3.5 plus PTFs.

  > **Note:** z/VM V4R4 supports the TS7700, but this level is out of service. An extended service contract is recommended for support.

- z/VSE™ 3.1.2

  With z/VSE, the TS7700 is transparent to host software. Version 3.1.2 or later is required.
- TPF 4.1 and z/TPF V1.1 and higher

  With the TS7700, both TPF releases are supported in both a Single and a Two-Cluster Grid environment with the appropriate software maintenance. The category reserve and release functions are not supported in a grid environment.

3.4.3 z/OS software environments

System-managed tape allows you to manage tape volumes and tape libraries according to a set of policies that determine the kind of service to be given to the data sets on the volume.

The automatic class selection (ACS) routines process every new tape allocation in the system-managed storage (SMS) address space. The production ACS routines are stored in the active control data set (ACDS). These routines allocate to each volume a set of classes that reflect your installation’s policies for the data on that volume. The ACS routines are invoked for every new allocation. Tape allocations are passed to the object access method (OAM), which uses its library control system (LCS) component to communicate with the Library Manager through the Asynchronous Operation Manager (AOM).

The Storage Class ACS routine determines whether a request is SMS-managed. If no Storage Class is assigned, the request is not SMS-managed, and allocation for non-specific mounts is made outside the tape library.

For SMS-managed requests, the Storage Group routine assigns the request to a Storage Group. The assigned Storage Group determines which Library Manager logical partitions in the tape library are to be used. Through the Storage Group construct, you can direct logical volumes to specific physical volume pools.

3.4.4 Sharing and partitioning considerations

When a System z host connected to an IBM TS3500 Tape Library contains both an IBM TS7700 and native TS1120 or 3592 Tape Drives, which are connected to the host through a TS1120 Model C06 or 3592 Model J70 controller, the Library Manager in the dedicated 3953 frame splits this logical library, which is a subset of the IBM 3584, into two logical library partitions:

- One partition for the TS7700 drives
- One partition for the native tape drives connected through the tape controller

To the host, each logical library partition looks like a separate physical library. There might be another TS7700 or VTS attached to this logical library partition, which will look like a third logical library to the host. Each logical library within a TS3500 library partition has its own full set of Library Manager categories that can overlap. The Library Manager maintains two sets of category tables and keeps the TS7700 and physical volumes separate.

A host can use category X'0001' for Media Type 1 for real and virtual volumes (based on DFSMS). Issuing a mount for category X'0001' to a TS7700 drive results in selecting a logical volume from the volumes that are owned by the TS7700 library. Issuing a mount from that category to a real tape drive results in selecting a volume from the volumes that are owned by the real side of the tape library.

Partitioning between multiple hosts

The virtual drives and virtual volumes in a TS7700 can be partitioned just like physical drives and real volumes. Any virtual volume can go to any physical stacked volume. The TS7700 places no restrictions on the use and management of those resources. You have the ability to partition your stacked media in up to 32 separate pools by assigning a Storage Group to a defined range of stacked volumes prior to insert time.

Sharing a TS7700 Virtualization Engine

A FICON-attached TS7700 supports two or four physical channels, each of which is capable of supporting 256 logical paths. Each logical path can address any of the 256 virtual devices in the TS7700.

We recommend that you use a FICON Director when connecting the TS7700 to more than one system.

The TS7700 places no limitations on the number of hosts that can use those channel paths, the types of hosts, or their operating system environments (as with any tape technologies that are supported in IBM tape libraries). An operating environment, however, through its implementation, does impose limits. z/OS DFSMS can support up to 32 systems or groups of systems.

Basically, anything that can be done with standalone drives in an IBM 3494 Tape Library or IBM TS3500 Tape Library, can be done with the virtual drives in a TS7700.

The TS7700 attaches to the host system or systems through two or four FICON channels. Each FICON channel provides 256 logical paths (starting with TS7700 release 1.4). With a four-FICON configuration that will result in a total of 1024 logical paths per TS7700.

You can use the following formula to calculate the number of logical paths required in an installation:

Number of logical paths = number of hosts x number of CU x number of channels

This formula assumes all hosts access all control units in the TS7700 with all channel paths.

3.5 Planning for logical and physical volumes

Before you define logical volumes to the TS7700, we recommend that you consider the total number of logical volumes that are required, the volume serial ranges to define, and the number of volumes within each range.

The VOLSERs must be unique throughout an SMSplex and throughout all storage hierarchies such as DASD, tape, and optical storage media.

The VOLSERs for logical volumes and physical volumes must be unique.

3.5.1 Logical volumes

You need to determine the number of logical volumes that are required to handle the workload that you are planning for the TS7700 Virtualization Engine. You need to consider the size of the logical volumes, the number of scratch volumes that you need per day, the time that is required for return to scratch processing, and how often scratch processing is performed.

Size of logical volumes

The TS7700 Virtualization Engine supports logical volumes with an uncompressed size of 400, 800, 1000, 2000, and 4000 MB. If the data can be compressed, the effective size will be larger. For example, if your data compresses with a 3:1 ratio, the effective maximum logical volume size for a 4000 MB logical volume is 12 000 MB.

Depending on the logical volume sizes that you choose, you might see the number of volumes that are required to store your data grow or shrink, depending on the media size from which you are converting. If you have data sets that fill native 3590 volumes, even with 4000 MB logical volumes, you will need more TS7700 logical volumes to store the data, which are stored as multi-volume data sets.

You can specify the 400 MB CST-emulated cartridges or the 800 MB with ECCST-emulated cartridges when adding volumes to the TS7700 Virtualization Engine. You can use these sizes directly, or you can use policy management to override them to provide for the 1000, 2000, or 4000 MB sizes.

You can set a logical volume size VOLSER and can change it dynamically using the DFSMS Data Class storage construct. The amount of data copied to the stacked cartridge is only the amount of data that is written to a logical volume. The choice between all available logical volume sizes does not affect the real space that is used in either the TS7700 cache or the stacked volume.

In general, unless you have a special need for CST emulation (400 MB), specify the ECCST media type when inserting volumes in the TS7700 Virtualization Engine.

Number of scratch volumes needed per day

As you run your daily production workload, you need enough logical volumes in scratch status to support the data that will be written to the TS7700 Virtualization Engine. This number could be hundreds or thousands of volumes, depending on your workload. You will likely want more than a day's worth of scratch volumes available at any point in time.

Time required for return to scratch processing

Return to scratch processing involves running a set of tape management tools that identify the logical volumes that no longer contain active data and then communicating with the

TS7700 Virtualization Engine to change the status of those volumes from private to scratch. The amount of time that the process takes depends on the type of tape management system being employed, as well as how busy the TS7700 Virtualization Engine is when it is processing the volume status change requests and whether a grid configuration is being used.

You should expect that during the day when the TS7700 Virtualization Engine is not handling a workload peak, that up to 5000 logical volumes per hour can be returned to scratch for a Single Cluster and up to 2500 logical volumes per hour can be returned to scratch in a Two-Cluster Grid configuration.

You can see elongated elapsed time in any Tape Management Systems return to scratch process, when you migrate to or install a multi-cluster configuration solution.

How often return to scratch processing is performed

If the number of logical volumes used on a daily basis is small (less than a few thousand), you can choose to only perform return to scratch processing every few days. A good rule of thumb is to plan for no more than a 4-hour time period to run return to scratch. By ensuring a nominal runtime of 4 hours, enough time exists during first shift to run the process twice should problems be encountered during the first attempt.

For example, assume that you use 2000 scratch volumes a night on a Two-Cluster Grid configuration. A 4-hour return to scratch period returns up to 10 000 logical volumes to scratch status. Running return to scratch processing every 5 days would just keep up with the rate at which the scratch logical volumes are being used. Therefore, running return to scratch every 3 or 4 days can provide some buffer in the number of available scratch logical volumes.

From a Tape Management System point of view, our best recommendation is to run the process, at least every day.

In planning for the number of logical volumes needed, first determine the number of private volumes that make up the current workload you are migrating. One way to determine this number is by looking at the amount of data on your current volumes and then matching that to the supported logical volume sizes. Match the volume sizes taking into account the compressibility of your data. If you do not know the average ratio, you can use a conservative value of 2:1.

Table 3-14 shows an example of mapping current volume sizes for data that compresses at 2:1 to supported logical volume sizes. (You can use the BatchMagic analysis tool for the current volume size information.)

Table 3-14 Volume sizes for data compressed at 2:1

Volume size in MB	Count	400 MB	800 MB	1000 MB	2000 MB	4000 MB
<50	15251	15251	0	0	0	0
50-200	5450	5450	0	0	0	0
200-400	4965	4965	0	0	0	0
400-800	5740	5740	0	0	0	0
800-1200	6280	0	6280	0	0	0
1200-1600	4430	0	4430	0	0	0
1600-2000	4100	0	0	4100	0	0
2000-2400	2652	0	0	0	2652	0

Volume size in MB	Count	400 MB	800 MB	1000 MB	2000 MB	4000 MB
2400-2800	2100	0	0	0	2100	0
2800-3200	820	0	0	0	820	0
3200-3600	357	0	0	0	357	0
3600-4000	295	0	0	0	295	0
4000-4400	653	0	0	0	0	653
4400-4800	512	0	0	0	0	512
>4800	1268	0	0	0	0	1268
Total	**54873**	**31406**	**10710**	**0**	**6224**	**2433**

Using Table 3-14 as example, you currently have 54 873 volumes and need 50 773 using the logical volume sizes supported.

Some of the original volumes are 3590 volumes and the rest are 3490 volumes.

If you choose to use only the 800 MB volume size, the total number that are needed will increase because some of the current volumes that contain more than 800 MB compressed need to expand to a multi-volume set. Taking that into account, for the above example the number of logical volumes that are required would increase by 22 180 (for each 1000 MB volume called out in the table, one additional 800 MB volume is needed for spanning; for each 2000 MB volume, two additional volumes are needed; and for each 4000 MB volume, four additional volumes are needed) for a total of 72 953.

Now that you know the number of volumes that you need for your current data, you can estimate the number of empty scratch logical volumes that you have to add. Based on your current operations, you can determine a nominal number of scratch volumes from your nightly use. If you have an existing VTS installed, you might have already determined this number and can set a scratch media threshold with that value through the Interactive Storage Management Facility (ISMF) Library Define panel.

You should then multiply that number by the value that will provide a sufficient buffer (typically 2×) and by the frequency with which you want to perform return to scratch processing.

You can use the following suggested formula to calculate the number of logical volumes that you need:

$$Vv = Cv + Tr + (Sc)(Si + 1)$$

Where:

Vv Total number of logical volumes needed

Cv Number of logical volumes needed for current data rounded up to the nearest 10 000

Tr The threshold value from the ISMF Library Define panel for the scratch threshold for the media type used (normally MEDIA2); usually set to equal the number of scratch volumes used per night.

Sc The number of empty scratch volumes used per night, rounded up to the nearest 500

Si The number of days between scratch processing (return to scratch) by the tape management system

For example, assuming the current volume requirements from Table 3-14 (using all the available volume sizes), using 2500 scratch volumes per night and performance of return to scratch processing every other day, you use the following formula to plan the number of logical volumes in the TS7700 Virtualization Engine:

75 000 (current, rounded up) + 2500 + (2500) (1+1) = 82 500 logical volumes

If you define more volumes than you need, you can always delete the additional volumes. Unused logical volumes do not consume space, but see 3.5.2, “Physical volumes” on page 110 for implications about having a large number of previously used but now scratch logical volumes.

The IBM System Storage TS7700 Virtualization Engine supports up to 1 000 000 logical volumes. This is the maximum number either in a standalone or in a grid configuration.

Restriction: Up to 10 000 logical volumes, including both a range of logical volumes and a quantity of logical volumes, can be inserted at one time. Attempting to insert more than 10 000 volumes will return an error.

Volume serial numbering

When you insert volumes, you provide starting and ending volume serial number range values.

The TS7700 Virtualization Engine determines how to establish increments of VOLSER values based on whether the character in a particular position is a number or a letter. For example, inserts starting with ABC000 and ending with ABC999 will add logical volumes with VOLSERs of ABC000, ABC001, ABC002, and so forth through ABC998, and ABC999 into the inventory of the TS7700 Virtualization Engine. You might find it helpful to plan for growth by reserving multiple ranges for each TS7700 Virtualization Engine that you expect to install.

If you have multiple partitions, it might be better to plan in advance which ranges to use in which partitions. For example, A* for first sysplex, B* for second sysplex, and so on. If you need more than a range, you can select A* and B* for the first sysplex, C* and D* for the second sysplex, and so on.

3.5.2 Physical volumes

You need to determine the number of physical volumes that are required to accommodate the workload that you are planning for the TS7700 Virtualization Engine.

You need to consider the number of logical volumes you define, average amount of data on a volume, average compression ratio achieved for the data, if the Selective Dual Copy function is to be used, whether the Delete Expired Volume Data setting is to be used, whether the Copy Export function is to be used, the reclaim threshold settings, scratch physical volumes, and number of physical volume pools.

Number of logical volumes

Because the data on a logical volume is stored on a physical volume, the number of logical volumes has a direct effect on the number of physical volumes required.

Average amount of data on a volume

The TS7700 Virtualization Engine stores only the amount of data that you write to a logical volume plus a small amount of metadata.

Average compression ratio achieved for the data

The data that a host writes to a logical volume might be compressible. The space that is required on a physical volume is after the results of compression. If you do not know the average number for your data, assume a conservative 2:1 ratio.

Selective Dual Copy function

If you use this function for some or all of your data, a second physical copy of the data is written to a physical volume.

For critical data that resides only on tape, you can currently make two copies of the data on physically separate tape volumes, either manually through additional job steps or through their applications. Within a TS7700, you might need to be able to control where a second copy of your data is placed such that it is not stacked on the same physical tape as the primary copy. Although this can be accomplished through the logical volume affinity functions, it simplifies the management of the tape data and better uses the host CPU resources to have a single command to the TS7700 subsystem to direct it to selectively make two copies of the data that is contained on a logical volume.

If you activate Dual Copy for a group of data or a specific pool, consider that all tasks and properties that are connected to this pool become duplicated:

- The number of reclamation tasks
- The number of physical drives used
- The number of cartridges used
- The number of writes to the cartridges (one from cache to the primary pool and another to the backup pool)

You have to plan for additional throughput and capacity. You do not need more logical volumes because the second copy uses an internal volume ID.

Number of Copy Export volumes

If you are planning to use the Copy Export function, you need to plan to have enough physical volumes for the Copy Export function and sufficient storage cells for the volumes in the library destined for Copy Export or in the copy export state. The Copy Export function allows for a maximum of 2000 physical volumes to be in the Copy Export state, which includes off-site volumes, the volumes still in the physical library that are in the copy export state, and the empty, filling, and full physical volumes that will eventually be set to the copy export state. After your copy export operations reach a steady state, approximately the same number of physical volumes are returned to the library for reclamation because there are those that are sent off-site as new members of the copy export set of volumes.

Delete Expired Volume Data setting

If the Delete Expired Volume Data setting on the Library Manager is not used, logical volumes occupy space on the physical volumes even after they have been returned to scratch. In that case, only when a logical volume is rewritten is the old data released to reduce the amount of active data on the physical volume.

With the Delete Expired Volume Data setting, the data associated with volumes that have been returned to scratch are deleted after a time period and their old data released. For example, assume that you have 20,000 logical volumes in scratch status at any point in time and that the average amount of data on a logical volume is 400 MB and that the data compresses at a 2:1 ratio. The space occupied by the data on those scratch volumes is 4 000 000 MBs or the equivalent of 14 JA cartridges (when using J1A Emulation mode). By using the Delete Expired Volume Data setting, you could reduce the number of cartridges required in this example by 14.

Reclaim threshold settings

The number of physical volumes also depends on the reclaim threshold percentage that you have specified for the minimum amount of active data on a volume before it becomes eligible for reclamation. The default is set to 10%. The reclaim threshold setting can have a large impact on the number of physical volumes required. Physical volumes will hold between the threshold value and 100% of data. On average, the percentage of active data on the physical volumes is (100%-10%)/2 or 45% (assuming a reclamation setting of 10%).

Having too low a setting results in more physical volumes being needed, and having too high a setting can impact the TS7700 Virtualization Engine's ability to perform host workload because it is using its resources to perform reclamation. You might need to experiment to find a threshold that matches your needs.

We suggest as a good start point to use 35%. We think that it will fit in most installations.

Number of physical volume pools

You should plan for at least 10 scratch physical volumes to be available in the common scratch pool.

For each physical volume pool you are planning on using, you should have at least three scratch physical volumes. These are in addition to the number of physical volumes calculated to hold the data on the logical volumes.

The suggested formula to calculate the number of physical volumes needed is:

$$Pv = (((Lv + Lc) \times Ls)\ Cr) / (Pc \times (((100-Rp)/100)/2))$$

Where:

Pv	Total number of physical volumes needed
Lv	Number of logical volumes defined
Lc	The number of logical volumes that will be dual copied
Ls	The average logical volume size in MB
Cr	Compression ratio of your data
Rp	The reclamation percentage
Pc	Capacity of a physical volume in MB

To this number, you would then add scratch physical volumes based on the common media pool and the number of physical volume pools you plan on using. For example, using the following assumptions:

Lv	82 500 (see "Number of logical volumes" on page 110)
Lc	10 000
Ls	400 MB
Cr	2
Rp	10
Pc	300 000 MB (capacity of a 3592 J1A written JA volume)
	500 000 MB (capacity of a TS1120 written JA volume)
	700 000 MB (capacity of a TS1120 written JB volume)
Common scratch pool	10
Volume Pools	5 (with 3 volumes per pool)

You need to plan on the following number of physical volumes in the TS7700 Virtualization Engine:

$$(((82{,}500 + 10{,}000) \times 400\ \text{MB})/2)/(300{,}000 \times (((100 - 10)/100)/2)) + 10 + 5 \times 3 = 162$$ physical volumes

If you insert more physical volumes in the TS7700 Virtualization Engine than you need, you can eject them at a later time.

Pooling considerations

Pooling can have an effect on the throughput of your TS7700. If you are using physical volume pooling at your site, consider the following possibilities:

1. The possible increase of concurrent TS1120 drive usage within the TS7700

   Depending on the number of pools and the amount of logical volume data being created per pool, you need to ensure that sufficient drives are available to handle:

   – Pre-migration of logical volumes from the Tape Volume Cache
   – Recall of logical volumes from physical stacked volumes to the Tape Volume Cache
   – The amount of logical volume data being dual copied

2. The reclamation process

   Reclamation is done at the pool level and each reclamation task will use two drives. To minimize the effects of the reclamation process, ensure you maintain a sufficient amount of physical TS7700 scratch cartridges so that the reclamation process is performed within the reclaim scheduled time.

3. An increase in the amount of cartridges being used
4. Library slot capacity
5. TS7700 processing/cache capacity

Out-of-Scratch condition for physical stacked volumes

It is highly recommended that you monitor the number of empty stacked volumes in a library. If the library is close to running out of a physical volume media type, actions should be taken to either expedite the reclaim of physical stacked volumes or add additional ones. With release 1.3, the Host Console Request function can be used to obtain the physical volumes counts within the TS7700. You can also use the Bulk Volume Information Retrieval (BVIR) function to obtain the physical media counts for each library. The information obtained includes the empty physical volume counts by media type for the common scratch pool and each defined pool.

Refer to “Out of Physical Volumes” on page 391 for more information about how to handle an out of scratch situation. For more information about BVIR, refer to 8.6.5, “Interpreting the BVIR Response Data” on page 449. For more information about the Host Console Request function, refer to 2.3.6, “Host Console Request” on page 51.

3.5.3 Data compression

When writing data to a virtual volume, the host compression definition is honored. Compression is turned on or off by the JCL parameter DCB=TRTCH=COMP (or NOCOMP), the Data Class parameter COMPACTION=YES|NO, or the COMPACT=YES|NO definition in the DEVSUPxx PARMLIB member. The TRTCH parameter overrides the Data Class definition and both override the PARMLIB definition.

We strongly recommend that you check your definitions to make sure that you specify compression when writing data to the TS7700 in order to achieve the optimum throughput.

3.5.4 Secure Data Erase

Expired data on a physical volume remains readable until the volume has been completely overwritten with new data. Some customers are concerned that a court order could expose them to liability and cost to be able to find an old version of a data volume. Another concern is security of old data.

TS7700 Release 1.3 adds physical volume erasure on a physical volume pool basis controlled by an additional reclamation policy. It utilizes the Outboard Policy Management (OPM) feature which is standard on the TS7700. With the secure data erase function, all reclaimed physical volumes in that pool are erased by writing a random pattern across the whole tape prior to being reused. In the case of a physical volume that has encrypted data, the erase might involve just “shredding” the encryption keys on the volume to accomplish the erasure. A physical cartridge is not available as a scratch cartridge as long as its data is not erased.

The secure data erase function is performed in one of two ways:

- For an unencrypted physical volume, or a physical volume that is encrypted, but its previous use was unencrypted, the erasure is performed by writing a random data pattern on the physical volume being reclaimed. In the case of the encrypted volume, the encrypted data keys are also erased.
- For an encrypted physical volume whose previous use was also encrypted, the physical volume is erased by erasing the encrypted data keys. By erasing the data keys, the data on the volume cannot be decrypted, rendering it “erased”. This is a much faster method because only the encrypted data keys need to be erased. This process takes only a few minutes compared to a multi-hour process for writing the random pattern across the whole tape.

Keep in mind, that if you choose this “erase” functionality, TS7700 needs a lot of time to “reclaim” every physical tape. So, TS7700 will need more time every day to do reclaim process.

As part of this data erase function, an additional reclaim policy is added. The policy specifies the number of days a physical volume can contain invalid logical volume data before the physical volume becomes eligible to be reclaimed. The data associated with a logical volume is considered invalidated as follows:

- A host has assigned the logical volume to a scratch category. The volume is subsequently selected for a scratch mount and data is written to the volume. The older version of the volume is now invalid.
- A host has assigned the logical volume to a scratch category that has the fast-ready attribute set, the category has a nonzero delete expired data parameter value, the parameter value has been exceeded, and the TS7700 has deleted the logical volume.
- A host has modified the contents of the volume. This could be a complete rewrite of the volume or appending to it. The new version of the logical volume will be migrated to a different physical location, and the older version is now invalid.

The TS7700 keeps track of the amount of active data on a physical volume. It starts at 100% when a volume becomes full. Although the granularity of the percent of full TS7700 tracks is 1/10%, it rounds down, so even one byte of inactive data will drop the percent to 99.9%. TS7700 keeps track of the time that the physical volume went from 100% full to less than 100% full by:

- Checking on an hourly basis for volumes in a pool with a non-zero setting
- Comparing this time against current time to determine if the volume is eligible for reclaim

This data erase function is enabled on a per pool basis. It is enabled when a nonzero value is specified for the data erase reclaim policy. When enabled, all physical volumes in the pool are erased as part of the reclaim process, independent of the reclaim policy under which the volume became eligible for reclaim.

Any physical volume that has a status of read-only is not subject to this function and is not designated for erasure as part of read-only recovery.

If you use the eject stacked volume function, no attempt is made to erase the data on the volume prior to ejecting the cartridge. The control of expired data on an ejected volume becomes a customer responsibility.

Volumes tagged for erasure cannot be moved to another pool until erased but they can be ejected from the library, because such a volume is usually removed for recovery actions.

The usage of the Move function of the LM will also cause a physical volume to be erased, even though the number of days specified has not yet elapsed. This includes returning borrowed volumes.

The TS7700 historical statistics are updated with the number of physical mounts for data erasure. The pool statistics are updated with the number of volumes waiting to be erased and the value for the days (number of days) until erasure reclaim policy.

As soon as you have decided to implement Secure Data Erase for a limited group of data separated on a dedicated pool, the number of additional reclamation tasks plus data erase tasks will increase. Less physical drives might be available even during times when you have inhibited reclamation.

The Inhibit Reclaim Schedule specification only partially applies to Secure Data Erase:

- No new cartridges are reclaimed during this time
- Cartridges already reclaimed could be erased during this time

This means that, although you do not allow reclamation during your peak hours to have all your drives available for recall and premigration, Secure Data Erase will not honor your settings and thus will run up to two concurrent erasure operations per physical drive type as long as there are physical volumes to be erased.

Because the first logical volume that expires triggers the physical volume to be erased, an almost full physical cartridge will be first reclaimed and second erased.

We highly recommend to group logical volumes that require secure erase after they are expired, in such a way that almost no unnecessary reclamation and subsequent erasure operations take place. Pooling by expiration date might help to reduce unnecessary reclamation. Although proper grouping reduces the amount of reclamation that needs to be done, it will not eliminate the erasure step.

3.6 Planning for encryption in the TS7700

The importance of data protection has become increasingly apparent with news reports of security breaches, loss and theft of personal and financial information, and government regulation. Encryption of backstore tapes helps control the risks of unauthorized data access without excessive security management burden or subsystem performance issues.

Encryption on the TS7700 is controlled on a storage pool basis. “Storage Group” and “Management Class” DFSMS constructs specified for logical tape volumes determine, through mapping in the Library Manager, which physical volume pools are used for the primary and backup (if used) copies of the logical volumes. The storage pools, originally created for management of physical media, have been enhanced to include encryption characteristics.

The encryption solution for tape virtualization consists of several components.

- Tape encryption solutions from IBM all use an Encryption Key Manager (EKM) as a central point from which all encryption key information is managed and served to the various subsystems. The EKM communicates with the TS7700 Virtualization Engine as well as tape libraries, control units, and open-systems device drivers.
- The TS1120 model E05 encryption-enabled tape drive, as the other common component to the various data encryption solutions, provides hardware that performs the cryptography function without reducing the data-transfer rate.
- The TS7700 Virtualization Engine provides the means to manage the use of encryption and what keys are used on a storage-pool basis. It also acts as a proxy between the tape drives and the EKMs, using Ethernet to communicate with the EKMs and Fibre Channel connections with the drives. Encryption support is enabled with Feature Code 9900.

TS7700 backend drive encryption was introduced with Release 1.2 of the TS7700. Although there are no library manager code changes to support the function, there are other functions of the TS7700 that require a compatible level of library manager code. There are no host software updates required for this function because the TS7700 controls all aspects of the encryption solution.

Tapes encrypted in the TS7700 backstore use a “wrapped key” model. The data on each cartridge is encrypted with a random 256-bit Advanced Encryption Standard (AES-256) Data Key (DK). The Data Key is stored on the cartridge in an encrypted, or *wrapped*, form. Four instances of these wrapped data keys are stored on each cartridge.

While the feature for encryption support is customer-installable, some of the prerequisites might require additional hardware installation or configuration by an IBM Service representative.

3.6.1 EKM

Your EKMs should be installed, configured, and operational before installing the encryption feature on the TS7700. You should also create the certificates and keys you plan to use for encrypting your backstore tape cartridges.

Although you can run with a single EKM, it is strongly suggested that you have two EKMs for use by the TS7700. Each EKM should have all the required keys in their respective keystores. The EKMs should have independent power and network connections to maximize the chances that at least one of them is reachable from the TS7700 when needed. If the TS7700 cannot contact either EKM when required, you might lose access to migrated logical volumes temporarily and cannot move logical volumes in encryption-enabled storage pools out of cache.

See *IBM System Storage Tape Encryption Solutions*, SG24-7320, and the *IBM Encryption Key Manager component for the Java platform Introduction, Planning, and User's Guide,* GA76-0418, for details about installing and configuring your EKMs.

Because the TS7700 maintains TCP/IP connections with the EKMs at all times, it is recommended that the EKM configuration file have the following setting to prevent the EKM from timing out on these always-on connections:

```
TransportListener.tcp.timeout = 0
```

3.6.2 IBM Tivoli Key Lifecycle Manager

The IBM Tivoli® Key Lifecycle Manager is an enhancement and follow-on to the Encryption Key Manager that you can use to encrypt data with the TS1130 and TS1120 tape drives. TS1130 tape drives are only supported starting with TS7700 R1.5. Similar to the Encryption Key Manager, the Tivoli Key Lifecycle Manager serves data keys to the tape drive. The first release of Tivoli Key Lifecycle Manager focuses on ease of use and provides a new graphical interface (GUI) to help with the installation and configuration of the key manager. It also allows you to create and manage the key encrypting keys (certificates). If you already use the existing Encryption Key Manager, you easily can migrate to the new Tivoli Key Lifecycle Manager. Host software, including the IOS proxy, has no direct knowledge of which key manager is used. Specification of either key manager is handled the same way, using the EKM keyword on existing IOS commands. For additional information about Tivoli Key Lifecycle Manager, refer to the IBM Tivoli Key Lifecycle Manager Information Center:

http://www-01.ibm.com/software/tivoli/products/key-lifecycle-mgr/

3.6.3 Tape drives

Because data encryption is performed on the tape drives themselves, TS1120 Model E05 encryption-capable tape drives must be attached to the TS7700. They also must be running in native (E05) mode rather than J1A emulation mode.

If you have 3592 Model J1A drives attached to the TS7700, they should be detached. The TS7700 does not allow a mixture of drive types to be used. The J1A drives can be redeployed in other subsystems or used as direct-attached drives for open-systems hosts. If you have a mixture of J1A and E05 drives attached to your TS7700 and cannot detach the J1A drives right away, you can proceed as long as you have a minimum of four encryption-capable E05 drives attached. Be aware, though, that the J1A drives will not be used by the TS7700 after the E05 drives are put into native mode.

All TS1120 tape drives with Feature Code 5592 or 9592 are encryption-capable. FC5595 must be ordered against each tape drive to enable encryption on the drive, and FC9900 must be ordered against the TS7740 to allow encryption for the TS7740.

3.6.4 TS7700 Virtualization Engine

The TS7700 must be running Release 1.2 or higher. Feature code 9900 must be installed to access the encryption settings.

The TS7700 must not be configured to force the TS1120 drives into “J1A” mode. This setting can only be changed by your service representative. If you need to update the microcode level, be sure the service representative checks and changes this setting if needed.

3.6.5 Installation

Feature 5595 on the drives provides explicit instructions on setting up the drives, and FC9900 on the TS7740 provides these instructions for activating the feature on the TS7740. Briefly, the installation steps are:

1. Determine which TS1120 E05 drives are attached to the TS7740.
2. Verify that the tape drives are encryption-capable.

3. Verify that the drives are configured as encryption-enabled by specifying the "System-Managed" encryption method.
4. Install the Feature Code 9900 License Key.
5. Set up the EKM IP addresses and port information.
6. Verify connection to the EKMs.

It is essential that you configure the tape drives for System-Managed encryption. The TS7740 uses the drives in this mode only and does not support Library-Managed or Application-Managed encryption.

When the TS7740 is using drives for encrypted physical tape volumes, it puts drives that are not properly enabled for encryption offline to the subsystem. 3494 and TS3500 library operators should be made aware to leave TS7740-attached drives in System-Managed encryption mode at all times so that drives are not taken offline.

3.7 Tape analysis and sizing the TS7740

In this section we document the process of using various tools to analyze your current tape environment and to size the TS7740 to meet your requirements. With this link you get access to a tools library that offers a bunch of jobs to analyze your current environment and a procedure to unload specific SMF records for a comprehensive sizing with *BatchMagic*, which has to be done by your IBM representative.

3.7.1 IBM tape tools

Most of the IBM tape tools are available to customers, but some, such as BatchMagic, are available only to IBM personnel. You can download the tools that are generally available using the following link:

`ftp://ftp.software.ibm.com/storage/tapetool/`

To files that you can download include text, PDF, and EXE files. First, you can open the OVERVIEW.PDF file to get a brief description of the different tool jobs. All jobs are found in the IBMTOOLS.EXE file. This file is self-extracting zipped file and, after you download it, it expands into the following separate files:

IBMJCL.BIN	JCL for current tape analysis tools.
IBMCNTL.BIN	Parameters needed for job execution
IBMLOAD.BIN	Load library for executable load modules
IBMPAT.BIN	Data pattern library, only needed if you run the QSAMDRVR utility

In this section, we describe the procedures that can help to detect weak points in the current tape environment. These weak points can negatively influence overall performance of the TS7740 by causing bad block size (that is, a block size smaller than 16 kilobytes) or small compression ratios.

Tape compression analysis for TS7740

By analyzing the Miscellaneous Data Records (MDR) from SYS1.LOGREC or the EREP history file, you can see how well current tape volumes are compressing.

The following job stream has been created to help analyze these records. See the installation procedure in member $$INDEX.

- EREPMDR: JCL to extract MDR records from EREP history file.
- TAPECOMP: A program that reads either SYS1.LOGREC or the EREP history file and produces reports on current compression ratios and MB transferred per hour.

MDR analysis for bad TS7740 block sizes

Again, by analyzing the MDR from SYS1.LOGREC or the EREP history file, you can identify tape volumes that are writing small blocks to the TS7740 and causing extended job runtimes.

The following job stream has been created to help analyze these records. See the installation procedure in member $$INDEX:

- EREPMDR: JCL to extract MDR records from EREP history file.
- BADBLKSZ: A program that reads either SYS1.LOGREC or the EREP history file, finds volumes writing small block sizes, then gets the job name and data set name from a Tape Management System copy.

Data collection and extraction

To correctly size the TS7740, you have to analyze your current workload. The SMF records that are required to perform the analysis are record types 14, 15, 21, and 30.

You should collect the stated SMF records for all z/OS systems that share the current tape configuration and that can have data migrated to the TS7740. The data collected should span one month (to cover any month-end processing peaks) or at least those days that represent the peak load in your current tape environment. You need to check in SYS1.PARMLIB in member SMF to see whether the required records are being collected. If they are not being collected, arrange for collection.

The steps shown in Figure 3-8 on page 120 are as follows:

Step 1: The TMS data and SMF data collection using the FORMCATS and SORTSMF procedure selects only the required tape processing-related SMF records and the TMS catalog information.

Step 2: The files created are compressed by BMPACKT and BMPACKS procedures.

Step 3: Download the packed files (ZIPPED format) to your PC and send them per e-mail to your IBM representative or burn them on a CD if the size is larger than allowed for transmission per e-mail.

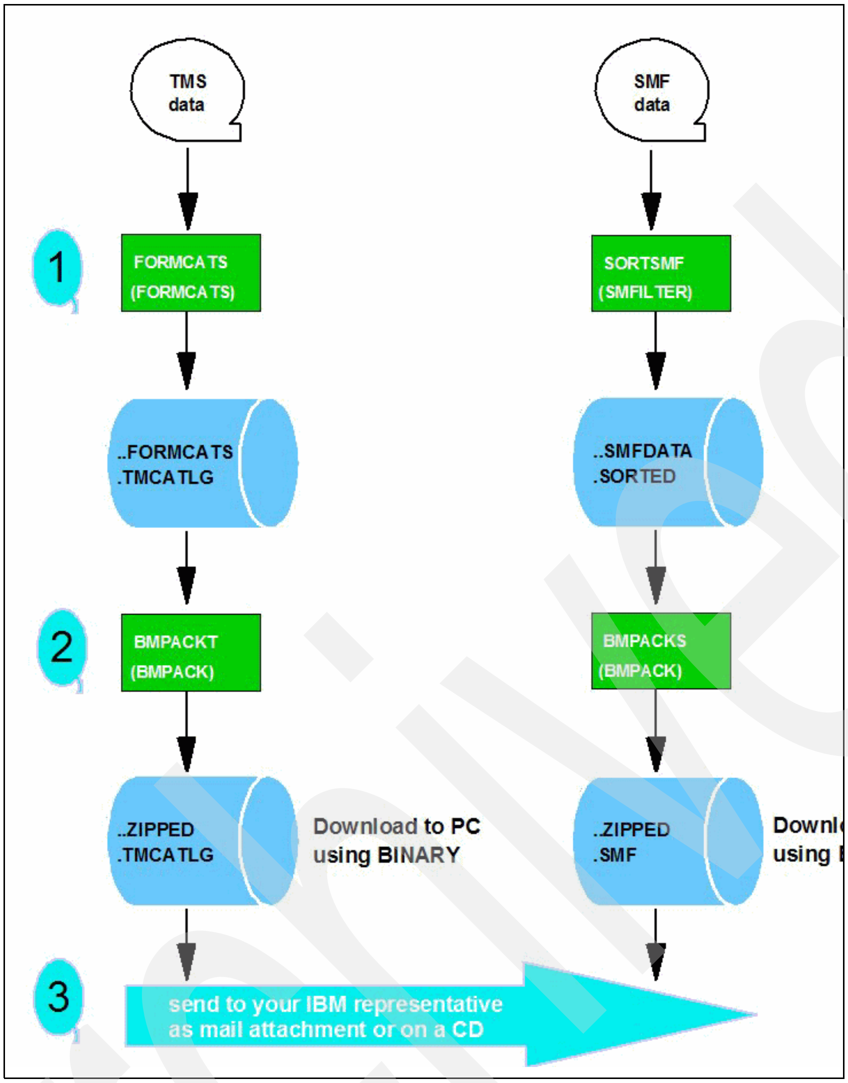

Figure 3-8 Unload process for TMS and SMF data

In addition to the extract file, the following information is useful for sizing the TS7740:

- Number of volumes in your current tape library

  This should include all your tapes (those in automated libraries, on shelves, and off site). If you provide the unloaded Tape Management Catalog data, you do not have to collect the number of volumes.

- Criteria for identifying volumes

  Because volumes that go off site, that are used as backup, or that should be ignored are of particular interest, their identification is important. Identifiers such as high-level qualifiers (HLQs), program names, or job names, must be documented.

- Number and type of tape control units installed

  This information provides a good understanding of the current configuration and will help you identify the reasons for any apparent workload bottlenecks.

- Number and type of tape devices installed

  This information will also help you identify the reasons for any apparent workload bottlenecks.

- Number and type of host channels attached to tape subsystems

  This information will also help you identify the reasons for any apparent workload bottlenecks.

3.7.2 BatchMagic

This tool gives a comprehensive view of the current tape environment as well as predictive modeling on workloads and technologies. The general methodology behind this tool involves analyzing customer SMF type 14, 15, 21, and 30 records, as well as data extracted from the customer tape management system. The TMS data is required only if you want to make a precise forecast of the cartridges to be ordered based on the current cartridge utilization that is stored in the TMS catalog.

A BatchMagic "run" involves data extraction, grouping data into workloads and then targeting workloads to individual or multiple IBM tape technologies. Batch Magic examines Tape management System catalogs and projects cartridges required with new technology and it models the operation of a TS7740 and 3592 drives and projects required resources. The reports from BatchMagic give a clear understanding of your current tape activities, and even more important, make projections for a TS7740 solution together with its major components such as 3592 drives, which cover your overall sustained and peak throughput requirements.

BatchMagic is specifically for IBM internal and Business Partner use.

3.7.3 Workload considerations

Like the prior generation of IBM Virtual Tape Server products, the TS7700 Virtualization Engine appears to the host systems as sixteen 3490E strings with a total of 256 devices attached. Any data that can reside on a 3480/3490/3590/3592, prior generations of VTS systems, or cartridges from other vendors can reside on the TS7740 Virtualization Engine. However, processing characteristics of workloads differ, so some data is more suited for the TS7700 Virtualization Engine than other data. This topic highlights some of the considerations to bear in mind when deciding what workload to place in the TS7700 Virtualization Engine.

- Throughput

  The TS7700 Virtualization Engine has a finite bandwidth capability, as does any other device attached to a host system, so you should plan to put workloads into the TS7700 Virtualization Engine that fit within its capabilities. With the 4 Gb FICON channels and large TS7700 Cache, there are few workloads that would not be suitable for the TS7700 Virtualization Engine based on throughput.

- Drive concurrency

  The TS7700 Virtualization Engine appears to the host operating system as 256 3490E drives. If there are periods of time during the day when your tape processing jobs are limited by drive availability, the TS7700 Virtualization Engine can help your processing considerably. The design of the TS7700 Virtualization Engine allows transparent access to multiple logical volumes on the same stacked physical volume, because access to the logical volumes is solely through the TS7700 Cache. If you need access to more than one logical volume on a physical volume, it is provided without requiring any user involvement, unlike some alternatives, such as stacking by using JCL.

- Cartridge capacity utilization

  One of the key benefits of the TS7700 Virtualization Engine is its ability to fully utilize the capacity of the 3592 cartridges independent of the data set sizes written and to manage that capacity effectively without host or user involvement. A logical volume can contain up to 4 GB of data (12 GB assuming data compressibility of 3:1) using the extended logical volume sizes. The actual size of a logical volume is only the amount of data written by the host. So even if an application only writes 20 MBs to a 4 GB volume, only the 20 MB is kept in the TS7700 Cache or on a physical volume. Very large data sets might gain little from the TS7700 Virtualization Engine's ability to stack data, so you might want to leave them on native 3590 or 3592 cartridges.

- Volume caching

  Often, one step of a job writes a tape volume and a subsequent step (or job) reads it. A major benefit can be gained using the TS7700 Virtualization Engine as the data is cached in the TS7700 Cache, which effectively removes the rewind time, the robotics time, and load or thread times for the mount.

  Figure 3-9 shows the effect that a TS7700 can have on a job and drive assignment as compared to a native drive. The figure is a freehand drawing—it is not to scale. It shows typical estimated elapsed times for elements that would make up the reading of data from a tape. When comparing the three timelines in Figure 3-9, notice that the TS7700 Cache hit time line does not include robotics, load, or thread time at the beginning of the time line, nor does it include any rewind or unload time at the end.

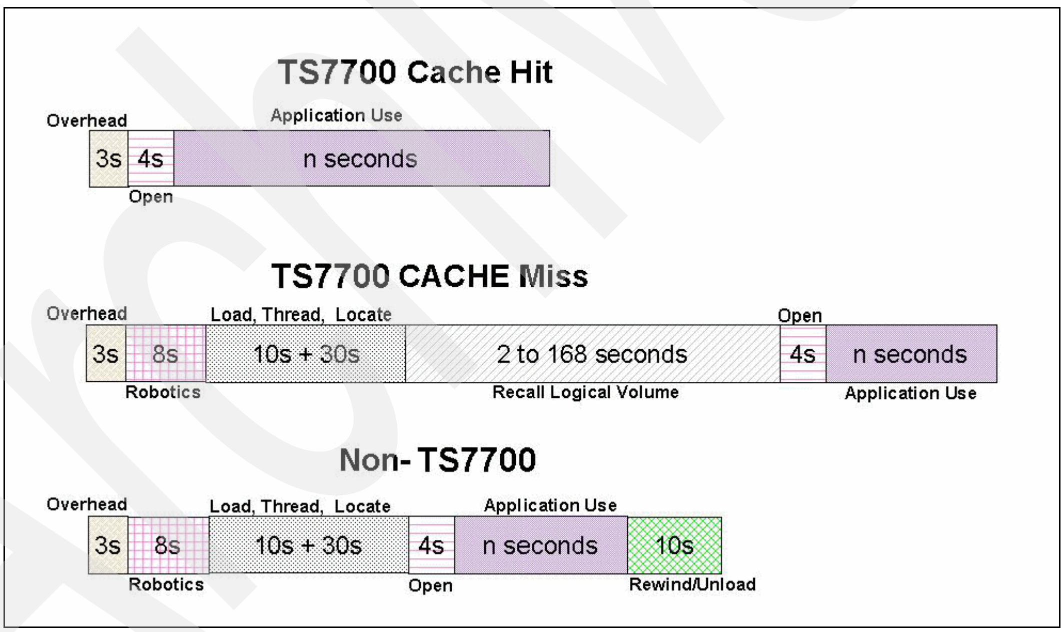

Figure 3-9 Tape processing time comparison (not to scale)

In this example, the TS7700 Cache hit results in a savings in tape processing elapsed time of 40 seconds.

The time reduction in the tape processing has two effects. First, it reduces the elapsed time of the job processing the tape. Second, it frees up a drive earlier so the next job needing a tape drive can access it sooner, as there is no rewind or unload and robotics time after closing the data set.

When a job attempts to read a volume that is not in the TS7700 Cache, the logical volume is recalled from a stacked physical volume back into the cache. When a recall is necessary, the time to access the data is greater than if it were already in the cache. The size of the cache and the use of the cache management policies can reduce the number of recalls. Too much recall activity can negatively impact overall throughput of the TS7700 Virtualization Engine.

- Scratch mount times

  When a program issues a scratch mount to write data, the TS7700 Virtualization Engine completes the mount request without having to recall the logical volume into the cache. For workloads that create a lot of tapes, this significantly reduces volume processing overhead times and improves batch window efficiencies. The effect of using the Fast-Ready attribute on TVC performance of scratch mounts is fast, because no physical mount is required. The performance for scratch mounts is the same as for TVC read hits, which are compared with non-TS7700 in Figure 3-9 on page 122.

- Disaster recovery

  The TS7700 Virtualization Engine's grid configuration is a perfect integrated solution for your disaster recovery data. The TS7700 clusters in a Multi Cluster Grid can be separated over long distances and interconnected using an I/P infrastructure to provide for automatic data replication. Data written to a local TS7700 Virtualization Engine is accessible at the remote TS7700 Virtualization Engine just as though it was created there. Flexible replication policies make it easy to tailor the replication of the data to your business needs.

  The Copy Export function, introduced with TS7700 release 1.3 for a Single Cluster Grid and release 1.4 for a Multi Cluster Grid, provides another disaster recovery method. The Copy Exported physical volumes can be used in an empty TS7700 to recover from a disaster. Refer to Chapter 2.3.4, “Copy Export” on page 46 for more details.

- Multifile volumes

  If you stack multiple files onto volumes today, by using JCL constructs or some other method, the reason you are stacking is most likely to better use cartridge capacity. Automatic utilization of physical cartridge capacity is one of the primary attributes of the TS7700 Virtualization Engine. Therefore, you should find that in many cases manual stacking of data sets onto volumes is no longer required. If you are planning a new application that would have used JCL to stack data sets onto a volume, the TS7700 Virtualization Engine makes this JCL step unnecessary. Multifile volumes moved to the TS7700 Virtualization Engine can also work without changing the stacking. However, the TS7700 Virtualization Engine recalls the complete logical volume to the TS7700 Cache if the volume is not in cache, rather than moving each file as you access it. Therefore, in some cases, it can be advantageous to let the TS7700 Virtualization Engine do the stacking automatically for you. It can save you manual management overhead and, in some cases, host CPU cycles, host channel bandwidth, DASD space, or a combination of all of these.

- Interchange or offsite storage

  As currently delivered, the TS7700 Virtualization Engine does not support a capability to remove a stacked volume to be used for interchange. Native 3490, 3590, or 3592 tapes are better suited to your data for interchange. The Copy Export function can be used for off-site storage of data for the purposes of disaster recovery. Refer to 2.3.4, “Copy Export” on page 46 for more details.

With the wide range of capabilities that the TS7700 virtualization Engine provides, unless your data sets are very large or require interchange, it is likely that the TS7700 Virtualization Engine is a very suitable place to put your data.

3.8 Education and training

If you have never had an IBM library installed before, plan education for your storage administrators as well as for the operating personnel. Education topics should include handling of the 3494 Web Specialist, TS3500 Specialist and the ETL Specialist, usage of the TSSC, and handling of the library itself. If you are familiar with the IBM 3494, the Library Manager in the 3953 Tape System will be familiar to you. However, there is an additional set of activities required to prepare the TS3500 logical libraries before you get to the setup of the Library Manager. Therefore, we recommend additional training for the usage of the TS3500 Tape Library Specialist, if appropriate, and TS7700 Management Interface.

The amount of education and training your staff requires on the TS7700 depends on a number of factors, including these:

- Are you installing the TS7700 in an existing IBM 3494 Tape Library environment?
- Are you installing the TS7700 in an existing IBM TS3500 Tape Library environment?
- Are you installing the grid TS7700 in an existing library?
- Are both the TS7700 and the library new to your site?
- Are you using BTLS or SMS tape to manage your library?
- Are the library and the TS7700 shared among multiple systems?
- Do you have existing tape drives at your site?

3.8.1 Adding an IBM TS7700 to an existing IBM 3494

If you are adding a TS7700 to an existing 3494 containing a TS7700, the training needed for your operators, system programmers, and storage administrators is minimal. Although the Library Manager posts operator interventions that are specific to the TS7700, the messages posted to the host are not new. The operator intervention and help pull-downs in the Library Manager have been updated to contain these TS7700-specific interventions and the actions necessary to resolve the conditions.

As an existing IBM 3494 we recommend that training for your operations staff include the following:

- IBM 3494 Library Manager-specific TS7700 functions
- ETL Specialist-specific TS7700 functions
- TS7700 Management Interface

We recommend that your storage administrators and system programmers receive the same training as the operations staff, plus:

- Software choices and how they affect the TS7700
- Disaster recovery considerations

3.8.2 Adding an IBM TS7700 to an existing IBM TS3500

If you are adding a TS7700 to an existing TS3500 containing a TS7700, the training needed for your operators, system programmers, and storage administrators is minimal. Although the Library Manager posts operator interventions that are specific to the TS7700, the messages posted to the host are not new. The operator intervention and help pull-downs in the Library Manager have been updated to contain these TS7700-specific interventions and the actions necessary to resolve the conditions.

As an existing IBM TS3500 we recommend that training for your operations staff include the following:

- IBM 3953 Library Manager-specific TS7700 functions
- ETL Specialist-specific TS7700 functions
- TS7700 Management Interface

We recommend that your storage administrators and system programmers receive the same training as the operations staff, plus:

- Software choices and how they affect the TS7700
- Disaster recovery considerations

3.8.3 New TS7700 in an existing 3494 without a VTS

If you are adding a new TS7700 to an existing 3494 but you have never had a VTS or TS7700 in the library, the education and training for your staff should include the TS7700 training items listed in 3.7, “Tape analysis and sizing the TS7740” on page 118, as well as the following topics:

- Hands-on training with the Web Specialist functions related to the TS7700
- Management Interface hands-on

We recommend that you review the following:

- *IBM TotalStorage Enterprise Automated Tape Library 3494 Operator Guide*, GA32-0449
- *IBM TotalStorage Enterprise Automated Tape Library 3494 Introduction and Planning Guide*, GA32-0448
- In this chapter, 3.7, “Tape analysis and sizing the TS7740” on page 118
- TS7700 Information Center

Education is a service offering of IBM Global Services. Also available are the services for library and TS7700 implementation described in the following sections.

3.8.4 New TS7700 and new TS3500 with IBM 3953 Library Manager

If you are installing a new TS3500/3953 with a TS7700, the education and training for your staff should include the TS7700 training items listed in 3.7, “Tape analysis and sizing the TS7740” on page 118, as well as the following topics:

- Role of the Library Manager and how the operator interacts with it
- Proper procedures for changing modes and states
- Proper procedures for entering and ejecting tape cartridges in the library
- Performing Manual Mode operations
- Hands-on training with the Web Specialist
- Management Interface hands-on

We recommend that you review the following:

- *IBM System Storage TS3500 Tape Library Introduction and Planning Guide*, GA32-0559
- *IBM System Storage TS3500 Tape Library Operator Guide*, GA32-0560
- In this chapter, 3.7, “Tape analysis and sizing the TS7740” on page 118
- TS7700 Information Center

Education is a service offering of IBM Global Services. Also available are the services for library and TS7700 implementation described in the following sections.

3.8.5 Sharing the TS7700 with other systems within a TS3500

In addition to the material in 3.7, “Tape analysis and sizing the TS7740” on page 118, refer to *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

3.8.6 Implementation services

A range of services is available to assist you with your TS7700. IBM can deliver end-to-end storage services to help you throughout all phases of the IT life cycle, including:

- Assessment

  Provides an analysis of the tape environment and an evaluation of potential savings and benefits of installing new technology, such as tape automation, virtual tape, and tape mounting management.

- Planning

  Assists in the collection of information required for tape analysis, analysis of the customer’s current environment, and the design of the ATL environment, including coding and testing of customized Data Facility Storage Management Subsystem (DFSMS) Automatic Class Selection routines.

- Implementation
  - TS7700 implementation provides technical consultation, software planning, and assistance and operation education to customers implementing an IBM TS7700. Options include Data Analysis and SMS Tape Design for analysis of tape data in preparation and design of a DFSMS tape solution, New Allocations for assistance and monitoring of tape data migration through new tape volume allocations, and Static Data for migration of existing data to a TS7700 or traditional automated tape library. (See also Tape Copy Services.)
  - Automated Tape Library (ATL) implementation provides technical consultation, software planning assistance, and operational education to customers implementing an ATL.
  - Tape Copy Service performs copying of data on existing media into an ATL. This service is generally performed subsequent to an Automated Library, TS7700, or grid implementation.

- Support

  SupportLine provides access to technical support professionals who are experts in all IBM tape products.

The people of IBM Global Services offer exceptional industry experience, resources, and capabilities across multiple product lines. Our service offerings are based on methodologies that have been time-tested and proven with hundreds of successful engagements in locations all over the world. Put your trust in a worldwide e-business leader. IBM Global Services offers comprehensive, integrated industry solutions that will deliver ongoing value to help keep you competitive in today’s networked economy.

For more information about storage services and IBM Global Services, contact your IBM sales representative, or visit:

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

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

IBM Integrated Technology Services

IBM services include business consulting, outsourcing, hosting services, applications, and other technology management tasks.

These services help you learn about, plan, install, manage, or optimize your IT infrastructure to be an on Demand business. They can help you integrate your high-speed networks, storage systems, application servers, wireless protocols, and an array of platforms, middleware, and communications software for IBM and many non-IBM offerings. IBM is your one-stop shop for IT support needs.

For details about available services, contact your IBM representative or visit:

http://www.ibm.com/services/

For details about available IBM Business Continuity and Recovery Services, contact your IBM representative or visit

http://www.ibm.com/services/continuity

For details about education offerings related to specific products, visit

http://www.ibm.com/services/learning/index.html

Select your country, and then select the product as the category.

3.9 TS7700 Functionality by release level

This section lists when the various functions of the TS7700 Virtualization Engine were made available as well as the minimum code levels required.

3.9.1 Release 1.0

This is the list of functionality in the original release of the TS7700 Virtualization Engine which became generally available on 29 September 2006.

- 3592 native J1A and E05 emulating J1A
- Attachment to TS3500 (3584) libraries
- 128 virtual drives
- 128 logical host paths per FICON adapter
- 500,000 virtual volumes
- Two or four 4 Gb FICON adapters
- 6 TB cache
- Two-Cluster Grid
- Requires 534.03 LM code
- Requires TS7700 microcode level 8.0.0.203.

3.9.2 Release 1.1

This is a list of the functionality added in the second release of the TS7700 Virtualization Engine which became generally available on 26 January 2007.

- 256 virtual drives
- Autonomic Ownership Takeover Manager
- 3592 E05 drives in native mode
- JB media
- 3 TB cache configuration

- 2 TB to 6 TB Cache Increments (FC5267)
- RPQ Q8B3409 - Fiber Ethernet Card for Grid IP Connection
- RPQ Q8B3411 (ordered on B10 or B20) to migrate standalone B10/B20 attached to 3584 to TS7740 (attached to same 3584)
- Requires 534.23 LM code
- Requires 8.0.1.24 TS7700 code
- For z/OS, software APAR OA19061 is required.

3.9.3 Release 1.2

This is a list of the functionality added in the third release of the TS7700 Virtualization Engine which became generally available on 09 March 2007.

- Out of Band Encryption by Pools
- Broadband Call Home (requires TSSC code 4.2.5)
- Requires 534.33 LM code
- Requires 8.2.0.33 TS7700 code

This is a list of the functionality added in a subsequent release of the third release for the TS7700 Virtualization Engine which became generally available on 18 May 2007.

- Attachment to 3494 libraries
- RPQ Q8B3411 (ordered on B10 or B20) to migrate B10/B20 PtP attached to 3584 to TS7740 Grid (attached to same 3584), B10/B20 standalone attached to 3494 to TS7740 standalone (attached to same 3494), and B10/B20 PtP attached to 3494 to TS7740 Grid (attached to same 3494)
- Requires 534.31 LM code
- Requires 8.2.0.27 TS7700 code

3.9.4 Release 1.3

This is a list of the functionality added in the fourth release of the TS7700 Virtualization Engine which became generally available on 31 August 2007:

- Three-Cluster Grid
- 1 000 000 virtual volumes
- Select B10/B20 VTS to TS7700 migration and merge scenarios are now Feature Codes (0522 and 0523)
- Select TS7700 Grid MESs
- Fiber Ethernet Card for Grid IP Connection is now a Feature Code (1031)
- Host Console Query
- Copy Export (standalone TS7700 only)
- Automatic Read-Only Recovery
- Secure Data Erasure
- Secure Data Erasure with Encryption
- Cache increments MES (2 TB - 6 TB in 1TB increments)
- Performance increments (100 MB/s - 600 MB/s in 100 MB/s increments)
- Two to four FICON card MES
- LW/SW and 4 km/10 km FICON feature conversions
- Requires 535.02 LM code
- Requires 8.3.0.106 TS7700 code
- For z/OS, software APAR OA20065 is required.

3.9.5 Release 1.4

This is a list of the functionality added in the fifth release of the TS770 Virtualization Engine which will become generally available on 23 November 2007:

- Copy Export for Grid
- Single cache drawer configuration
- 256 logical host paths per FICON adapter
- Additional Cache increment MES (1 TB - 6 TB in 1 TB increments)
- Requires 535.xx LM code
- Requires 8.4.x.x TS7700 code

3.9.6 Release 1.4a

This is a list of the functionality added in the fifth release of the TS770 Virtualization Engine which will become generally available on 25 April 2008:

- Dynamic Grid Network Balancing
- Host Copy Control
- Removal of a cluster from a grid
- Cluster Cleanup
- Upgrading a single cache unit configuration to a two cache unit configuration (FC5269)

4

Hardware implementation

In this chapter we describe the hardware-related implementation steps for the IBM System Storage TS7700 Virtualization Engine. Because the TS7700 Virtualization Engine works only with a Tape Library and a Library Manager, it must be installed together with either the IBM 3494 Tape Library or the IBM System Storage TS3500 Tape Library and the IBM System Storage 3953 Tape System. The IBM 3953 Tape System consists of an IBM 3953 Frame Model F05 and one or two IBM 3953 Library Manager models L05.

> **Note:** For this chapter, we use the term "Tape Library" when the subject applies to both the 3494 Tape Library and TS3500/3953 Tape Library. We use the specific library name when a subject is specific to a particular tape library. We use the term "Library Manager" when the subject applies to both the 3494 and 3953 Library Manager. We use the specific Library Manager name when the subject is specific to a particular Library Manager.

You can install the TS7700 Virtualization Engine together with your existing tape library as long as the existing Library Manager has only one existing IBM TotalStorage Virtual Tape Server (VTS) Model B10 or B20 attached, or one IBM TS7700 Virtualization Engine. You can also install a new TS3500/3953 and a new TS7700 Virtualization Engine at the same time. You can also add a TS7700 to an existing 3494 Tape Library by using existing frames or adding new D frames. Either way, you must have a 3494 or TS3500/3953 if you want to take advantage of the unique performance and functions that the IBM TS7700 Virtualization Engine offers.

We cover all implementation steps in this chapter that relate to the setup of the following:

- TS3500 Tape Library
- IBM 3953 Library Manager
- IBM 3494 Tape Library
- TS7700 Virtualization Engine

For details about host software implementation, refer to Chapter 5, "Software implementation" on page 193 where we describe the Hardware Configuration Definition (HCD) steps on the host as well as operating system-related definitions.

4.1 TS7700 implementation and installation considerations

The following sections discuss the implementation and installation tasks to set up the TS7700 Virtualization Engine.

4.1.1 Implementation tasks

The TS7700 Virtualization Engine implementation can be logically separated into three major sections:

1. TS7700 Virtualization Engine and Tape Library setup

   You use the TS7700, TS3500 (if applicable), and the Library Manager interfaces for these setup steps:

   – You define the physical properties of the TS7700 partition, such as physical tape drives and cartridges, using the TS3500 Tape Library Specialist, the Web browser interface to the TS3500 if you are using a TS3500.

   – You define logical volumes, management policies, and volume categories using the Library Manager console or the Enterprise Tape Library (ETL) Specialist, which is the Web browser interface of the Library Manager.

   – You use the TS7700 Management Interface (MI) to define specific customer settings, such as encryption, and to insert logical volumes into the TS7700.

   We describe these implementation steps in detail in this chapter.

2. Hardware I/O configuration definition

   This section relates to the system generation. It consists of processes such as FICON channel attachment to the host, HCD/IOCP definitions, and Missing Interrupt Handler (MIH) settings. This activity can be done prior to the physical hardware installation and be part of the pre-installation planning and activity.

   We describe these installation steps in Chapter 5, “Software implementation” on page 193.

3. TS7700 Virtualization Engine software definition

   Here you define the new tape library to the individual host operating system.

   In a System z environment with DFSMS/MVS, this includes updating DFSMS Automatic Class Selection (ACS) routines, Object Access Method (OAM), and your tape management system during this phase. You also define Data Class (DC), Management Class (MC), Storage Class (SC), and sTorage Group (SG) constructs and selection policies, which are passed to the TS7700.

   We describe these installation steps in Chapter 5, “Software implementation” on page 193.

These three groups of implementation tasks can be done in parallel or sequentially—the HCD and host definitions can be completed before or after the actual hardware installation.

4.1.2 Installation tasks

The tasks outlined in this section are specific to the simultaneous installation of a TS3500/3953 and a TS7700 Virtualization Engine. If you are installing a TS7700 in an existing 3494 or TS3500/3953 environment, some of these tasks might not apply to you:

- Hardware related activities (completed by your IBM System Service Representative):
  - Install the IBM TS3500 Tape Library and IBM 3953 Library Manager (if applicable).
  - Install any native drives that will not be TS7700 controlled.
  - Install the TS7700 Frame and corresponding D2x Frame(s) in the 3494 or TS3500 Tape Library.
- Host-specific activities (define drives to the host):
  - z/OS
  - z/VM and VM/ESA®
  - z/VSE
  - TPF and z/TPF
- Software-specific items:
  - Apply maintenance for the IBM 3494 or IBM TS3500/3953.
  - Apply maintenance for the TS7700.
  - Verify or update exits for the tape management system (if applicable) and define logical volumes to it.
  - Update the host definitions:
    - SYS1.PARMLIB members.
    - OAM parameters.
    - Define the TCDB.
    - Define or update SMS classes and groups.
    - Define or update ACS routines.
    - Install applicable OAM exits.
  - DFSMS/VM RMS:

    Consult *z/VM V5R2.0 DFSMS/VM Removable Media Services,* SC24-6090, and 5.4, “Software implementation in z/VM and z/VSE” on page 223.
- TS7700-specific installation items:
  - Define the TS7700 environment, using the Library Manager.
  - Define the logical VOLSER ranges for TS7700-owned logical volumes through the TS7700 Management Interface (MI).
  - Define the physical VOLSER ranges for TS7700-owned physical volumes to the Library Manager and TS3500 Tape Library (if applicable).
  - Insert TS7700-owned physical volumes in the Tape Library.

These tasks will be discussed further, including the recommended order of events, later in this chapter.

After your TS7700 Virtualization Engine is installed on the 3494 Tape Library or TS3500 Tape Library/3953 Library Manager, carry out these post-installation tasks:

- Schedule and complete operator training.
- Schedule and complete storage administrator training.

4.2 IBM TS3500 Tape Library definitions

We now describe how to implement the IBM TS3500 Tape Library in a System z environment. Use this section if you are implementing the TS7700 in a TS3500 library. Your IBM Systems Service Representative (SSR) performs the hardware installation of the library, its associated tape control units, the IBM 3953 Library Manager, and the frames. This does not require customer involvement other than the appropriate planning (see Chapter 3, "Pre-installation planning and sizing" on page 77 for details). The following topics are covered in this section:

- Defining a Logical Library
- Cartridge Assignment Policies
- Eight-Character-VOLSER support

After the SSRs have physically installed the library hardware, you can use the TS3500 Tape Library Specialist to set up the logical library, which then is managed by the IBM 3953 Library Manager, which is attached to the System z host.

Note that the steps described in the following section relate to the installation of a new IBM TS3500 Tape Library with all the required features, such as ALMS, already installed. If you are attaching an existing IBM TS3500 Tape Library—which is already attached to Open Systems hosts—to System z hosts as well, refer to *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789, for additional actions that might be required.

4.2.1 Defining a logical library

The IBM TS3500 Tape Library Specialist is required to define a logical library and perform the tasks below. Therefore, you should make sure that it is set up properly and working. For access using a standard-based Web browser, an IP address must be configured, which will be done initially using the IBM SSR during hardware installation at the TS3500 operator panel.

The Advanced Library Management System (ALMS) is required to define a logical library partition in the TS3500.

Ensure ALMS is enabled

You can check the status of ALMS with the TS3500 Specialist as shown in Figure 4-1.

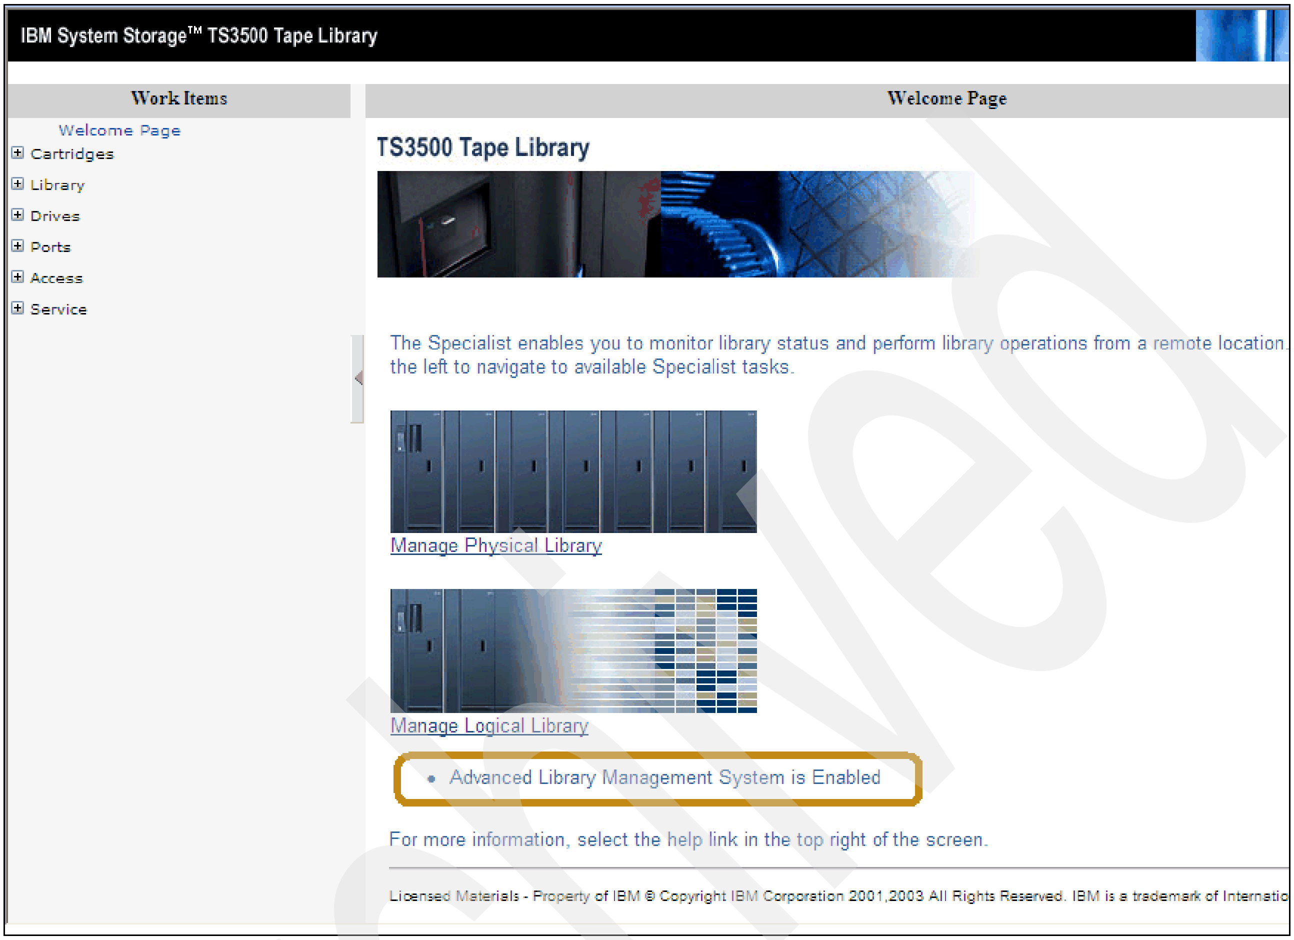

Figure 4-1 IBM TS3500 Specialist Welcome Page

As you can see at the bottom of the panel, ALMS is enabled for this IBM TS3500 Tape Library.

In addition, when ALMS is enabled, the work items under Manage Library show Disable ALMS. When ALMS is disabled, the work items display Enable ALMS. If ALMS is not enabled, select this work item. From the next panel, enable ALMS.

When enabling ALMS in an already partitioned IBM TS3500, the cartridges in the library will be correctly assigned to their respective logical libraries. When enabling ALMS in a single partition IBM TS3500 Tape Library, cartridges already residing inside the library also remain in this single partition and no longer become unassigned.

Prior to enabling ALMS, the ALMS License Key has to be entered using the TS3500 Tape Library operator panel, because ALMS is a chargeable feature.

Define a logical library with ALMS

This function is only valid and available if ALMS is enabled.

> **Note:** You can create or remove a logical library from the TS3500 Tape Library by using the Tape Library Specialist Web interface but not by using the operator panel at the IBM 3953 Library Manager.

From the main section of TS3500 Specialist Welcome Page, select **Manage Logical Library**. Alternatively, you can use the work items on the left side of the panel to navigate to the required panel by selecting **Library** → **Logical Libraries** as shown in Figure 4-2. From the Select Action drop-down menu, select **Create** and press **Go**.

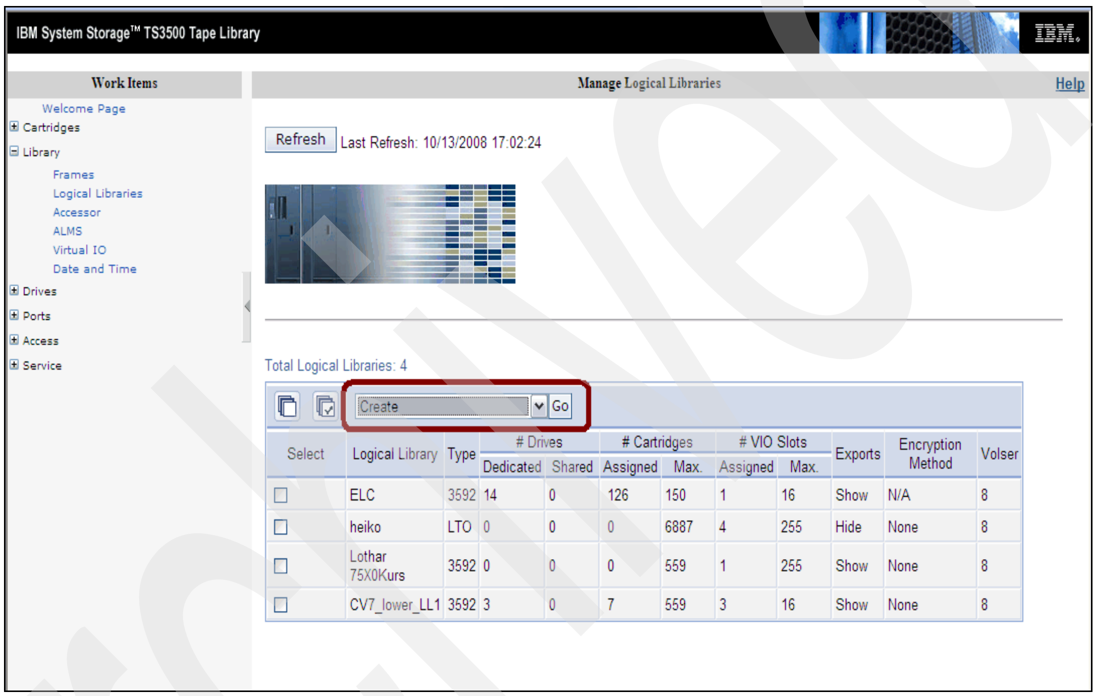

Figure 4-2 Create Logical Library starting panel

This causes an additional panel, Create Logical Library, to display; both panels are shown in Figure 4-3.

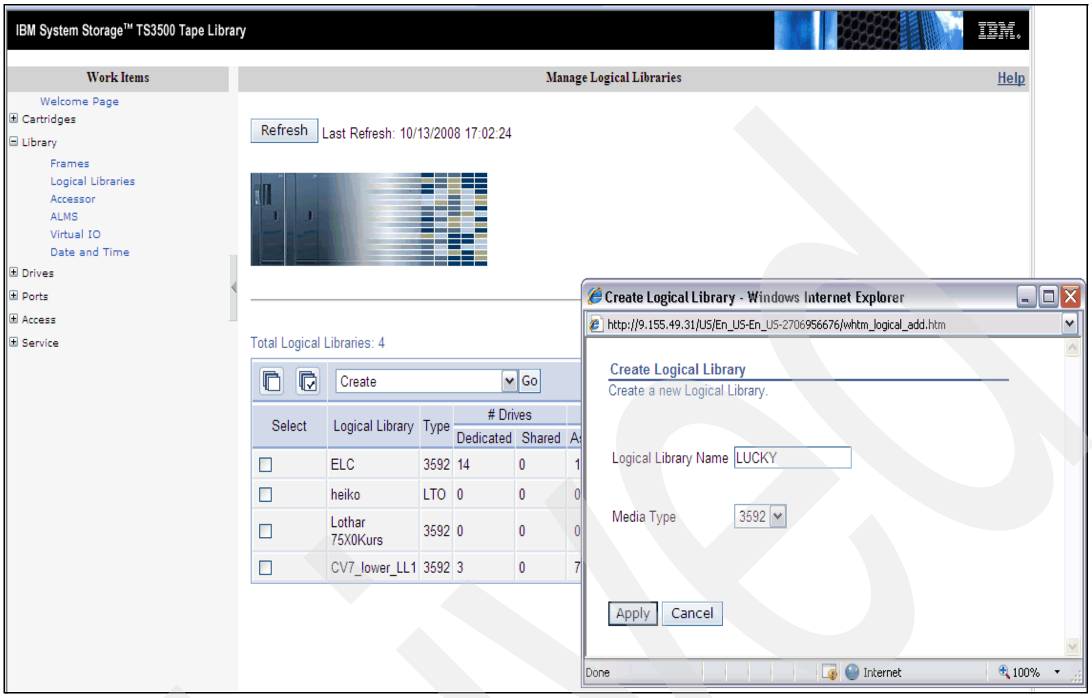

Figure 4-3 Create Logical Library pop-up windows

When you have defined your new logical library name and the Media Type 3592, click **Apply**, which creates the logical library and lists it on the Manage Logical Libraries panel.

After the logical library is created, you can use the work items on the left side of the panel to navigate to the required panel by selecting **Library** → **Logical Libraries** as shown in Figure 4-4. From the Select Action drop-down menu select **Detail** and press **Go**. You can see the element address range in the windows for the Logical Library Details. The starting element address needs to be recorded and needs to be given to the IBM System Service Representative for base installation and configuration of the IBM 3953 Library Manager. Because each created logical library will get a unique starting element address (except those logical libraries that existed before ALMS was enabled), this is the method by which the communication is synchronized between the TS3500 Tape Library and the IBM 3953 Library Manager.

The starting element address of each newly created logical library starts one element higher, such as:

- Logical Library 1: Starting SCSI element address is 1025
- Logical Library 2: Starting SCSI element address is 1026
- Logical Library 3: Starting SCSI element address is 1027

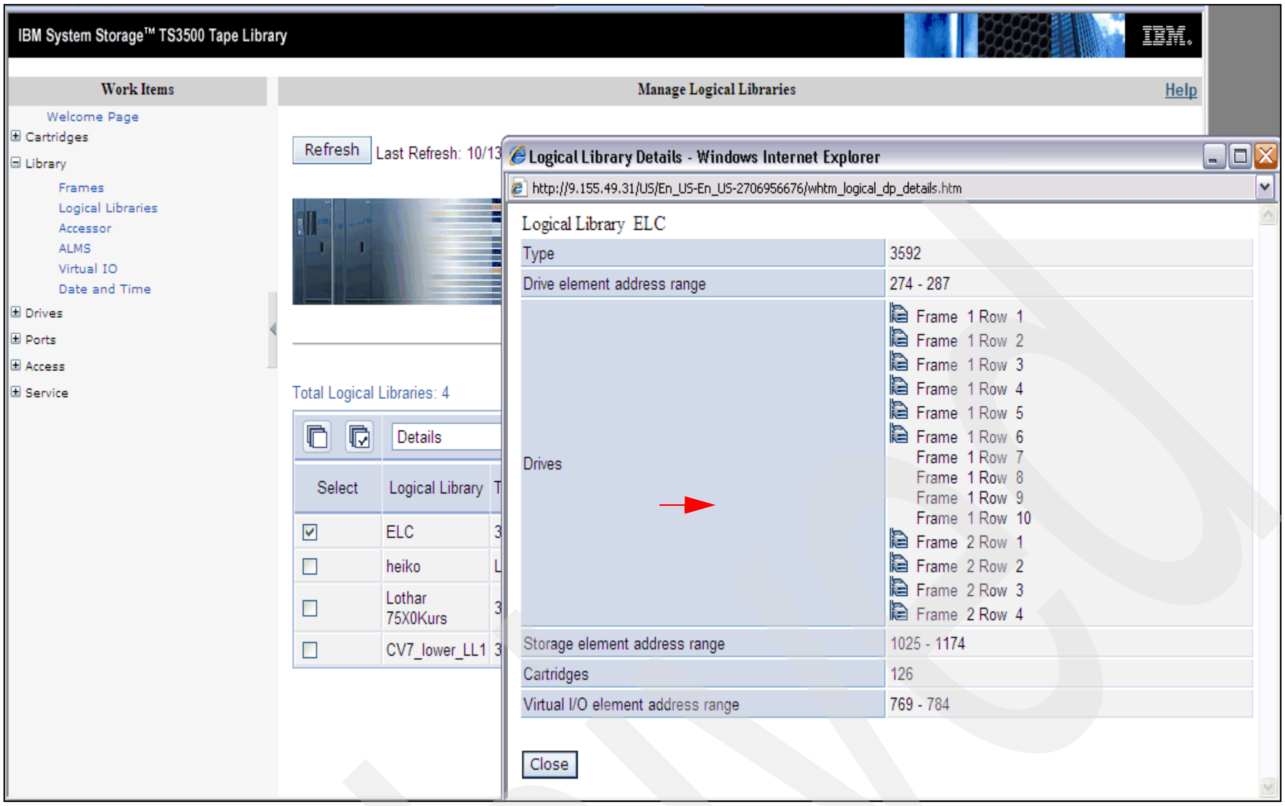

Figure 4-4 Recording of the starting SCSI element address of a Logical Library

Define cartridge slots

Define the maximum number of cartridge slots for the logical library. If multiple logical libraries are defined, you can define the maximum number of cartridge slots for each logical library. This allows a logical library to grow without a reconfiguration each time you want to add empty slots. To define the cartridge slots, click **Maximum Cartridges** on the Manage Logical Libraries panel, after you define the logical library. Figure 4-2 on page 136 shows a maximum number of 237 cartridges specified for logical library 1-01.

Add drives to the logical library

From the Manage Logical Libraries panel shown in Figure 4-2 on page 136, click **Drive Assignment**. Alternatively, you can use the work items on the left side of the panel to navigate to the requested Web page by selecting **Drives** → **Drive Assignment**.

Both these links take you to a filtering panel from where you can select to have the drives displayed **By Drive** or **by Logical Library** and, upon your selection, to a panel that allows you to add or remove a drive from a library configuration. It also enables you to add, remove, and share drives in a logical library, and change a control path.

Figure 4-5 shows the drive assignment panel of a logical library that has all drives assigned. Unassigned drives would appear in the Unassigned column with the box checked, so to assign them you could check the appropriate drive box under the Logical Library name and click **Apply**.

The small gray question mark at the top, right corner of the panel shown in Figure 4-5 provides you with extended Help information, which contains detailed explanation of all the fields and functions of the panel. The other TS3500 Specialist panels provide similar help support.

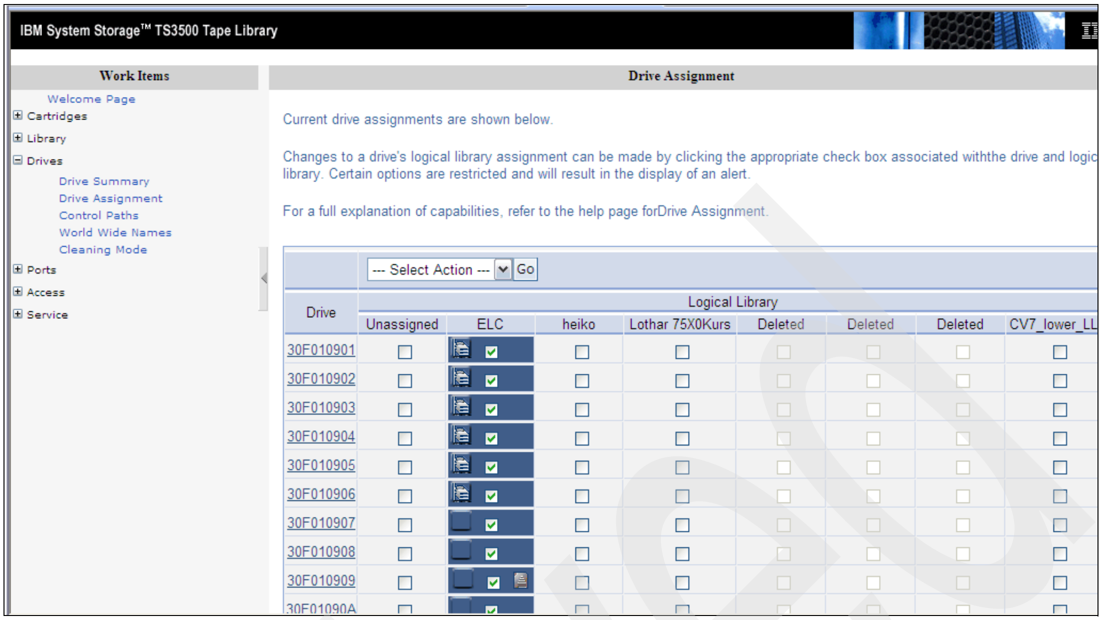

Figure 4-5 Drive Assignment panel

In a multi-platform environment, you see more logical libraries on the panel shown in Figure 4-5, and you can reassign physical tape drives from one logical library to another. You can easily do this for the Open Systems environment where the tape drives attach directly to the host systems without a tape controller or VTS/TS7700.

In a System z environment, a tape drive always attaches to one tape controller or VTS/TS7700 only. If you reassign a tape drive from an IBM 3953 Library Manager partition to an Open Systems partition temporarily, you must also physically detach the tape drive from the TS7700 or tape controller first and then attach the tape drive to the Open Systems host. Only IBM Systems Service Representatives (SSR) should perform these tasks to protect your tape operation from unplanned outages. The reconfiguration might require host software updates or TS7700 reconfiguration in order to not seriously interfere with your System z production work.

Important: You should not attempt to temporarily assign tape drives owned by an IBM 3953 Library Manager partition to another logical library.

In a System z environment, use the Drive Assignment panel only to:

- *Initially* assign the tape drives to the IBM 3953 Library Manager partition.
- Assign additional tape drives after they have been attached to the TS7700 or a tape controller.
- Remove physical tape drives from the configuration, *before* they are physically detached from the TS7700 or tape controller.
- Never disable ALMS at the TS3500 Tape Library after it has been enabled for System z host support and System z tape drive attachment.

Define control path drives

Each 3592-C06 or 3592-J70 controller and TS7700 must have at least one control path drive defined, and should have the optimum number of four control path drives. Distribute the control path drives over more than one IBM TS3500 frame to avoid single points of failure.

> **Important:** We recommend that you define four control path drives for the TS7700 Virtualization Engine.

In a logical library, you can designate any empty, dedicated drive to become a control path drive. A drive that is loaded with a cartridge cannot become a control path until you remove the cartridge. Similarly, any drive that is a control path cannot be disabled until you remove the cartridge that it contains.

The definition of the control path drive is specified on the Drive Assignment panel shown in Figure 4-5 on page 139. The first four drives of the logical library and then drives seven to ten are identified as control path drives by the symbol on the left side of the drive box. You can change the control path drive definition by selecting or deselecting this symbol.

Note that you do not have to and cannot define cleaning policies for the IBM 3592 tape drives installed inside the IBM TS3500. If you try to do so by selecting the work item **Cleaning Mode**, the panel shown in Figure 4-6 displays.

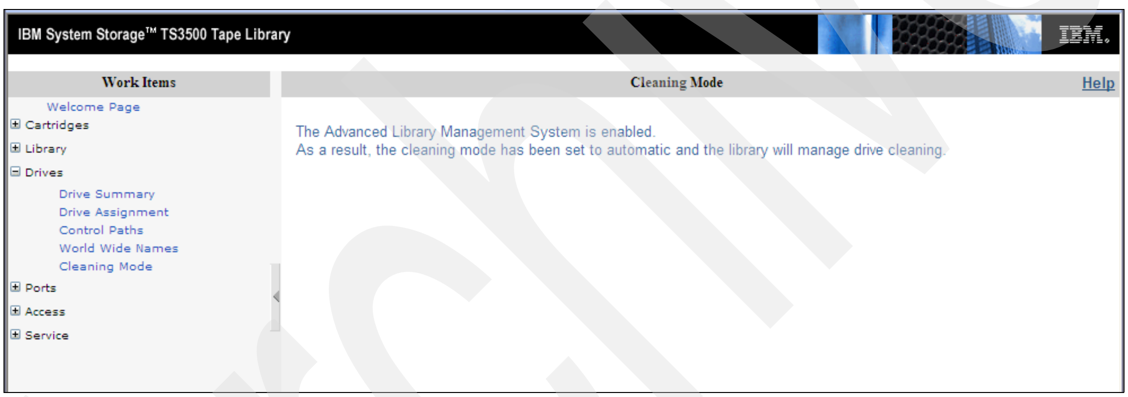

Figure 4-6 Cleaning Mode panel

When ALMS is enabled, a tape drive is cleaned automatically when it requests cleaning.

4.2.2 Defining Cartridge Assignment Policies

The Cartridge Assignment Policy (CAP) of the TS3500 Tape Library allows you to assign ranges of physical cartridge volume serial numbers to specific logical libraries. If you have previously established a Cartridge Assignment Policy and place a cartridge with a VOLSER that matches that range into the I/O station, the library automatically assigns that cartridge to the appropriate logical library.

Select **Cartridge Assignment Policy** from the Manage Cartridges work items to add, change, and remove policies. The maximum quantity of Cartridge Assignment Policies for the entire IBM TS3500 library is 300. Figure 4-7 shows two VOLSER ranges defined for logical library “1-01”.

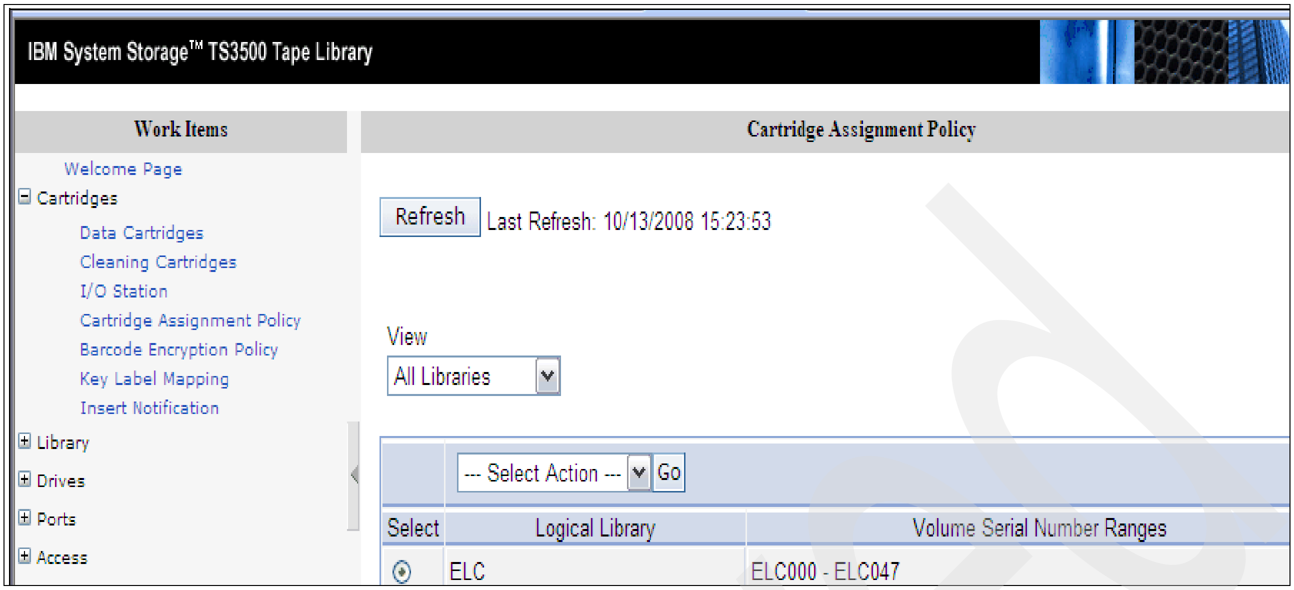

Figure 4-7 Cartridge Assignment Policy panel

The IBM TS3500 Tape Library allows duplicate VOLSER ranges for different media types only. For example, Logical Library 1 and Logical Library 2 contain LTO media, and Logical Library 3 contains IBM 3592 media. Logical Library 1 has a Cartridge Assignment Policy of ABC100-ABC200. The library will reject an attempt to add a Cartridge Assignment Policy of ABC000-ABC300 to Logical Library 2 because the media type is the same (both LTO). However, the library will allow an attempt to add a Cartridge Assignment Policy of ABC000-ABC300 to Logical Library 3 because the media (3592) is different.

In an SMS-managed z/OS environment, all VOLSER identifiers across all storage hierarchies are required to be unique. We strongly recommend that you follow the same rules across host platforms as well whether or not you are sharing an IBM TS3500 between System z and Open Systems hosts.

> **Note:** The Cartridge Assignment Policy does not reassign an already assigned tape cartridge.

4.2.3 Inserting System z data volumes

You can insert two types of volumes into the Library Manager logical partition: physical volumes and logical volumes. Physical volumes might be used as native cartridges, or as backend cartridges (stacked tapes) for the TS7700 Virtualization Engine.

Perform the following steps to add physical cartridges:

1. Define Cartridge Assignment Policies at the *IBM TS3500 Tape Library level* using ALMS through the Web Specialist. These ensure that all System z VOLSER ranges are recognized by the library hardware to be assigned to the correct IBM TS3500 logical library partition, before you begin any 3953 Library Manager definitions.
2. Assign the physical volumes to the correct *Library Manager* partition (native or TS7700). Only VOLSER ranges for TS7700 backend cartridges can be defined. If a cartridge is inserted, and no definition in the 3953 Library Manager exists, it will be treated as a native cartridge. If a host accepts that cartridge, it will be assigned to its scratch volume category. If no host accepts that cartridge, the cartridge remains in the *insert category.*

3. *Physically insert* volumes into the library using the I/O station or by opening the library and placing cartridges in empty storage cells. Cartridges are assigned to the Library Manager partitions, according to the definitions.

These procedures ensure that TS7700 backend cartridges will never be assigned to a host by accident.

The assignment process during insert processing is illustrated in Figure 4-8.

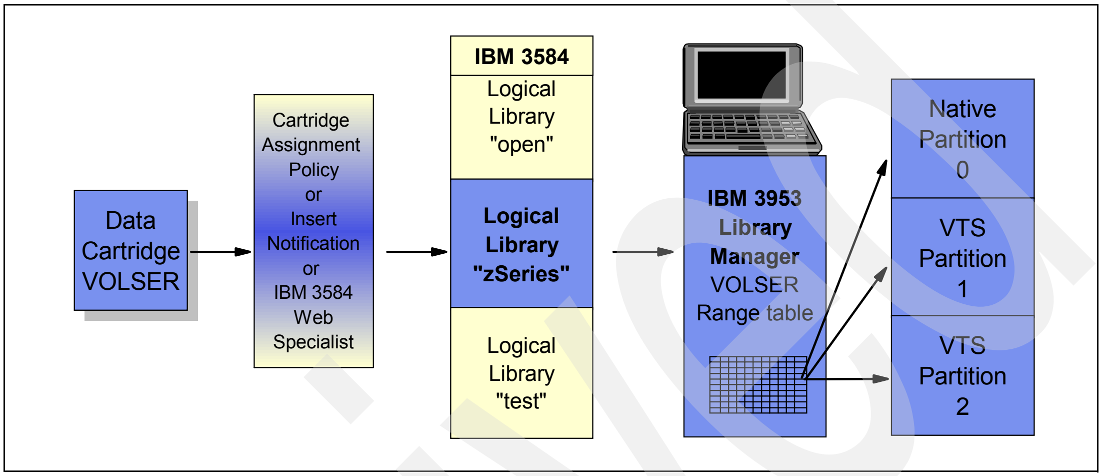

Figure 4-8 Volume assignment to logical libraries

Inserting physical volumes into the IBM TS3500 Tape Library

There are two methods for inserting physical volumes into the IBM TS3500 Tape Library:

- Opening the library doors and directly inserting into tape library storage cells
- Using the TS3500 Tape Library I/O station

Inserting directly into storage cells

The IBM TS3500 Cartridge Assignment Policy defines which volumes are assigned to which logical library partition; if the VOLSER is included in the System z range, then it will be assigned to the TS3500 logical library partition.

After the doors on the library are closed and the tape library has performed inventory, the upload of the inventory to the 3953 Tape System is processed before the IBM TS3500 Tape Library reach the READY state. The 3953 Library Manager updates its database accordingly.

Note: The inventory is performed only on the frame where the door is opened and not on the frames to either side. If you insert cartridges into a frame adjacent to the frame that you opened, then you must perform a manual inventory of the adjacent frame, using the operator panel on the IBM TS3500 Tape Library itself.

Inserting using the I/O station

The TS3500 Tape Library detects volumes in the I/O station, and then moves the volumes to empty cells.

The IBM TS3500 Cartridge Assignment Policy (CAP) defines which volumes are assigned to which logical library; if the VOLSER is included in the System z range then it will be assigned to the IBM TS3500 logical library partition. If any VOLSER is not in a range defined by the CAP, then the operator identifies a System z logical library as the destination using the Insert Notification process.

Under certain conditions, cartridges will not be assigned to a logical library partition in the TS3500 Tape Library or in the 3953 Library Manager.

Attention: Unassigned cartridges can exists in the TS3500 Tape Library, as well as in the 3593 Library Manager. But “unassigned” have different meanings and need different actions from the operator in each system.

Unassigned volumes in the TS3500 Tape Library

If a volume does not match the definitions in the CAP and during the Insert Notification process no owner was specified, the cartridge remains unassigned in the TS3500 Tape Library. You can then assign the cartridges to the logical library partitions using the IBM TS3500 Tape Library Specialist Web interface. You must access the ETL and request an upload of the inventory to make sure that the now assigned volumes are transferred to the 3593 Library Manager. This is because the manual assignment of the cartridges will not automatically process an upload of the inventory to the 3953 Library Manager.

Unassigned volumes in the 3953 Library Manager

If a volume did match the CAP, and is already assigned to a specific logical library partition, it will be uploaded to the assigned 3953 Library Manager. However, if the following conditions are true, the Library Manager will put the cartridge in the *Unassigned* category (FF16):

1. An IBM TotalStorage Virtual Tape Server (VTS) is installed attached to the same IBM 3953, which supports Export and Import operations. In this case, the convenience I/O station is per default in IMPORT mode.
2. A cartridge is entered through the convenience I/O station, and is assigned to that 3953.

The cartridge will be put in the Unassigned category, until the operator assigns it manually to the insert category. If the cartridges were directly assigned to the Insert category, the Import/Export cartridge can be overwritten.

The ETL Specialist Web interface is used to assign unassigned volumes to the Insert or Import category.

Assign cartridges in the IBM TS3500 to a logical library partition

This procedure is only necessary if a cartridge was inserted, but no Cartridge Assignment Policy (CAP) was provided in advance. Then you must assign the cartridge manually to a logical library in the TS3500 Tape Library.

This function is only valid and available if the Advanced Library Management System (ALMS) is enabled on the TS3500 Tape Library. You can assign a data cartridge to a logical library by using the TS3500 Tape Library Specialist Web interface, but not by using the TS3500 Operator Panel. Also, it is normal for a cleaning cartridge to be unassigned if ALMS or cleaning mode is enabled. If cleaning mode is disabled, use the operator panel or the Tape Library Specialist Web interface to assign the cleaning cartridge as part of an import operation.

To assign a data cartridge to a logical library in the TS3500 Tape Library, perform the following steps:

1. Type the Ethernet IP address or the library URL on the URL line and press Enter. The Welcome Page displays.
2. Select **Cartridges** → **Data Cartridges**. The Data Cartridges window displays.
3. Select the logical library to which the cartridge is currently assigned and select how you want the cartridge range to be sorted. (The library can sort the cartridge by volume serial

number, SCSI element address, or frame, column, and row location.) Select **Search**. The Cartridges window displays all ranges for the logical library that you specified.

4. Select the range that contains the data cartridge that you want to assign.
5. Select the data cartridge, then select **Assign**.
6. Select the logical library partition to which you want to assign the data cartridge.
7. Select **Finish** to complete the function.

Note: An inventory upload at the IBM 3953 Library Manager is required because inventory does not automatically update when cartridges are assigned to logical libraries through the Tape Library Specialist Web interface of the library.

Inserting cleaning cartridges

Each drive in the TS3500 Tape Library needs cleaning. The TS1120/3592 drives request the cleaning by themselves if a cleaning is needed. The cleaning in the IBM TS3500 does not involve the 3953 Library Manager at all. It is only the responsibility of the TS3500 Tape Library itself. However, you must enable automatic cleaning, and provide the necessary cleaning cartridges.

Important: If ALMS is installed, cleaning is automatically enabled, and cannot be disabled.

The process to insert cleaning cartridges varies depending on the setup of the IBM TS3500 Tape Library.

With ALMS enabled, you can use the IBM TS3500 Specialist or the library's Operator Panel to insert a cleaner cartridge.

To insert a cleaner cartridge using the IBM TS3500 Specialist, perform the following steps:

1. Open the door of the I/O station and insert the cartridge so that the bar code label faces the interior of the library and the write-protect switch is on the right.
2. Close the door of the I/O station.
3. Type the Ethernet IP address on the URL line of the browser and press Enter. The Welcome Page displays.
4. Select **Manage Cartridges** → **I/O Station**. The I/O Station panel displays.
5. Follow the instructions on the panel.

4.3 TS7700 definitions from the 3494 and 3953 Library Manager

The tasks you perform on the 3494 or 3953 Library Manager are related to the management of the specific Library Manager partitions. The tasks are dependent on whether only tape controllers with native drives are attached or whether one or two VTS/TS7700 subsystems are attached as well. For both the VTS/TS7700 systems and the native tape drives for a System z host, you have to define VOLSER ranges on the Library Manager. You will need to define VOLSER ranges as well on the IBM TS3500 Tape Library Specialist if you are attaching to a TS3500/3953.

In addition, after your IBM service representative installs the TS7700 Virtualization Engine, you must define the environment so that you can use all the TS7700 high-performance functions. All TS7700 definition procedures can be performed from the Library Manager console. Most of these functions can also be performed using the Library Manager Web browser interface, the Enterprise Tape Library (ETL) Specialist.

In this section, we show the 3953 panels but not the 3494 panels unless the 3494 panels are significantly different.

For further details about operating an IBM 3953 Library Manager and the IBM TS3500 Tape Library, refer to the following documentation:

- *IBM TotalStorage TS3500 Operator Guide,* GA32-0468
- *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473

For further details about operating an IBM 3494 Tape Library, refer to *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

4.3.1 Library Manager setup tasks

The procedure for defining the TS7700 definitions from the Library Manager includes defining initial VOLSER ranges and management policies for the TS7700 operation and the 3592 cartridge insertion for TS7700 stacked volume preparation. You must perform the following tasks at the Library Manager:

- Define stacked volume ranges.
- Define logical volume ranges through the TS7700 Management Interface.
- Define Fast Ready categories.
- Define TS7700 Virtualization Engine management policies.
- Define Inhibit reclaim schedule.
- Set reclaim policies.
- Define Free storage threshold.
- SNMP Traps definition.
- Operator Intervention definitions.

These tasks are most likely to be performed after the system IPL has been completed or the library is visible to the host. This section details the tasks to be performed at the Library Manager, not the order of events. For System z implementation of the TS7700 Virtualization Engine, the Library Manager definitions and the System z software definitions are intertwined.

4.3.2 IBM 3494/3953 Library Manager interfaces

Although a physical tape library logically looks like two or three separate libraries to the host, the libraries share the same Library Manager and database. Therefore, the VOLSERs of both the virtual and physical volumes in the same physical tape library must be unique. Before you can insert physical 3592 cartridges for TS7700 Virtualization Engine's use into a library, you must define them on the Library Manager (LM). You must also define a beginning set of VOLSER ranges for your Virtual Volumes.

Most of the tasks can be performed either through the Library Manager console or through the Enterprise Tape Library (ETL) Specialist, the Web-based browser interface to the Library Manager. Some configuration tasks, for example inserting of logical volumes, have moved to the TS7700 Management Interface (MI). We show the ETL Specialist panels and the MI panels for those tasks that can be performed using these interfaces. Figure 4-9 shows the ETL Specialist Home Page for the 3953 Library Manager. The 3494 Web Specialist panel is similar and is shown after the 3953 panel.

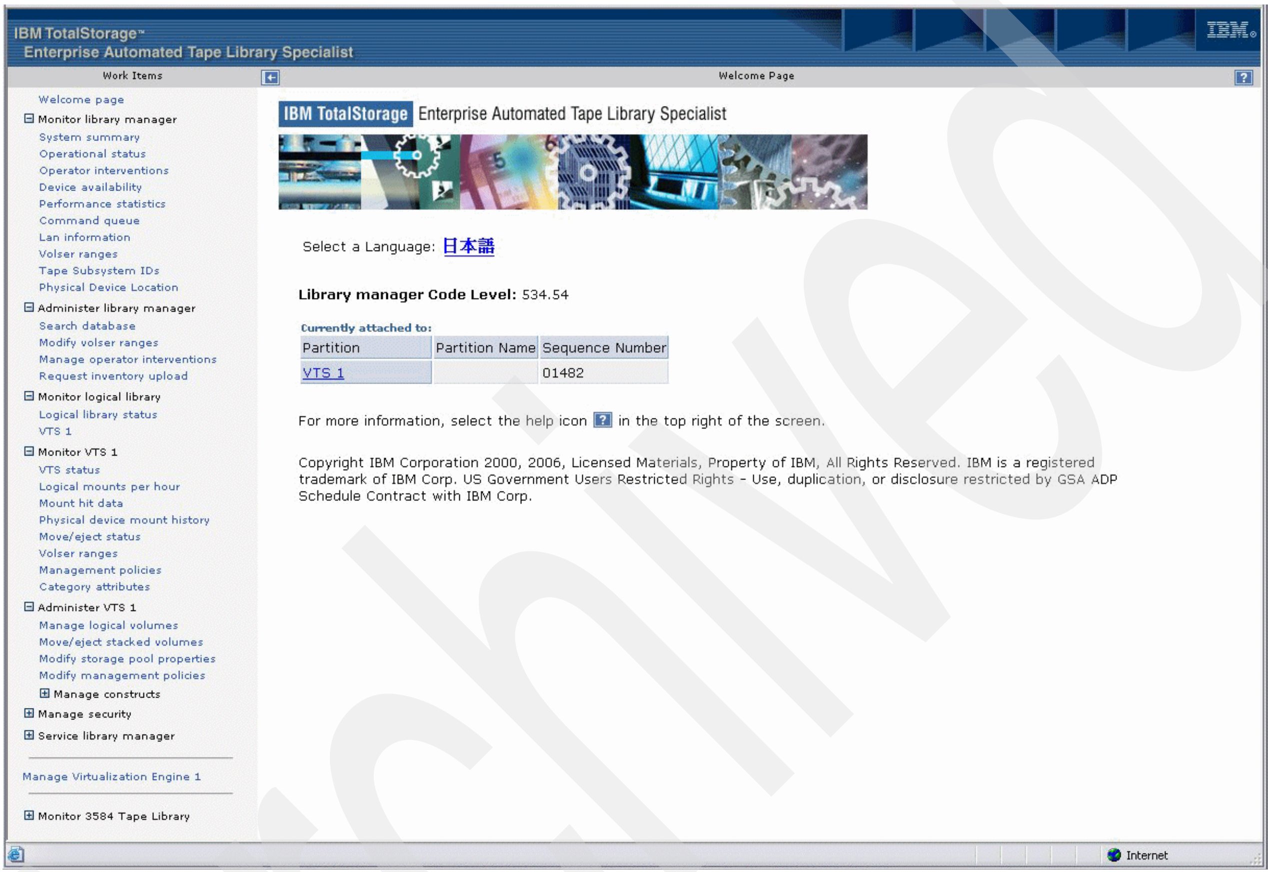

Figure 4-9 IBM 3953 Enterprise Tape Library (ETL) Specialist

Figure 4-10 shows the 3494 Web Specialist welcome panel. As you can see it is similar to the corresponding 3953 panel.

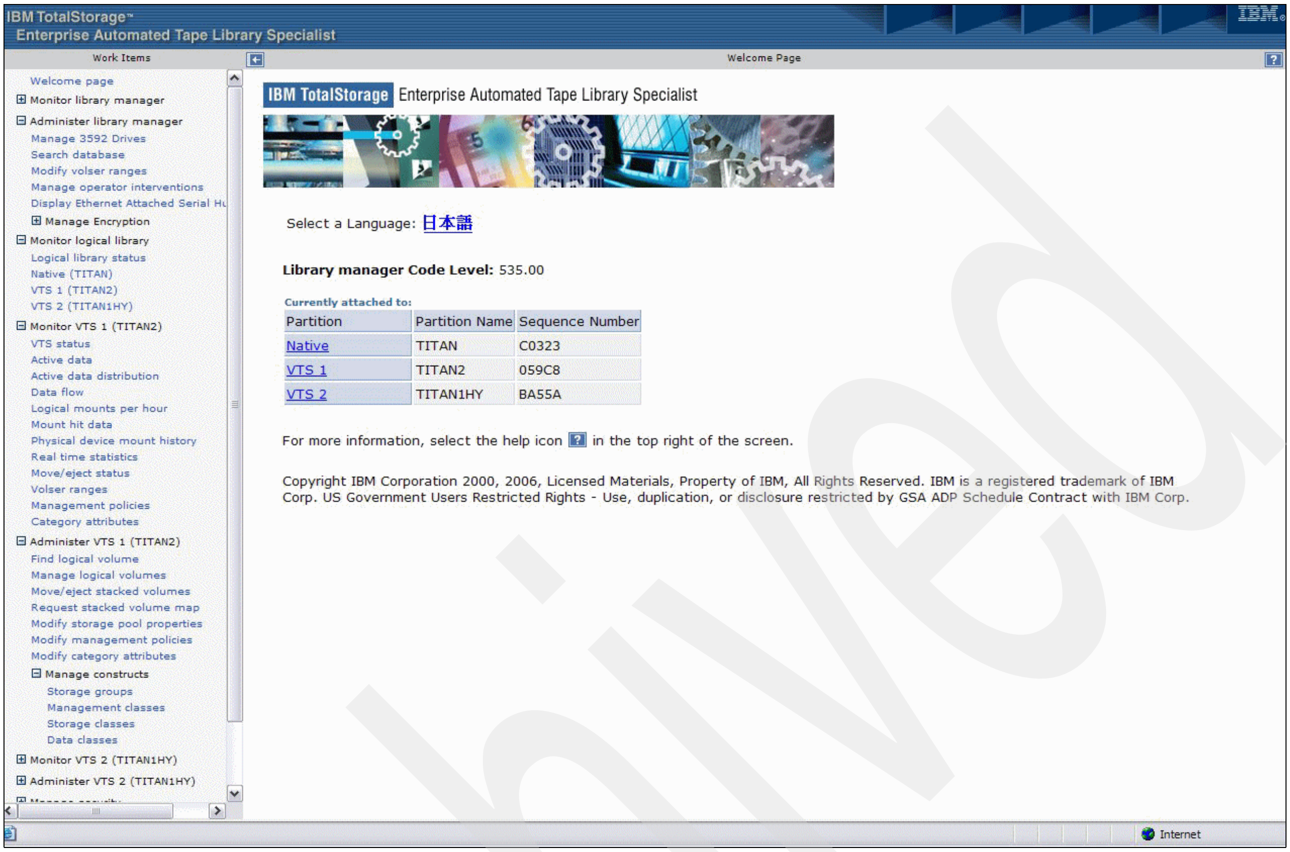

Figure 4-10 3494 Tape Library Web Specialist

We show the Library Manager console panels only if these panels are the only interface to perform a certain task. Figure 4-11 shows the IBM 3953 Library Manager console panel.

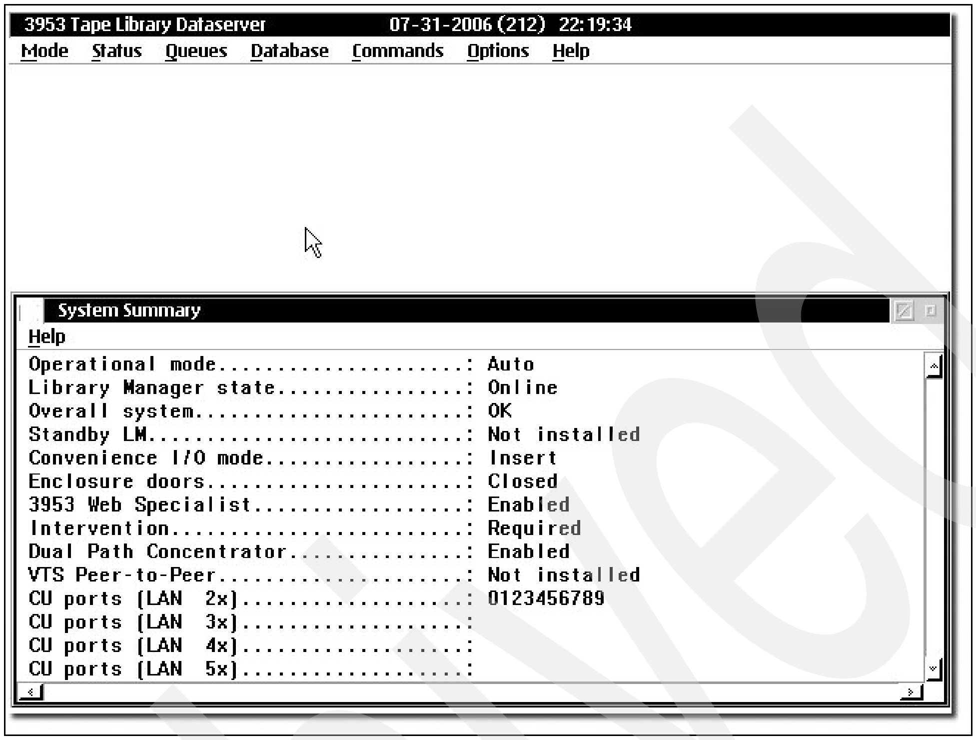

Figure 4-11 The IBM 3953 Library Manager console

From the Library Manager Console, click **Commands**, then **System management**. The System management pop-up window has the following selections:

- **Volser range for media types**: Use this panel to define the physical VOLSER range of TS7700 Virtualization 3592 cartridges to use for the TS7700. These 3592 cartridges are used for stacked volumes in the TS7700.

  On this panel you can enter up to 256 VOLSER ranges and associated media types. The VOLSER ranges are used to help determine a VOLSER's media type when it is inserted into the library. Volser ranges are used only for physical volumes.

- **Insert VTS logical volumes**: Use this panel only for the IBM TotalStorage Virtual Tape Server to insert logical volumes into a library. The logical volumes for a TS7700 Virtualization Engine are inserted through the TS7700 Management Interlace (MI) which is described in “TS7700 definition using the TS7740 Management Interface” on page 174. If you try here to insert logical volumes for the TS7700 you will get an error message as show in Figure 4-15 on page 154.

- **Delete VTS logical volumes**: Use this panel only for the IBM TotalStorage Virtual Tape Server to delete logical volumes that have not been checked into a host’s tape management system. You can use this window to delete only logical volumes that are in the Insert Category (FF00) from here. With the TS7700 Virtualization Engine this task has also been moved to the TS7700 Management Interlace (MI) which is described in “TS7700 definition using the TS7740 Management Interface” on page 174.

- **Set VTS category attributes**: Use this panel to assign the Fast Ready attribute to categories as well as other fast ready category attributes such as expire time and expire hold.

- **Set VTS management policies**: Use this panel to enter the Inhibit reclaim schedule policies and the *Reclaim threshold percentage* and *Free storage* threshold values.
- **Manage insert volumes**: Use this panel to reevaluate the physical volumes in the Insert category for 3592 native use. By redefining the VOLSER ranges, you can move the volumes to the Insert categories for the TS7700. You can also eject the volumes from the tape library.
- **Manage constructs and pools**: Use this panel to access to multiple panels that allow you to manage the storage management constructs and stacked volume pool properties, move/eject stacked volumes, manage logical volumes, and transfer LM administrative data.

You can find a complete description of the 3953 Library Manager console panels in the following documentation:

- *IBM TotalStorage TS3500 Operator Guide,* GA32-0468
- *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473

You can find a complete description of the 3494 Library Manager console panels in *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

4.3.3 Verify the library sequence number and distributed library name

You need to ensure that the composite library sequence number defined through the Interactive Storage Management Facility (ISMF) is *not* the same as the LIBRARY-ID given to the TS7700 by the IBM service representative during the teach/configuration operation.

As explained before, even a Single Cluster Grid (standalone) TS7700 Virtualization Engine requires definition of a Composite LIBRARY-ID and of a Distributed LIBRARY-ID. The Composite LIBRARY-ID is used for the host definition of the TS7700 logical library; the Distributed LIBRARY-ID is used to link to the hardware aspects of the TS7700, such as displaying scratch stacked volumes.

Check the library sequence number on the Operational Status panel on the ETL Specialist Monitor Library Manager selection. Scroll down until you see the field, *Library Sequence Numbers NNNNN*, as shown in Figure 4-12 on page 150.

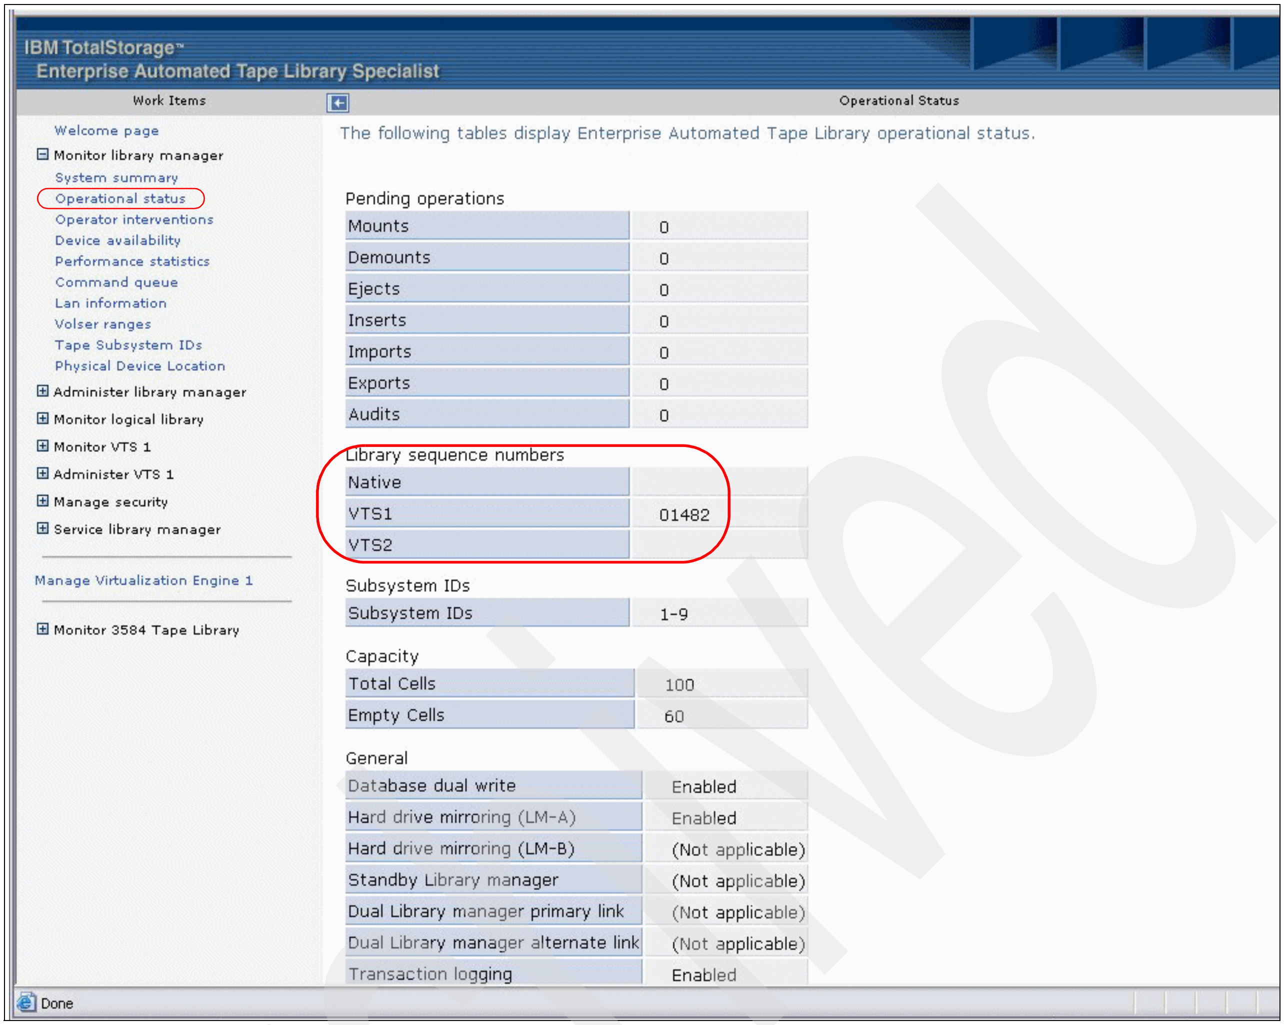

Figure 4-12 TS7700 Operational Status panel using the ETL Web Specialist

Even if there are no static connections between Distributed LIBRARY-ID and the Distributed library name—a name is optional—you can specify a name for each Distributed LIBRARY-ID in the 3953 Library Manager. The same name can be used in the System z host software configuration as a name for the Distributed LIBRARY-IDs defined there. This helps to identify and administer each of the configured Distributed libraries.

The names from a Library Manager view can be entered and updated any time through the Library Manager console using **Configure** → **Update Library Definitions** as shown in Figure 4-13.

The Library Manager access to the service functions must be enabled. To do so, you should check with your IBM service representative.

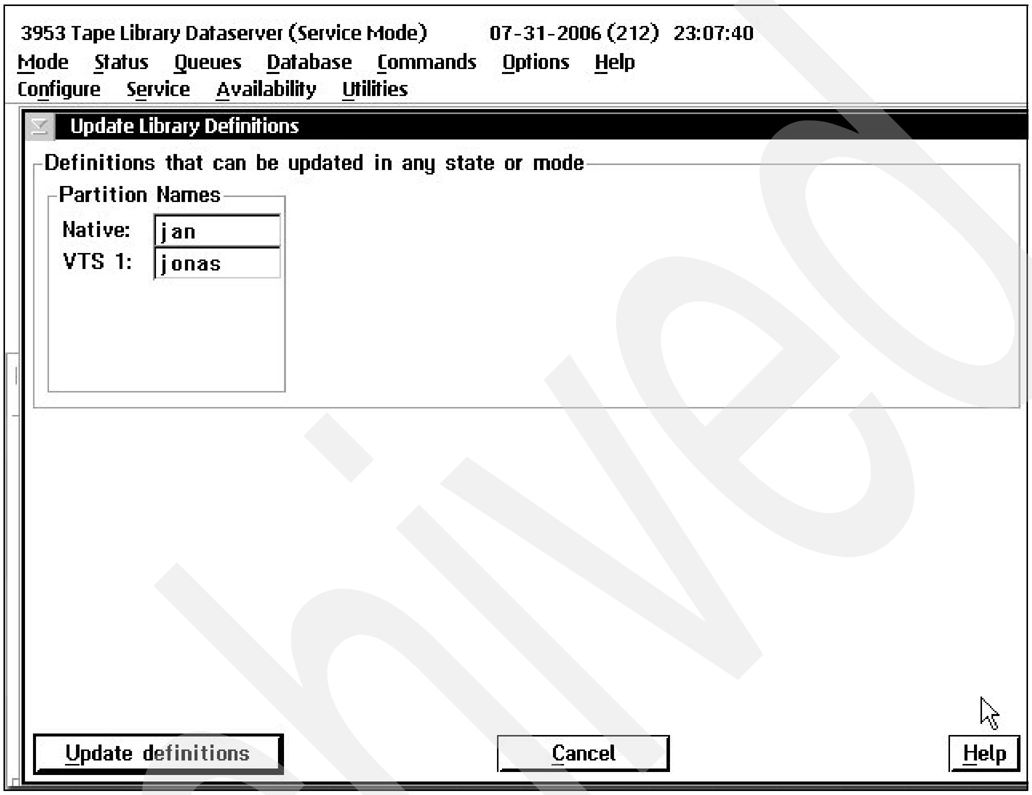

Figure 4-13 Updating the Distributed library names

Note: You should not use Distributed library names starting with a *V* because on the z/OS host, a library name cannot start with a *V*.

On the Library Manager console, the Distributed library names are called *Partition Names*.

4.3.4 Define VOLSER ranges for physical (stacked) volumes

When a cartridge is assigned to a Library Manager partition, the Library Manager uses the VOLSER ranges defined in its VOLSER Range Table to direct the cartridge to the proper partition (Native, VTS1/VE-1, or VTS2/VE-2), and assign the proper Library Manager category. We recommend that you define the proper policies in the VOLSER Range Table *before* inserting the cartridges into the Tape Library.

Important:

- When using a TS3500 Tape Library, you must assign Cartridge Assignment Policies (CAP) at the library hardware level before using the library with System z hosts.
- When using a TS3500 Tape Library, the TS7700 and native physical volumes must fall within ranges that are assigned with CAP to System z host logical libraries in the IBM TS3500.
- The policies for the System z host cartridge ranges must match any host tape management policies.

You can use the ETL Specialist to define the VOLSER ranges for the native and TS7700 partitions. From the ETL Specialist Welcome Page, under **Administer Library Manager**, select **Modify VOLSER ranges**. The panel shown in Figure 4-14 displays, and there you can define new VOLSER ranges or update existing VOLSER ranges.

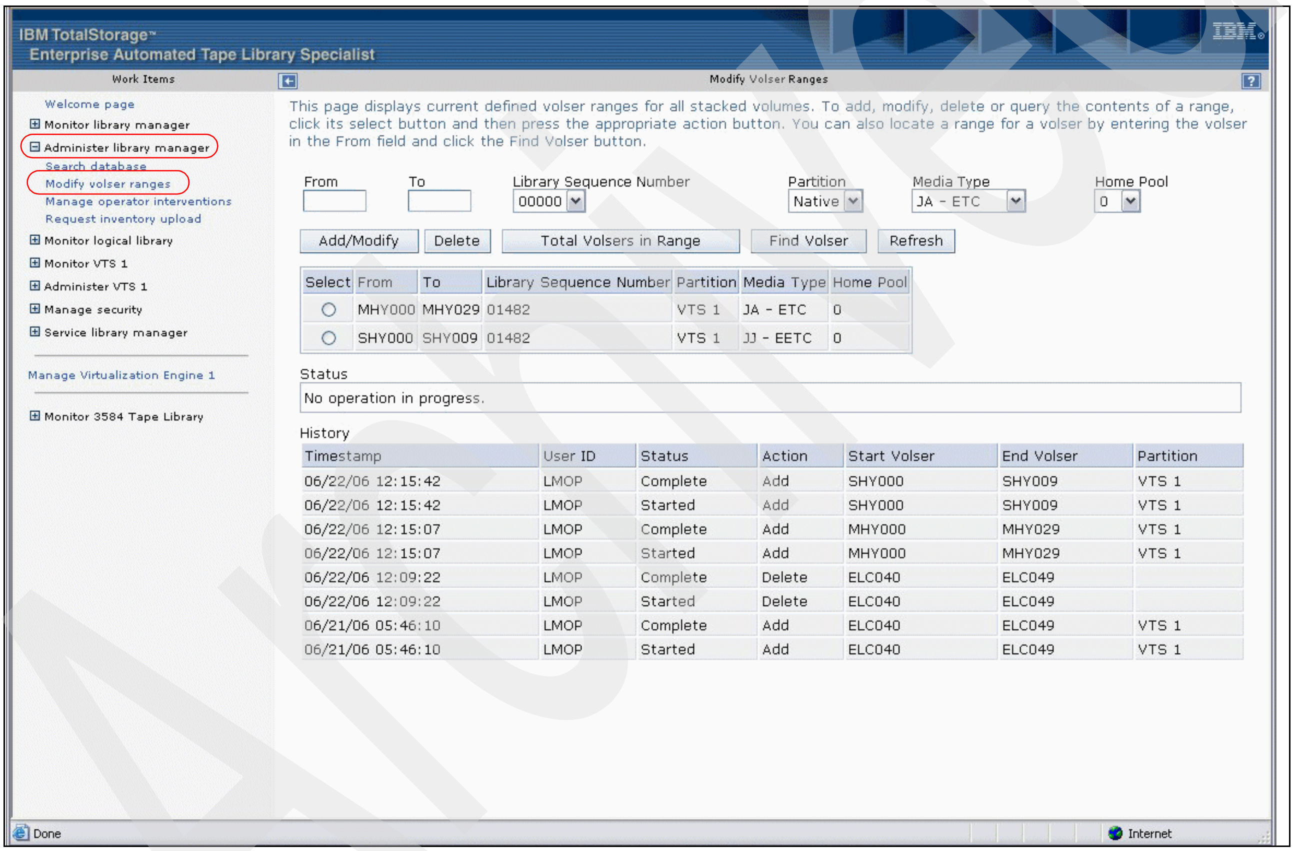

Figure 4-14 Modify VOLSER Ranges panel

The VOLSER entry fields must contain six characters. The characters can be letters, numerals, or a space. The two VOLSERs must be entered in the same format. *Corresponding characters* in each VOLSER must both be either alphabetic or numeric. For example, AAA998 and AAB004 are of the same form, but AA9998 and AAB004 are not. The VOLSERs that fall within a range are determined as follows: The VOLSER range is increased such that alphabetic characters are increased alphabetically, and numeric characters are increased numerically. For example, VOLSER range ABC000–ABD999 would result in a range of 2,000 VOLSERs (ABC000–ABC999 and ABD000–ABD999).

In Figure 4-7 on page 141, we showed how to define the Cartridge Assignment Policies (CAP) for one logical library in a TS3500. You must define the same VOLSER ranges again as shown in Figure 4-14. Here they determine to which logical library within the IBM 3953 Library Manager partition the cartridges will be assigned.

Note: The VOLSER ranges you define on the IBM TS3500 Tape Library and on the IBM 3953 Library Manager or on the 3494 Library Manager apply to physical cartridges only. You should not define VOLSER ranges for logical volumes for a TS7700. You must define logical volumes using the TS7700 Management Interface. See 4.5.8, “Insert logical volumes using the TS7700 Management Interface” on page 182 for more information.

For the TS7700, no additional definitions are required at the hardware level other than setting up the proper VOLSER ranges at the TS3500 library.

The definition of VOLSER ranges for the native partition is optional. The Library Manager logic assigns the VOLSER to the native partition when it matches a range in the range table or when it does not match any range. VOLSERs destined for a TS7700 partition must have a range defined. In other words, if a VOLSER does not match a VOLSER range, the Library Manager assigns it per default to the native partition.

Note: Although you could now enter cartridges into the 3494 or TS3500 library, we recommend that you first complete the required definitions at the host as described in subsequent sections of this chapter, before you insert any physical cartridges into the Tape Library.

In previous generations of the TS7700 (as the IBM TotalStorage Virtual Tape Server - VTS) the logical volumes were defined and inserted through the Library Manager console or the IBM ETL Specialist.

This process has changed for the TS7700 Virtualization Engine, where now this task is performed through the TS7700 Management Interface (MI), which is described in detail later in 4.5.8, “Insert logical volumes using the TS7700 Management Interface” on page 182.

If you still try to enter a logical volume range using the Library Manager console for a TS7700 subsystem, you will get a window with a user warning that this operation is not supported for a TS7700 Virtualization Engine, as shown in Figure 4-15.

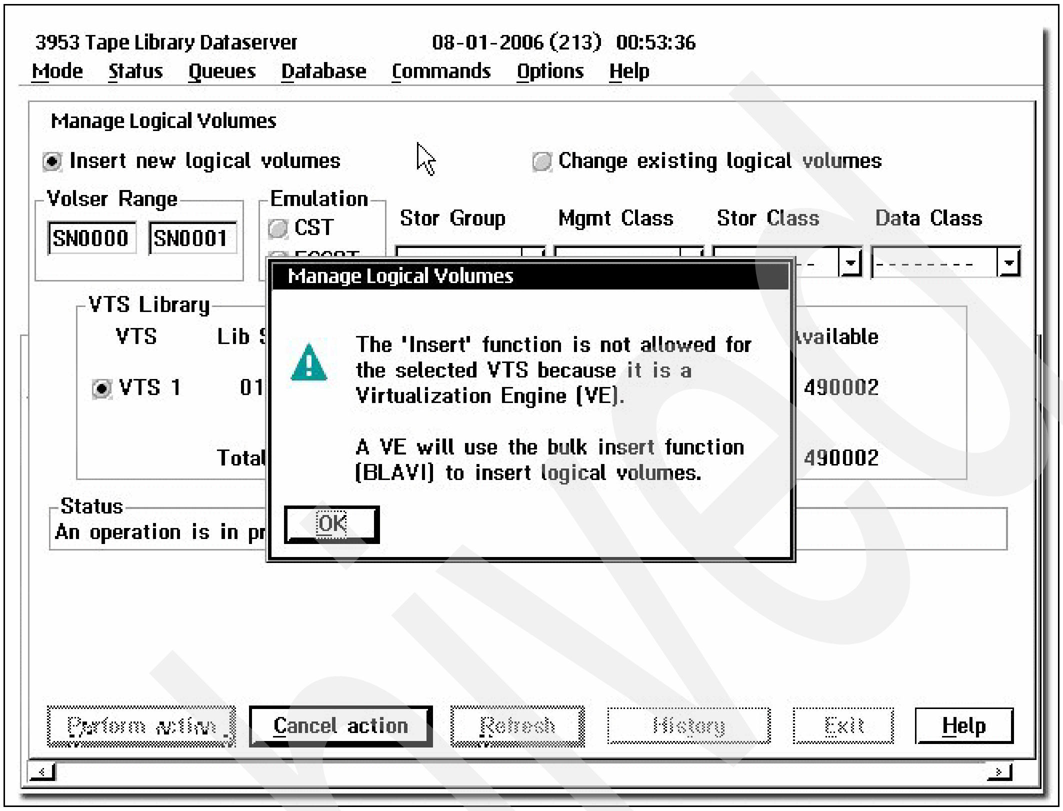

Figure 4-15 User warning for logical volume inserting using Library Manager console

4.3.5 Define Fast Ready categories

To take advantage of the scratch mount performance of the TS7700 Virtualization Engine and to prevent recalls for scratch mounts, you need to assign the Fast Ready attribute to the categories used by the host for scratch volumes.

The MOUNT FROM CATEGORY command, as used for scratch mounts, is not exclusively used for scratch mounts; therefore, the TS7700 cannot assume that any MOUNT FROM CATEGORY is for a scratch volume.

The Fast Ready attribute provides a definition of a category to supply scratch mounts. The Fast Ready definition is done through the Library Manager ETL Web Specialist. Figure 4-16 shows the Define Fast Ready Categories panel.

The actual category hexadecimal number depends on the software environment and on the definitions in the SYS1.PARMLIB member DEVSUPxx for library partitioning. Also, the DEVSUPxx member must be referenced in IEASYSxx to be activated.

Note: Any category up to and including the X'FEFF' can be overridden in DEVSUPxx, but be aware that categories can be overridden only if the TS7700 will not be accessed by the owning operating systems, which are described in Appendix B, “Library Manager volume categories” on page 531.

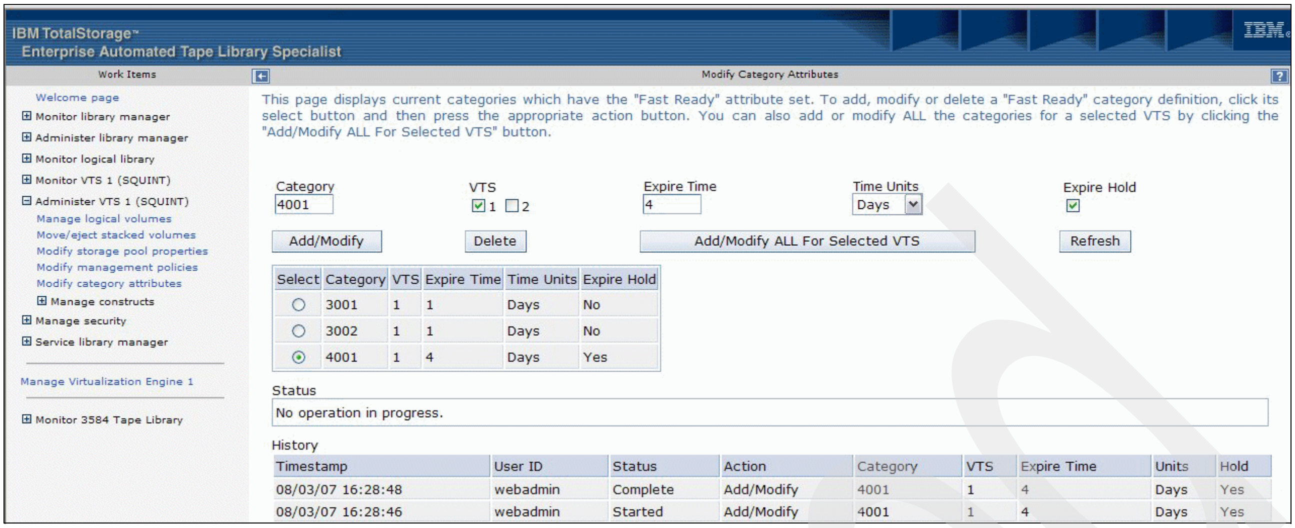

Figure 4-16 Set TS7700/VTS Fast Ready category attributes panel

In this panel you add the category to the list of categories in the tape library that have the Fast Ready attribute set.

To modify a category's Fast Ready attribute, from the ETL Specialist Welcome Page, under Administer VTS 1 (or VTS 2), select **Modify category attributes**. From the panel select the button next to the desired category, make any desired changes to the fields, then click **Add/Modify**.

To define a category as Fast Ready, on the Define Fast Ready Categories window enter the 4-digit hexadecimal category number, select the desired TS7700, set the expire field, and select **Add/Modify**. You should specify the scratch mount categories of MEDIA1 and MEDIA2 as the TS7700 Fast Ready when you use z/OS DFSMS. These scratch mount categories are 0001 and 0002, respectively, if you use the z/OS DFSMS default values. Figure 4-16 on page 155 shows non-default category values of 3001, 3002, and 4001.

When defining a Fast Ready category, you can also set up an expire time as well as further define the expire time as an expire hold time. Refer to 4.3.6, "Define logical volume Expiration Time" on page 156 for details concerning this aspect of the Fast Ready categories.

Note: We recommend that you add a comment to DEVSUPnn to make sure that the Library Manager Fast Ready categories are updated whenever the category values in DEVSUPnn are changed. They need to be in sync at all times.

Refer to Appendix B, "Library Manager volume categories" on page 531 for the scratch mount category for each software platform. In addition to the z/OS DFSMS default value for the scratch mount category, you can define your own scratch category to the Library Manager. In this case, you should also add your own scratch mount category to the Fast Ready category list.

To delete a category from the Fast Ready list, highlight the category in the list box and select **Delete category**. You are prompted to confirm the delete operation. You can click **Help** to display the Set Category Attributes help panel. After all updates have been completed, select **Cancel** to exit this option. Such an action would only serve to increase the mount wait time for scratch volumes within the particular category. This should only be considered as a "cleanup" activity after a category is to be removed from service.

4.3.6 Define logical volume Expiration Time

You define the expiration time from the Library Manager panel shown in Figure 4-16 on page 155. If the Delete Expired Volume Data setting on the Library Manager is not used, logical volumes occupy space on the physical volumes even after they have been returned to scratch. In that case, only when a logical volume is rewritten is the old data released to reduce the amount of active data on the physical volume. With the Delete Expired Volume Data setting, the data associated with volumes that have been returned to scratch are deleted after a time period and their old data released.

For example, assume that you have 20 000 logical volumes in scratch status at any point in time and that the average amount of data on a logical volume is 400 MB and that the data compresses at a 2:1 ratio. The space occupied by the data on those scratch volumes is 4 000 000 MB or the equivalent of 14 3592 JA cartridges. By using the Delete Expired Volume Data setting, you can reduce the number of cartridges required in this example by 14.

The parameter Expire Time specifies the amount of time in hours, days, or weeks. The data will continue to be managed by the TS7700 after a logical volume is returned to scratch before the data associated with the logical volume is deleted. A minimum of 24 and a maximum of 32 767 hours (approximately 194 weeks) can be specified. Specifying a value of zero means that the data associated with the volume is to be managed as it was prior to the addition of this option, meaning that it is never deleted. In essence, specifying a value (other than zero) provides a “grace period” from when the logical volume is returned to scratch until its associated data is eligible for deletion. A separate Expire Time can be set for each category defined as fast-ready.

Establishing the Expire Time for a volume occurs as a result of specific events or actions. The possible events or actions and their effect on the expire time of a volume are as follows:

- A volume is mounted

  The data associated with a logical volume will not be deleted, even if it is eligible, if the volume is mounted. Its Expire Time is set to zero, meaning it will not be deleted. It will be re-evaluated for deletion when its category is subsequently assigned.

- A volume's category is changed

  Whenever a volume is assigned to a category, including assignment to the same category it is currently in, it is re-evaluated for deletion.

- To a category with a non-zero Expire Time

  If the category has a non-zero Expire Time, the volume's data is eligible for deletion after the specified time period, even if its previous category had a different non-zero Expire Time.

- To a category with an Expire Time of zero

  If the volume's previous category had a non-zero Expire Time or even if the volume was already eligible for deletion (but has not yet been selected to be deleted) and the category it is assigned to has an Expire Time of zero, the volume's data is no longer eligible for deletion. Its Expire Time is set to zero.

- A category’s Expire Time is changed

  If a user changes the Expire Time value through the fast-ready definition panel on the Library Manager, the volumes assigned to the category are re-evaluated for deletion.

- Non-zero to zero

  If the Expire Time is changed from a non-zero value to zero, volumes assigned to the category that currently have a non-zero Expire Time are reset to an Expire Time of zero. If a volume was already eligible for deletion, but had not been selected for deletion, the volume's data is no longer eligible for deletion.

- Zero to non-zero

  Volumes currently assigned to the category continue to have an Expire Time of zero. Volumes subsequently assigned to the category will have the specified non-zero Expire Time.

- Non-zero to non-zero

  Volumes maintain their current Expire Time. Volumes subsequently assigned to the category will have the updated non-zero Expire Time.

After a volume's Expire Time has been reached, it is eligible for deletion. Not all data eligible for deletion will be deleted in the hour it is first eligible. Once an hour, on the half hour, the Library Manager selects up to 500 eligible volumes for data deletion. The volumes are selected based on the time that they became eligible, with the oldest ones being selected first. Up to 500 eligible volumes for the TS7700 in the library are selected first.

When the Library Manager database is backed up, the Expire Time for all volumes is set to zero in the backup file. The Expire Times in the database itself are not changed.

You can view the status of a volume's Expire Time using the query database functions that are provided through the Library Manager.

The Hold function, which you can enable by setting the check mark, essentially "locks" the volume from being used during the grace period. This means that it will not be selected to satisfy either a scratch or specific mount. In addition, it cannot be moved out of the category without IBM service intervention. This gives the guarantee that within the specified grace period a scratch logical volume can be recovered easily by just changing the category back to private from the host. For more information, see Chapter 7, "Operation" on page 303.

4.3.7 Define TS7700 management policies

You define TS7700 management policies for reclamation and free storage thresholds through the ETL Specialist. These policies include:

- Reclaim Threshold Percentage and Inhibit Reclaim Schedule (see Figure 4-17 on page 160) and Free Storage Threshold (see Figure 4-18 on page 162), which are defined in the Modify Management Policy panel.
- Days without Access, Age of Last Data Written, Days without Data Inactivation, Days before Secure Data Erase, and Maximum Active Data (see Figure 4-21 on page 164), which are defined in the Modify Stacked Volume Pool Properties panel.

Reclamation and reconciliation

To minimize the effect of TS7700 internal processes like space reclamation on your tape operation, you can inhibit space reclamation for certain periods of time and adjust reclamation thresholds through the Library Manager console and ETL Specialist.

Over time, more and more logical volume copies on a stacked volume become obsolete and the stacked volume contains less and less active data. The storage management software monitors the amount of active data on stacked volumes. It marks the cartridge eligible for reclamation when the percentage set by the Library Manager Reclaim Threshold Percentage value is met. During reclamation, the active data is copied to another stacked volume, leaving the source volume to be used as a scratch stacked volume by the storage management software in the TS7700 Virtualization Engine.

Note: Each reclamation task uses two tape devices, a source and a target. The movement of active data is a tape-to-tape copy function that does not use the TVC.

All volumes with active data that hit any of the specified reclamation policies is eligible for space reclamation.

Multiple Reclamation processes can run in parallel. The maximum would be the maximum amount of physical installed tape drives divided by two. Whenever a new reclamation task is started, the stacked volume with the least amount of active data is picked.

Reclamation enablement

To minimize any impact on TS7700 activity, the storage management software monitors resource utilization in the TS7700 and enables or disables reclaim as appropriate. You can optionally prevent reclamation activity at specific times of day by specifying an Inhibit Reclaim Schedule on the Library Manager Console. However, the TS7700 determines whether reclamation is to be enabled or disabled once an hour depending on the number of available scratch cartridges and will ignore the Inhibit Reclaim Schedule, if the TS7700 is close to running out of scratch cartridges.

Using the Bulk Volume Information Retrieval (BVIR) process, you can run the query for PHYSICAL MEDIA POOLS to monitor the amount of active data on stacked volumes to help you plan for a reasonable and effective Reclamation Threshold Percentage. You can also use the Host Console Request function to obtain the physical volume counts.

Even though reclamation is enabled, there might not always be stacked volumes going through the process all the time. Other conditions must be met, such as stacked volumes that meet one of the reclaim policies and drives available to mount the stacked volumes.

Reclaim for a volume is stopped by the TS7700 internal management functions if a tape drive is needed for a recall or copy (because these are of a higher priority) or a logical volume is needed for recall off a source or target tape being used in the reclaim process. If this happens, reclamation is stopped for this physical tape after the current logical volume move is complete.

Reclamation is enabled or disabled according to a set of rules. These rules are required to enable the TS7700 to compensate for the pooling structure and scratch thresholds. Pooling is enabled as a standard feature of the TS7700, even if you are only using one pool. Reclamation can occur on multiple volume pools at the same time, as well as processing multiple tasks for the same pool. One of the reclamation methods selects the volumes for processing based on the percentage of active data. For example: If the reclaim threshold was set to 30% generically across all volume pools, the TS7700 would select all the stacked volumes from 0% to 29% of remaining active data. The reclaim tasks would then process the volumes from least full (0%) to most full (29%) up to the defined reclaim threshold of 30%.

Individual pools can have different reclaim policies set. The number of pools can also influence the reclamation process, because the Library Manager always evaluates the stacked media starting with Pool 1.

The scratch count for physical cartridges also affects reclamation. The *scratch state* of pools is assessed as follows:

1. A pool enters a *Low scratch state* when it has access to less than 50 but two or more stacked volumes.
2. A pool enters a *Panic scratch state* when it has access to less than two empty stacked volumes.

Access to includes any borrowing capability; that means that if the pool is configured for borrowing, and if there are more than 50 cartridges in the Common Scratch Pool (CSP), the pool will not enter the Low scratch state.

Whether borrowing is configured or not, as long each pool has two scratch cartridges, the Panic Reclamation mode is not entered. Panic Reclamation mode is entered when a pool has less than two scratch cartridges and no more scratch cartridges can be borrowed from any other pool defined for borrowing. Borrowing is described in "Physical Volume Pooling" on page 40.

Pools in either scratch state (Low or Panic state) get priority for Reclamation. Table 4-1 summarizes the thresholds.

Table 4-1 Reclamation priority table

Priority	Condition	Reclaim schedule honored	Active data threshold % honored	Number of concurrent reclaims	Comments
1	Pool in Panic scratch state	No	No	1, regardless of idle drives	
2	Priority move	Yes or No	No	1, regardless of idle drives	If a volume is within 10 days of a Secure Data Erasure (SDE) and still has active data on it, it will be reclaimed at this priority. An SDE priority move will honor the inhibit reclaim schedule. For an LM operator initiated priority move, the option to honor the inhibit reclaim schedule is given to the operator.
3	Pool in Low scratch state	Yes	Yes	1, regardless of idle drives	Volumes subject to reclaim due to Maximum Active Data, Days Without Access, Age of Last Data Written, and Days Without Data Inactivation will use priority 3 or 4 reclamation.
4	Normal reclaim	Yes	Yes, pick from all eligible pools	(# idle drives / 2) minus 1 6 drv: 2 max 12 drv: 5 max	Volumes subject to reclaim due to Maximum Active Data, Days Without Access, Age of Last Data Written, and Days Without Data Inactivation will use priority 3 or 4 reclamation.

Note: A physical drive is considered *idle* when there has been no activity for the previous ten minutes.

The Inhibit Reclaim Schedule is not honored by the Secure Data Erase function for a volume that has no active data.

Inhibit reclamation

The Inhibit Reclaim Schedule defines when the TS7700 should refrain from reclaim operations. Reclaim operations require physical drives. Therefore, drives are used for reclaim operations at the same time others are used to recall data to satisfy mount requests. During times of heavy mount activity, it might be desirable to make all of the physical drives available for recall operations. If these periods of heavy mount activity are predictable, you can use the Inhibit Reclaim Schedule to inhibit reclaim operations for the heavy mount activity periods. You can add up to 14 entries to the schedule.

To define the Inhibit Reclaim Schedule, select **Modify management policies** from the Administer VTSx work items at the ETL Specialist. See Figure 4-17.

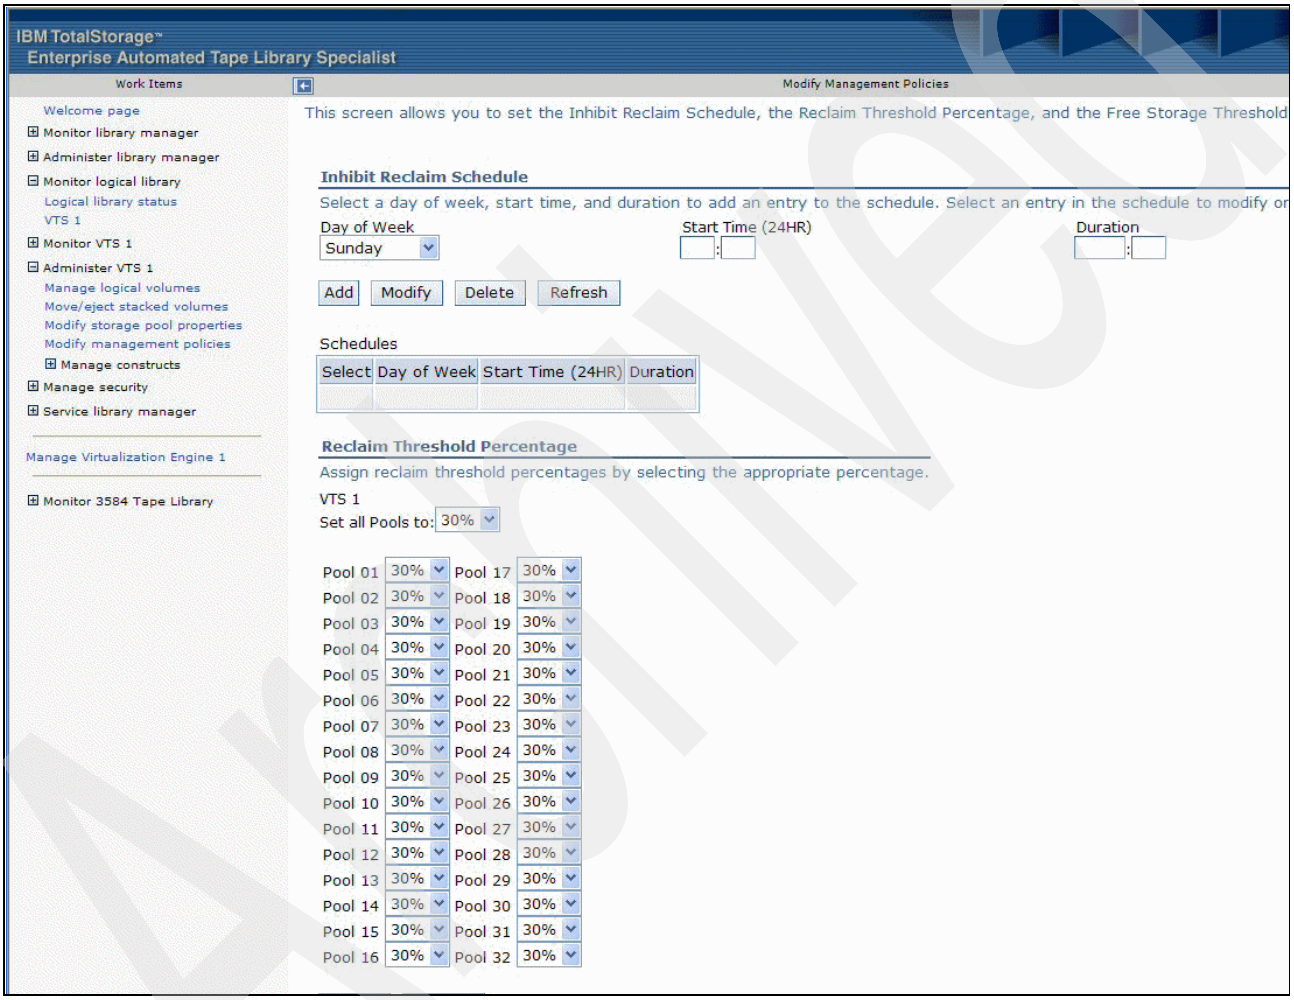

Figure 4-17 Modify management policies for the TS7700 using the ETL Specialist, Part 1

Note: Enter your local times in this panel. You also need to enter the time that the Library Manager needs to be configured.

Reclaim Threshold Percentage

The Reclaim Threshold Percentage (Figure 4-17) is used to identify when a cartridge is to be made available for reclamation. Each stacked volume has some amount of active data and some amount of invalidated data which has been deleted from the active volume list. If the percentage of active data in stacked volume is less than the percentage specified in this panel, the stacked volume is available to reclaim. During the reclamation process all of the active data from the original stacked volume is moved to another stacked volume. After all active data is moved from the original stacked volume, its category is set to scratch. This makes it available for reuse immediately.

The Reclaim Threshold Percentage is initially set at 10%. We recommend that you start with this value and slowly raise it by 5% increments, if you need to. As a general rule, try not to go above 30% to 40%. It is better to add additional stacked volumes rather than raise this value. The higher this number is, the longer it takes the TS7700 to reclaim a stacked volume, because more data must be copied from one cartridge to the other.

You can set a pool's Reclaim Threshold Percentage to 0%, if you do not want it to be a reason to reclaim a volume or want to prevent reclaim from occurring (assuming that all of the other reclaim criteria are also set to 0).

Be aware that a multiple of two drives are involved in the reclamation process and because of this resource usage, you should not specify high percentages. BVIR reports can help you adjust the percentages. See 8.6, "Bulk Volume Information Retrieval (BVIR)" on page 442 for details.

Free Storage Threshold percentage

The *Free Storage Threshold (GB)* provides a warning when the TS7700 is running low on free storage, the capacity of all the empty stacked volumes in the TS7700. A threshold is provided for each TS7700 installed in the library and is entered in GB. The default value is 600 GB. If the free storage drops below the threshold (alarm level), the Library Manager signals an intervention-required condition to notify you to add more stacked volumes. Refer to the appropriate Library Manager Operator's Guide section under "VTS Management Policies" for additional information and Free Storage Threshold tables.

Figure 4-18 shows the ETL Specialist panel to define the threshold.

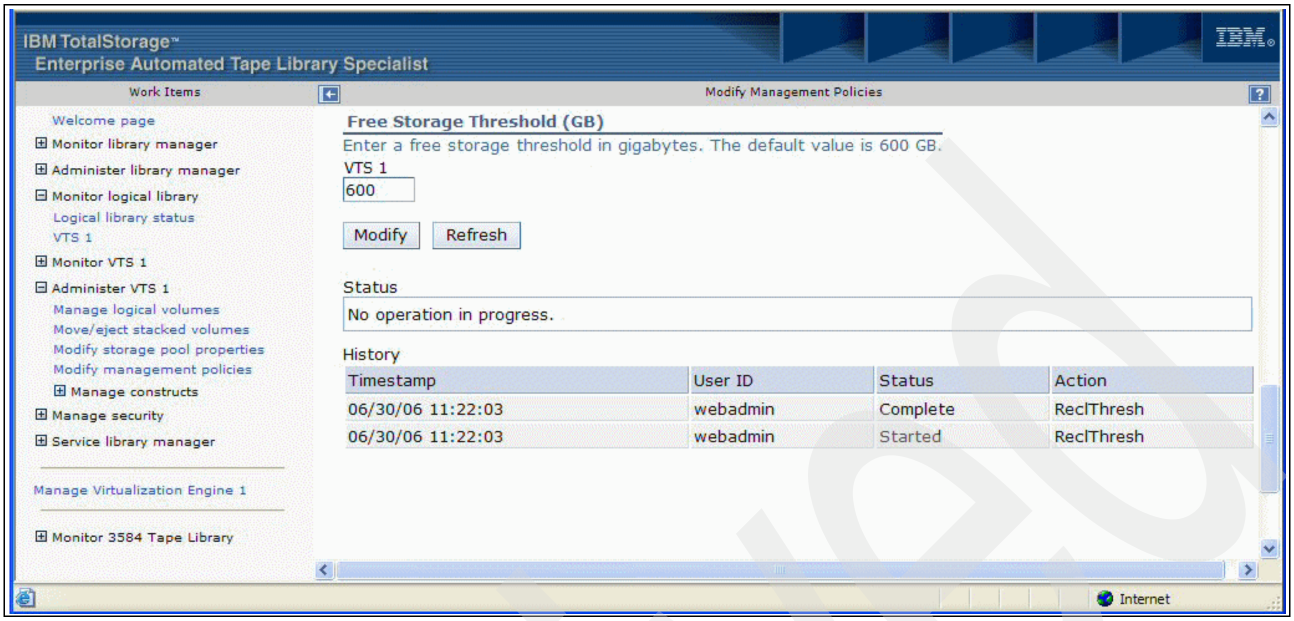

Figure 4-18 Modify management policies for the TS7700 using the ETL Specialist, Part 2

4.3.8 Operator interventions

To ensure all raised Operator Interventions in the Library Manager are getting reported to the attached System z hosts, you need to verify through the **Manage operator interventions** Web panel that a check mark is set for sending the information up to the hosts. Be sure to click the Submit button after making your selection. This is shown in Figure 4-19.

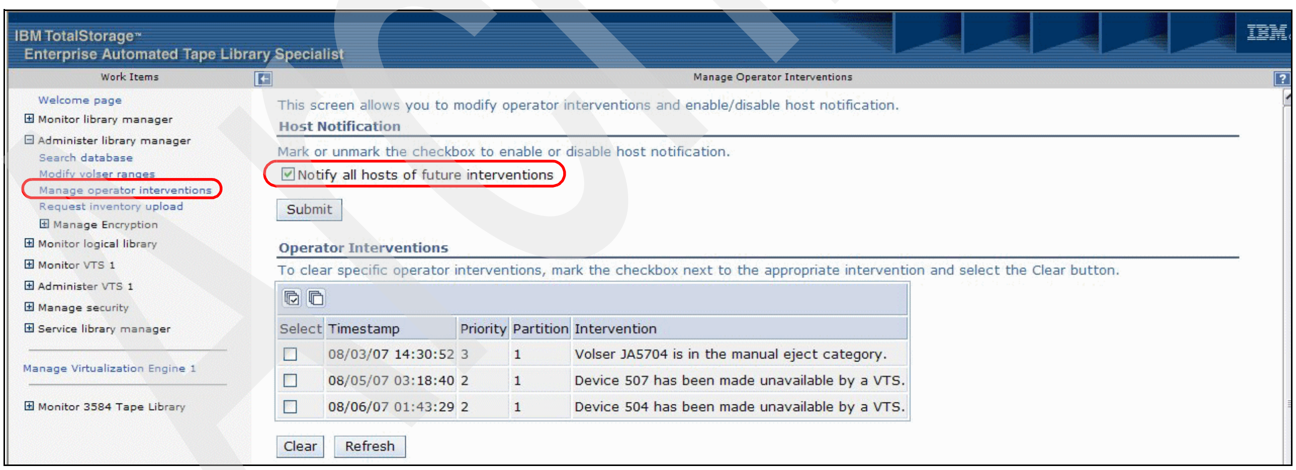

Figure 4-19 Send interventions to host consoles

4.4 Define storage pools and storage constructs

Outboard Policy Management was introduced in 2002 as Advanced Policy Management in the previous TS7700 generation, the IBM TotalStorage Virtual Tape Server (VTS). Today, each TS7700 comes by default with this policy management capability enabled.

You can use the ETL Specialist for the necessary definitions. Alternatively, you can make these definitions from the Library Manager console. For details about using the Library Manager panels, refer to:

- *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473
- *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632

To exploit the outboard management functions, define these four types of constructs:

- Storage Group (SG)
- Management Class (MC)
- Storage Class (SC)
- Data Class (DC)

These construct names are passed down from the z/OS host and stored with the logical volume. Depending on the actions defined in the constructs on the Library Manager, these actions will be performed as part of mount processing. The actions defined for each of the constructs are performed by the TS7700.

From the Administer VTS 1 or Administer VTS 2 selection you can choose the different administration functions, as shown in Figure 4-20.

Figure 4-20 Manage Constructs and Pools using ETL Specialist panel

Note that all ETL Specialist panels that allow modifications of volumes or definitions require that you enter a user ID and password.

4.4.1 Defining stacked volume pool properties

If you want to use physical volume pooling and separate the stacked cartridges into different volume pools, define the stacked volume pool properties. In addition, you can specify reclaim parameters for the physical cartridges of a specific pool. Both the Library Manager and the TS7700 management Interface are used for pool property definitions.

> **Note:** The default Storage Pool is Pool 1. If you choose not to use multiple pools, no further actions are required. All private stacked volumes will reside in Pool 1.

The Library Manager Stacked Volume Pool Properties panel allows you to modify several of the pool properties. To do this, perform the following steps:

1. From the Work items of the ETL Specialist panel, click **Administer VTSx → Modify storage pool properties.** You will see a list of 32 Stacked Volume Pools, as shown in Figure 4-22.

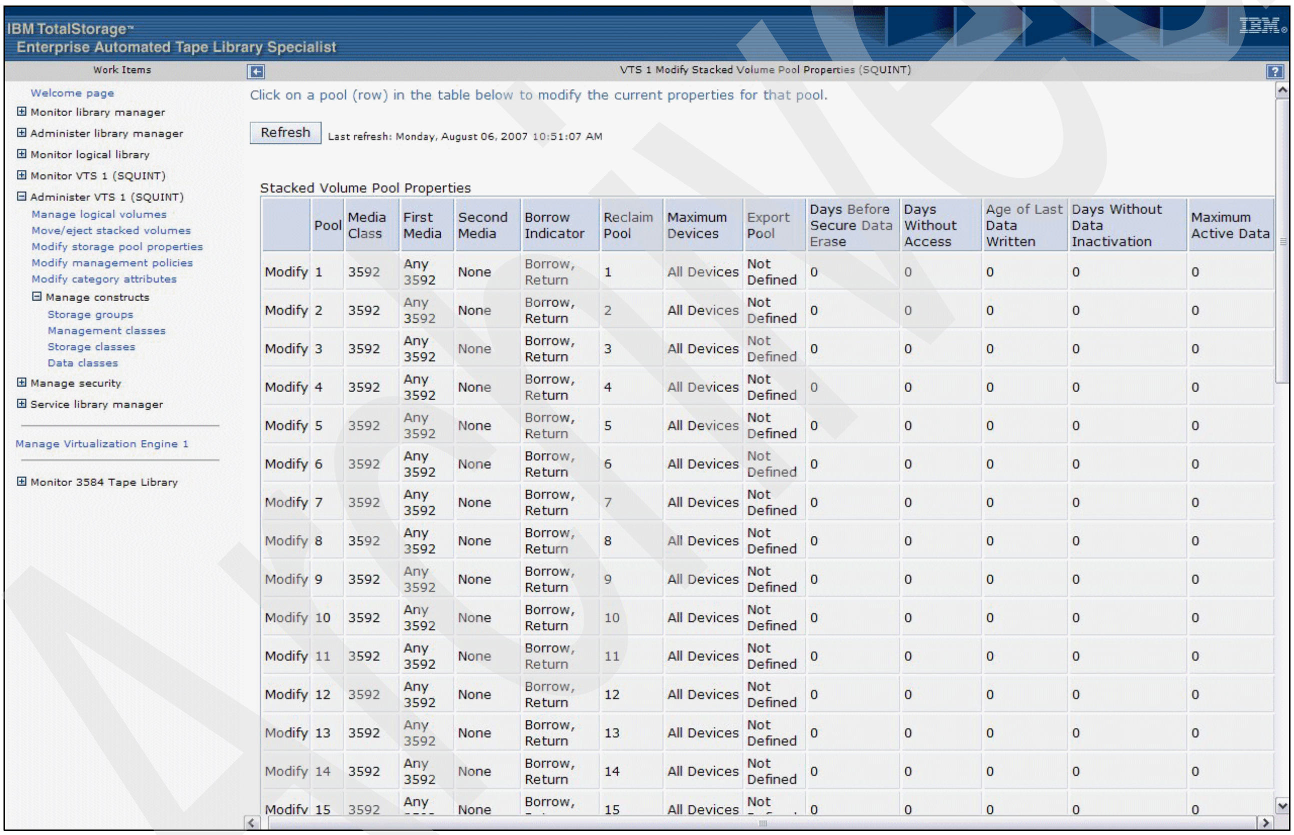

	Pool	Media Class	First Media	Second Media	Borrow Indicator	Reclaim Pool	Maximum Devices	Export Pool	Days Before Secure Data Erase	Days Without Access	Age of Last Data Written	Days Without Data Inactivation	Maximum Active Data
Modify	1	3592	Any 3592	None	Borrow, Return	1	All Devices	Not Defined	0	0	0	0	0
Modify	2	3592	Any 3592	None	Borrow, Return	2	All Devices	Not Defined	0	0	0	0	0
Modify	3	3592	Any 3592	None	Borrow, Return	3	All Devices	Not Defined	0	0	0	0	0
Modify	4	3592	Any 3592	None	Borrow, Return	4	All Devices	Not Defined	0	0	0	0	0
Modify	5	3592	Any 3592	None	Borrow, Return	5	All Devices	Not Defined	0	0	0	0	0
Modify	6	3592	Any 3592	None	Borrow, Return	6	All Devices	Not Defined	0	0	0	0	0
Modify	7	3592	Any 3592	None	Borrow, Return	7	All Devices	Not Defined	0	0	0	0	0
Modify	8	3592	Any 3592	None	Borrow, Return	8	All Devices	Not Defined	0	0	0	0	0
Modify	9	3592	Any 3592	None	Borrow, Return	9	All Devices	Not Defined	0	0	0	0	0
Modify	10	3592	Any 3592	None	Borrow, Return	10	All Devices	Not Defined	0	0	0	0	0
Modify	11	3592	Any 3592	None	Borrow, Return	11	All Devices	Not Defined	0	0	0	0	0
Modify	12	3592	Any 3592	None	Borrow, Return	12	All Devices	Not Defined	0	0	0	0	0
Modify	13	3592	Any 3592	None	Borrow, Return	13	All Devices	Not Defined	0	0	0	0	0
Modify	14	3592	Any 3592	None	Borrow, Return	14	All Devices	Not Defined	0	0	0	0	0
Modify	15	3592	Any	None	Borrow,	15	All Devices	Not	0	0	0	0	0

Figure 4-21 IBM ETL Specialist panel to modify storage pool properties

2. Click **Modify** in front of the storage pool you want to change. The panel shown in Figure 4-22 will be presented.

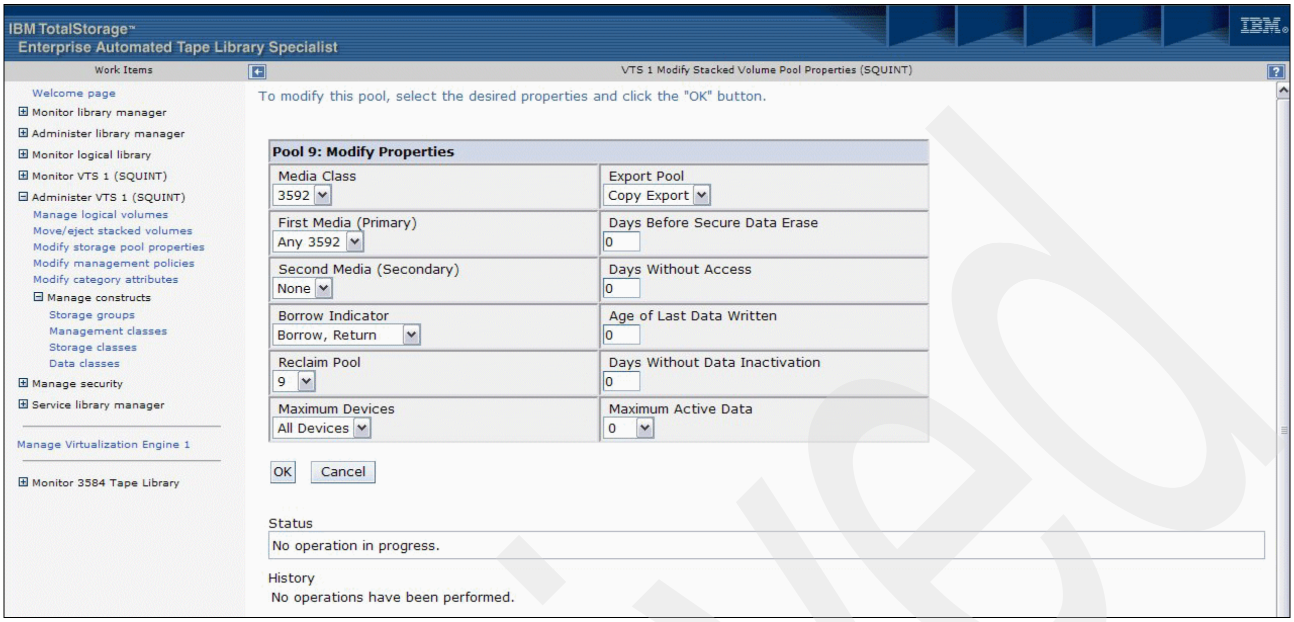

Figure 4-22 Modify Stacked Volume Pool Properties update panel from ETL Specialist

3. In the Pool Properties for TS7700 section, modify the fields as needed:

Media Class
Shows the device type 3590 or 3592. In a TS7700, only TS1120 Tape Drives can be installed. However, only one device type is supported per pool and its supported media.

First Media
Shows the preferred media for the Media Class specified for this pool. If the TS1120 tape drives are running in E05 native mode, for a Media Class of 3592, this field can contain *JA, JB, JJ,* or *Any 3592*. If you have J1A drives attached to the TS7700, or if the TS1120 drives are operating in J1A Emulation mode, the First Media can contain *only JA, JJ,* or *JA/JJ.*

Second Media
Contains the secondary media for this pool. The Second Media can only be set to *None* for a TS7700, if the First Media contains *Any 3592*. If a specific media type has been selected for the First Media, you can select any of the remaining media types for the Second Media.

Borrow Ind
Defines whether and how the pool is populated with scratch cartridges. *Borrow, Return* enables borrowing from the Common Scratch Pool When volumes become scratch, they are returned to the CSP. *Borrow, Keep* enables borrowing from the CSP. When volumes become scratch, they remain in the borrowing pool and are not returned. *No Borrow, Return* does not allow borrowing from the CSP. When any volumes become scratch, they are returned to the CSP (this setting can be used to empty pools). *No Borrow, Keep* does not allow borrowing from the CSP. When any volumes become scratch, they remain in the pool and are not returned.

Reclaim Pool
Lists the pool to which active logical volumes will be assigned when stacked volumes of this pool are reclaimed by the TS7700. The stacked volume itself is treated according to the Borrow Indicator.

Maximum Devices Defines the maximum number of physical tape drive that this pool can use for pre-migration.

Export Pool Select Not Defined or Copy Export for this pool. In order for a logical volume to be exported it must be written to a stacked volume whose pool is designated as a Move Export Pool or a Copy Export Pool.

> **Note:** The pool that is selected as an Export Pool cannot be designated as a Primary Pool by a Storage Group construct. If a pool that is designated a Primary Pool is selected as an Export Pool, an error message is displayed.

Days Before Secure Data Erase

Enter the Days Before Secure Data Erase. This guarantees that the logical volume data that has expired is erased and cannot be recovered by any reasonable means after the “days” value. A value of zero (0) turns off the use of this criterion.

Days Without Access

Can be 0 to 365 days. If a stacked volume with active data on it has not been accessed because of a recall for the specified number of days, the volume becomes eligible for reclaim. A value of zero deactivates this policy.

Age of Last Data Written

Can be 0 to 365 days. If a stacked volume with active data on it has not been written to for the specified number of days, the volume becomes eligible for reclaim. A value of zero deactivates this policy.

Days Without Data Inactivation

Can be 0 to 365 days. If a customer specified period of time has elapsed since the last decrease in the amount of active data on a volume and the amount of data falls below the specified threshold, the volume becomes eligible for reclaim. A value of zero deactivates this policy.

Maximum Active Data (%)

Can be 0 to 95% and defines the threshold for Days without Data Inactivation. A value of zero deactivates these policies.

You can use one or a multiple of the reclamation policies listed above. All policies for a pool are taken into account independently. You should also consider that:

- Each pool can have a different set of policies.
- When dual copy is used, properties for primary and secondary pool should be defined.
- Reclamation workload will increase during data migration.

The History table at the bottom of the panel shown in Figure 4-22 on page 165 is displayed on all panels of the ETL Specialist that allow modification of parameters. It is also displayed on those panels of the Library Manager console that allow modifications of library or TS7700 settings, and it serves multiple purposes:

- Indicate actions that are currently in progress or have already completed.
- Coordinate remote users (ETL Specialist and Library Manager console operator)
- Notify the current user if another user has performed the same kind of action while the current user is preparing to perform the same or a similar action.

Encryption is defined using the TS7700 Management Interface. Refer to 4.5.4, “Review and define pool encryption settings” on page 178 for details.

4.4.2 Creating Storage Groups

On the z/OS host, the Storage Group construct determines into which tape library a logical volume is written. Within the TS7700, the Storage Group construct allows you to define the Storage Pool you want the logical volume written to.

Even before you define the first Storage Group, there is always at least one Storage Group present: the Default Storage Group that is identified by eight dashes (--------). This Storage Group cannot be deleted, but you can modify it to point to a different Storage Pool, for example.

If you assign a logical volume a Storage Group at the host, and if this Storage Group has not been defined on the Library Manager before, this Storage Group will be created on the Library Manager using the specifications of the Default Storage Group. You can define up to 256 Storage Groups including the Default.

To define a new Storage Group, or modify or delete an existing Storage Group, click **Administer VTS** → **Manage constructs** → **Storage Groups.** Figure 4-23 shows the Storage Group Definition panel of the ETL Specialist.

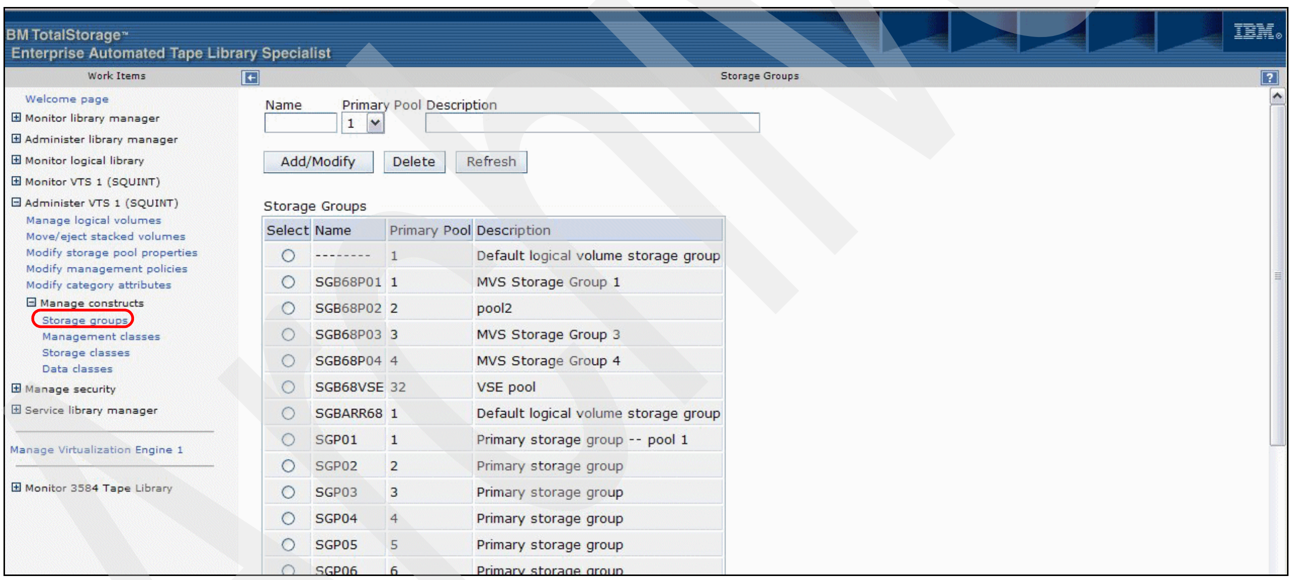

Figure 4-23 Library Manager panel to manage Storage Groups

To create a Storage Group, proceed as follows:

1. Enter a 1to 8 character alphanumeric name in the *Name* field. This name must be unique within the Storage Group construct names. Use the same name as your host-defined DFSMS Storage Group name.
2. Enter the number of the pool you want to use for this Storage Group in the Primary Pool field. You can choose from any of the 32 storage pools, but you cannot enter 0 for the Common Scratch Pool. The default Primary Pool is Pool 1.
3. Optionally, enter a short description in the Description field.
4. Select **Add/Modify**.

To modify a Storage Group, select from the current Storage Groups presented in the list box. Use the mouse or keyboard to highlight the Storage Group you want to modify. Modify the primary pool or description, and select **Add/Modify**.

To delete a Storage Group, select from the current Storage Groups presented in the list box. Use the mouse or keyboard to highlight the Storage Group to delete, then select **Delete**.

Inadvertently deleting a Storage Group will have no impact on the accessibility of the logical volumes. This is because at allocation time, constructs are assigned to the logical volume. These constructs are stored in the TS7700 database. There can be an effect if the deleted Storage Group pointed to a different primary storage pool than the default storage pool. When the logical volume is closed, the TS7700 will query the logical volumes primary pool and the Library Manager will return the primary pool associated with the default Storage Group.

The purpose of the History Table at the bottom of the ETL Specialist panel shown in Figure 4-23 is to:

- Indicate actions that are currently in progress or have already completed.
- Coordinate remote users (ETL Specialist and Library Manager console operator).
- Notify the current user if another user has performed the same kind of action while the current user is preparing to perform the same or a similar action.

> **Tip:** In an environment with multiple pools in use, we recommend not to assign a Storage Group to Pool 1. If you then encounter physical cartridges and logical volumes in Pool 1, you have an indication that your Storage Group definitions on the host and on the Library Manager might not be matching.

4.4.3 Creating Management Classes

You can define, through the Management Class, whether you want to have a dual copy of a logical volume within the same TS7700. In a grid configuration, you will most likely choose to copy logical volumes over to the other TS7700s instead of creating a second copy in the same TS7700. In a Single Cluster configuration, however, you might want to protect against media failures by using the dual copy capability. The second copy of a volume can be in a pool that is designated as an Export Copy pool. Refer to Chapter 2.3.4, “Copy Export” on page 46 for more information.

If you want to have dual copies of selected logical volumes, you must use at least two Storage Pools, because the copies cannot be written to the same Storage Pool as the original logical volumes.

A Default Management Class is always available. It is identified by eight dashes (--------) and cannot be deleted. If you assign a logical volume a Management Class at the host, and if this Management Class has not been defined on the Library Manager before, it will be created on the Library Manager using the specifications of the Default Management Class. You can define up to 256 Management Classes including the Default.

To define a new Management Class, or modify or delete an existing Management Class, click **Administer VTS → Manage constructs → Management Class.** Figure 4-24 shows the Management Class Definition panel of the ETL Specialist.

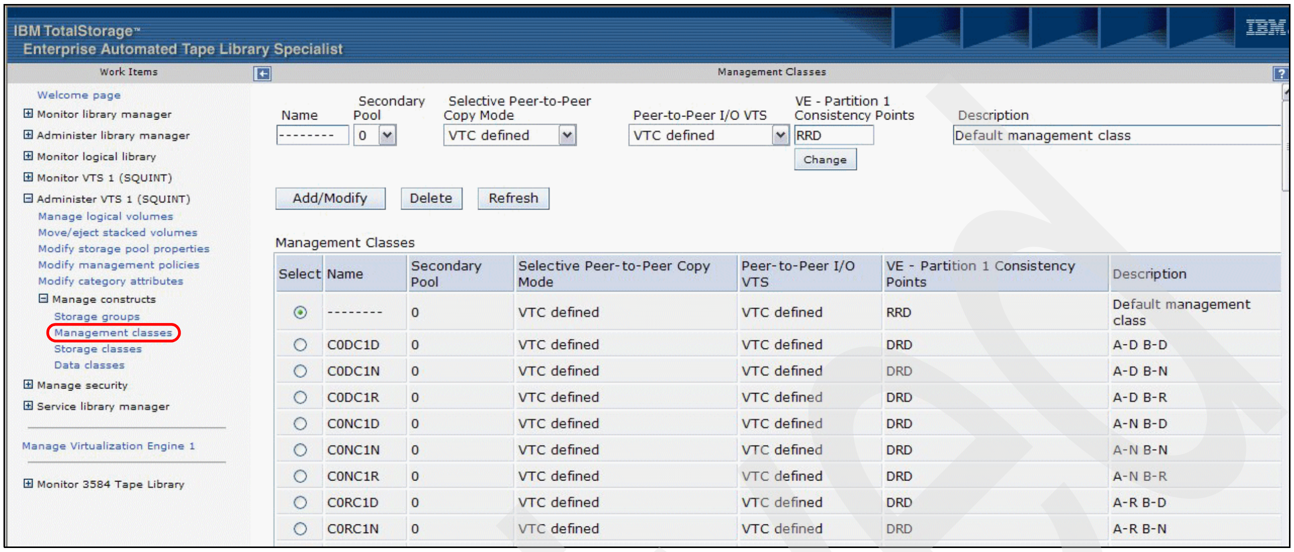

Figure 4-24 Manage Management Classes panel

To add a Management Class, proceed as follows:

1. Enter a 1 to 8 character name in the Name field. This name must be unique in the Management Class construct names.
2. Specify dual copy pool, if you want to create a dual copy within the same TS7700. Select one of the following:
   – Specify the secondary pool number (1-32). This will determine in which physical pool the copy of the logical volumes will reside.
   – If '**00**' is selected, no secondary copy will be made.
3. Enter a short description in the Description field.
4. Click **Add/Modify**.

> **Important:** The settings for PTP Copy Control and PTP I/O VTS only apply for IBM TotalStorage Peer-to-Peer Virtual Tape Server (PtP VTS).
>
> For implementation of an IBM TS7700 Virtualization Engine, those fields should be left at their defaults saying "VTC defined".

To modify a Management Class, select one from the current Management Classes presented in the list box. Use the mouse or keyboard to highlight the Management Class you want to modify. Modify pools, copy consistency points if you have a Multi Cluster Grid, or description. Select **Add/Modify**.

For a TS7700 Multi Cluster Grid you will also need to define the copy consistency points. Using the TS7700 Virtualization Engine in a Multi Cluster Grid is explained in 4.6, "TS7700 Multi Cluster Grid definitions" on page 183.

To delete a Management Class, select it from the current Management Classes presented in the list box. Use the mouse or keyboard to highlight the Management Class you want to delete. Select **Delete**.

The purpose of the history table is to:

- Indicate actions that are currently in progress or have already completed.
- Coordinate remote users (ETL Specialist and Library Manager console operator).
- Notify the current user if another user has performed the same kind of action while the current user is preparing to perform the same or similar action.

Note: Depending on its use, dual copy might have an impact on overall throughput in your TS7700. Prior to utilizing this function of Outboard Policy Management, ensure that your TS7700 is configured correctly. To help identify any impact, contact your IBM representative, who can use a modelling tool. BatchMagic has been updated to accommodate the new functions.

4.4.4 Creating Storage Classes

Using the Storage Class construct, you can influence when a logical volume is to be removed from cache.

A Default Storage Class is always available. It is identified by eight dashes (--------) and cannot be deleted. If you assign a logical volume a Storage Class at the host, and if this Storage Class has not been defined on the Library Manager before, it will be created on the Library Manager using the specifications of the Default Storage Class. You can define up to 256 Storage Classes including the Default.

To define a new Storage Class, or modify or delete an existing Storage Class, click **Administer VTS** → **Manage constructs** → **Storage Class.** Figure 4-25 shows the Storage Class Definition panel of the ETL Specialist.

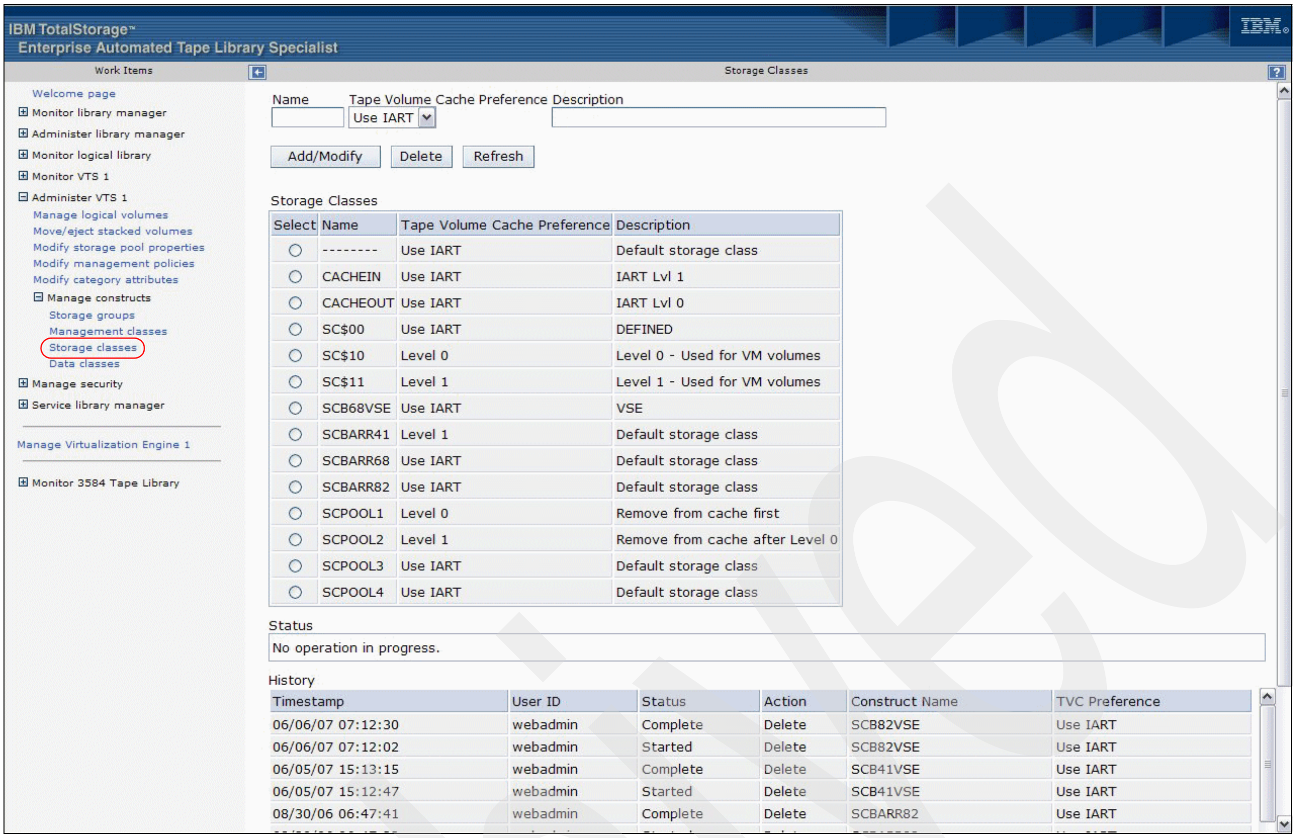

Figure 4-25 Storage Classes panel from the ETL Specialist

To add a Storage Class, proceed as follows:

1. Enter a 1 to 8 character name in the Name field. Define the same name as host-defined DFSMS Storage Class construct. The name must be unique within the Storage Class construct names.
2. Determine the Tape Volume Cache Preference. Three options are available:

   IART Instructs TS7700 to honor the host-supplied Initial Access Response Time (IART) value.

   Level 0 Instructs TS7700 that it can remove these volumes from Tape Volume Cache (TVC) as soon as they are copied to tape. There are actually two processes going on. If space is needed, the largest PG0 volumes are preferred to be removed over PG1. During idle time, the smallest PG0 volumes are removed.

   > **Note:** We recommend that you not use Preference Group 0 in a Three-Cluster Grid. The use of Preference Group 1 or IART is fine for a Three-Cluster Grid.

   Level 1 Instructs TS7700 to remove volumes from TVC after the copy has been made, but only if space is needed in the TVC. Use the Least Recently Used (LRU) algorithm.

3. Enter a short description of the Storage Class in the Description field.
4. Click **Add/Modify**.

To modify a Storage Class, select from the current Storage Classes presented in the list box. Use the mouse or keyboard to highlight the Storage Class you want to modify. Modify the TVC preference or description. Select **Add/Modify**.

To delete a Storage Class, select from the current Storage Classes presented in the list box. Use the mouse or keyboard to highlight the Storage Class you want to delete. Select **Delete**.

The purpose of the history table is to:

- Indicate actions that are currently in progress or have already completed.
- Coordinate remote users (ETL Specialist and LM operator).
- Notify the current user if another user has performed the same kind of action while the current user is preparing to perform the same or similar action.

4.4.5 Creating Data Classes

From a z/OS perspective, for SMS managed tape, the DFSMS Data Class defines:

- Media type parameters
- Recording technology parameters
- Compaction parameters
- Performance Scaling
- Segmentation

For the TS7700, only the Media type, Recording technology, and Compaction parameters are used. The use of larger logical volume sizes is controlled through Data Class.

A Default Data Class is always available. It is identified by eight dashes (--------) and cannot be deleted. If you assign a logical volume a Data Class at the host, and if this Data Class has not been defined on the Library Manager before, it will be created on the Library Manager using the specifications of the Default Data Class. You can define up to 256 Data Classes including the Default.

To define a new Data Class, or modify or delete an existing Data Class, click **Administer VTS → Manage constructs → Data Class.** Figure 4-26 on page 173 shows the Data Class Definition panel of the ETL Specialist.

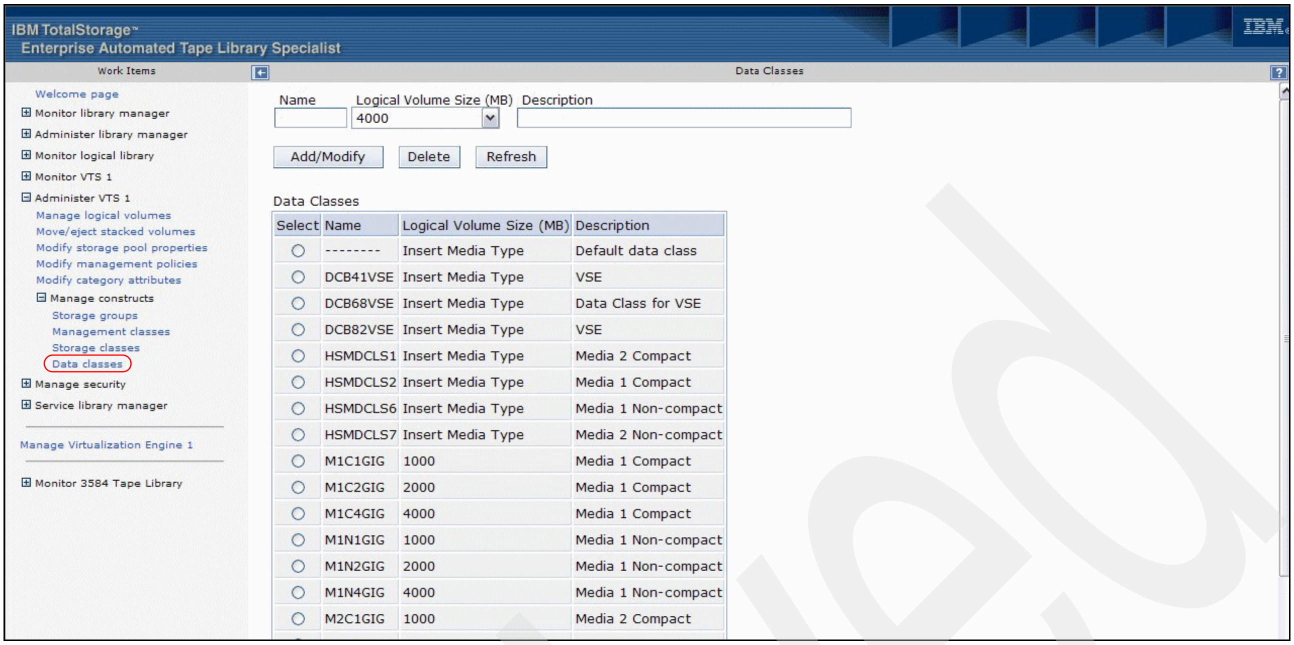

Figure 4-26 Manage Data Classes panel

To add a Data Class, proceed as follows:

1. Enter a 1 to 8 character name in the Name field. The name must be unique in the Data Class construct names.
2. Enter the Logical Volume Size in MB you want to use with this Data Class.
   You can select:

1000	For 1000 MB logical volumes
2000	For 2000 MB logical volumes
4000	For 4000 MB logical volumes
Insert Media Type	To use the media type (MEDIA1 or MEDIA2) that was defined at the time of logical volume insert.

3. Enter a short description in the Description field.
4. Click **Add/Modify**.

To modify a Data Class, select from the current Data Classes presented in the list box. Use the mouse or keyboard to highlight the Data Class you want to modify. Modify the description and click **Add/Modify**.

To delete a Data Class, select from the current Data Class presented in the list box. Use the mouse or keyboard to highlight the Data Class you want to delete and click **Delete**.

The purpose of the history table is to:

- Indicate actions that are currently in progress or have already completed.
- Coordinate remote users (ETL Specialist and Library Manager console operator).
- Notify the current user if another user has performed the same kind of action while the current user is preparing to perform the same or a similar action.

4.5 Setup of the TS7700 Virtualization Engine

In this section, we introduce the definitions to be made for the TS7700 Virtualization Engine using its TS7740 Management Interface (MI), a Web browser interface that allows you to perform these major tasks:

- Definitions made for you by the IBM Service Representative (SSR) during installation of the IBM TS7700 Virtualization Engine
- Definitions through the TS7740 Virtualization Engine Management Interface (MI)
- Insertion of logical volumes through the IBM TS7740 Virtualization Engine MI

4.5.1 TS7700 definition using the TS7740 Management Interface

The Management Interface (MI) for the TS7700 Virtualization Engine complies with industry standards that make it easier to manage devices from different manufacturers. It is a Web-based GUI that you access using any standard Web browser by entering the TS7740 Virtualization Engine IP address. The setup for the initial IP address is part of the hardware installation that is done by the IBM System Service Representative (SSR). Later on you can change or modify those IP addresses through the Management Interface at any time.

The implementation of storage management standards allows you to manage storage devices from different manufacturers.

To log on to the TS7700 Management Interface you need to enter a user ID and a password, which comes from manufacturing as the default values “admin” and “admin”. You can change and modify these values as well through the Management Interface later at any time. Using the admin/admin login, you can add and administer other users.

Figure 4-27 shows the welcome display after logon to a Single Cluster IBM TS7700 Virtualization Engine.

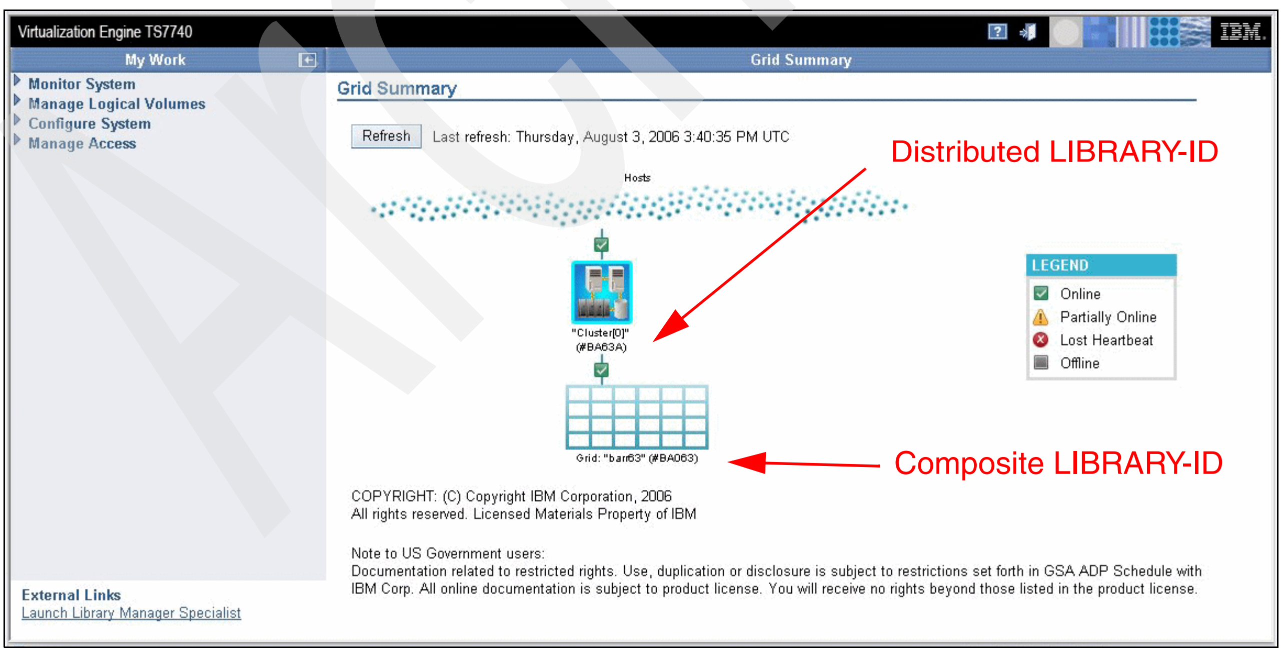

Figure 4-27 Welcome panel of the TS7700 Virtualization Engine Management Interface

As you can see in the example in Figure 4-27, even if it is a Single Cluster Grid TS7700 Virtualization Engine, it is always displayed as though it were part of a Multi Cluster Grid installation. You can see the configured Composite LIBRARY-ID and the configured Distributed LIBRARY-ID, as well as the configured cluster and grid names which are the Distributed and Composite Library names defined to DFSMS.

The next example (Figure 4-28) shows a grid summary panel of an IBM TS7700 Virtualization Engine configured in a Multi Cluster Grid environment.

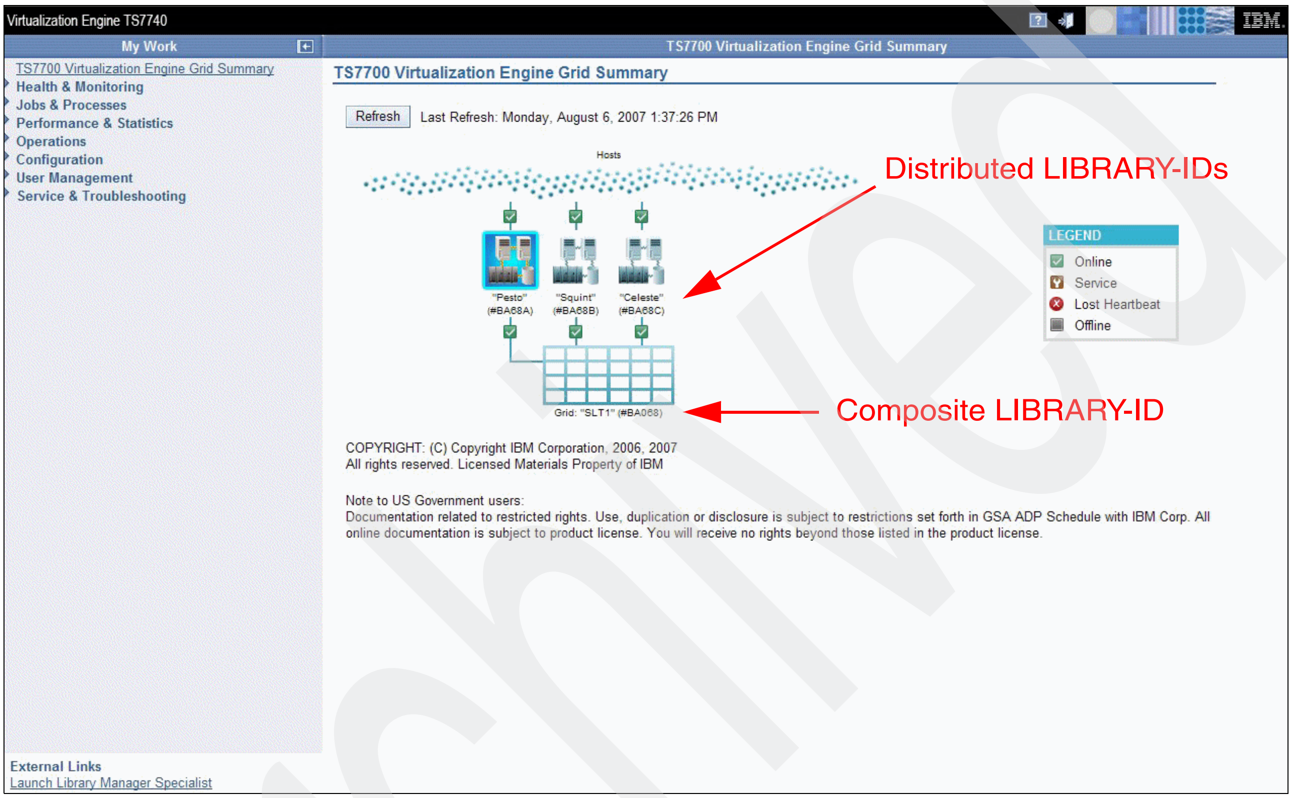

Figure 4-28 Welcome panel (Grid Summary) of the TS7700 VE Management Interface

4.5.2 Cache Enablement and Performance Increment license key entry

The amount of disk cache capacity and performance capability are enabled using feature license keys. You will receive feature license keys for the features that you have ordered. Each of these feature increments allow you to tailor the subsystem to meet your disk cache and performance needs.

You use the Feature Licenses panel to activate feature licenses in the IBM Virtualization Engine TS7700. This panel is shown in Figure 4-29 on page 176. Select an action of **Activate New Feature License** from the list and click **Go**. The license key entry field panel appears as shown in Figure 4-30 on page 176. Then enter the license key into the fields provided and select **Activate**.

To remove a license key, select the feature license to be removed, select the action of **Remove Selected feature License** from the list, and click **Go**.

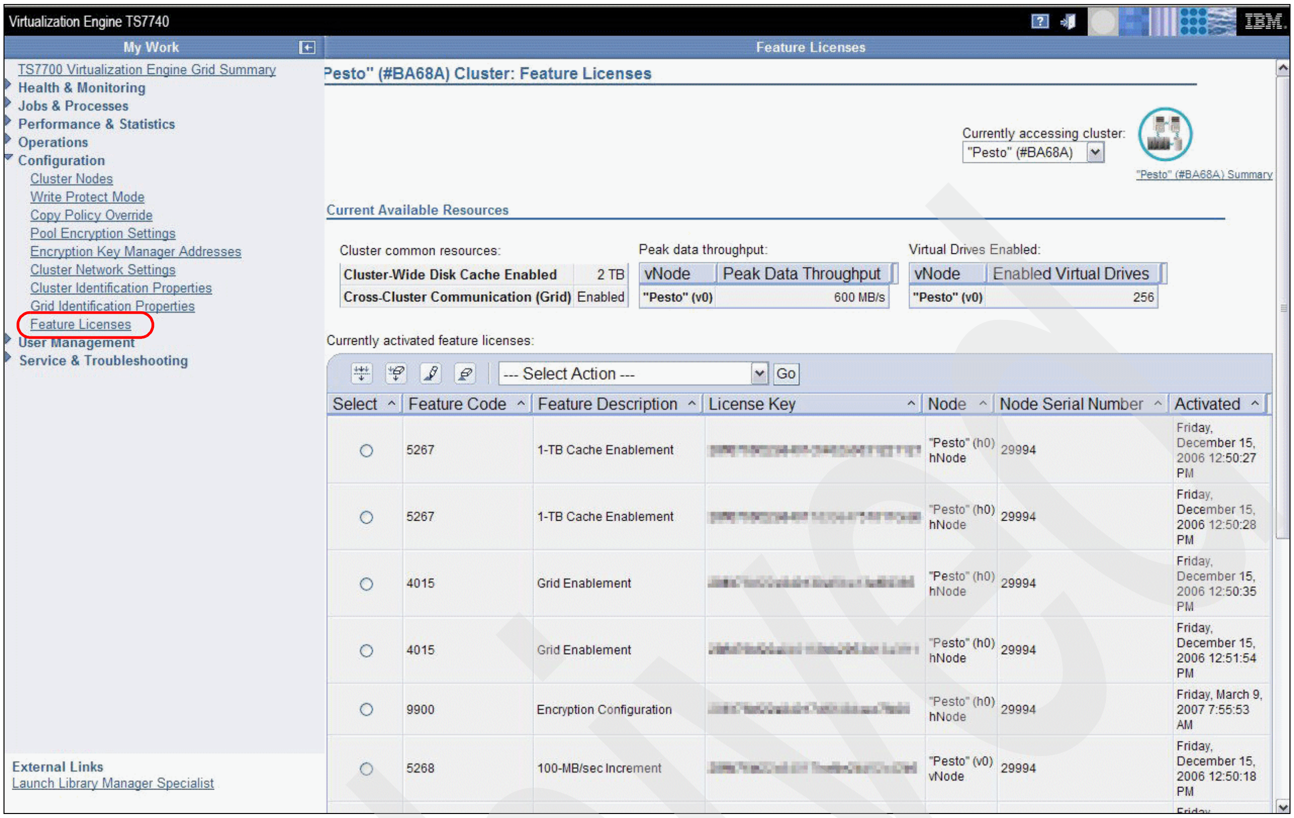

Figure 4-29 Feature Licenses Panel

> **Note:** Removal of a Cache Enablement feature is not allowed because it would severely impact host performance.
>
> A Performance Increment can be removed; this will impact performance immediately.

When you select the action of **Activate New Feature License**, the Feature License entry panel appears as shown in Figure 4-30. When you enter a valid feature license key, the feature will be activated.

> **Note:** Performance Increments become active immediately. Cache Increments become active within 30 minutes.

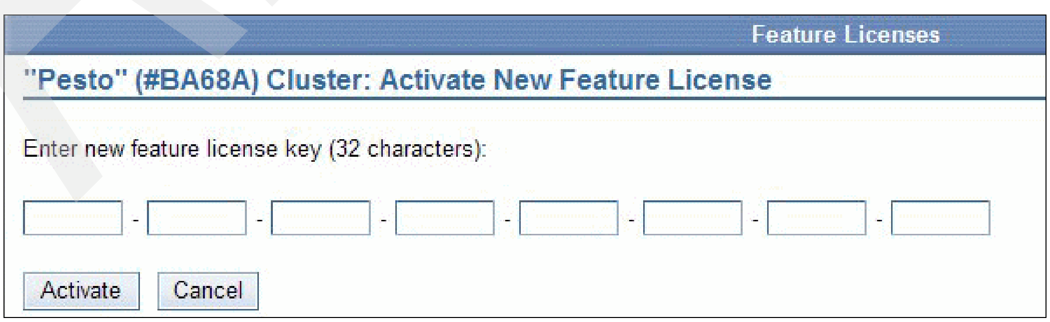

Figure 4-30 Feature License Entry Panel

4.5.3 Define Copy Policy Override Settings

With the IBM TS7700 Virtualization Engine, you can define and set the optional override settings that influence the selection of the I/O Tape Volume Cache (TVC) and replication responses. The settings are specific to a cluster in a Multi Cluster Grid configuration, meaning that each cluster can have different settings if desired. The settings take effect for any mount requests received after the settings were saved. Mounts already in progress are not affected by a change in the settings. You can define and set the following settings:

- Prefer local cache for fast ready mount requests
- Prefer local cache for non-fast ready mount requests
- Force volumes mounted on this cluster to be copied to the local cache
- Allow fewer RUN consistent copies before reporting RUN command complete

You can view and modify these settings from the TS7700 Management Interface when you select **Copy Policy Override Settings** from the Configure System pull-down, as shown in Figure 4-31.

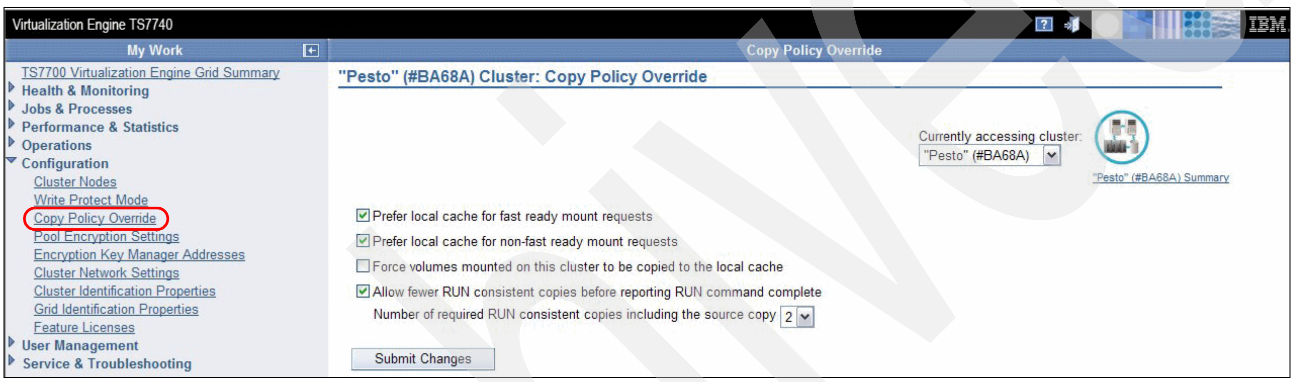

Figure 4-31 MI panel Copy Policy Override settings

The settings you can select in the MI panel shown in Figure 4-31 are:

- Prefer local cache for fast ready mount requests

  A fast ready mount selects a local copy as long as a copy is available and a cluster copy consistency point is not specified as No Copy in the Management Class for the mount. The cluster is not required to have a valid copy of the data.

- Prefer local cache for non-fast ready mount requests

  This override causes the local cluster to satisfy the mount request as long as the cluster is available and the cluster has a valid copy of the data, even if that data is only resident on physical tape. If the local cluster does not have a valid copy of the data, then default cluster selection criteria applies.

- Force volumes mounted on this cluster to be copied to the local cache

  For a non-fast ready mount, this override causes a copy to be performed on the local cluster as part of mount processing. For a fast ready mount, this setting has the effect of overriding the specified Management Class with a copy consistency point of Rewind/Unload for the cluster. This does not change the definition of the Management Class, but serves to influence the replication policy.

- Allow fewer RUN consistent copies before reporting RUN command complete

  If selected, the value entered at **Number of required RUN consistent copies including the source copy** will be used to determine the number of copies to override before the Rewind/Unload operation reports as complete. If this option is not selected, the management class definitions are to be used explicitly. Thus, the number of RUN copies can be from one to the number of clusters in the grid.

Note: In a Geographically Dispersed Parallel Sysplex (GDPS) all three Copy Policy Override Settings must be selected on each cluster, to ensure that wherever the GDPS primary site is, this TS7700 cluster is preferred for all I/O operations.

In case the TS7700 cluster of the GDPS primary site fails, you have to perform the following recovery actions:

1. Vary virtual devices from a remote TS7700 cluster online from the primary site of the GDPS host.
2. Manually invoke, through the TS7700 Management Interface, a Read/Write Ownership Takeover unless Automated Ownership Takeover Manager (AOTM) has already transferred ownership.

4.5.4 Review and define pool encryption settings

Use this page for viewing and modifying storage pool encryption settings on the TS7700 Virtualization Engine. The TS7700 Engine allows encryption by storage pool. If you are planning to enable encryption for one or more pools, use the panel shown in Figure 4-32 to set the encryption settings. You can modify the settings for several pools at the same time by selecting multiple pools. All of the selected pools will be modified to the same settings. You modify a pool's encryption settings by selecting one or more storage pools, selecting **Modify Encryption** from the list and selecting **Go**.

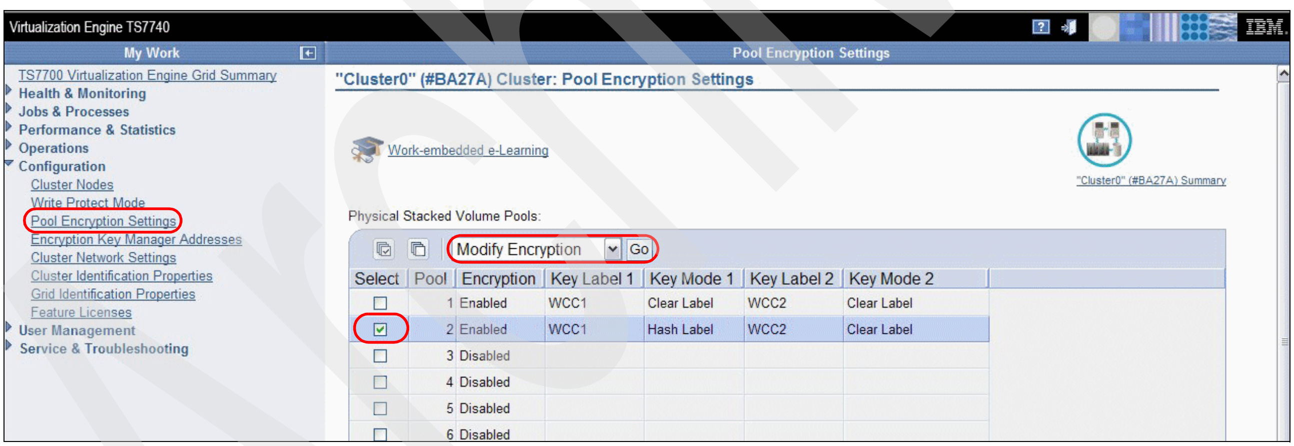

Figure 4-32 Pool Encryption Settings

The Pool Encryption Settings Modify panel displays, as shown in Figure 4-33.

Note: Pool encryption defaults to Disabled.

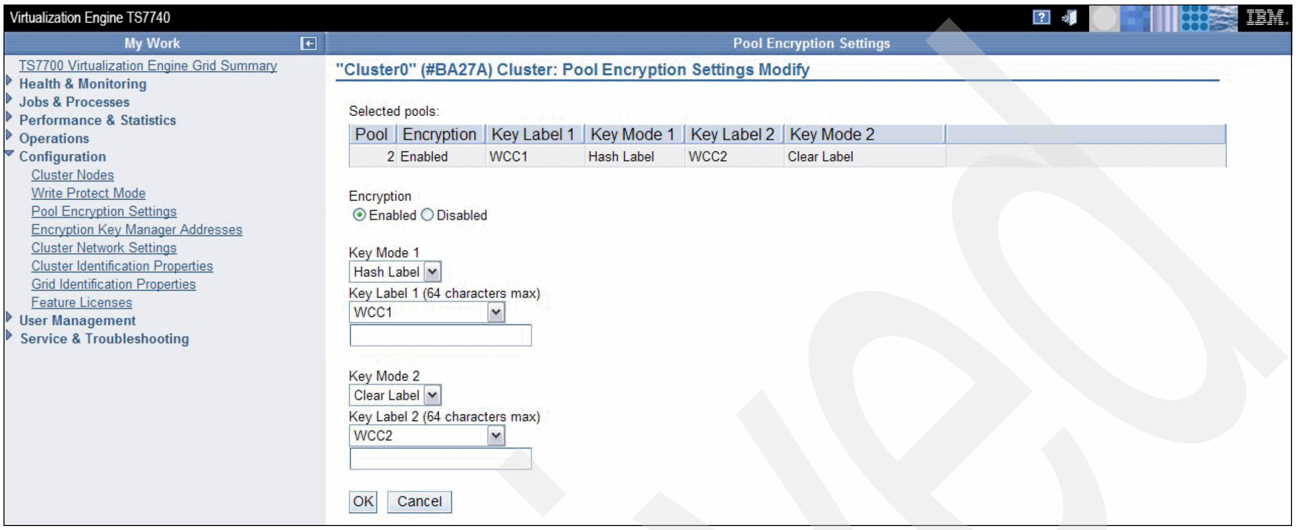

Figure 4-33 Pool Encryption Settings Modify

Referring to Figure 4-33, the selected pools are displayed at the top of the panel. Make the desired modifications and click **OK**.

You can then modify the pool encryption settings as follows:

Encryption Enabled or disabled for the pool.

Key Label 1/2 An alias for the key encrypting key (KEK) used by the encryption key manager. A user can select up to two key labels and the maximum length of each label can be 64 characters. Key Label 2 will be enabled only if Key Label 1 is enabled. A previously selected label can be used by selecting one from the previously selected key labels drop-down.

Key Mode 1/2 The method by which the Encryption Key Manager (EKM) identifies the public/private keys that were used to encrypt it.

Clear Label The key label points to an externally encoded data key (EEDK). It is possible that the same key can be in one keystore with one label and be imported into another keystore with a different label. This can cause problems across sites, such as a TS7700 Grid with a main site and a disaster recovery site. In the case that there are inconsistent label names, the disaster recovery site would be unable to decrypt.

Hash Label A numerical value is calculated from the public key, independent of any key label that can be assigned to the key. The hash is valid across multiple EKM/keystores even if different key labels are assigned. This mode avoids the inconsistency problems that occur with the Clear Label mode.

4.5.5 Review and define Encryption Key Manager addresses

If you are using encryption for the TS7700, you will need to define where the Encryption Key Managers (EKMs) are located. EKMs are defined with an IP address or domain name, and a port number. Each cluster in a grid that is using encryption must set up the EKM addresses. Select Encryption Key Manager Addresses from the MI menu. The EKM address page as shown in Figure 4-34 is presented.

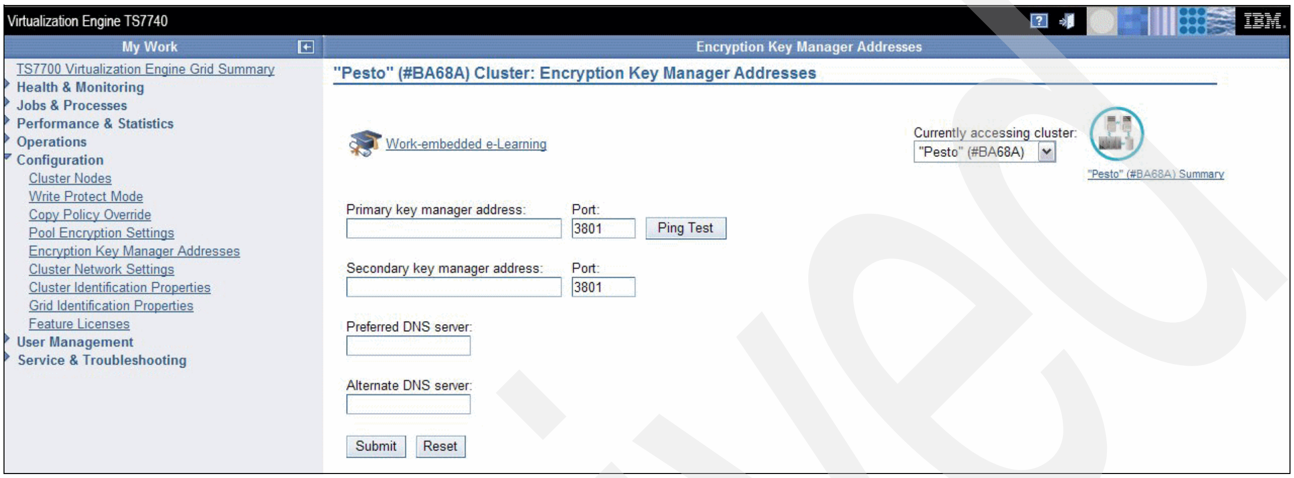

Figure 4-34 Encryption Key Manager addresses

Modify the following settings to configure the TS7700 connection to one or two Encryption Key Managers.

> **Note:** We highly recommend that two EKMs be set up on separate machines to provide redundancy. Connection to an EKM is required to read encrypted data.

Primary key manager address
: Enter the primary EKM IP address or domain name.

Secondary key manager address
: Enter the secondary EKM IP address or domain name.

Port
: Enter the port number to be used for each of the EKM addresses or domain name. The default port address setting is 3801.

Preferred DNS server
: Enter the primary Domain Name Server (DNS). DNS addresses are only needed if you specify a symbolic domain name for one of the key manager addresses rather than a numeric IP address. If you need to specify a DNS, it is highly recommended that you specify both a primary and an alternate so you do not lose access to your EKM due to one of the DNS servers being down or inaccessible.

Alternate DNS server
: Enter the secondary Domain Name Server. The secondary DNS is used in case the preferred DNS server is unavailable. If a preferred DNS server is specified, it is recommended to specify an alternate DNS as well.

4.5.6 Review and define cluster identification properties

Use this page for viewing and altering cluster identification properties for the IBM TS7700 Virtualization Engine. Select **Cluster Identification Properties** from the Configure System pull-down, as shown in Figure 4-35.

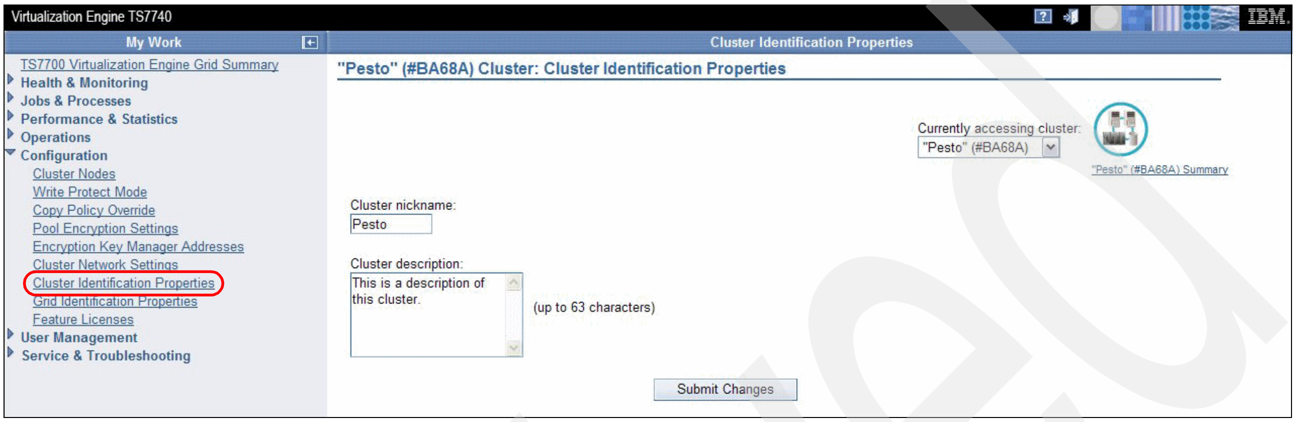

Figure 4-35 MI Cluster Identification Properties panel

In the MI panel shown in Figure 4-35, you can update the following fields:

- Cluster nickname: An alphanumeric, 8 character name for the cluster
- Cluster description: A short description of the cluster of up to 63 characters

To change cluster identification properties, edit the available fields and click **Submit Changes**.

> **Note:** We suggest you make the cluster nickname the same as the distributed library name that was defined through DFSMS.

4.5.7 Review and modify grid identification properties

Use this page for viewing and altering grid identification properties for the IBM TS7700 Virtualization Engine. Select **Grid Identification Properties** from the Configure System pull-down, as shown in Figure 4-36.

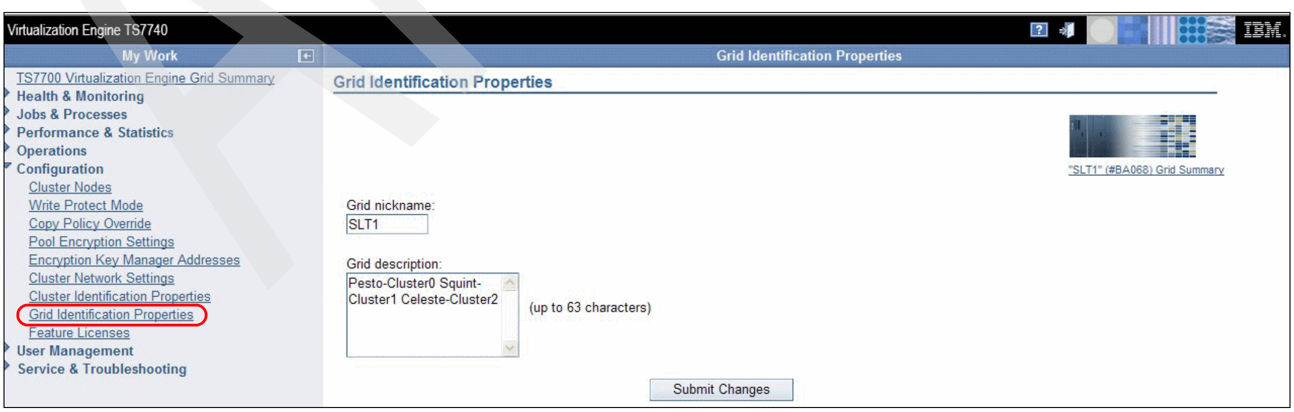

Figure 4-36 Grid Identification Properties page from the TS7700 Management Interface

In the MI panel shown in Figure 4-36, you can update the following fields:

- Grid nickname: An alphanumeric, 8 character name for the cluster
- Grid description: A short description of the cluster of up to 63 characters

To change grid identification properties, edit the available fields and click **Submit Changes**.

> **Note:** We suggest you make the grid name the same as the composite library name that was defined through DFSMS.

4.5.8 Insert logical volumes using the TS7700 Management Interface

Use this page to insert logical volumes to the TS7700 Virtualization Engine. Select **Insert Logical Volumes** as shown in Figure 4-37.

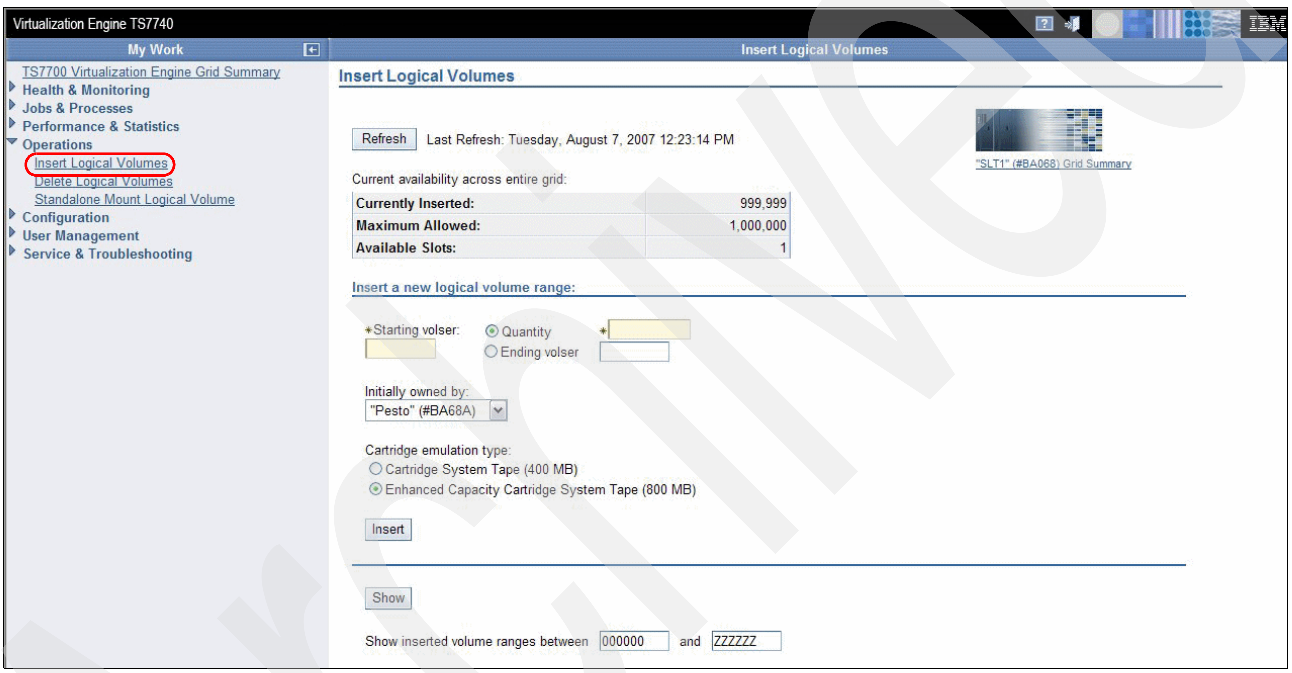

Figure 4-37 Insert Logical Volumes page from the TS7700 Management Interface

Specify the following fields to insert logical volumes:

Starting VOLSER The first logical volume to be inserted. The range for inserting logical volumes begins with this VOLSER identifier.

Quantity Select this option to insert a set amount of logical volumes starting with the starting VOLSER. The adjoining text field is where the quantity of logical volumes to be inserted is entered.

Ending VOLSER Select this option to insert a range of logical volumes. The adjoining text field is where the ending VOLSER identifier is entered.

Cartridge emulation type
Media type of the logical volume. Possible values are:

- Cartridge System Tape (400 MB)
- Enhanced Capacity Cartridge System Tape (800 MB)

Select Cluster Used to select the initial owning cluster of the logical volume.

Click **Insert** to insert the logical volumes with the selected options.

Note: You can insert up to 10 000 logical volumes at one time. This applies to both inserting a range of logical volumes and inserting a quantity of logical volumes.

4.6 TS7700 Multi Cluster Grid definitions

In this chapter we describe and illustrate what additional configuration tasks have to be done, or need to be considered, for an installation of the TS7700 Virtualization Engine in a Multi Cluster Grid configuration. In the previous product generation of the TS7700 Virtualization Engine, this kind of installation was called a Peer-to-Peer (PtP) Virtual Tape Server (VTS) configuration.

As shown in Figure 4-38, each cluster within a grid configuration has its own Cluster ID. This Cluster ID is configured during initial hardware installation. Each cluster within a grid configuration is typically attached to its own host. The grid communication is done through two 1 GB Ethernet connections between the TS7700 Virtualization Engines.

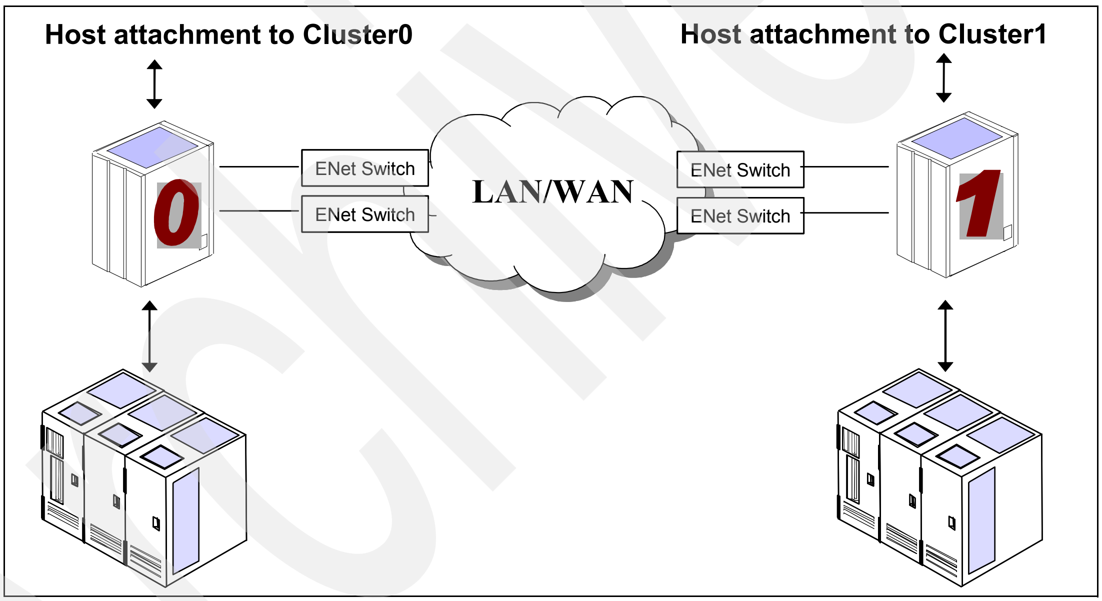

Figure 4-38 TS7700 in a Two-Cluster Grid configuration

4.6.1 Define Grid copy mode control

To define Grid copy mode control, you must set the data consistency point, which defines when a copy of data becomes consistent with the original.

The data consistency point is defined in the Management Class construct definition on the Library Manager. You can perform this task only for an existing grid system; in a single cluster configuration, you have fewer options during Management Class definition. Refer to 4.4.3, “Creating Management Classes” on page 168 for information about creating a Management Class for a Single Cluster Grid configuration.

The Management Class panel (Figure 4-39) is displayed through the ETL Specialist when you select **Management classes** from the Manage constructs pull-down of the Administer VTSx work items. There you find one or two additional menu options to open a separate window

(Figure 4-40) to select the Data Consistency Point for each attached TS7700. In this example the TS7700 Virtual Engine (named here *VE - Partition 1*) is part of a Multi Cluster Grid configuration. This additional pull-down menu is only displayed if a TS7700 is part of a Multi Cluster Grid environment.

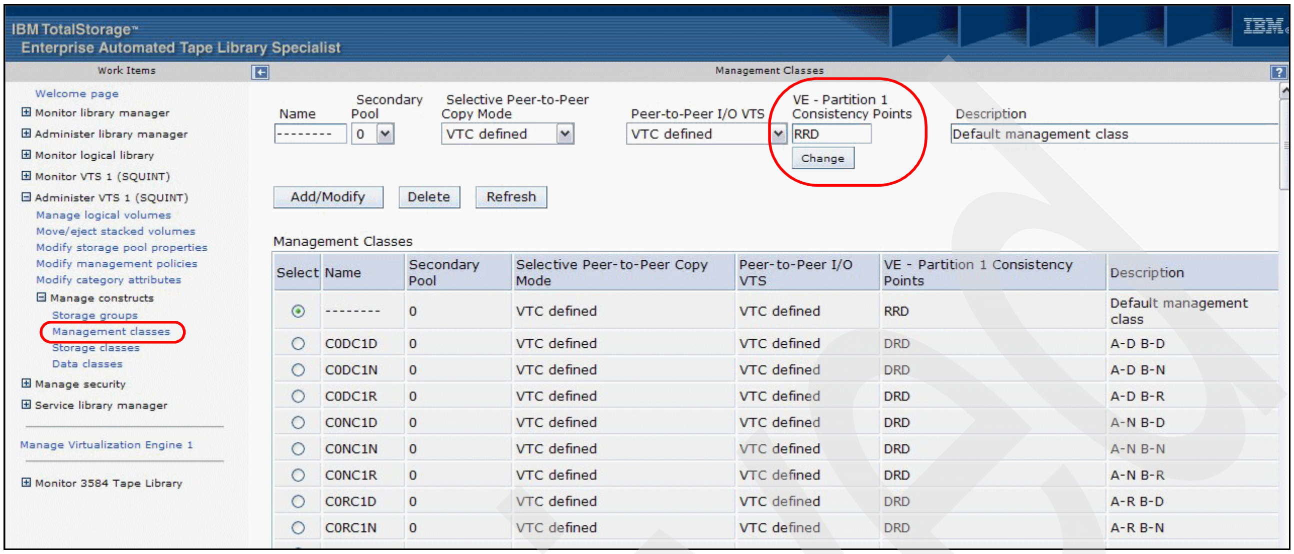

Figure 4-39 Manage Management Classes panel with pull-down for Data Consistency Points

As shown in Figure 4-40, you can choose between three consistency points per cluster:

Rewind/Unload (RUN) time
A valid version of the logical volume has been copied to this cluster as part of the volume unload processing.

Deferred Copy (after RUN)
A replication of the modified logical volume is made to this cluster after the volume had been unloaded.

No Copy
No copy is made to this cluster.

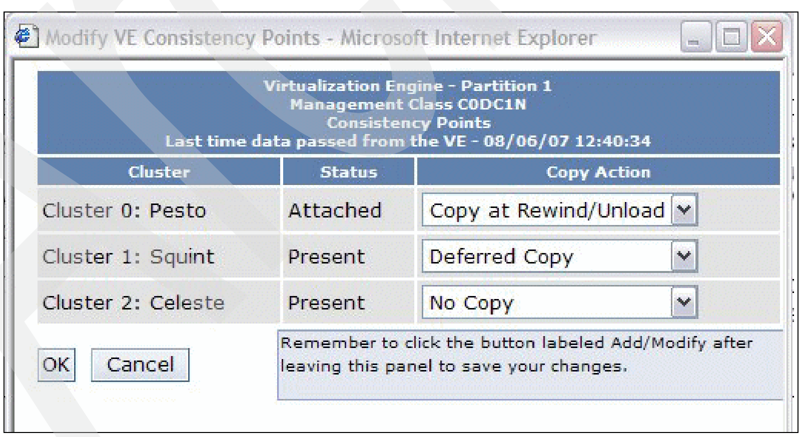

Figure 4-40 Modify VE Consistency Points through ETL Specialist

Table 4-2 provides an overview of the possible Cluster Copy Data Consistency Points for the Management Class in a Two-Cluster Grid, and their consequences. The table uses the following terminology:

- RUN = Copy at Rewind/Unload
- DEF = Deferred Copy
- NC = No Copy

Table 4-2 Possible settings of the Data Consistency Points combinations - Two-Cluster Grid

Cluster 0	Cluster 1	Consequence
RUN	RUN	Use this setting to have both clusters maintaining a consistent copy when the Rewind/Unload is acknowledged back to the host. The grid will manage the utilization of each cluster.
RUN	DEF	This ensures that Cluster 0's site will have a valid copy of the logical volume at job completion. Cluster 1 can have a valid copy at job completion, or sometime after job completion, depending on the virtual drive address selected and the override settings.
DEF	RUN	This ensures that Cluster 1's site will have a valid copy of the logical volume at job completion. Cluster 0 can have a valid copy at job completion, or sometime after job completion, depending on the virtual drive address selected and the override settings.
DEF	DEF	Use this setting if you do not care to which cluster the initial creation of the virtual volume is directed. A copy will be made as soon as possible after Rewind/Unload is acknowledged back to the host.
RUN	NC	Use this setting to specify Cluster 0 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 1. Note, the “force” override and virtual device address can have the system create a copy at Cluster 1.
NC	RUN	Use this setting to specify Cluster 1 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 0.
DEF	NC	Use this setting to specify Cluster 0 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 1.
NC	DEF	Use this setting to specify Cluster 1 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 0.
NC	NC	This setting is not supported, because it implies that the volumes should not be consistent anywhere. You will receive an error message when trying to specify this combination on the Library Manager.

Note: A Single Cluster Grid (standalone TS7700) always uses Rewind/Unload (RUN) as the Data Consistency Point.

A Three-Cluster Grid has approximately the same consequences as a Two-Cluster Grid. Table 4-3 provides a few select consistency point variations that are important for the Three-Cluster Grid configuration. The table assumes Clusters 0 and 1 are production clusters and Cluster 2 is a disaster recovery cluster. The table uses the following terminology:

- RUN = Copy at Rewind/Unload
- DEF = Deferred Copy
- NC = No Copy

Table 4-3 Possible settings of the Data Consistency Points combinations - Three Way Grid

Cluster 0	Cluster 1	Cluster 2	Consequence
RUN	RUN	DEF	Use this setting to have the production clusters maintaining a consistent copy when the Rewind/Unload is acknowledged back to the host. The grid will manage the utilization of each cluster. Cluster 2 can have a valid copy at job completion, or sometime after job completion, depending on the virtual drive address selected and the override settings.
RUN	RUN	NC	Use this setting to have the production clusters maintaining a consistent copy when the Rewind/Unload is acknowledged back to the host. The grid will manage the utilization of each cluster. No copy will be made to Cluster 2. Use this setting when a high availability copy is desired but a remote disaster recovery copy is not needed.
RUN	NC	DEF	Use this setting to specify Cluster 0 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 1. Cluster 2 can have a valid copy at job completion, or sometime after job completion, depending on the virtual drive address selected and the override settings. Note, the “force” override and virtual device address can have the system create a copy at Cluster 1. Use these settings when a high availability copy is not required but a remote disaster recovery is needed.
NC	RUN	DEF	Use this setting to specify Cluster 1 to have the initial valid copy when the Rewind/Unload is acknowledged back to the host. No copy will be made to Cluster 0. Cluster 2 can have a valid copy at job completion, or sometime after job completion, depending on the virtual drive address selected and the override settings. Note, the “force” override and virtual device address might have the system create a copy at Cluster 0. Use these settings when a high availability copy is not required but a remote disaster recovery is needed.
NC	NC	RUN	Use this setting to specify that only Cluster 2 is to have a valid copy. Clusters 0 and 1 will not have a copy. This setting is useful for disaster recovery testing on Cluster 2 so that test data is not created on the production clusters.

Because the Library Manager supports attachment of two TS7700/VTS subsystems, you can attach two TS7700 Virtualization Engines to one Library Manager. Each TS7700 subsystem can be part of a Multi Cluster Grid configuration. However, the two TS7700s attached to the same Library Manager cannot be in the same Multi Cluster Grid.

The Data Consistency Points for the TS7700 Multi Cluster Grid copy control are defined through the Management Class definition for each cluster. These settings can be different for each cluster—so it is possible to define asymmetric definitions.

For a Two-Cluster Grid, let us assume that Cluster 0 has defined RUN/DEF and Cluster 1 has defined, for the same Management Class, DEF/RUN. Although this setting sounds contradictory, it might make sense under some circumstances: If a deferred copy is sufficient, and the local cluster's Tape Volume Cache is supposed to be used as the I/O VTC, these settings will ensure that the initial copy is always created locally before a deferred copy is done to the remote cluster.

4.6.2 Management construct definitions and transfer

It is essential to understand that the outboard policy management construct definitions that you make at the Library Manager are on a Library Manager Level only. Because the TS7700 Virtualization Engines of a Multi Cluster Grid configuration are managed by different Library Managers, you need to make the construct definitions on *both* Library Managers, or you can transfer the definitions from one Library Manager to the other.

This function cannot be performed through the ETL Specialist—you must do it from the Library Manager console. To do so, on the Library Manager console, select **Commands** → **Systems Management** → **Manage Constructs and Pools** → **Transfer (Backup/Restore) LM Administrative Data.**

The panel shown in Figure 4-41 displays, and you can select that data that you want to transfer to the other Library Manager. You can also use this function to create a backup copy of your APM definitions before or after you make significant changes. These backup file sets are stored on any preformatted standard 3.5 inch floppy diskette or can be stored on your hard drive.

When you select to copy the data to your hard drive, the files are stored in the `C:\LM` directory with a file extension of .XFR. You can use FTP to electronically copy these files to another Library Manager. Be sure to copy all of the .XFR files to the C:\LM directory on the other Library Manager. After FTPing the files, select the **Restore from hard drive** button, then click **Start transfer**.

In the example shown in Figure 4-41, **Backup from VTS 1** has been selected. To restore the LM Administrative Data on the other VTS, click the buttons on the right side of the panel, and select the corresponding Restore Options.

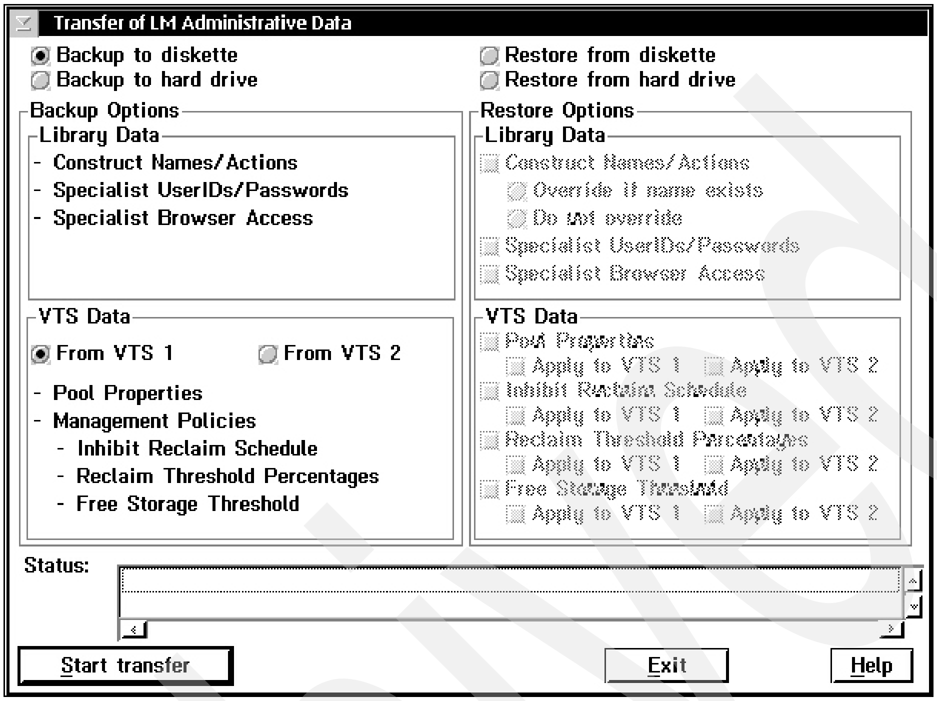

Figure 4-41 Transfer Library Manager Administrative Data

Important: In a TS7700 Multi Cluster Gird configuration, you must define outboard policy management constructs and pool definitions on all Library Managers.

4.6.3 Data management settings

These definitions for the TS7700 Virtualization Engine are optional. The IBM Service Representative configures these settings during installation of the TS7700 or at any time later using the TS7700 service menus. These data management settings are:

- Copy files preferred to reside in cache
- Recalls preferred for cache removal

Copy files preferred to reside in cache

Normally, the Tape Volume Cache (TVC) in both TS7700 Virtualization Engines in a Multi Cluster Grid is managed as one to increase the likelihood that a needed volume will be in cache. By default, the volume on the TS7700 selected for I/O operations is preferred to stay in cache on that TS7700, where the copy made on the other TS7700 is preferred to be removed from cache:

- Preferred to stay in cache means that when space is needed for new volumes, the oldest volumes are removed first. This algorithm is called the Least Recently Used (LRU) algorithm. This is also referred to as Preference Group 1 (PG1).
- Preferred to be removed from cache means that when space is needed for new volumes, the largest volumes are removed first, regardless of when they were written to the cache. This is also referred to as Preference Group 0 (PG0).

For a TS7700 Virtualization Engine running in a dual production Multi Cluster Grid configuration, both TS7700s are being selected as the I/O TVCs and will have the original volumes (newly created or modified) preferred in cache while the copies to the other TS7700 will be preferred to be removed from cache. The result is that each TS7700 TVC is filled with unique, newly created or modified volumes, thereby roughly doubling the amount of cache as seen by the host.

For a TS7700 Virtualization Engine running in a Multi Cluster Grid configuration used for business continuance, particularly when all I/O is preferenced to the local Tape Volume Cache (TVC), this default management method might not be desired. In the case where the remote site of the Multi Cluster Grid is used for recovery, the recovery time is minimized by having most of the needed volumes already in cache. What is really needed is to have the most recent copy volumes remain in the cache, not being preferred out of cache.

Based on customer requirements, the IBM Service Representative (SSR) can set or modify this control through the TS7700 Virtualization Engine service menu for the remote TS7700.

- Default is `off`.
- When `off`, copy files are managed as Preference Group 0 volumes (prefer out of cache first by largest size).
- When `on`, copy files are managed based on the Storage Class construct definition.

Recalls preferred for cache removal

Normally, a volume recalled into cache is managed as though it were newly created or modified, because it resides in the TS7700 selected for I/O operations on the volume. A recalled volume will displace other volumes in cache.

In the case where the remote TS7700 is used for recovery, the recovery time is minimized by having most of the needed volumes in cache. However, it is not likely that all of the volumes to restore will be resident in the cache and that some amount of recalls will be required. Unless you can explicitly control the sequence of volumes to be restored, it is likely that recalled volumes will displace cached volumes that have not yet been restored from, resulting in further recalls at a later time in the recovery process.

When the restore completes from a recalled volume, that volume is no longer needed. Such volumes need to be removed from the cache after they are accessed so that they minimally displace other volumes in the cache.

Based on customer requirements, the IBM Service Representative (SSR) can set or modify this control through the TS7700 Virtualization Engine service menu of the remote TS7700.

- Default is `off`.
- When `off`, recalls are managed as Preference Group 1 volumes (LRU).
- When `on`, recalls are managed as Preference Group 0 volumes (prefer out of cache first by largest size).

This control is independent of and not affected by cache management controlled through the Storage Class SMS construct as well. Storage Class cache management affects only how the volume is managed in the I/O Tape Volume Cache (TVC).

4.6.4 Delete expired volume data in a Multi Cluster Grid environment

With this optional function, the data associated with logical volumes assigned to a category with the Fast-Ready attribute set (scratch categories), are deleted permanently from the TS7700 inventory after a specified time period.

The function is provided to help you manage the amount of physical cartridge space occupied by expired data.

If the Delete Expired Volume Data function is used, the next time the host mounts a logical volume that was deleted by this function, the TS7700 generates an image of a newly initialized volume. Although most tape management systems tolerate a newly initialized volume for a scratch mount, it is recommended that you verify this with your tape management system provider before implementing the Delete Expired Volume Data function.

4.6.5 Implementing Outboard Policy Management for non-z/OS hosts

Outboard Policy Management and its constructs are exploited only in DFSMS host environments where OAM has knowledge of the construct names and dynamically assigns and resets them. z/VM, z/VSE, TPF, z/TPF, and other hosts do not have knowledge of the construct names and therefore cannot change them. In addition, non-z/OS hosts use multiple Library Manager (LM) categories for scratch volumes and, therefore, can use multiple logical scratch pools on the Library Manager, as shown in Table 4-4.

Table 4-4 Scratch pools and Library Manager volume categories

Host SW	LM Scratch Categories	# Scratch pools	LM Private Cat.
VM (+ VM/VSE)	X’0080’ - X’008F’	16	X’FFFF’
BTLS	X’0FF2’ - X’0FF8’, X’0FFF’	8	X’FFFF
Native VSE	X’00A0’ - X’00BF’	32	X’FFFF
Unassigned (can be used by Open Systems)	X’012F’ - X’0FF1’ X’0FF9’ - X’0FFE’ X’F00F’ - X’FEFF’		

Note: In a TPF environment manipulation of construct names for volumes can occur when they are moved from scratch through a user exit. The user exit allows the construct names and clone volser to be altered. If the exit is not implemented, TPF does not alter the construct names.

TPF use of categories is very flexible and is difficult to contain in the above table. TPF allows each drive to be assigned a scratch category. Concerning private categories, each TPF has their own category that volumes are assigned to when they are mounted.

For more information refer to the zTPF Information Center:

http://publib.boulder.ibm.com/infocenter/tpfhelp/current/index.jsp

Because the hosts do not know about constructs, they ignore static construct assignment, and the assignment is kept even when the logical volume is returned to scratch. Static assignment means that at insert time of logical volumes, they are assigned construct names as well. Construct names can also be assigned later at any time.

Note: In a z/OS environment, OAM controls the construct assignment and will reset any static assignment made before at the Library Manager level using the ETL Specialist or the Library Manager console. Construct assignments are also reset to blank, when a logical volume is returned to scratch.

From the ETL Specialist, select **Manage Logical Volumes** from the Administer VTSx work items to obtain the panel shown in Figure 4-42, which allows you to change existing logical volumes.

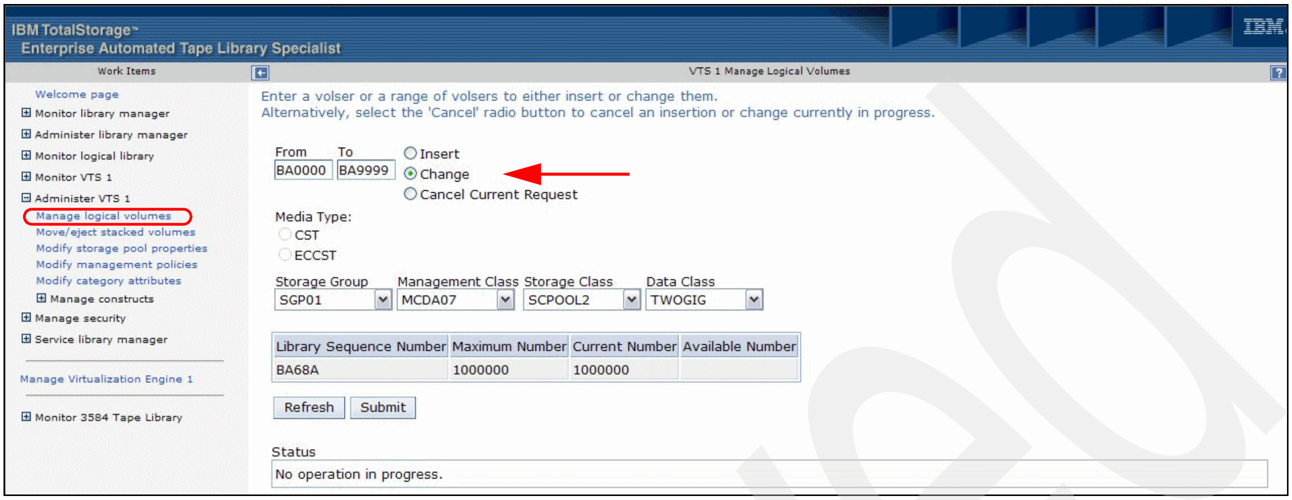

Figure 4-42 Change existing logical volumes from the ETL Specialist

To implement Outboard Policy Management for non-z/OS hosts attached to an IBM TS7700 VE, proceed as follows:

1. Define your pools and constructs as described earlier in this chapter.
2. Then insert your logical volumes in groups through the TS7700 Management Interface (MI) as described in 4.5.8, “Insert logical volumes using the TS7700 Management Interface” on page 182.
3. After that go to the IBM Library Manager and assign the required static construct names to the logical volume ranges through the “change existing logical volume” function. We recommend that you define groups of logical volumes with the same construct names assigned and, during insert processing, direct them to different LM volume categories so that all volumes in one LM volume category have identical constructs assigned.

Host control is given through utilization of the appropriate scratch pool. By requesting a scratch mount from a specific scratch category, the actions defined for the constructs assigned to the logical volumes in this category are executed at Rewind/Unload of the logical volume.

You should make sure that, when the logical volume is returned to scratch, it is returned to the correct Library Manager volume category.

5

Software implementation

In this chapter we describe how to implement the IBM System Storage TS7700 on the Series z host system. The tasks you have to complete include the definition of the hardware to the hosts through the Hardware Configuration Definition (HCD) dialog boxes and other, operating system-specific tasks depending on the System z software you are using.

In a z/OS environment, DFSMS provides automated system management of tape libraries. In addition to the implementation steps described in this chapter, you can find a step-by-step implementation cookbook in Appendix F, “TS7700 Virtualization Engine implementation, step-by-step” on page 591.

5.1 Host implementation considerations

From a host perspective, the TS7700 as a subsystem supports 16 tape control units, each with 16 IBM 3490E tape drives, for a total of 256 tape drives in a Single Cluster configuration. In a Two-Cluster Grid, 32 tape control units provide 512 virtual tape drives, and 48 control tape control units with a total of 768 tape drives are provided in a Three-Cluster Grid. Each TS7700 can be attached physically through two or four FICON channels.

The host does not know whether it is dealing with “real” 3490E tape drives or with the virtual 3490E tape drives of the TS7700. Therefore, the TS7700 with virtual 3490E tape drives is defined just like multiple physical IBM 3490-A10 controllers with 16 addresses through the Hardware Configuration Definition (HCD) dialog boxes.

Before you can use the TS7700, you need to define it to the System z host through HCD. Because the virtual tape drives of the TS7700 are library resident, you must define them through the HCD dialog, specifying LIBRARY=YES. If FICON directors are being installed for the first time, the directors themselves also have to be defined in the IOCP and HCD input/output definition file (IODF). See the subsequent sections for more information about HCD definitions.

In a z/OS environment, you then must define the TS7700 logical library to SMS. Update the SMS constructs and ACS routines to direct mounts to the TS7700. See 5.3, “TS7700 software definitions for z/OS” on page 209 for more details on the implementation in a System z environment. You might need to update Missing Interrupt Handler (MIH) values as well, as described in 5.2.5, “Set values for the Missing Interrupt Handler” on page 208.

The software implementation steps for z/VM and z/VSE are described in 5.4, “Software implementation in z/VM and z/VSE” on page 223. For TPF-related implementation details, refer to 5.5, “Software implementation in TPF” on page 229.

After defining the TS7700 to a system by whichever method, verify that the devices can be brought online. Also plan to update the expected IPL configuration so that this is automatic and coincides with the implementation through the 3494 or 3953 Library Manager and host software.

5.1.1 Sharing the TS7700 in multiple hosts

Each logical library has its own library sequence number, which is used to define the logical library to the host. Each logical library, a Composite Library in the case of the TS7700, looks like a separate library to the host. A TS7700 can be shared by multiple System z systems, VM VSE, and TPF systems.

Sharing can be achieved in two different ways: By logically dividing it into different partitions (partitioning) or by allowing all attached systems to sequentially access all physical as well as logical volumes (sharing).

Sharing of an IBM Automated Tape Library means that all attached hosts have the same access to all volumes in the tape library. To achieve this true sharing you need to share the host control data sets, that is, the tape management system inventory, the catalog information, and the tape configuration database (TCDB), among the attached hosts. In a non-SMS environment, all systems must share the ICF catalog that contains the BTLS inventory.

In general, these requirements can be met only in a single-platform environment. In this configuration only one global tape volume scratch pool is available.

5.1.2 Partitioning the TS7700 between multiple hosts

Partitioning is the solution if you need to dedicate the use of volume ranges to certain systems or complexes or different host platforms. Dividing one or more libraries into logical libraries is the easiest way to allow different hosts to access them. Each host or complex owns its own set of drives and volumes, which another system or complex cannot access without manual intervention. Each system knows only about its part of the library.

Partitioning is also appropriate for the attachment to a z/OS logical partition (LPAR) for testing. If there is a need for running a test environment with a date different from the actual date, as it was the case during Y2K tests, you should have a separate TCDB and tape management system inventory for the test complex. For details about sharing and partitioning a library, refer to *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

5.1.3 Logical path considerations

The TS7700 attaches to the host system or systems through two or four FICON channels. Each FICON channel provides 256 logical paths. With a four FICON configuration, that will result in a total of 1024 logical paths per TS7700.

You can use the following formula to calculate the number of logical paths that are required in an installation:

Number of logical paths = number of LPARs x number of CUs

This formula assumes all LPARs access all control units in the TS7700 with all channel paths.

For example, if one LPAR has 16 CUs defined, then you are using 16 logical paths of the 256 logical paths that are available. Table 5-1 shows the number of logical paths that are used in different scenarios.

Table 5-1 Logical paths

Number of CUs defined	Number of LPARs	Logical Paths Used
16	8	128
16	16	256
8	32	256

If you might have more than 16 LAPRs, we recommend that you do not define more than 8 CUs per LPAR.

For more information about the planning and implementation of FICON channels and operating in FICON native (FC) mode, see *FICON Native Implementation and Reference Guide*, SG24-6266. That book also discusses the FICON and Fibre Channel architectures, terminology, and supported topologies.

Define one tape control unit (CU) in the HCD dialog for every 16 virtual devices. Up to eight channel paths can be defined to each control unit. A logical path can be thought of as a 3-element entity:

- A host port
- A TS7700 port
- A logical control unit in the TS7700

Note: A reduction in the number of physical paths will reduce the throughput capability of the TS7700 and the number of available logical paths. A reduction in control units will reduce the number of virtual devices available for any individual host.

5.1.4 Library names, library IDs, and port IDs

Library names, library IDs, and port IDs are defined in the hardware but also used in the host definition of the TS7700. Thus, several definitions must be made in HCD and others must be made in SMS. You can use the Library names and IDs needed in a z/OS implementation listed in Table 5-2 as an example, and you can create a similar table during your planning phase to use in later steps.

Table 5-2 Sample of library names and ID's needed in a TS7740 implementation

TS7700 virtual names	SMS Name	LIB-ID	HCD defined?	SMS defined?
IBMC1 (Composite)	IBMC1	77401	Yes	Yes
IBMD1TU (Distributed Tucson)	IBMD1TU	11312	No	Yes
IBMD1PH (Distributed Phoenix)	IBMD1PH	11307	No	Yes
IBMD1SJ (Distributed San Jose)	IBMD1SJ	11300	No	Yes

Distributed Library name and Composite Library name

Those two names are entered by your IBM Service Representative (SSR) during installation and are arbitrary. The Distributed Library name is entered through the IBM 3953 or 3494 Library Manager (LM) console and the Composite Library name (also called Domain Name) during the TS7700 initial configuration and cannot start with the letter *V*.

Both names are not directly connected to the settings you specify on the host site during Interactive Storage Management Facility (ISMF) Library definition, but it is useful to associate the same names at the host for the same Distributed and Composite LIBRARY-IDs, to make administration easier.

Note: We recommend that the Distributed and Cluster library names entered at the host, the Library Manager, and the TS7700 Management Interface be the same. Although they do not have to be the same, it will simplify management of the subsystem.

LIBRARY-ID and LIBPORT-ID

LIBRARY-ID and LIBPORT-ID are z/OS HCD parameters that allow HCD to provide the logical library configuration information that is normally obtained by the operating system at IPL time. If the devices are unavailable during IPL, the HCD information allows the logical tape devices to be varied online (when they subsequently become available to the system) without reactivating the IODF.

Note: We strongly recommend that you specify the LIBRARY-ID and LIBPORT-ID in your HCD/IOCP definitions. It reduces the likelihood of having to reactivate the IODF when the library is not available at IPL, as well as providing enhanced error recovery in some cases. It can protect you from having to IPL when you make changes to your I/O configuration.

Distributed LIBRARY-ID

Each of the three IBM 3953 Library Manager partitions (Native,TS7700-1/VTS-1, and TS7700-2/VTS-2) is assigned a unique, five-digit hexadecimal number, called a *Sequence Number* by the IBM Service Representative (SSR) during hardware installation. This number is the *Distributed LIBRARY-ID*. It is an arbitrary number that you can select.

We recommend that you use the last five digits of the hardware serial number if it only consists of hexadecimal characters. For the native part of a TS3500/3953 library, this hardware serial number is the last five digits of the serial number of the 3952-F05 frame. For each TS7700 Distributed LIBRARY-ID, it is the last five digits of the serial number of the TS7700 itself.

If you are installing a new Multi Cluster Grid configuration, you might consider choosing LIBRARY-IDs that clearly identify the cluster and the grid. For example, the Distributed LIBRARY-IDs of your first Three-Cluster Grid configuration might be:

Cluster 0:	BA01A
Cluster 1:	BA01B
Cluster 2:	BA01C

The Composite Library-ID for this Three-Cluster Grid can then be BA010.

Note: Although both ways are correct, we recommended that you use the hardware number ID for distributed library because that number is used only in the HCD definition step. Because the name that displays in all system messages is the SMS library name, the important name to distinguish libraries is the SMS name.

Composite LIBRARY-ID

The composite LIBRARY-ID defined by your IBM Service Representative (SSR) is arbitrary as well. It is entered during hardware installation at the TS7700 (as well as the Composite Library name). We have no specific recommendation for what to use for this five hex-character sequence number. You might want to consider using a name out of the hex-characters A, B, C, D, E, F that makes it easily recognizable as being associated with a Composite Library. These sequence numbers must match the LIBRARY-IDs used in the HCD library definitions and the LIBRARY-IDs listed in the ISMF Tape Library Define panels.

Note: In a Single Cluster configuration, do not use the same sequence number for the Distributed and the Composite Library.

LIBPORT-ID

The LIBPORT-ID reflects the order in which the tape control units are connected to the Library Manager and provides the tape drive pool ID, which is transparent and only used by allocation and JES3.

DEVICE SERVICES QTAPE command

In an existing installation, you can use the DEVSERV QTAPE system command to determine what to specify. All tape drives (logical or physical) connected to a given logical control unit (LCU, CU) have the same LIBPORT-ID. Therefore, you only have to issue the DS QT command once per control unit (for any logical device number in that string of 16).

The command syntax is:

`DS QT,`*`devnum`*`,1,RDC`

Where:

`DS`	Device service
`QT`	Query tape
`devnum`	Device address
`1`	Number of devices to be displayed
`RDC`	Read device characteristics

Figure 5-1 details the output of a DS QT system command.

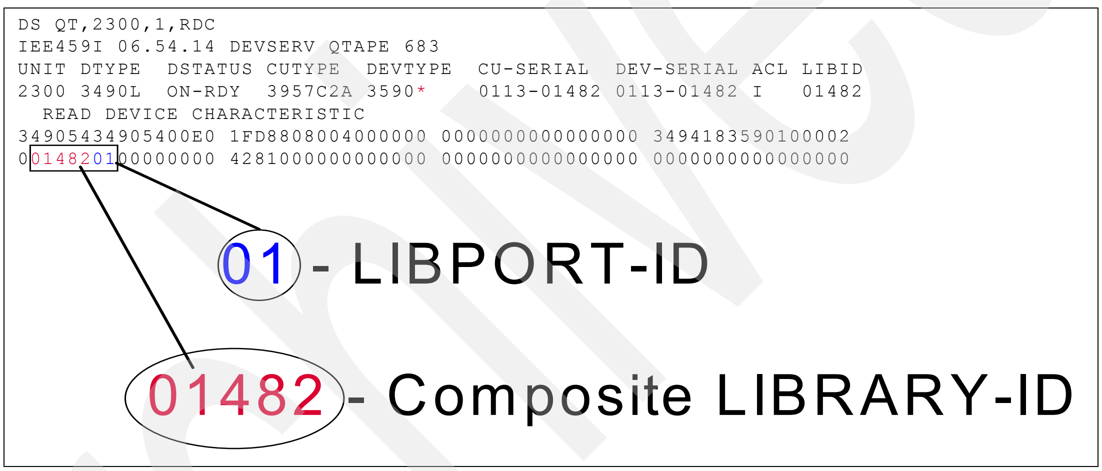

Figure 5-1 Sample DEVSERV QT command output

From the DEVSERV Q T command in Figure 5-1 you can derive LIBRARY-ID from the Composite Library and LIBPORT-ID. The real device type of the physical devices is unknown to the host and DEVSERV always shows 3590 as DEVTYPE.

> **Note:** You can get the real device type located in the Distributed Library from the Host Console request command:
>
> `LIBRARY REQUEST,`*`<Distributed Library Name>`*`,PDRIVE`
>
> The short form of LIBRARY REQUEST is LI REQ.

5.2 Hardware Configuration Definition

In this section we describe the process of defining the TS7700 through Hardware Configuration Definition (HCD) panels. The most important points to observe are as follows:

- HCD definitions are required.
- Per TS7700, up to sixteen 3490 tape control units can to be defined, with 16 x 3490E drives each.
- Keep the link address blank when no FICON director is used.
- Specify LIBRARY=YES when using system-managed tape.

We recommend that the z/OS system administrator, system programmer, hardware planner or similar person complete the information in Table 5-3 with the definitions that are needed to create the HCD. In general, they will know the information that is required, such as available CU numbers and channel paths.

Table 5-3 HCD definitions table

Chpid	CU	CUADD	Link	Devices	ADD	LIB-ID	Libport
		0			00-0F	77401	01
		1			00-0F	77401	02
		2			00-0F	77401	03
		3			00-0F	77401	04
		4			00-0F	77401	05
		5			00-0F	77401	06
		6			00-0F	77401	07
		7			00-0F	77401	08
		8			00-0F	77401	09
		9			00-0F	77401	0A
		A			00-0F	77401	0B
		B			00-0F	77401	0C
		C			00-0F	77401	0D
		D			00-0F	77401	0E
		E			00-0F	77401	0F
		F			00-0F	77401	10

5.2.1 Defining devices through HCD

You can define up to 16 control units with 16 devices each per cluster in the grid configuration. Use CUADD=0 through CUADD=7 and LIBPORT-IDs of 01 through 08 for the first eight control units as shown in Table 5-4.

Table 5-4 CUADD and LIBPORT-ID for the first set of 256 virtual devices

CU	1	2	3	4	5	6	7	8
CUADD	0	1	2	3	4	5	6	7
LIBPORT-ID	01	02	03	04	05	06	07	08

For the ninth to sixteenth control units, use CUADD=8 through CUADD=F and LIBPORT-IDs of 09 through 10. See Table 5-5.

Table 5-5 CUADD and LIBPORT-ID for the second set of virtual devices

CU	9	10	11	12	13	14	15	16
CUADD	8	9	A	B	C	D	E	F
LIBPORT-ID	09	0A	0B	0C	0D	0E	0F	10

Figure 5-2 and Figure 5-3 on page 201 show the two important panels for specifying a tape control unit.

```
------------------------ Add Control Unit ---------------------
CBDPCU10

Specify or revise the following values.
Control unit number  . . . . 0440  +

Control unit type  . . . . . 3490          +

Serial number  . . . . . . . __________
Description  . . . . . . . . ________________________________

Connected to switches  . . . 01  01  01  01  __  __  __  __  +
Ports  . . . . . . . . . . . D6  D7  D8  D9  __  __  __  __  +

If connected to a switch, select whether to have CHPIDs/link
addresses and unit address range proposed.

Auto-assign  . . . . . . . . 2   1.  Yes
                                 2.  No

F1=Help     F2=Split    F4=Prompt   F5=Reset    F9=Swap    F12=Cancel
```

Figure 5-2 Adding the first TS7700 control unit through HCD, Part 1

Specify the control unit number and the type (here, 3490) as shown in Figure 5-2, then press Enter. The panel shown in Figure 5-3 displays for you to choose the processor to which to connect the control unit.

```
-------------------------- Add Control Unit -------------------------
CBDPCU12

Specify or revise the following values.
Control unit number  . : 0440        Type . . . . . . : 3490
Processor ID . . . . . : PROC1       This is the main processor

Channel path IDs . . . . 40   50   60   70   __   __   __   __  +
Link address . . . . . . D6   D7   D8   D9   __   __   __   __  +

Unit address . . . . . . 00   __   __   __   __   __   __   __  +
Number of units  . . . . 16  ___  ___  ___  ___  ___  ___  ___

Logical address  . . . . 0  + (same as CUADD)

Protocol . . . . . . . . __  + (D,S or S4)
I/O concurrency level  . 2  + (1, 2 or 3)


 F1=Help     F2=Split    F4=Prompt   F5=Reset    F9=Swap    F12=Cancel
```

Figure 5-3   Adding the first TS7700 control unit through HCD, Part 2

Note that because the TS7700 is not being attached through FICON directors, the link address fields would be blank. You should define the second to eighth TS7700 tape control units, specifying the logical unit address (CUADD) = 0 to F in the CBDPCU12 panel, as shown in Figure 5-3.

To define the TS7700 virtual drives, you need to go to the Device List Panel either from the Main Panel by entering `1` and then `5`, or by means of the Control Unit List Panel by using action **s**. To add the TS7700 virtual drives, press PF11. The panel shown in Figure 5-4 is then displayed.

```
------------------------------- Add Device ---------------------------
CBDPDV10

Specify or revise the following values.
Device number  . . . . . . . . 0A40  (0000 - FFFF)
Number of devices  . . . . . . 16__
Device type  . . . . . . . . . 3490_________  +

Serial number  . . . . . . . . __________
Description  . . . . . . . . . ________________________________

Connected to CUs . . 0440  ____  ____  ____  ____  ____  ____  ____  +


F1=Help     F2=Split    F4=Prompt   F5=Reset    F9=Swap    F12=Cancel
```

Figure 5-4   Adding the first 16 drives through HCD

After entering the required information and specifying to which processors and operating systems the devices are connected, the panel in Figure 5-5 displays, where you can update the device parameters.

```
 CBDPDV13 Define Device Parameters / Features Row 1 of 6
Command ===> __________________________________________ Scroll ===> PAGE
Specify or revise the values below.
Configuration ID . : AB          MVS operating system
Device number  . . : 0440        Number of devices :16
Device type  . . . : 3490

Parameter /
Feature      Value  P Req.  Description
OFFLINE      Yes            Device considered online or offline at IPL
DYNAMIC      Yes            Device supports dynamic configuration
LOCANY       No             UCB can reside in 31 bit storage
LIBRARY      Yes            Device supports auto tape library
AUTOSWITCH   No             Device is automatically switchable
LIBRARY-ID   A123A          5 digit library serial number
LIBPORT-ID   01             2 digit library string ID (port number)
ALTCTRL      No             Separate physical control unit path
SHARABLE     No             Device is Sharable between systems
COMPACT      Yes            Compaction
***************************** Bottom of data ****************************
 F1=Help        F2=Split     F4=Prompt       F5=Reset        F7=Backward
 F8=Forward     F9=Swap     F12=Cancel      F22=Command
```

Figure 5-5 Define Device Parameters HCD panel

Notes:

If you are defining drives that are installed in a system-managed IBM Tape Library, either an IBM 3494 or an IBM 3953 attached to an IBM TS3500 Tape Library, you must specify LIBRARY=YES.

If more than one System z host will be sharing the virtual drives in the TS7700, specify SHARABLE=YES. This will force OFFLINE to YES. It is up to the installation to ensure proper serialization from all attached hosts.

We strongly recommend that you specify LIBRAYRY-ID and LIBPORT-ID. Use the Composite LIBRARY-ID of the TS7700 for the HCD definition.

To define the next 16 TS7700 3490E virtual drives, you need to go back to the Device List Panel from the Main Panel by entering `1` and then `5`. To add the TS7700 virtual drives, press PF11. The CBDPDV10 panel displays again.

5.2.2 Activate the I/O configuration

There are differences in the process of activating the IODF for a new tape library or for changes made to an existing tape library that cause changes in the device addresses, if you want to avoid an IPL. As an alternative to the procedures described below, you can always IPL the system.

Installing a new tape library

If you are installing a TS7700 for the first time, from a host software definition point of view this is the installation of a new library. When you are activating the IODF for a new tape library, the following steps are required to get the tape library or TS7700 ONLINE without IPLing your systems:

1. Activate the IODF.
2. MVS VARY ONLINE the devices in the library. This will create some of the control blocks. You will see the message

   `IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED`
3. Do the final ACTIVATE. This is required to build the Eligible Device Table (EDT) for MVS Allocation.

You can check details using the DEVSERV QTAPE command, which provides information about Unit Control Block (UCB), UCB prefix, UCB common extension, Device Class Extension (DCE), and Read Device Characteristics (RDC) data and Read Configuration Data (RCD) data, which are data buffers acquired directly from the device.

> **Note:** If you are just adding additional device address ranges to an existing TS7700, you can use the same process as for a new tape library.

Modifications to an existing tape library

When you are modifying an existing tape library so that existing device addresses are to be changed, follow this process flow:

1. Activate an IODF deleting all devices from the library.
2. Activate an IODF that defines all of the devices of the modified library.
3. MVS VARY ONLINE the devices in the library. This will create some of the control blocks. You will see this message:

   `IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED`
4. Do the final ACTIVATE.

Alternatively, you can use DS QL,nnnnn,DELETE (where nnnnn is the LIBID) to delete the library's dynamic control blocks. If you have IODEF with LIBID and LIBPORT coded already, proceed as follows:

1. Use QLIB LIST to display that the INACTIVE control blocks have been deleted.
2. Use ACTIVATE IODF to redefine the devices.
3. Use QLIB LIST to display that the ACTIVE control blocks are properly defined.

If LIBRARY-ID (LIBID) and LIBPORT-ID are not coded, perform the following steps:

1. MVS VARY ONLINE the devices in the library. This will create some control blocks, and you will see this message:

   `IEA437I TAPE LIBRARY DEVICE(DDD), ACTIVATE IODF=XX, IS REQUIRED`
2. Use ACTIVATE IODF to redefine the devices.
3. Use QLIB LIST to verify that the ACTIVE control blocks are properly defined.

5.2.3 HCD considerations for Multi Cluster Grid operation

Each TS7700 presents 256 virtual device images for a total of 768 virtual devices if used in a Three-Cluster Grid configuration. Each TS7700 has 256 virtual devices with 16 logical control units (LCU 0-F). The host generates the physical paths to each site separately, so the host sees one Composite Library image, and three Distributed libraries. An example of a TS7700 Three-Cluster Grid configuration is shown in Figure 5-6.

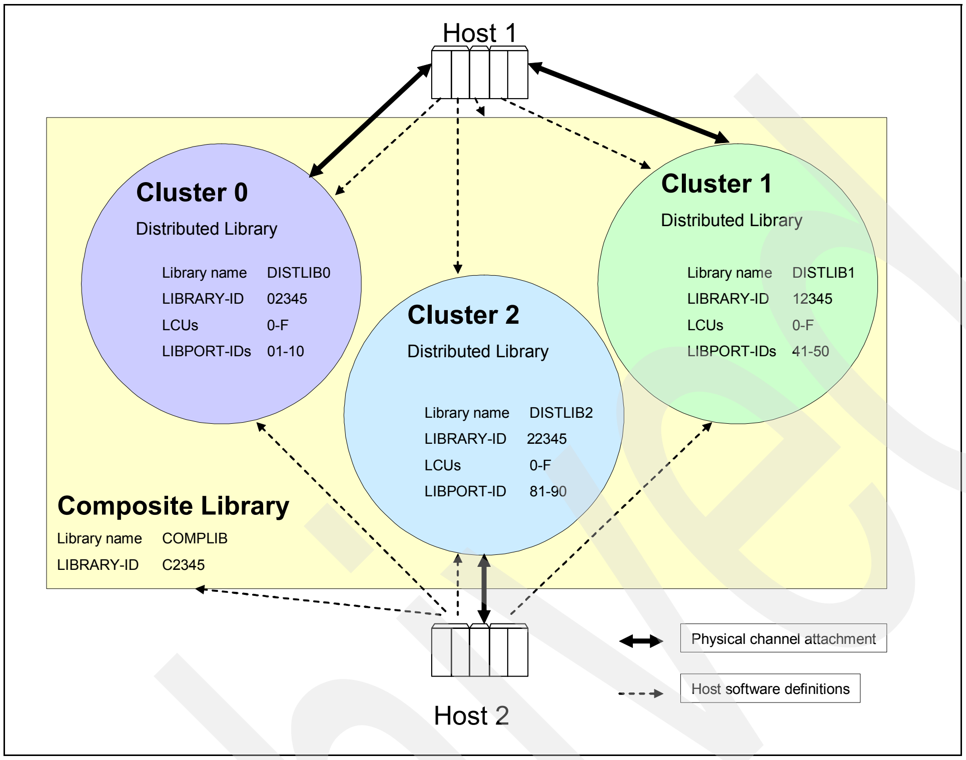

Figure 5-6 Example of a TS7700 Three-Cluster Grid configuration

With Figure 5-6 as example you define that Host 1 has physical connections to Cluster 0 and Cluster 1. Cluster 2 and Host 2 is probably far away in a Disaster Recovery center and Host 2 has only physical connections to Cluster 2. You then define that Host 1 has 2*256 3490 units and Host 2 has 256 units. In HCD you use LIBRARY-ID from the Composite Library (C2345). In SMS on both Host 1 and Host 2 all four libraries are defined as three Distributed Libraries and one Composite Library. Internally, the three clusters are connected with a customer-supplied IP-network.

Logical control units and physical paths are defined on a vNode/gNode boundary, similar to the Virtual Tape Controllers (VTCs) in the previous generation of the PtP VTS. All of them are part of the same Composite Library image presented to the host. Table 5-6 shows the possible subsystem IDs (LIBPORT-IDs) for each cluster in a Three-Cluster Grid configuration.

Table 5-6 LIBPORT-ID for each cluster in a Three-Cluster Grid configuration

Cluster	Logical Control Units (LCUs)	LIBPORT-IDs (hex)
0	0-7	01-08
	8-F	09-10
1	0-7	41-48
	8-F	49-50
2	0-7	81-88
	8-F	89-90

This definition is essentially the same as for a Single Cluster Grid configuration. The only major and important difference is that you need to specify the range 41-50 for the LIBPORT-IDs of Cluster 1 and range 81-90 for Cluster 2.

Note: The virtual device control for each cluster in a Dual Cluster Grid configuration is managed by the host. That means the host randomly picks a device from each cluster for I/O operation by the host device allocation algorithm. If now the remote cluster is attached through limited bandwidth connection, it might affect the I/O performance, because the remote cluster might be selected as the I/O cluster, even if the data is residing inside the TVC of the local cluster.

To avoid those situations it is recommended—whether you have one host attached to both clusters in a Dual Cluster Grid or two hosts attached redundantly to both clusters in a Dual Cluster Grid—to vary the remote virtual devices from each host's point of view offline for normal operation. Only in case of disaster recovery should those remote virtual drives be varied online from the host.

In your installation there can be considerations about use of FICON switches. Policies are based on requirements for availability and accessibility. Your local policies might require that all FICON equipment must be attached through two FICON switches, half of the connections on each?

If you have a setup with two centers and FICON switch equipment in both sites attached to the host, you can also use cascading FICON attachment to your tape setup.

Note: For the latest information about supported FICON directors and TS7700 Licensed Internal Code levels, see:

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ116133

For an example of how to do Cascading FICON attachment with two sites, see:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2844

5.2.4 Display and control your settings

In a z/OS environment, you can use the DISPLAY SMS and DEVSERV QUERY LIBRARY commands to check some portions of your definitions.

Note: The following examples were obtained from a Three-Cluster Grid with 768 virtual devices.

DISPLAY SMS command

You can use the DISPLAY SMS command to display and check the settings in the DEVSUPxx member for the scratch categories:

```
DISPLAY SMS,LIBRARY(xxxxxxxx),DETAILS
```

Note that the scratch count of MEDIA2 does not necessarily match the number of scratch volumes of your tape management system when you use the Expire Hold function in the TS7700. OAM displays the scratch count it receives from the 3953 or 3494 Library Manager.

Example 5-1 shows the sample output of a DISPLAY SMS,LIBRARY command against the Composite Library and, in this case, for a TS7700 the OPM function is always supported by this type of library. The display shows that MEDIA2 uses category 0002 and MEDIA1 uses category 0001.

Example 5-1 Display SMS,LIB from a Composite Library

```
D SMS,LIB(COMPLIB),DETAIL
F OAM,D,LIB,COMPLIB,L=ST6T10-Z
CBR1110I OAM LIBRARY STATUS: 141
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY SCRTCH  ON OP
LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS   VOLS
COMPLIB   VCL  3957-V06  768  768  287      0      0 368298  Y  Y
----------------------------------------------------------------------
MEDIA      SCRATCH       SCRATCH       SCRATCH
TYPE         COUNT     THRESHOLD      CATEGORY
MEDIA1       170345            0          0001
MEDIA2       197953            0          0002
----------------------------------------------------------------------
DISTRIBUTED LIBRARIES:  DISTLIB0    DISTLIB1    DISTLIB2
----------------------------------------------------------------------
LIBRARY ID:  10001
OPERATIONAL STATE:  AUTOMATED
ERROR CATEGORY SCRATCH COUNT:             33
CORRUPTED TOKEN VOLUME COUNT:              0
----------------------------------------------------------------------
LIBRARY SUPPORTS IMPORT/EXPORT.
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.
```

Example 5-2 shows the sample output of a DISPLAY SMS,LIBRARY command against the Distributed Library.

Example 5-2 Display SMS,LIB output from a Distributed Library

```
D SMS,LIB(DISTLIB1),DETAIL
F OAM,D,LIB,DISTLIB1,L=ST6T10-Z
CBR1110I OAM LIBRARY STATUS: 062
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY SCRTCH  ON OP
LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS   VOLS
DISTLIB1  VDL 3957-V06     0    0    0   1348    819      0  Y  Y
----------------------------------------------------------------------
COMPOSITE LIBRARY:      COMPLIB
----------------------------------------------------------------------
LIBRARY ID:  10001
OPERATIONAL STATE:  AUTOMATED
SCRATCH STACKED VOLUME COUNT:            222
PRIVATE STACKED VOLUME COUNT:            108
----------------------------------------------------------------------
LIBRARY SUPPORTS IMPORT/EXPORT.
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.
CONVENIENCE I/O STATION INSTALLED.
CONVENIENCE I/O STATION IN OUTPUT MODE.
BULK INPUT/OUTPUT NOT CONFIGURED.
```

DEVSERV QUERY LIBRARY command

The DEVSERV QUERY LIBRARY or DS QL command should always be used to query your Library configuration before and after an activate of your IODF when you do changes for the library.

Example 5-3 shows an example of how you would list all libraries using the DS QL command. Those LIBRARY-IDs (LIBIDs) marked with asterisks (*) are actually attached to the host.

Example 5-3 DEVSERV QLIB,LIST

```
DS QL,LIST
IEE459I 14.57.36 DEVSERV QLIB 708
The following libids are defined in the   ACTIVE configuration:
*C0323 *BA094 *BA055 *CA002 *BA012 *BA091  BA049 *BA022 *BA044 *059C8
*BA095 *BA048  BA092 *BA010 *BA008 *BA060 *BA036 *11975 *B0009  BA069
 CA022  C0159  11974 *C0076  BA009 *CA003  BA056  12087  BA066  BA035
 BA071  21252  BA072  BA042  BA046  BA063  BA041  BA040  BA061  BA070
 BA047  BA034  BA033  BA013  BA096  BA067

NOTE: asterisks indicate libraries that are actually attached to the host.
```

Example 5-4 shows a detailed list of one single Library using the DS QL,<library-id>, DETAIL command. Check that no duplicate port IDs are listed and that each port has 16 devices. This is the correct output for a TS7700.

Example 5-4 Sample output of the DEVSERV QLIB,<library-id>,ACTIVE command

```
DS QL,BA094,ACTIVE
IEE459I 15.02.46 DEVSERV QLIB 774
The following are defined in the   ACTIVE configuration:
LIBID  PORTID       DEVICES
BA094  04           293B* 293C* 293D* 293E* 293F* 2930* 2931* 2932*
                    2933* 2934* 2935* 2936* 2937* 2938* 2939* 293A*
       02           291B* 291C* 291D* 291E* 291F* 2910* 2911* 2912*
                    2913* 2914* 2915* 2916* 2917* 2918* 2919* 291A*
       06           295B* 295C* 295D* 295E* 295F* 2950* 2951* 2952*
                    2953* 2954* 2955* 2956* 2957* 2958* 2959* 295A*
       08           297B* 297C* 297D* 297E* 297F* 2970* 2971* 2972*
                    2973* 2974* 2975* 2976* 2977* 2978* 2979* 297A*
       03           2920* 2921* 2922* 2923* 2924* 2925* 2926* 2927*
                    2928* 2929* 292A* 292B* 292C* 292D* 292E* 292F*
       01           2900* 2901* 2902* 2903* 2904* 2905* 2906* 2907*
                    2908* 2909* 290A* 290B* 290C* 290D* 290E* 290F*
       05           2940* 2941* 2942* 2943* 2944* 2945* 2946* 2947*
                    2948* 2949* 294A* 294B* 294C* 294D* 294E* 294F*
       07           2960* 2961* 2962* 2963* 2964* 2965* 2966* 2967*
                    2968* 2969* 296A* 296B* 296C* 296D* 296E* 296F*

NOTE: asterisks indicate devices that are currently attached.
```

You can display the command syntax with DS QL,?. For a complete description of the QLIB command, see Appendix H, “DEVSERV QLIB command” on page 623 or refer to *z/OS V1R9.0 MVS System Commands*, SA22-7627.

5.2.5 Set values for the Missing Interrupt Handler

The TS7700 emulates 3490E devices and does not automatically communicate the Missing Interrupt Handler (MIH) time-out values to the host operating system in the Read Configuration Data Channel Control Word (CCW).

> **Important:** An MIH value of 45 minutes is recommended for the virtual devices in a Multi Cluster Grid when a copy consistency for the remote clusters is set to RUN.

You must specify the MIH time-out value for IBM 3490E devices. The value applies only to the virtual 3490E drives and not to the real IBM TS1120/3592 drives that the TS7700 manages in the backend. Remember that the host only knows about logical 3490E devices. Table 5-7 summarizes recommended minimum values that might need to be increased, depending on specific operational factors.

Table 5-7 Tape device MIH values

Tape device	MIH
3480, 3490 with less than 8 devices per CU or low usage	3 minutes
3480, 3490 with 8 devices per CU or heavy usage	5 minutes
3490E or 3480, 3490 with ESCON	10 minutes
3490E with ECST	20 minutes
TS7700 Single Cluster Grid with 3490E emulation drives	20 minutes
TS7700 Multi Cluster Grid with 3490E emulation drives	45 minutes

You should specify the MIH values in PARMLIB member IECIOSxx. Alternatively, you can also set the MIH values through the System z operator command, SETIOS. This setting will be available until manually changed or until the system is initialized. Use the following statements in PARMLIB, or manual commands to display and set your MIH values:

1. You can specify the MIH value in the IECIOSxx PARMLIB member:

   ```
   MIH DEV=(0A40-0A7F),TIME=45:00
   ```

2. To manually specify MIH values for emulated 3490E tape drives, use:

   ```
   SETIOS MIH,DEV=(0A40-0A7F),TIME=45:00
   ```

   – To display the new settings:

     ```
     D IOS,MIH,DEV=0A40
     ```

   – To check the current MIH time:

     ```
     D IOS,MIH,TIME=TAPE
     ```

The settings of the SETIOS and the MIH values in the IECIOSxx member change the value for the primary time-outs, but you cannot change the secondary time-outs. Those are delivered by the self-describing values from the device itself.

You can find more information about MIH settings in *MVS Initialization and Tuning Reference*, SA22-7592.

When specifying time intervals, consider the following possibilities:

- The MIH detects a missing interrupt condition within 1 second of the time interval that you specify.
- If the time interval is too short, a false missing interrupt can occur and cause early termination of the channel program. For example, if a 30-second interval is specified for a tape drive, a rewind might not complete before the MIH detects a missing interrupt.
- If the time interval is too long, a job or system could hang because the MIH has not yet detected a missing interrupt. For example, if a 15 minute time interval is specified for a tape drive used as an IMS log tape, the MIH could delay IMS for 15 minutes because of MIH detection.

During IPL (if the device is defined to be ONLINE) or during the VARY ONLINE process, some devices can present their own MIH time-out values, through the *primary/secondary* MIH timing enhancement contained in the self-describing data for the device. The *primary* MIH time-out value is used for most I/O commands, but the *secondary* MIH time-out value can be used for special operations such as long-busy conditions or long running I/O operations.

Any time a user specifically sets a device or device class to have an MIH time-out value that is different from the IBM-supplied default for the device class, that value will override the device-established primary MIH time value. This implies that if an MIH time value that is equal to the MIH default for the device class is explicitly requested, IOS will not override the device-established primary MIH time value. To override the device-established primary MIH time value, you must explicitly set a time value that is not equal to the MIH default for the device class.

Note that overriding the device-supplied primary MIH time-out value might adversely affect MIH recovery processing for the device or device class.

Refer to the specific device's reference manuals to determine if the device supports self-describing MIH time values.

5.3 TS7700 software definitions for z/OS

In this section, we describe the software definition considerations for implementing the TS7700 in z/OS, VM/ESA, and z/VSE environments. From a software point of view, the TS7700 is the same as an IBM 3494 Enterprise Tape Library with IBM 3490E tape drives.

The TS7700 must be defined as a new tape library with emulated IBM 3490E tape drives from the host system. Refer to *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632, for more information about how to define this configuration.

The software support levels required to support a TS3500 Tape Library, or 3953 Library Manager configuration with TS3500, are explained in 3.4.2, “Software requirements” on page 104.

Tape management systems

From the host perspective, a TS7700 is a single subsystem whether it is configured as a Single Cluster Grid configuration or within a Multi Cluster Grid configuration. The tape management system sees only the Composite Library and logical drives. There is no difference from the tape management system’s point of view between a Multi Cluster Grid TS7700 installation (Peer-to-Peer) and a Single Cluster Grid installation (standalone).

Entering logical volumes after hardware changes

During cartridge entry processing, even if the library is online and operational to a given host, at least one device needs to be online (or have been online) to that host for the library to be able to send the cartridge entry attention interrupt to that host. If the library is online and operational, yet there are no online devices to a given host, that host will not receive the attention interrupt from the library unless a device had previously been VARYed online. To get around this, ensure that at least one device is online (or had been online) to each host or use the LIBRARY RESET,CBRUXENT command to initiate cartridge entry processing from the host. This is especially important if you only have one host attached to the library that owns the volumes being entered. In general, after you have entered volumes into the library, if you do not see the expected CBR36xxI cartridge entry messages being issued, you can use the LIBRARY RESET,CBRUXENT command from MVS to initiate cartridge entry processing. The LIBRARY RESET,CBRUXENT command causes the host to ask for any volumes in the insert category.

5.3.1 z/OS and DFSMS/MVS SMS-managed tape

To define the TS7700 to DFSMS, use the ISMF panels to create a new definition of the TS7700 logical tape library to be recognized from the host. This definition is done in the same way for a new installation of a Single Cluster Grid as it is done for a new installation of a Multi Cluster Grid TS7700, except for definitions of the Distributed Library or libraries:

- In a Single Cluster configuration, you define one Composite and one Distributed Library.
- In a Multi Cluster configuration, you define one Composite Library and two or three distributed libraries.

To use the TS7700, at least one Storage Group must be created to allow the TS7700 logical tape library virtual drives to be allocated by the ACS routines. Because all of the logical drives and volumes are associated with the Composite Library, only the Composite Library can be defined in the Storage Group. The distributed libraries must not be defined in the Storage Group.

For a complete discussion of the host software implementation tasks for IBM tape libraries, refer to:

- *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427
- *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789

The steps to define the TS7700 Tape Library in a z/OS environment are:

1. Modify the SYS1.PARMLIB member (such as IEFSSNxx, IGDSMSxx, LOADxx, DEVSUPxx, and COMMANDxx).
2. Create the Tape Configuration Database (TCDB), an ICF catalog type VOLCAT.
3. IDCAMS IMPORT CONNECT of the TCDB to the other system is required when tape library sharing capability is used.
4. Add the procedure to start the OAM address space.

5. Define the tape library as a DFSMS resource. Define the Composite Library and one or more distributed libraries. Remember that library names cannot start with a “V”. Figure 5-7 shows the definition of a Composite Library.

```
Panel   Utilities   Scroll    Help
------------------------------------------------------------------------------
                                TAPE LIBRARY DEFINE              Page 1 of 2
Command ===>_

SCDS Name . : SCDS.TEMP.PRIMARY
Library Name : P2PLIB1

To Define Library, Specify:
  Description ===> TS7700 Grid Composite library
              ===>
  Library ID . . . . . . . . . . .1234C       (00001 to FFFFF)
  Console Name . . . . . . . . . .LIB1CON
  Entry Default Data Class . . . .DCATLDS
  Entry Default Use Attribute . . S           (P=PRIVATE or S=SCRATCH)
  Eject Default . . . . . . . . . K           (P=PURGE or K=KEEP)

  Media Type:         Scratch Threshold
    Media1 . . . . 100          Media3 . . . . 0   (0 to 999999)
    Media2 . . . . 400          Media4 . . . . 0   (0 to 999999)

Use ENTER to Perform Verification; Use DOWN Command to View Next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure 5-7 Composite Library definition

Figure 5-8 shows a sample panel to define one of the distributed libraries.

```
Panel   Utilities   Scroll    Help
------------------------------------------------------------------------------
                                TAPE LIBRARY DEFINE              Page 1 of 2
Command ===>_

SCDS Name . : SCDS.TEMP.PRIMARY
Library Name : P2PLIBA

To Define Library, Specify:
  Description ===> TS7700 Distributed library A
              ===>
  Library ID . . . . . . . . . . .12349       (00001 to FFFFF)
  Console Name . . . . . . . . . .
  Entry Default Data Class . . . .
  Entry Default Use Attribute . .             (P=PRIVATE or S=SCRATCH)
  Eject Default . . . . . . . . .             (P=PURGE or K=KEEP)

  Media Type:         Scratch Threshold
    Media1 . . . . 0          Media3 . . . . 0   (0 to 999999)
    Media2 . . . . 0          Media4 . . . . 0   (0 to 999999)

Use ENTER to Perform Verification; Use DOWN Command to View Next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure 5-8 Distributed Library definition

Note: Library ID is the only field that applies for the distributed libraries; all other fields can be blank or left as the default.

6. Create or update the Data Classes (DCs), Storage Classes (SCs), and Management Classes (MCs) for the TS7700. Make sure that these defined construct names are the same as those you have defined at the Library Managers, especially in a grid configuration, because outboard policy management is being used for Multi Cluster Grid copy control.
7. Create the Storage Groups (SGs) for the TS7700. Make sure that these defined construct names are the same as those you have defined at the Library Manager.

   The Composite Library must be defined in the Storage Group. Do *not* define the distributed libraries in the Storage Group.

   Note: At OAM address space initialization, if a distributed library is defined to a Storage Group, the warning message CBR3017I is issued indicating that the distributed library is incorrectly defined to the Storage Group.

8. Create ACS routines to assign the constructs. Translate, test, and validate ACS routines.
9. Activate the new SCDS.
10. SCDS activation will initiate an OAM restart if parameter RESTART=YES is specified in the OAM startup procedure in PROCLIB. IF RESTART=NO is used, you must issue an OAM restart command manually with command F OAM,RESTART. Example 5-5 shows a sample OAM procedure where RESTART=NO is used.

    Example 5-5 OAM start-up procedure from PROCLIB

```
//OAM PROC
//IEFPROC  EXEC PGM=CBROAM,REGION=OM,
// PARM=('OSMC=NO,APLAN=CBROAM,OAM=60,MAXS=2,RESTART=NO')
//SYSABEND DD SYSOUT=A
```

11. Vary the Composite Library and distributed libraries online.
12. Vary the TS7700 virtual drives online.

For more detailed information about defining a tape subsystem in a DFSMS environment, refer to:

- *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.
- *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632

Note: The CBRXLCS FUNC=PTPDATA and FUNC=PTPMC programming interfaces are not supported in the TS7700. The command is accepted but is treated as a no-op.

5.3.2 Implementing Copy Export

Copy Export is the function that supports moving physical volumes with data critical for continuing business operations outside a library to an offsite vault. You can also use the Copy Export function as a possibility for disaster testing. It supports moving the logical copy of the original to another location and keep the original logical copy within the TS7700 available for normal production.

In the following sections, we describe step by step how to implement and execute Copy Export. For more details and error messages related to the Copy Export function, refer to *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide* and search for TS7700. The guide is available on the Techdocs Web site:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs

Overview of Copy Export

From a host perspective it is managed based on the definitions of the assigned Management Class, that direct the logical tapes onto a dedicated secondary volume pool, as shown in Figure 5-9.

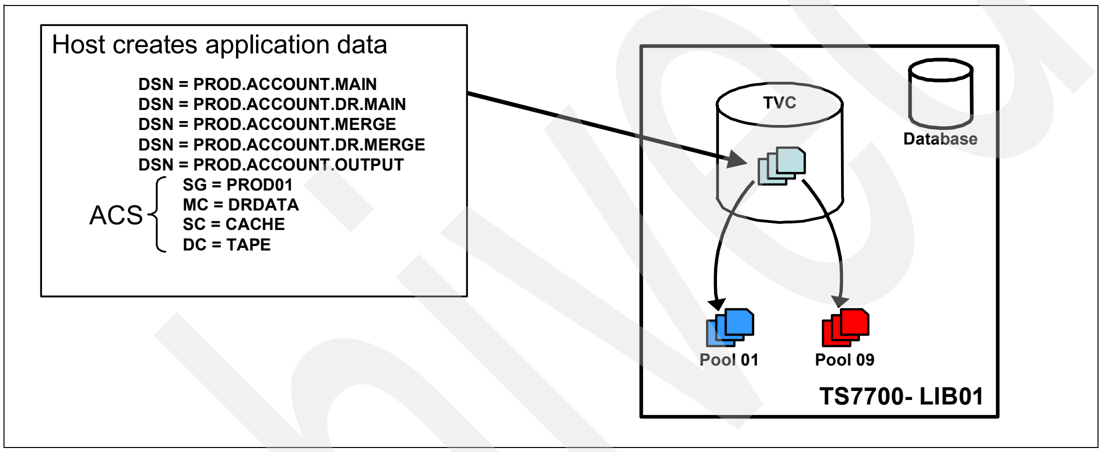

Figure 5-9 Copy Export for assigning different volume pools

In Figure 5-9, MC DRDATA is used to identify the data to be taken offsite. The outboard definitions in the LM direct the primary copy of a logical volume to the pool associated with Storage Group PROD01(Pool 01 in our example) and one copy of the logical volume to the pool defined in Management Class DRDATA (Pool 09 in our example). The stacked volumes of Pool 09 can then be ejected and moved to a secure location.

Copy Export Recovery is the function to use when the stacked volumes have to be recovered. Recovery must be into an empty TS7700. You cannot merge the exported copies into an existing, running TS7700 or into a Multi Cluster Grid.

Note: You can specify either Selective Dual Copy or Copy Export for a logical volume, but not both. Selective Dual Copy and Copy Export are defined on a per storage pool basis at the Library Manager.

To implement the Copy Export function based on creating two copies of one logical volume in one TS7700, where one is to reside in the library and another will be used for Export Copy, these steps are required:

- Enable the host setup to use the Copy Export function.
- Validations before activating the Copy Export function.
- Initiate and evaluate the Copy Export operation.
- Return the Export stacked volumes to the Library.

Enable the host setup to use Copy Export

The description can apply to a Single Cluster as well as a Multi Cluster Grid and focus on describing the actions directly related to Copy Export. You can find more details about how to define MC and ACS routines on the host in Appendix F, "TS7700 Virtualization Engine implementation, step-by-step" on page 591.

We assume that you already have implemented the SMS environment to direct logical volumes to a TS7700 and that you plan to use the Management Class DRDATA for the logical volumes to be exported. The additional steps to enable the host to use the Copy Export function are:

1. Define the Management Class named DRDATA on the host using ISMF.
2. Modify the Management Class ACS routine. Create a FILTLIST that includes the data set names you want to have MC DRDATA assigned.
3. Translate the Management Class ACS routine.
4. Activate the new SCDS.
5. Define MC DRDATA on the LM and relate it to the physical pool you are planning to use for Copy Export (Pool 09 in our example). If used in a Multi Cluster Grid the setting of DRDATA should only apply to the cluster you will use for Copy Export.
6. From the ETL Specialist panel or the Library Manager console, define the physical pool for Copy Export as an Export Pool. Figure 5-10 shows the result of defining Pool 09. The reclaim policy of Pool 09 applies to physical volumes that are residing inside the library and those that have been removed from the library.

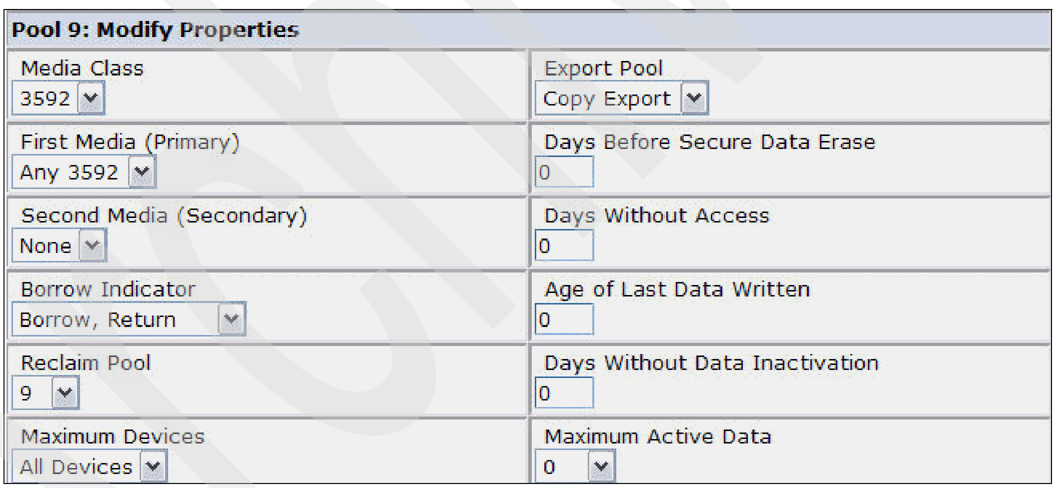

Figure 5-10 Pool properties for Copy Export Pool 09

Validations before activating the Copy Export function

When the logical volumes are to be exported, there are some general validations that should be done. Before you initiate the operation, you should check to ensure that the TS7700 has the needed physical drives and scratch physical volume resources and that it is not near the limit of the number of physical volumes that can have a status of copy exported (2000 physical volumes). Depending on your production environment, you might want to automate these validation steps.

The validation needed is as follows:

1. If you had migrated from a B10 or B20 VTS to the TS7700 using the outboard migration method, you might have data that is still in the older VTS format. The TS7700 cannot export data in the old format, so you should check to see whether any of the data to export was written with the old format.
2. Validate that the TS7700 has at least four available physical tape drives. You can use the Library Request host console command that specifies the PDRIVE request. This returns the status of all physical drives attached to the TS7700. If there are fewer than your required numbers of physical drives available, you must call for service to repair drives prior to performing the copy export operation. See Example 5-6 for the output of the PDRIVE request.

Example 5-6 Data returned by the PDRIVE request

```
LI REQ,BARR68A,PDRIVE
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,PDRIVE.
CBR1280I LIBRARY BARR68A REQUEST. 896
KEYWORDS: PDRIVE
---------------------------------------------------------------------
PHYSICAL DRIVES V1
   SERIAL NUM     TYPE  MODE AVAIL  ROLE  POOL    PVOL    LVOL
 000007878161 3592E05E  E05E Y      RCLS    01  S70470  Z09866
 000007878175 3592E05E  E05E Y      MIGR    01  JA8149  Z04381
 000001365176 3592E05E  E05E Y      RCLT    01  S70421  Z09866
 000001365177 3592E05E  E05E Y      MIGR    01  JA8145  Z08629
 000001365137 3592E05E  E05E Y      RCLS    03  310112  XC4487
 000001365181 3592E05E  E05E Y      IDLE    01  310451
 000007878194 3592E05E       Y      IDLE    00
 000007878312 3592E05E  E05E Y      RCLT    03  S70479  XC4487
```

The response shown in Example 5-6 shows the following information:

– Eight drives are defined. Their serial numbers (SN) are shown in the left column.

– All drives are encryption-capable as indicated by TYPE 3592E05E. TS1120 tape drives that are not encryption-capable will show as 3592E05.

– All eight drives are available (AVAIL=Y).

– MODE indicates the format the drive is currently operating in. The information is only available when a tape is mounted on the drive.

– ROLE describes what the drive is doing at the moment, for example RCLS is reclaim source and RCLT is reclaim target. In our example, we can see that two reclamation and two premigration operations are running:

  - Logical volume Z09866 is being reclaimed from physical volume S70470 mounted on drive with SN-7878161 to physical volume S70421 mounted on drive with sn-1365176. Both stacked volumes reside in Pool 01.
  - Logical volume XC4487 is being reclaimed from physical volume 310112 mounted on drive sn-1365137 to physical volume S70479 mounted on drive sn-7878312. Both stacked volumes reside in Pool 03.
  - Logical volume JA8149 is being written to physical volume Z04381 mounted on drive sn-7878175. Logical volume JA8145 is being written to physical volume Z08629 mounted on drive sn-1365177.

– Six drives are in use and two are IDLE, meaning ready for use. Serial number 1365181 is IDLE but a physical volume is still mounted.

Note: The Copy Export operation requires a single drive to write the TS7700 database to the volumes being exported. Be sure to consider this when analyzing workload and drive utilization. Refer to Chapter 8, “Performance and monitoring” on page 409 for more information about workload and drive utilization.

3. Check that the pool to be exported has sufficient scratch physical volumes and that the TS7700 is under the 2000 volume limit for copy exported volumes in all pools. You can use the Library Request host console command that specifies the POOLCNT request. See Example 5-7 for the response to the POOLCNT command.

Example 5-7 Data returned from POOLCNT command

```
LI REQ,BARR68A,POOLCNT
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,POOLCNT.
CBR1280I LIBRARY BARR68A REQUEST. 919
KEYWORDS: POOLCNT
--------------------------------------------------------------------
PHYSICAL MEDIA COUNTS V1
    POOL MEDIA  EMPTY  FILLING   FULL  ERASE    ROR  UNAVAIL  CXPT
       0    JA    164
       0    JJ     38
       1    JA      2        6     12      0      0        1     0
       9    JJ      0        4     22      0      0        0    45
```

Pool 00 is the Common Scratch Pool. Pool 9 is the one used for Copy Export. Example 5-7 shows the command POOLCNT; the response that is listed per pool is:

– The media type used for each pool
– The number of empty physical volumes
– The number of physical volumes in the filling state
– The number of full volumes
– The number of physical volumes that have been reclaimed, but need to be erased
– The number of physical volumes in read-only recovery state
– The number of volumes unavailable or in a destroyed state (1 in Pool 1)
– The number of physical volumes in the copy exported state (45 in Pool 9)

You should determine when you usually want to start the Copy Export operation. Thresholds could be the number of physical scratch volumes or other values that you define. These thresholds could even be automated by creating a program that interprets the output from the Library Request commands PDRIVE and POOLCNT, and acts based on the required numbers.

For more information about the Library Request Command, refer to 7.3.3, “Host Console Request” on page 375.

4. Create an Export List volume that provides the TS7700 with information about which data to export, as well as options to use during the operation.

   If you use a Multi Cluster Grid, be sure to create the Export List volume only on the same TS7700 that is used for Copy Export, but not on the same physical volume pool as used for Copy Export. If more than one TS7700 in a Multi Cluster Grid configuration contains the Export List volume, the Copy Export operation will fail.

Figure 5-11 shows the setting of a management class on the LM for the Export List volume in a Three-Cluster Grid configuration. RNN means one copy locally at RUN (R) and no copy (NN) on the two other clusters.

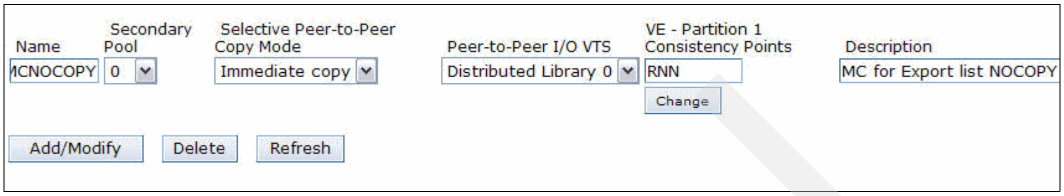

Figure 5-11 Management Class settings for the Export List volume

5. Then create an Export List volume as shown in Example 5-8.

Example 5-8 Sample JCL to create an Export List volume

```
//****************************************
//* FILE 1: EXPORT LIST
//****************************************
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=HILEVELQ.EXPLIST,MGMTCLAS=MCNOCOPY,
// UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(1,SL),
// VOL=(,RETAIN),
// DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1 DD *
EXPORT LIST 03
EXPORT PARAMETERS PHYSICAL POOL TO EXPORT:09
OPTIONS1,COPY,EJECT
/*
//****************************************
//* FILE 2: RESERVED FILE
//****************************************
//STEP2 EXEC PGM=IEBGENER,COND=(4,LT)
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=HILEVELQ.RESERVED,MGMTCLAS=MCNOCOPY,
// UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(2,SL),
// VOL=(,RETAIN,REF=*.STEP1.SYSUT2),
// DCB=*.STEP1.SYSUT2
//SYSUT1 DD *
RESERVED FILE
/*
//****************************************
//* FILE 3: EXPORT STATUS FILE
//****************************************
//STEP3 EXEC PGM=IEBGENER,COND=(4,LT)
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=HILEVELQ.EXPSTATS,MGMTCLAS=MCNOCOPY,
// UNIT=VTS1,DISP=(NEW,CATLG),LABEL=(3,SL),
// VOL=(,,REF=*.STEP1.SYSUT2),
// DCB=*.STEP1.SYSUT2
//SYSUT1 DD *
EXPORT STATUS 01
/*
```

The information required in the Export List file is, like for BVIR, provided by writing a logical volume that fulfills these requirements:

– That logical volume must have a standard label and contain three files:

  - An Export List file as created in Example 5-8 in STEP1. Notice that we want to export Pool 09. Option EJECT in record 2 tells the TS7700 to eject the stacked volumes upon completion. With OPTIONS1,COPY the physical volumes will be placed in the export-hold category for later handling by an operator.
  - A Reserved file as created in STEP2 in Example 5-8. This file is reserved for future use.
  - An Export Status file as created in STEP3. In this file, the information is stored from the Copy Export operation. This file is essential to keep because it contains information related to the result of the Export process.

– All records must be 80 bytes in length.

– The Export List file must be written without compression. Therefore, you need to assign a Data Class that specifies COMPRESSION=NO or you can overwrite the Data Class specification by coding TRTCH=NOCOMP in the JCL.

– Make sure that the files are assigned a Management Class that specifies that only the local TS7700 has a copy of the logical volume. You can either have the ACS routines assign this Management Class, or you can specify it in the JCL as we have done with MGMTCLAS=MCNOCOPY in Example 5-8 on page 217. We recommend that these files have the same expiration dates as the longest of the logical volumes you export because they must be kept for reference.

Initiate and evaluate the Copy Export operation

We can now initiate the Copy Export operation based on the logical volume created using the JCL in Example 5-8 on page 217:

- The Copy Export operation is initiated by issuing the LIBRARY EXPORT command. In this command, *logical volser* is a variable and is the logical volume used in creating the Export List file volume.

  ```
  LIBRARY EXPORT,logical volser
  ```

- The host will send a command to the Composite Library and from there it will be routed to the TS7700 where the Export List VOLSER resides.
- The executing TS7700 validates the request, checking for required resources, and if all is accepted the Copy Export continue.
- Logical volumes related to Pool 09 that still only reside in cache will delay the process. They will be copied to physical volumes in pool 9 as part of the Copy Export execution.
- Messages informing about the progress will be sent to the system console. All messages are in the format

  ```
  CBR3750I Message from library library-name: message text.
  ```

Table 5-8 shows examples of messages. Refer to the White Paper *WP101092 TS7700 Copy Export V1.0* on Techdocs for the most recent set of messages. Search for TS7700 from the Techdocs page:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs

Table 5-8 Message examples for Copy Export

Message Description	Action Needed
`E0000 EXPORT OPERATION STARTED FOR EXPORT LIST VOLUME XXXXXX` This message is generated when the TS7700 begins the copy export operation.	None
`E0005 ALL EXPORT PROCESSING COMPLETED FOR EXPORT LIST VOLUME XXXXXX` This message is generated when the TS7700 completes an export operation.	None
`E0006 STACKED VOLUME YYYYYY FROM LLLLLLLL IN EXPORT-HOLD` This message is generated during Copy Export operations when an exported stacked volume YYYYYY has been assigned to the eject-hold category. The LLLLLLLL field is replaced with the distributed library name of the TS7700 performing the export operation.	Use the Library Manager console dialog box, Manage Export-Hold Volumes, to move the exported stacked volumes to the Eject category to have them placed in the convenience I/O station.
`E0013 EXPORT PROCESSING SUSPENDED, WAITING FOR SCRATCH VOLUME` This message is generated every five minutes when the TS7700 needs a scratch stacked volume to continue export processing and there are none available.	Make one or more physical scratch volumes available to the TS7700 performing the export operation. If the TS7700 does not get access to a scratch stacked volume in 60 minutes, the operation is terminated.
`E0014 EXPORT PROCESSING RESUMED, SCRATCH VOLUME MADE AVAILABLE` This message is generated when, after the export operation was suspended because no scratch stacked volumes were available, scratch stacked volumes are again available and the export operation can continue.	None
`E0015 EXPORT PROCESSING TERMINATED, WAITING FOR SCRATCH VOLUME` This message is generated when the TS7700 has terminated the export operation because scratch stacked volumes were not made available to the TS7700 within 60 minutes of the first E0013 message.	Operator should make more TS7700 stacked volumes available, perform analysis of the Status file on the Export List File volume, and reissue the Export operation.
`E0016 COPYING LOGICAL EXPORT VOLUMES FROM CACHE TO STACKED VOLUMES` This message is generated when the TS7700 begins, and every 10 minutes during, the process of copying logical volumes that are only resident in the tape volume cache to physical volumes in the specified secondary physical volume pool.	None
`E0017 COMPLETED COPY OF LOGICAL EXPORT VOLUMES TO STACKED VOLUMES` This message is generated when the TS7700 has completed the copy of all needed logical volumes from cache to physical volumes in the specified secondary physical volume pool.	None

Message Description	Action Needed
`E0018 EXPORT TERMINATED, EXCESSIVE TIME FOR COPY TO STACKED VOLUMES` The export process has been terminated because one or more cache resident-only logical volumes needed for the export could not be copied to physical volumes in the specified secondary physical volume pool within a 10-hour period from the beginning of the export operation.	Call for IBM support
`E0019 EXPORT PROCESSING STARTED FOR POOL XX` This message is generated when the TS7700 has started export processing for the specified secondary physical volume pool XX.	None
`E0020 EXPORT PROCESSING COMPLETED FOR POOL XX` This message is generated when the TS7700 has completed processing for the specified secondary physical volume pool XX.	None
`E0022 EXPORT RECOVERY STARTED` The export operation has been interrupted by a TS7700 error or a power off condition. When the TS7700 is restarted, it will attempt recovery of the operation.	None
`E0023 EXPORT RECOVERY COMPLETED` The recovery attempt for interruption of an export operation has been completed.	Perform analysis of the Status file on the Export List File volume and reissue the Export operation, if necessary.
`E0024 XXXXXX LOGICAL VOLUME WITH INVALID COPY ON LLLLLLLL` This message is generated when the TS7700 performing the export operation has determined that one or more (XXXXXX) logical volumes that are associated with the secondary storage pool specified in the Export List file do not have a valid copy resident on the TS7700. The LLLLLLLL field is replaced by the distributed library name of the TS7700 performing the export operation. The export operation continues with the valid copies.	When the export operation completes, perform analysis of the Status file on the Export List file volume to determine the logical volumes that were not exported. Ensure that they have completed their copy operations and then perform another export operation.
`R0000 RECLAIM SUCCESSFULL FOR EXPORTED STACKED VOLUME YYYYYY` This message is generated when the TS7700 has successfully completed reclaim processing for an exported stacked volume exported during a previous copy export operation. Note: A copy exported physical volume can become eligible for reclaim based on the reclaim policies defined for its secondary physical volume pool or through the host console request command.	The physical volume no longer contains active data and can be returned from its offsite location for re-use.
`R0001 RECLAIM UNSUCCESSFUL FOR EXPORTED STACKED VOLUME YYYYYY` This message is generated when the TS7700 has attempted to reclaim an exported stacked volume exported during a previous Copy Export operation, but one or more primary copies of the active logical volumes were unavailable to complete the operation. Note: A copy exported physical volume can become eligible for reclaim based on the reclaim policies defined for its secondary physical volume pool or through the host console request command.	Obtain a list of the active logical volumes on the exported stacked volume and determine why the primary copies of them are not available to complete the reclaim.

After a successful completion all physical tapes related to Pool 09 (in our example) are ejected. The operator can empty the I/O station and transport the tapes to another location.

Note: We describe how to perform Copy Export Recovery for disaster recovery or disaster recovery testing in Chapter 9, “Disaster recovery and failover scenarios” on page 489.

Return the exported stacked volumes to the library

You need to set up procedures for when exported physical volumes can return to the library.

Normally the physical volumes must return when all logical volumes are expired. To determine when exported volumes have expired, it is simplest to use the Bulk Volume Information Retrieval facility. Use BVIR with *Physical Volume Status Pool xx* as input. With BVIR you can retrieve information for all physical volumes in a pool. See Chapter 8, “Performance and monitoring” on page 409 for more information about how to use BVIR.

Another reason for bringing back a physical volume applies to a Single Cluster Grid, where the original physical volume is destroyed, meaning the Exported Copy would be the only available copy available.

You insert the volumes through the I/O station and then assign all the inserted volumes to the Insert category. See Figure 5-12.

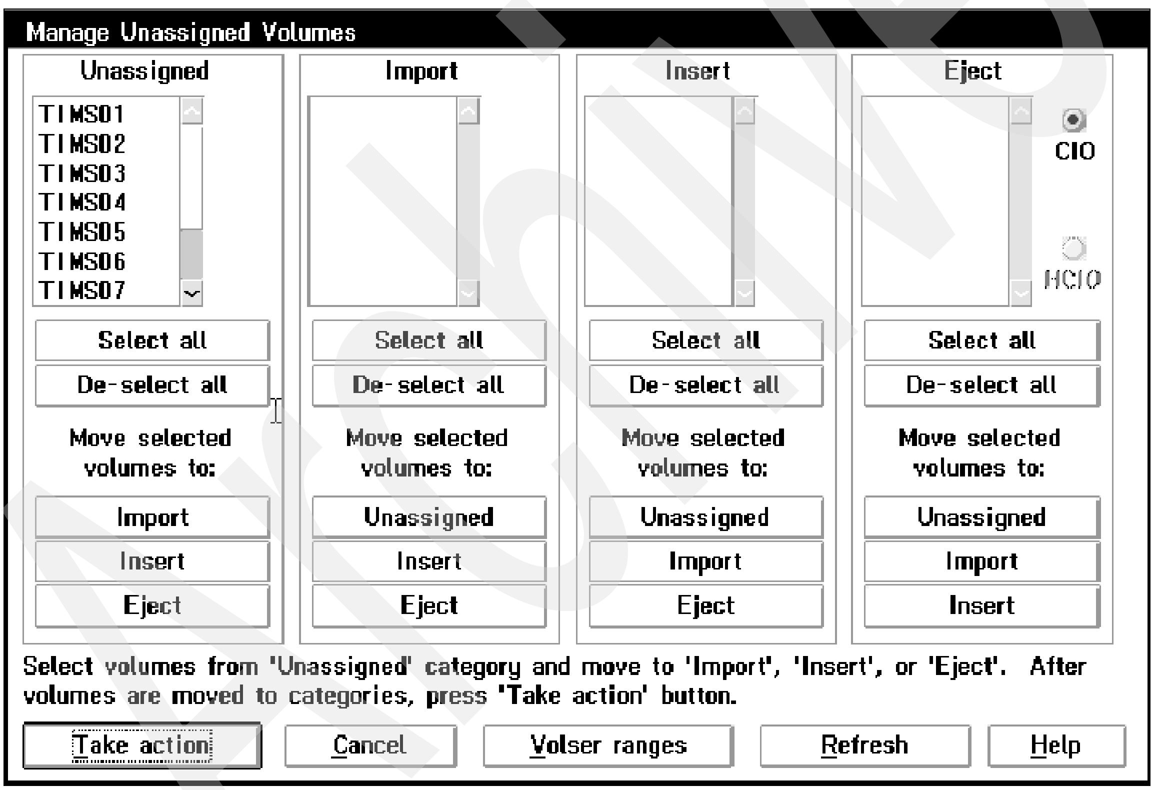

Figure 5-12 Insert of Export Physical Volumes

The selected volumes are then added to the insert category and the TS7700 will see that they are volumes it owns and return them to the pool they were in. The state of the volumes is changed to Read-Write and if they are empty and without active logical volumes they will have an empty status as well. When Insertion of the volumes is complete, the entire Export Copy operation for a single physical volume life cycle is done. For more information about managing unassigned volumes refer to “Unassigned volumes in the TS3500 Tape Library” on page 143 and “Unassigned volumes in the 3953 Library Manager” on page 143.

Tape Encryption is supported for stacked volumes used for Copy Export as described in 5.3.3, “Implementing Tape Encryption for stacked volumes” on page 222. You need the TS1120 drives to be encryption-enabled, the tape library and Library Manager on an LIC level that supports Tape Encryption, and the Encryption Key Manager available at the disaster recovery site in order to be able to process encrypted stacked volumes that have been exported.

5.3.3 Implementing Tape Encryption for stacked volumes

Encryption of backstore tapes helps control the risks of unauthorized data access without excessive security management burden or subsystem performance issues. Tape Encryption solutions from IBM all use an Encryption Key Manager (EKM) as a central point for key management. The EKM communicates with the TS7700 Virtualization Engine. Encryption on the TS7700 is controlled outboard on a physical stacked volume pool basis and is set on the TS7700 Management Interface (MI). As you set up encryption for all logical volumes in a specific pool, you can use it in addition to the Copy Export function as well, where physical volumes are moved to another location.

From the host you assign volumes to a Storage Group, and that SG is assigned on the LM to a specific Physical Volume Pool and MC. The same pool is connected through the MI to where you configure the key used for encryption.

> **Note:** Support for encryption based on the assigned SMS Data Class applies to native tape devices such as the TS1120, but it does not apply to virtual devices in a TS7700.

Let us use the example from 5.3.2, “Implementing Copy Export” on page 212 and do encryption on Pool 09 used there. It is suitable to use encryption when moving physical volumes offsite, based on the risk related to transport of physical volumes from one location to another.

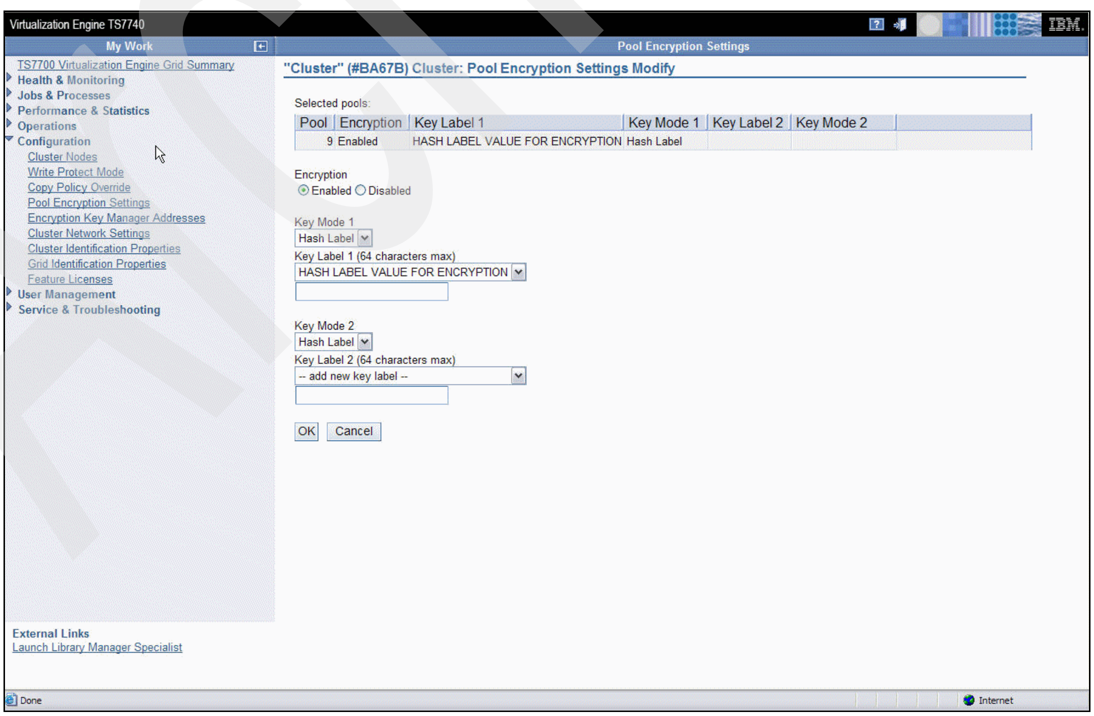

Figure 5-13 Pool Encryption settings

Figure 5-13 shows that encryption is enabled for Pool 09 and also shows what Key Labels are used to encrypt the data key on all volumes.

For details on encryption setting, refer to 4.5, “Setup of the TS7700 Virtualization Engine” on page 174.

5.4 Software implementation in z/VM and z/VSE

This section explains how to implement and run the TS7700 under z/VM and z/VSE. It covers basics for software requirements, implementation, customization, and platform-specific considerations about operations and monitoring. For more detailed information see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

5.4.1 General support information

Not all IBM Tape Libraries and TS7700 solutions are supported in all operating systems. Table 5-9 shows a summary of several supported tape solutions for non-z/OS environments. For information about support for TPF, refer to 5.5, “Software implementation in TPF” on page 229.

Table 5-9 Supported tape solutions for non-z/OS platforms in System z environments

Platform / Tape System	IBM 3494 or TS3500 Tape Library	IBM TS7700	3592 drives
z/VM native	Yes	Yes[1]	Yes
z/VSE 3.1 native z/VSE 4.1 native	Yes	Yes	Yes
z/VSE 3.1 under z/VM z/VSE 4.1 under z/VM	Yes	Yes[1]	Yes
zTPF	Yes	Yes	Yes

1 With restrictions.

Even if z/VM and z/VSE can use the TS7700, there are some restrictions you should consider.

Restrictions in all TS7700 environments

Neither z/VM nor z/VSE are able to provide SMS constructs to the Library Manager. Therefore, the functions provided by outboard policy management for the TS77000 cannot be used on these platforms.

However, a possibility is provided to use dedicated physical pools in a TS7700 environment. After insert processing of virtual volumes, you can define a default construct to the volume range as described in 4.6.5, “Implementing Outboard Policy Management for non-z/OS hosts” on page 190.

TS7700 Multi Cluster Grid environments

As you can see from Table 5-9, TS7700 is supported in a z/VM, or a z/VSE under z/VM environment. In both operating systems, there is no difference from the host point of view between a Multi Cluster Grid and a Single Cluster Grid TS7700 configuration. Neither z/VM nor z/VSE under z/VM know about the Composite Library and can monitor the entire environment.

z/VM and z/VSE are not able to handle unsolicited messages and actions (for example, intervention required and hardware messages); therefore, a PtP VTS requires an active z/OS in order to support VM and VM/VSE. This requirement has been eliminated with the TS7700 and its Management Interface and communication with the Library Manager.

5.4.2 z/VM native support using DFSMS/VM

DFSMS/VM Function Level 221 (FL221) is the only way for a z/VM system to communicate with an TS7700. DFSMS/VM FL221 is part of VM/ESA and z/VM. The RMS function of DFSMS/VM FL221 provides TS7700 support in VM/ESA environments at Version 1 Release 2 and all higher levels, as described in *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141.

Tape management

Although the RMS functions themselves do not include tape management system services, such as inventory management and label verification, RMS functions are designed to interface with a tape management system that can perform these functions. You can find additional information about third-party tape management systems that support the IBM TS7700 in the VM/ESA and z/VM environment in the following documents:

- *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-46322
- *Lights Out! Advanced Tape Automation Using VM/ESA*, GG24-4347

Figure 5-14 shows the z/VM native support for the TS7700. This is also true for VM/ESA.

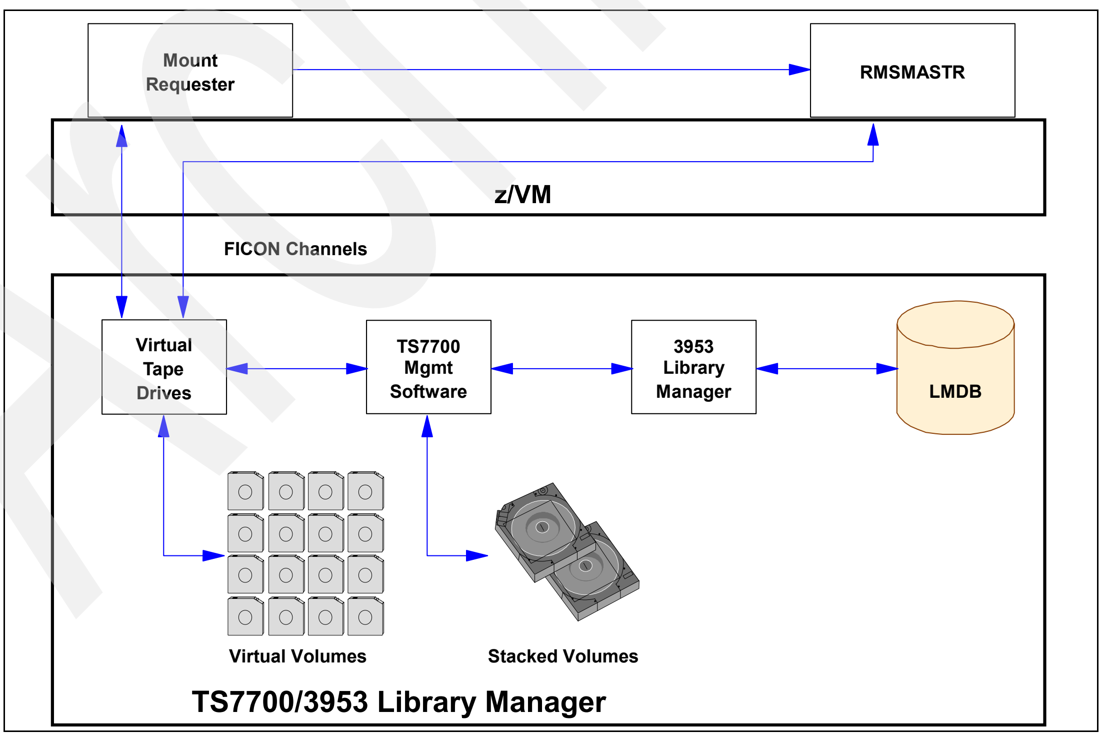

Figure 5-14 TS7700 in a native z/VM environment using DFSMS/VM

When you use the TS7700 in a VM environment, consider that many VM applications or system utilities use specific mounts for scratch volumes, so every time a mount request is issued from the host, the TS7700 has to recall the requested logical volume from the stacked cartridge if it is not already in the TVC. This can lead to performance degradation when writing data in a VM environment. In addition, VM backups usually require off-site movement, so the TS7700 is not the best candidate for this data.

DFSMS/VM

After you have defined the new TS7700 tape library through HCD, you must define the TS7700 to DFSMS/VM if the VM system is to use the TS7700 directly. You define the TS7700 tape library through the DFSMS/VM DGTVCNTL DATA control file. Also, you define the available tape drives though the RMCONFIG DATA configuration file.

You have the removable media services (RMS) as a component of DFSMS/VM. To allow RMS to perform automatic insert bulk processing, you must create the RMBnnnnn DATA file in the VMSYS:DFSMS CONTROL directory, where *nnnnn* is the five character tape library sequence number that is assigned to the TS7700 during hardware installation.

Note: The outboard policy management functions are currently not supported with z/VM.

For details on implementing of DFSMS/VM and RMS, refer to *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141. If the TS7700 is shared by your VM system and other systems, additional considerations apply. Refer to *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-4409, for further information.

5.4.3 Native z/VSE

Native support is provided for the Single Cluster Grid TS7700 configuration in z/VSE Version 3.1.2 and 4.1, which supports all IBM TS1120/3592-J1A configurations (J70, C06, and TS7700) without APARs in all automation offerings including IBM 3494 Enterprise Tape Library and IBM TS3500/3953 configurations. A Multi Cluster Grid is not supported under native VSE.

Note: The outboard policy management functions are currently not supported with z/VSE.

z/VSE supports the IBM TS3500/3953 natively through its Tape Library Support (TLS). In addition to the old tape library support, an added function allows the tape library to be supported through the /390 channel command interface commands, thus eliminating any XPCC/APPC communication protocol as required with the old interface. The external interface (LIBSERV JCL and LIBSERV macro) remains unchanged.

Define library support

First define which type of support you are using by specifying the SYS ATL statement. You can define:

TLS Tape Library Support, which provides full support

VSE LCDD, which does not support TS1120/3592 (only IBM 3490E and 3590). and therefore not the IBM TS3500

VM VM Guest Support

For native support under VSE, select TLS.

Define tape libraries

Define your tape library or libraries. This is done through a batch job, as shown in Example 5-9.

Example 5-9 Define libraries

```
* $$ JOB JNM=TLSDEF,CLASS=0,DISP=D
* $$ LST CLASS=A
// JOB TLSDEF
// EXEC LIBR,PARM='MSHP'
   ACCESS S=IJSYSRS.SYSLIB
   CATALOG TLSDEF.PROC REPLACE=YES
    LIBRARY_ID TAPELIB1 SCRDEF=SCRATCH00 INSERT=SCRATCH00             --- default library
    LIBRARY_ID TAPELIB2                * SECOND LIB DEF
    DEVICE_LIST TAPELIB1 460:463       *  DRIVES  460 TO 463
    DEVICE_LIST TAPELIB2 580:582       *  DRIVES  580 TO 582
    QUERY_INV_LISTS LIB=TLSINV         * MASTER INVENTORY FILES
    MANAGE_INV_LISTS LIB=TLSMAN        * MANAGE FROM MASTER
 /+
```

LIBSERV

The communication from the host to the Library Manager is through the LIBSERV JCL or macro interface. Example 5-10 shows a sample job using LIBSERV to mount volume 123456 for write on device address 480 and, in a second step, to release the drive again.

Example 5-10 Sample LIBSERV JCL

```
$$ JOB JNM=BACKUP,CLASS=0,DISP=D
$$ JOB BACKUP
// ASSGN SYS005,480
// LIBSERV MOUNT,UNIT=480,VOL=123456/W
// EXEC LIBR
BACKUP S=IJSYSRS.SYSLIB TAPE=480
/*
// LIBSERV RELEASE,UNIT=480
/&
$$ EOJ
```

LIBSERV provides the following functions:

Query all libraries for a volume	LIBSERV AQUERY,VOL=123456
Mount from category	LIBSERV CMOUNT,UNIT=480,SRCCAT=SCRATCH01
Mount a specific volume	LIBSERV MOUNT,UNIT=480,VOL=123456
Demount a volume	LIBSERV RELEASE,UNIT=480
Query count of volumes	LIBSERV CQUERY,LIB=TAPELIB1,SRCCAT= SCRATCH01
Query device	LIBSERV DQUERY,UNIT=480
Query inventory of library	LIBSERV IQUERY,LIB=TAPELIB1,SRCCAT=SCRATCH01
Query library	LIBSERV LQUERY,LIB=TAPELIB1
Manage inventory	LIBSERV MINVENT,MEMNAME=ALL,TGTCAT=SCRATCH01
Change category	LIBSERV SETVCAT,VOL=123456,TGTCAT=SCRATCH01
Query library for a volume	LIBSERV SQUERY,VOL=123456,LIB=TAPELIB1

For additional information, refer to *z/VSE System Administration Guide*, SC33-8224 and to *z/VSE System Macros Reference*, SC33-8230.

5.4.4 VM/ESA and z/VM guest support

In the following sections, we discuss two host environments that allow you to use an IBM TS7700 while running as a guest host system under VM/ESA or z/VM.

> **Note:** When z/OS is installed as a VM/ESA or z/VM guest on a virtual machine, you must specify the following statement in the virtual machine directory entry for the VM user ID under which the z/OS guest operating system is IPLed:
>
> STDEVOPT LIBRARY CTL

z/OS guests

It is possible for the environments described in 5.3.1, “z/OS and DFSMS/MVS SMS-managed tape” on page 210 to operate when z/OS is running as a guest of VM/ESA Release 2 or higher or z/VM Release 3.1. The considerations are the same as when z/OS runs natively without VM/ESA.

In this environment, additional software products are not required.

The STDEVOPT statement specifies the optional storage device management functions available to a virtual machine. The LIBRARY operand with CTL tells the control program that the virtual machine is authorized to issue tape library commands to an IBM Automated Tape Library Dataserver. If the CTL parameter is not explicitly coded, the default of NOCTL is used. NOCTL specifies that the virtual machine is not authorized to issue commands to a tape library, and this results in an I/O error (command reject) when MVS tries to issue a command to the library. For further information about the STDEVOPT statement, refer to *z/VM V5R1.0 DFSMS/VM Planning Guide*, SC24-6089 and *z/VM V5R2.0 Running Guest Operating Systems*, SC24-6115.

z/VSE guests

Some VSE tape management systems require VSE Guest Server (VGS) support as well as DFSMS/VM RMS for communication with the Library Manager of the TS7700 library.

If the VGS is required, define the LIBCONFIG file on the VGS service machine’s A disk. This file simply cross-references the z/VSE guest's tape library names with the names that DFSMS/VM uses. To enable z/VSE guest exploitation of inventory support functions through the LIBSERV-VGS interface, the LIBRCMS part must be installed on the VM system.

If VGS is to service inventory requests for multiple z/VSE guests, you must edit the LIBRCMS SRVNAMES cross-reference file. This file enables the inventory support server to access Librarian files on the correct VSE guest machine. See 5.4.5, “z/VSE as a z/VM guest using a VSE Guest Server (VGS)” on page 227. For further information, refer to *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-4409.

The CA DYNAM/TM-VSE does not use the VGS machine.

5.4.5 z/VSE as a z/VM guest using a VSE Guest Server (VGS)

When a z/VSE guest machine uses a tape drive in the TS7700, the virtual tape drive must be attached to that machine and the virtual tape volume must be mounted on the drive. Because, as a virtual machine, z/VSE cannot communicate with the Library Manager to request a tape mount, RMSMASTR (a VM/ESA machine) must attach the tape drive and mount the volume. z/VSE or VM/ESA cannot use RMSMASTR directly, however, because RMS functions run only in CMS mode.

Therefore, some z/VSE guest scenarios use the CMS service machine, called the VSE Guest Server (VGS), to communicate with RMSMASTR. VGS uses the standard facilities of RMS to interact with the 3953 Library Manager and the virtual drives of the TS7700. Figure 5-15 shows the flow and connections of a TS7700 in a z/VSE environment under VM.

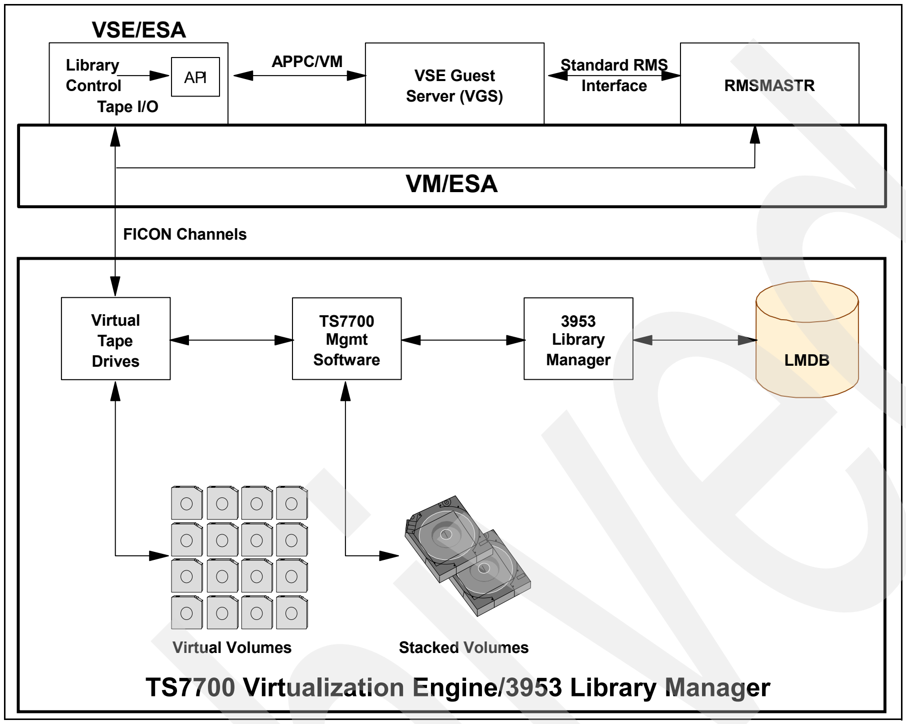

Figure 5-15 TS7700 in a z/VSE environment as a VM guest

Tape management

As with the VM/ESA native environment (see 5.4.2, “z/VM native support using DFSMS/VM” on page 224), the tape management system is responsible for keeping an inventory of volumes in the IBM TS7700 that belong to z/VSE. Some vendor tape management support scenarios do not use VGS. Instead, they communicate directly with RMSMASTR through CSL calls. Figure 5-16 shows the case of CA-DYNAM/T VSE.

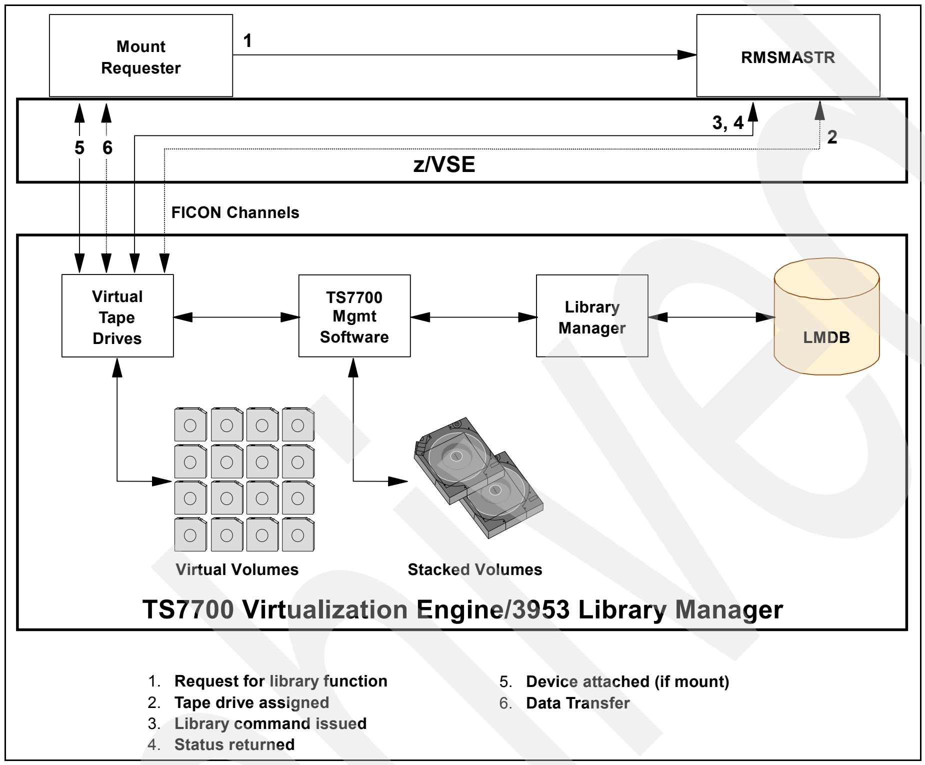

Figure 5-16 TS7700 in a z/VSE environment as a VM guest (no VGS)

VSE uses OEM tape management products that support scratch mounts, so if you are using VSE under VM, you have the benefit of using the fast-ready attribute for the VSE Library Manager scratch category.

For more information about z/VSE, refer to *z/VSE V4R1.0 Administration,* SC33-8304.

5.5 Software implementation in TPF

This section describes the support for an IBM TS7700 in a native Transaction Processing Facility (TPF) environment with TPF 4.1 or z/TPF. The TPF control program and several new and modified TPF E-type programs support the TS7700. The support is limited to a command-based interface.

Because TPF does not have a tape management system nor a tape catalog system, customers usually rely on z/OS to manage this function. In a TPF environment most tape data is passed between the systems. In general, 90 percent of the tapes are created on TPF and read on z/OS. The remaining 10 percent are usually created on z/OS and read on TPF.

We recommend that you use the normal z/OS and TS7700 installation process.

Specifics to TPF and z/OS with a shared TS7700

From the virtual drive side TPF must be allocated certain drive addresses. This will depend upon what tape function are need on TPF and will vary with the customer set. So the TS7700 will have tape addresses allocated to multiple TPF and z/OS systems.

Using the TS7700 with TPF

TPF uses virtual volumes from the z/OS scratch pools and shares the Library Manager scratch categories with z/OS. The z/OS host performs the insert processing for these virtual volumes and continues to manage them based on the input obtained from TPF. TPF has a set of commands (ztplf) that make it possible to load the volumes in TPF-allocated virtual drives.

After a volume is loaded into a TPF drive, you should have an automated solution in place that passes the volume serial number (VOLSER), the tape data set name, and the expiration date over to z/OS to have it processed automatically.

On z/OS you need to update the tape management system's catalog and the TCDB. This will allow z/OS to process virtual volumes that have been created by TPF. When the TPF-written volumes have been added to the z/OS tape management system catalog and the TCDB, normal expiration processing applies. When the data on a virtual volume has expired, and the volume is returned to scratch, the Library Manager and the TS7700 information are updated according to volume information maintained in z/OS.

Tapes that are created on z/OS and read into TPF

Tapes that are created on z/OS and read into TPF will use the same z/OS process for creating tapes. Now when TPF wants to read this z/OS-created tape, it does a specific mount of the tape VSN into a TPF-allocated drive using the TPF (ztplf) commands.

TS7700 performance for TPF

You can use the normal TPF Data Collection and Reduction reports that summarize read and write activity to the TPF allocated drive.

For TS7700-specific performance you should be using the normal TS7700 statistics that are off-loaded to z/OS through the TS7700 Bulk Volume Information Retrieval (BVIR) function.

Support of large virtual volumes for TPF (2 GB and 4 GB)

TPF itself does not use functions such as Data Class to control the cartridge volume size for specific cartridges. User exits provide the ability to set construct names for a volume. If you are not using the user exits, you can set the default in the TS7700/ Library Manager to the size for all virtual cartridges using the function for support as described in 4.6.5, “Implementing Outboard Policy Management for non-z/OS hosts” on page 190.

Here are some points in summary to consider when implementing an TS7700 in a TPF environment:

- Reserving a tape category does not prevent another host from using that category. It is the user's responsibility to monitor the use of reserved categories.
- Automatic insert processing is not provided in TPF.
- There is currently no IBM tape management system for TPF.

Advanced Policy Management is supported in TPF through a user exit. The exit is called any time a volume is loaded into a drive. At that time the user can specify, through the TPF user exit, whether the volume should inherit the attributes of an existing volume using the clone VOLSER attribute or the code can elect to specifically set any or all of the storage group, management class, storage class, or data class construct names. If the exit is not coded then the volume attributes remain unchanged as the volume is used by TPF. For the two levels of

TPF two different APARs are required for this support. For TPF 4.1 the APAR number is PJ31643 and for zTPF 1.1 the APAR is PJ31394.

Library Manager interface

TPF's only operator interface to the TS7700 is a TPF functional message, ZTPLF. The various ZTPLF functions provided allow the operator to manipulate the tapes in the library as operational procedures require. These functions include Reserve, Release, Move, Query, Load, Unload and Fill. Refer to *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

Control data sets

The TPF host does not keep a record of the volumes in the TS7700 Tape Library or manage the tape volumes in it. You can use the QUERY command to obtain information about the tape volumes held in the IBM TS3500/3952 Tape Library.

SIM and MIM presentation

SIMs and MIMs report hardware-related problems to the operating system. Refer to the *Statistical Analysis and Reporting System User Guide*, which you can access on the Web at:

http://www-1.ibm.com/support/docview.wss?uid=ssg1S7000247

SIM and MIM are presented in TPF by the messages CEFR0354I, CEFR0355W, CEFR0356W, CEFR0357E, CEFR0347W, CDFR0348W, and CDFR0349E, as well as EREP reports.

6

Migration

Although building on the proven concepts of several generations of IBM tape virtualization products, the Virtual Tape Servers, the TS7700 represents a new virtualization technology. Like most technology migrations, the upgrade to the TS7700 might require that the data is copied from other media into the TS7700. An outboard migration is available for specific configurations which we describe in this chapter.

We provide guidance on how to migrate your tape data into the TS7700 Virtualization Engine, covering the following topics:

- Migration overview
- Hardware migration
- Migration tools and methods
- Migration of DFSMShsm data
- Migration of Tivoli Storage Manager data
- Migration of database backups and logs

6.1 Migration overview

Migration of VTS Model B10 or B20 hardware to a TS7700 Virtualization Engine, also called *outboard VTS migration*, provides an upgrade path for existing B10 or B20 VTS models to a TS7700 Virtualization Engine. We provide different migration paths, depending on the existing B10/B20 VTS configuration and on the TS7700 target configuration which can be a Single Cluster Grid, a Two-Cluster Grid, or a Three-Cluster Grid configuration.

When migrating data from VTS to TS7700, and the TS7700 is connected to the existing library and the physical cartridges are not moved, this process is called *Merge*. If the data migration is performed from VTS to the TS7700 and if the TS7700 is connected to another library and the physical cartridges are moved from the existing library to the other, this process is called *Move*.

The replacement of the B10/B20 VTS with the TS7700 can be accomplished in different ways. Table 6-1 gives an overview of the possible migration scenarios for VTS and other migrations.

Table 6-1 TS7700 migration scenario

From → To	**TS7700 Single Grid**	**TS7700 Two-Cluster**	**TS7700 Three-Cluster**
B10/B20 VTS Standalone	See 6.2.1, "Standalone VTS to Single Cluster Grid" on page 244	Install second TS7700 and create Two-Cluster Grid configuration	See 6.2.4, "Three-Cluster Grid" on page 267
Two Standalone VTSs	See 6.2.2, "Merge two VTSs to a Single Cluster Grid TS7700" on page 252	Install second TS7700 and create Two-Cluster Grid configuration	See 6.2.4, "Three-Cluster Grid" on page 267
B10/B20 Peer-to-Peer VTS	N/A	See 6.2.3, "PtP VTS to Two-Cluster Grid" on page 260	See 6.2.4, "Three-Cluster Grid" on page 267
Native Drives and non-IBM tape hardware	See 6.3, "Moving data in and out of the TS7700" on page 276	See 6.3, "Moving data in and out of the TS7700" on page 276	See 6.3, "Moving data in and out of the TS7700" on page 276

Migration from native tape drives of any type to the TS7700 always requires host involvement to copy the data into the TS7700. See the sections starting from 6.3, "Moving data in and out of the TS7700" on page 276 for more information about the methods that you can use.

Similar to the previous B18 to B20 VTS migration, the hardware migration scenarios have the following aspects:

- Software changes in SMS, HCD, TCDB and RMM
- Migration of the Library Manager database from one existing tape library to another
- Physical swap of the B10/B20 VTS to the TS7700 hardware
- Migration of the database from B10/B20 VTS to theTS7700 Virtualization Engine

6.1.1 Migration planning

In the following sections, we describe the major considerations for migration from VTS Model B10 or B20 to a TS7700 Virtualization Engine and provide you with important information for the migration.

Replacing of the hardware is one but not the least of the items we consider here. These are the main aspects you should consider:

- Planning costs and benefits

  You will need to determine the best migration method for you. The various methods have temporary infrastructure implications. You can maintain the existing channel path configuration or duplicate the channel paths during the migration to help minimize system outage during the migration. Additional channel paths means the need for additional FICON director ports, additional channel ports on the host system, the need for more addresses in the unit control block (UCB), and more FICON cables.

- Planning software

  You can plan to verify any software prerequisite needed for the migration, prepare a new dynamic HCD generation and definitions, new DFSMS definitions, SMS, TCDB, RMM and what is needed depending on the scenario we are considering.

- Planning environment

  You can decide to use additional channel paths for disaster recovery or to minimize the system outages. Then new cables will be laid out in the environment. New Matrix configurations for the FICON Directors you might plan to update or reconfigure.

- Planning activities

  Plan to perform the hardware migration activities at a time when you minimize the impact on the production activities. Consider that the hardware migrations from B10/B20 VTS to TS7700 are disruptive, and therefore, with the assistance of the IBM System Service Representative, plan all the aspects of the activities for minimizing system outage.

> **Note:** When you start writing data into the TS7700 after the migration, the TS7700 will start normal background operations, such as reclamation and Secure Data Erase. As soon as there are changes to the logical or physical volumes there is no way to back out of the migration and go back to using the B10/B20 VTS.

Migration VTS rollout

With TS7700 R1.3 and R1.4, various hardware migration scenarios will be supported that allow migration without having to copy existing VTS data into the TS7700. See Table 6-2 for a summary of these migration scenarios and their availability.

With TS7700 R1.4a introduces new migration scenarios that include removal of a cluster from a Grid. These scenarios apply only if you are already in a TS7700 Multi Cluster Grid.

Table 6-2 Migration rollout and availability plan

Source Configuration	Target Configuration	Availability
B10/B20 PtP VTS in TS3500	TS7700 Grid in TS3500	Available[a]
B10/B20 PtP VTS in an IBM 3494	TS7700 Single or Multi Cluster Grid in an IBM 3494	Available[a]
B10/B20 Standalone or PtP VTS in an IBM 3494	TS7700 Single or Multi Cluster Grid in TS3500	Available[a]
Two Standalone B10/B20 VTSs attached to TS3500	TS7700 Single Cluster Grid attached to the same TS3500	Available with TS7700 R1.3[a]
Two Standalone B10/B20 VTSs attached to an IBM 3494	TS7700 Single Cluster Grid attached to different TS3500	Available with TS7700 R1.3[a]

Source Configuration	Target Configuration	Availability
Two-Cluster Grid and New TS7700	TS7700 Three-Cluster Grid	Available with TS7700 R1.3
Single Cluster Grid and new TS7700	TS7700 Two-Cluster Grid configuration	Available with TS7700 R1.3
Standalone B10/B20 VTSs attached to a TS3500	TS7700 Single Cluster Grid attached to different TS3500	Available with TS7700 R1.4[a]
Two Standalone B10/B20 VTSs attached to a TS3500	TS7700 Single Cluster Grid attached to different TS3500	Available with TS7700 R1.4[a]
Two Standalone B10/B20 VTSs attached to an IBM 3494	TS7700 Single Cluster Grid attached to the same IBM 3494	Available with TS7700 R1.4[a]
Existing Single Cluster Grid TS7700 and Existing Two-Cluster Grid TS7700 configuration	Merge to form TS7700 Three-Cluster Grid configuration	Available with TS7700 R1.4
Two existing Single Cluster Grid configurations	Merge to form TS7700 Two-Cluster Grid configuration	Available with TS7700 R1.4
TS7700 Three-Cluster Grid	Two-Cluster Grid and One TS7700 Single Cluster	Available with TS7700 R1.4ab
TS7700 Two-Cluster Grid configuration	Two TS7700 Single Cluster Grid	Available with TS7700 R1.4a[b]

a. Currently the upgrades from VTS Model B10 or B20 to TS7700 are available as a MES process ordering one of the following FCs to the B10 or B20 that will be replaced with a TS7700:
FC0522, Move VTS Database, if the new TS7700 is in the same Tape Library where VTS is installed
FC0523, Move VTS Database and Cartridges, if the new TS7700 will be installed in a different Tape Library

b. Removal of a Cluster from a Grid is a new function available in TS7700 R1.4a and requires a Cluster Cleanup procedure to be applied to the removed Cluster.

TS7700 Multi-Cluster Grid considerations

A TS7700 Grid refers to one, two, or three physically separate TS7700 clusters connected by means of a *customer-supplied* TCP/IP network.

The TCP/IP infrastructure connecting a TS7700 Multi Cluster Grid with two or three clusters is known as the *Grid Network*. The term *grid* refers to the code and functionality that provides replication and management of logical volumes and their attributes in cluster configurations. A Multi Cluster Grid provides remote logical volume replication and can be used to provide disaster recovery and high availability solutions. A disaster recovery solution is achieved when multiple clusters are geographically distant from one another.

You can make the following system configuration upgrades to an existing, on-site TS7700 cluster or TS7700 Grid:

- Two-Cluster TS7700 Grid

  You can add a new, empty TS7700 Cluster to an existing TS7700 cluster to create a TS7700 Two-Cluster Grid. With Release 1.4, you can combine two existing TS7700 clusters to create a Two-Cluster TS7700 Grid.

- Three-Cluster TS7700 Grid

  You can add a new, empty TS7700 cluster to an existing Two-Cluster TS7700 Grid to create a Three-Cluster TS7700 Grid. With Release 1.4, you can add an existing TS7700 cluster to an existing Two-Cluster TS7700 Grid to create a Three-Cluster TS7700 Grid.

> **Note:** For any of those upgrades, you must order FC4015, Grid Enablement, against the TS7700 Server.

Removal of a Cluster from a Grid and Cluster Cleanup procedure

You can remove a TS7740 Cluster from a TS7700 Grid concurrently with operations on the remaining TS7740 Clusters as long as you do not inhibit inserts, ejects, or exports. You must determine ahead of the removal process how to address volumes having only a copy consistency point at the TS7740 Cluster being removed. Prior to the cluster's removal, you must eject these volumes, move them to the scratch category, or copy them to another cluster using Management Class mount and demount operations.

You can clean up a standalone TS7740 Cluster or one that has been removed from a TS7700 Grid to achieve a state similar to new.

- TS7700 Two-Cluster Grid

  With TS7700 R1.4a, removal of one cluster is supported so that you will have a Single Cluster Grid remaining. Using the Cluster Cleanup feature, the TS7700 that has been removed is returned to a state similar to the state the machine is shipped from manufacturing.

  You can remove one TS7700 Cluster from the existing TS7700 Two-Cluster Grid. With Release 1.4a, it is possible and you will have one TS7700 Single Cluster Grid remaining, and the removed TS7700 will be returned to a state similar to one that is received from manufacturing with the feature codes still in place through the Cluster Cleanup procedure.

- TS7700 Three-Cluster Grid

  You can remove one TS7700 Cluster from the existing TS7700 Three-Cluster Grid. With Release 1.4a, you can remove a Cluster from a Grid and continue working with a Two-Cluster TS7700 Grid. The removed TS7700 will be returned to a state similar to one that is received from manufacturing with the feature codes still in place through the Cluster Cleanup procedure.

> **Note:** For the Removal of a Cluster of a Grid, you must order FC4016, Remove Cluster from a Grid, and FC4017, Cluster Cleanup, against the TS7700 that is going to be removed.

For a migration from a B10/B20 VTS to a TS7700 Multi-Cluster Grid configuration using the customer-supplied TCP/IP network, we recommend that you have the network prepared at the time the migration starts.

The TS7700 Virtualization Engine provides two independent 1 Gbit copper (RJ-45) or shortwave fiber Ethernet links for Grid Network connectivity. It is recommended that you connect each through an independent WAN interconnection to be protected from a single point of failure that would disrupt service to both WAN paths from a node.

VTS migration procedures and times

Migration procedures involve the replacement of an existing B10 or B20 VTS with a TS7700 Virtualization Engine. The replacement can be accomplished in two different ways:

- The database on the existing VTS is moved to a new TS7700. In this case the TS7700 is attached to the same tape library as the VTS that is being replaced. You must order FC0522 (*Migrate a VTS Database*) to perform this operation.
- The database on the existing VTS is moved to a new TS7700. In this case the TS7700 is attached to a different library than the VTS that is being replaced. All of the physical tapes associated with the VTS partition are also moved to the new library. You must order FC0523 (*Move VTS Database and Cartridges*) to perform this operation.

> **Note:** FC0522, Migrate VTS Database, and FC0523, Move VTS Database and Cartridges are available until 05 December 2008. After this date, the upgrade path from VTS to TS7700 will be available as an IBM Global Technology Service offering.

All the migration procedures are disruptive including those that involve existing PtP VTSs. You should plan for a downtime of between 8 and 12.5 hours, depending upon the migration type. During this time, you will not have access to the data in the VTS that is being migrated. You should plan systems outage to perform the necessary steps because the outboard migration is not concurrent with the system production activities.

If the VTS shares an IBM 3953 or an IBM 3494 Library Manager with another VTS or TS7700, or with native drives, those subsystems will also be affected because the Library Manager must be re-taught that it is now attached to a TS7700. Those other subsystems will not be impacted for the entire outage time, but only for the amount of time needed to vary the Library Manager offline, perform the re-teach, and return it again to the online status.

If you have multiple VTSs installed, we recommend that you reduce workload and if possible shift new workload away from the VTS that is next to be migrated a day prior to the migration. This gives VTS time to copy all of the logical volumes in its TVC to the physical tape. This can reduce the outage time needed for the migration.

To minimize the impact of the migration, consider temporarily redirecting your tape workload to a different library and VTS, TS7700, or native drives. This can be easily done by changing the ACS routines and enables you to write tape data. However, you will not be able to read the logical volumes in the VTS being migrated.

IBM 3494 Tape Library attachment

With the release of the TS7700 microcode level R1.2 (8.2.x.x), IBM 3494 Tape Library attachment is now available for the TS7700 Virtualization Engine. The attachment of the TS7700 Virtualization Engine to the 3494 Tape Library involves the following connections:

- Power control

  The 3494 Tape Library contains power sequencing and control equipment that remotely turns the power to the TS7700 Virtualization Engine on and off. This is the normal mode of power control for the TS7700 Virtualization Engine.
- Library Manager

  The 3494 Tape Library is one or two internal Library Managers. Communication and control signals travel among the Library Manager, the TS7700 Virtualization Engine, and the tape drives within the 3494 Tape Library through a local area network (LAN). Data between the host and the TS7700 Virtualization Engine travels along FICON connections, while data travels between the TS7700 Virtualization Engine and the tape drives within the 3494 Tape Library along a Fibre Channel.

- Tape Drives

  A VTS supports 3590, 3592-J1A, and TS1120 (3592-E05) Tape Drives. The TS7700 Virtualization Engine can only attach to 3592 Tape Drives. Up to 12 3592 Tape Drives in a single model D22 frame in the 3494 Tape Library can be attached to the TS7700 Virtualization Engine. A maximum of two TS7700s can be connected to a single 3494 Tape Library.

> **Note:** If you are planning to migrate a B20 VTS with 3590 tape drives to a TS7700, you must first upgrade the 3590 drives to TS1120 drives and migrate all data to the new media. For more information about upgrading tape drives, see:
>
> - *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229
> - *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115

The following feature codes are required on the 3494 Tape Library for connection to a TS7700 Virtualization Engine:

- FC5124 – Second Disk Drive for the Library Manager
- FC5246 – Dual Control Path Facility
- FC5047 or FC9047 – The LAN PCI Library Manager
- FC9013 on LM for TS7700 Attached
- FC9218 on TS7700 to indicate 3494 attachment

Microcode levels

To support the upgrade, the B10/B20 VTS must be at least at microcode release 2.32.747.xx (R7.4+7). That code level supports a method to make a copy of the VTS database that can be converted by the TS7700 into its DB2® database. Also, because the TS7700 does not support 2 Gbit Fibre Channel switches to the 3592 or to TS1120 drives, configurations that have 2 Gbit switches must have the switches replaced with 4 Gbit switches as part of the hardware change. We describe this in detail in 6.2.1, “Standalone VTS to Single Cluster Grid” on page 244 and in 6.2.3, “PtP VTS to Two-Cluster Grid” on page 260.

The Library Manager microcode of the IBM 3953 Enhanced Library Controller (ELC™) or the IBM 3494 Tape Library associated with the upgrade, must be at the level that supports the TS7700 Virtualization Engine implementation. We recommend that you plan this upgrade separately and before the installation of the migration.

DFSMS definitions

Because a TS7700 Virtualization Engine, even as a Single Cluster Grid system, looks to the host as a peer type system with a Distributed and a Composite LIBRARY-ID, the DFSMS library definitions will need to be updated as part of the migration. We recommend that the current library name (8 character name defined at DFSMS) be retained as the Composite Library name and a new name be defined for the underlying Distributed Library.

Keeping the same library name for the Composite Library avoids having to change the library name associated with all of the logical volumes in the TCDB and the tape management system catalog. It might also eliminate the need to change the storage group definitions.

HCD definitions

We recommend that you prepare the IODF in advance with the LIBRARY-ID and LIBPORT-IDs of the Composite Library. We provide an example of Distributed LIBRARY-ID, the Logical Control Unit and the LIBPORT-ID IDs relationship. See Table 6-3 for reference. For details refer to 5.1.3, “Logical path considerations” on page 195.

Table 6-3 LIBPORT ID

Distributed Library	Logical Control Unit	Libport/Subsystem IDs
Cluster 0	0-F	01-10
Cluster 1	0-F	41-50
Cluster 2	0-F	81-90

You might need additional FICON director ports unless the B10/B20 ports can be reused. Because the TS7700 only supports FICON attachment, reuse of existing paths is not possible for ESCON-attached B10/B20 VTSs.

6.1.2 Hardware processes

Here we cover general considerations regarding the hardware processes involved in the hardware migration from the B10/B20 VTS to the TS7700 Virtualization Engine.

Data format conversion

The physical data formats used by the B10/B20 VTS and the TS7700 are different. The TS7700 format was designed to allow the physical tapes to be directly exportable in the future. The TS7700 can read the old VTS format, but will only write data in the new format.

The conversion of the physical data format is handled completely transparent to the customer’s applications. Data written in the old data format is automatically converted to the new data format whenever the data is accessed. The access can be because of a host recall of the data or as part of reclamation processing for a physical volume. Other than to be ready for the export functions, there is no need to convert the data as the data is completely accessible in the old format. Figure 6-1 on page 241 illustrates this concept.

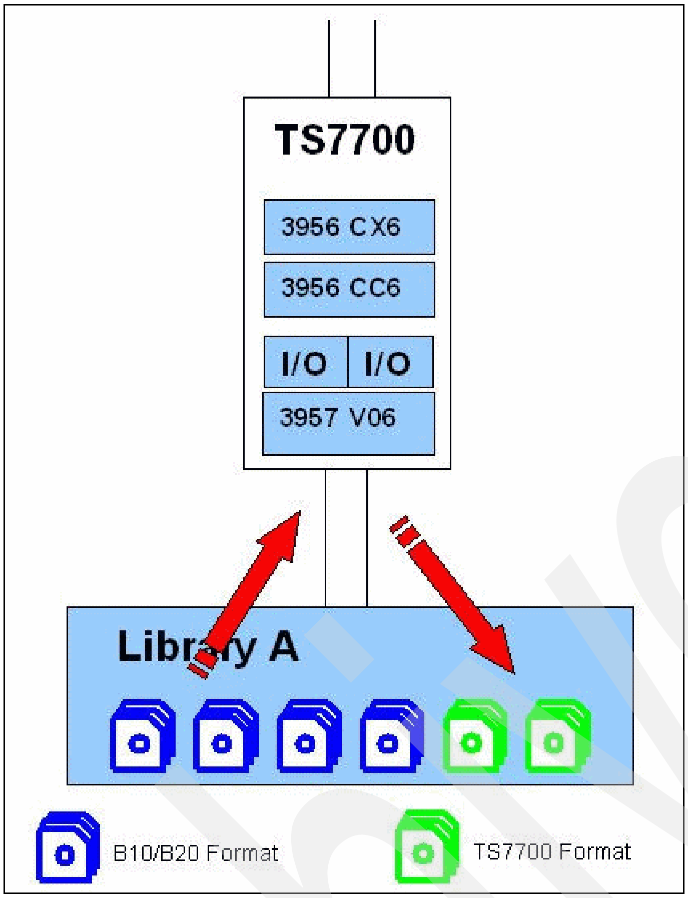

Figure 6-1 Conversion of the physical data format

If you want, you can expedite the conversion by changing the reclamation criteria for the pool the old format volumes are in and defining a new pool for the new format data (including new writes from the host). The target pool for the old data would be the new pool. You should consider adding extra physical scratch volumes prior to the migration to be sure that there will be enough scratch media on the new TS7700 configuration.

Back up VTS data

This procedure is used to extract data from the existing B10/B20 VTS to be exported on the TS7700 using outboard migration. This procedure is an IBM System Service Representative (SSR) initiated task. The process we show creates a number of files that will be used by the TS7700 to build its DB2 database and to determine the configuration of the new machine.

The backup process does the following:

- Migrates all data from B10/B20 VTS cache to tape
- Reconciles the B10/B20 VTS database
- Extracts data from the B10/B20 VTS database to comma-separated files
- Gathers logical volume data from the previously active Library Manager

See Figure 6-2 on page 242 for reference.

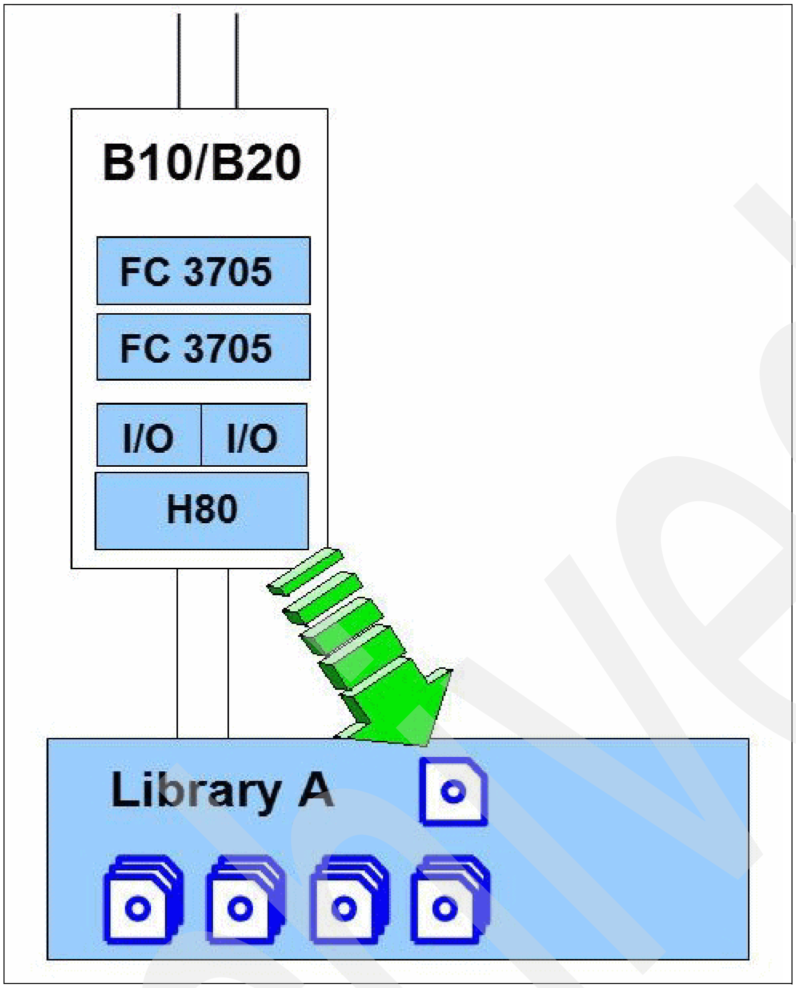

Figure 6-2 B10/B20 VTS data backup

Restore VTS data

This procedure is used to export data into the TS7700 using outboard migration. This is an IBM SSR initiated task. The process we show restores the data from the previously created backup tapes to the TS7700. The VTS files are post-processed, and a series of DB2 import files are created.

The restore process performs the following tasks:

- Converts the B10/B20 VTS data into importable DB2 files
- Creates an empty database if one does not exist
- Cleans up existing databases if they are not needed
- Merges the data with the existing data if needed

See Figure 6-3 on page 243 for reference.

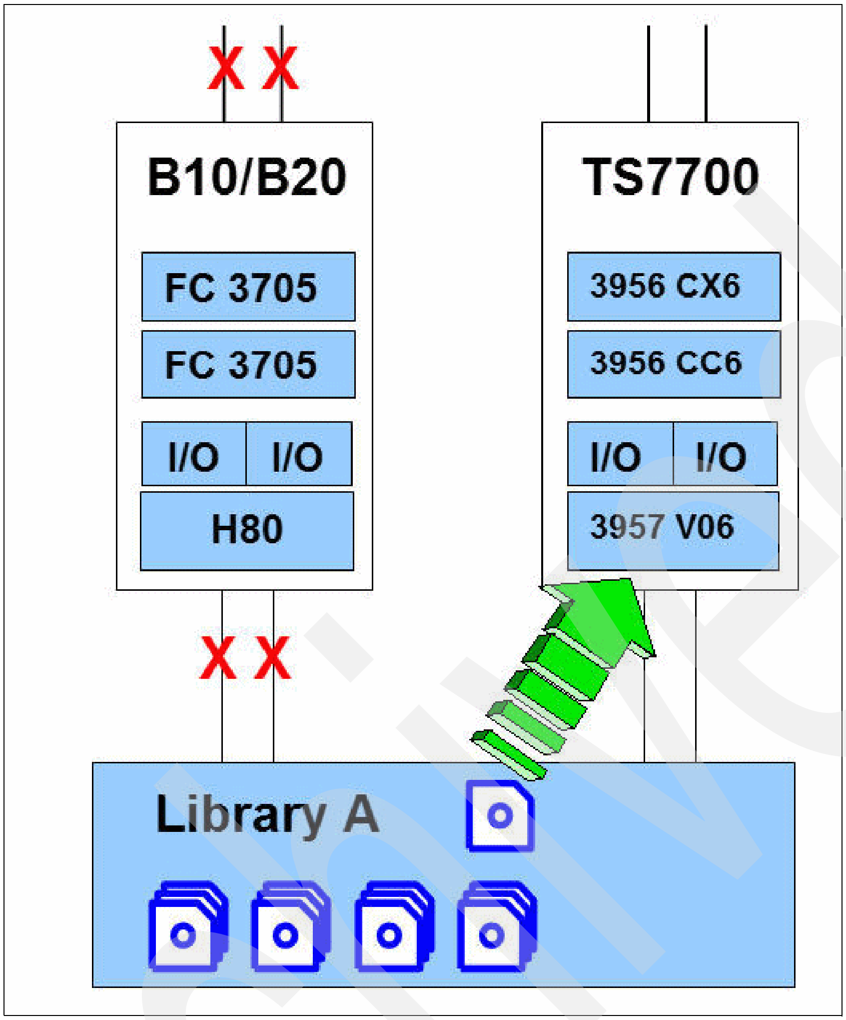

Figure 6-3 TS7700 Virtualization Engine data restore

6.2 TS7700 hardware migration scenarios

In this section we describe in detail the hardware migration from a B10/B20 VTS to the TS7700 Virtualization Engine. We provide various migration scenarios, depending on the existing VTS configuration and the final TS7700 Virtual Engine configuration, which can be a Single Cluster Grid, a Two-Cluster Grid, or a Three-Cluster Grid configuration. Refer to Table 6-2 on page 235, which shows the matrix of which migration steps apply to which scenario.

Note: In this section we describe various scenarios, but we cannot cover every single configuration. If your migration scenario is not described, it should be possible to combine information from the listed scenarios and derive information from them for your own scenario.

6.2.1 Standalone VTS to Single Cluster Grid

In this section we describe the data and hardware migration from a Standalone VTS B10 or B20 to the TS7700 Virtualization Engine connected to an IBM 3494 Tape Library or to a TS3500 Tape Library using the 3953 Library Manager.

We discuss two scenarios:

- The TS7700 connects to the same tape library as the VTS (see Figure 6-4)
- The TS7700 connects to a different library (see Figure 6-7 on page 250)

Existing tape library

If you will be migrating the data from a Standalone B10/B20 VTS to a Single Cluster Grid, where your existing VTS is connected to an existing tape library, either IBM 3494 or TS3500, use this procedure. Figure 6-5 on page 245 gives an overview of the hardware migration.

Figure 6-4 gives you an overview of the migration process that we describe in this section.

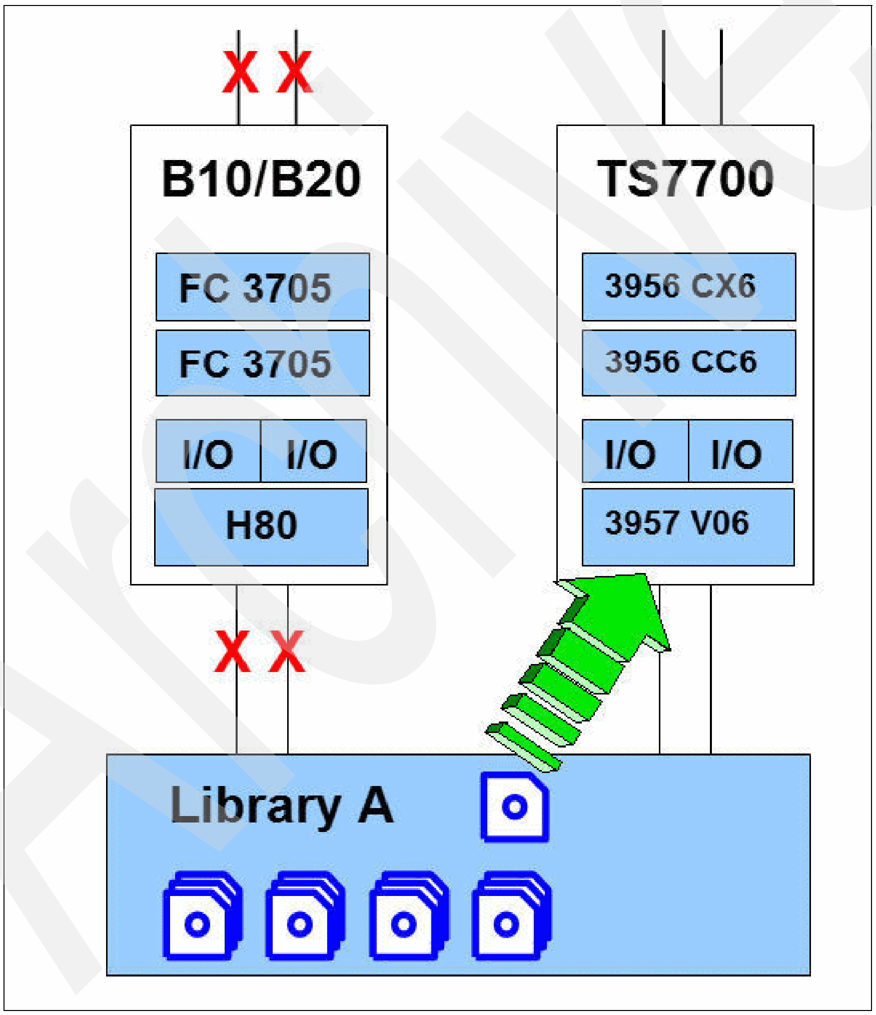

Figure 6-4 Standalone VTS hardware migration

Note: This migration does not change the tape drive configuration. Therefore, in a 3494 Tape Library, the VTS to be replaced must already have 3592 tape drives installed, and all data must have already been migrated to 3592 cartridges.

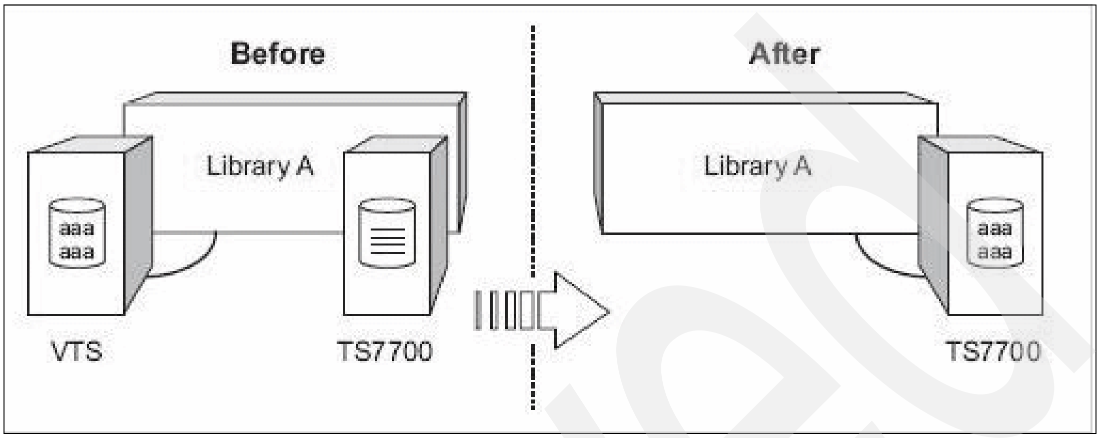

Figure 6-5 Standalone VTS to TS7700 Single Grid in the same tape library

These are the required migration steps:

1. Start the installation of the TS7700 hardware a few days prior to the outage window.
2. Validate whether there is any conflicting information in the TCDB, the tape management system catalog, and the Library Manager database. This is only needed if you plan to change the library name as described in Step 4. Example 6-1 offers a sample job for using the RMM utility EDGUTIL for verification.

   Example 6-1 Verify information in RMM CDS, Library Manager database, and TCDB

   ```
   //EDGUTIL  EXEC PGM=EDGUTIL,PARM='VERIFY(ALL,VOLCAT)'
   //SYSPRINT DD SYSOUT=*
   //MASTER   DD DSN=your.rmm.database.name,DISP=SHR
   //VCINOUT  DD UNIT=3390,SPACE=(CYL,(900,500))
   ```

   As a result of running EDGUTIL, you will get information about all volumes with conflicting information. We recommend that you resolve discrepancies before the migration. For more information about this utility, refer to *z/OS V1R8.0 DFSMSrmm Implementation and Customization Guide,* SC26-7405. The job should be run before the migration starts.

3. Stop all host activity.
   a. If there is another VTS or TS7700 system available to the host during the migration, you can change the ACS routines to direct allocations to that other system.
   b. Complete or cancel all host jobs for the VTS.
   c. Vary offline all device addresses associated with the VTS for all attached hosts.
   d. Vary the VTS to be migrated offline to all hosts.
   e. Vary the channel paths to the VTS offline.
4. Prepare the software changes. All of the following steps can be done concurrent with the hardware changes that follow (see Step 6 on page 249 and the following).
   a. Define the library names to SMS.

      When you define the Single Cluster Grid TS7700, you need to define a Composite Library as well as one Distributed Library in SMS. We recommend that you reuse the existing library name of the VTS as the Composite Library name as summarized in the top half of Figure 6-6. Reusing the library name has the advantage that you do not have to update the library name in the volume records in the TCDB and the tape management systems catalog. In this case, you just update the existing library definition “MyLib” and define the new distributed library “Distlib”.

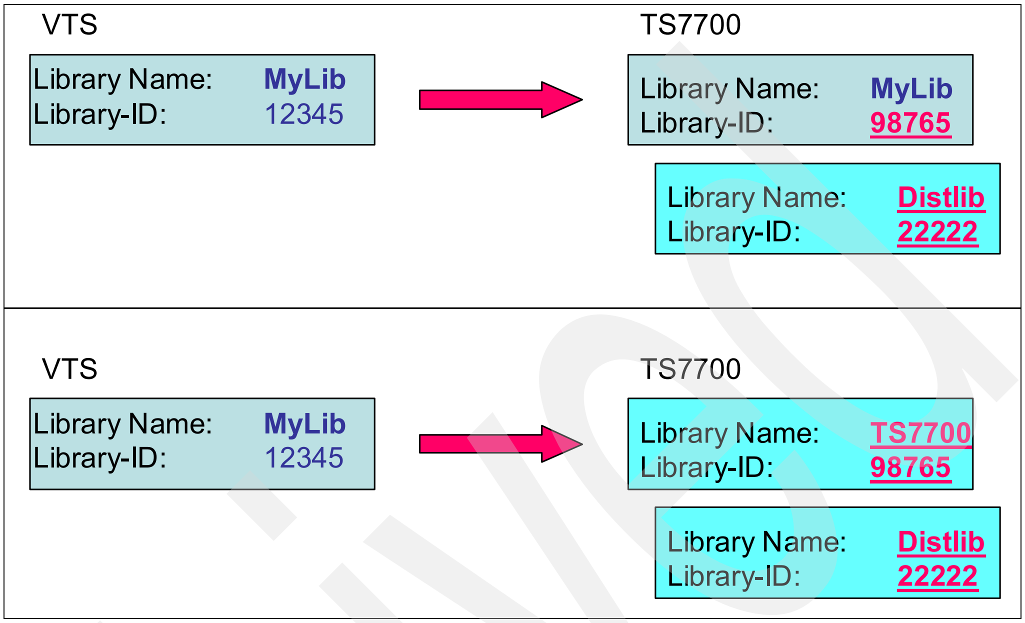

Figure 6-6 Defining library names

If you do not keep existing library names as shown in the bottom half of Figure 6-6, you have to change the library names in all volume entries in the TCDB and your tape management system catalog to the new names before you can use the volumes. Delete the old library definition “MyLib” and define the two new libraries “TS7700” and “Distlib” in SMS through Interactive Storage Management Facility (ISMF). Remember to write the new Library-IDs in the definitions. If a new Composite Library name is used, then update existing Storage Group definitions to relate to that new library name. When you delete a library in the SCDS through ISMF, the change needs to be reflected in the TCDB as shown in Example 6-2

Example 6-2 Delete a library in the TCDB

```
//STEP1      EXEC   PGM=IDCAMS
//SYSPRINT   DD     SYSOUT=*
//SYSIN      DD     *
 DELETE (vtsname) -
 LIBRARYENTRY
```

b. Changes to HCD channel definitions

   You need to plan appropriately if you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or might switch from ESCON to FICON with this process. If you define the devices as offline in HCD and use any product for device sharing, you need to define the new addresses to that product as well.

c. Changes to LCU definitions

   Changes are needed because you change the Library-ID. If the VTS is a model with 64 or 128 logical units, you also need to define more LCUs to enable the use of 256 logical units supported on a Single Cluster Grid. See Chapter 5, “Software implementation” on page 193 for more information.

d. Missing Interrupt Handler values

   If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.5, “Set values for the Missing Interrupt Handler” on page 208.

e. Complete the software definitions.

   If you have redirected workload to other systems during the migration, you can need to change the ACS routines back. If you are planning to use new SMS constructs and implement functions provided through outboard policy management, you should define them on the host now.

   Activate the IODF and SMS definitions and issue an OAM restart if not done after the SMS activation. For more information, see Chapter 5, “Software implementation” on page 193.

> **Note:** If the new SCDS is activated ahead of the new library being ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:
>
> `CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.`

5. Change volume entries in the TCDB

   If the Composite Library name remains the same as the previous name, no changes are required and you can continue with Step 6 on page 249.

   If you change the Library name, you must change all volume entries in the TCDB to relate to the new name for every single volume. You can get a list of all volumes in the TCDB by running a job like the one shown in Example 6-3.

   Example 6-3 JCL to list all entries in the TCDB

```
//STEP1    EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 LISTC VOLUMEENTRIES(V*) LIBRARY(vtsname)
```

   Update of all volume entries can be a time-consuming process. Because of the time used, you can change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible and after that take all the rest. However, we do not recommend that you start using the TS7700 before all TCDB volume records have been updated with the new library name.

   See Example 6-4 for a sample of the JCL to change one scratch volume and one private volume.

   Example 6-4 JCL for changing the TCDB to a new TS7700

```
//*****************************************************************
//**** Change TCDB for a scratch volume to a new TS7700       ****
//*****************************************************************
//TCDBSCR EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLENTRY LIBRARYNAME(TSname) USEATTRIBUTE(SCRATCH)
```

```
//*****************************************************************
//**** Change TCDB entry for a private volume to a new TS7700 ****
//**** Also change sgname to the one used (same as on the VTS)****
//*****************************************************************
//TCDBPRIV EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLUMEENTRY LIBRARYNAME(TSname) -
       USEATTRIBUTE(PRIVATE) STORAGEGROUP(sgname)
```

If you are using RMM as your tape management system, a command to RMM for each volume is also needed. The command is shown in Example 6-5.

Example 6-5 JCL for changing volumes in DFSMSrmm to a new TS7700

```
//PROCESS  EXEC PGM=IKJEFT01,DYNAMNBR=25,
//             TIME=100
//ISPLOG    DD DUMMY
//SYSPRINT  DD SYSOUT=*
//SYSTSPRT  DD SYSOUT=*
//SYSTSIN   DD *
 RMM CV volser LOCATION(TSname)
```

Even if you specify the FORCE parameter, it will only take effect if necessary. This parameter requires that you are authorized for a specific RACF® Facility class named STGADMIN.EDG.FORCE. You should verify that you have the required authorization.

Example 6-6 provides a sample REXX EXEC for updating the library name for volumes in the TCDB.

Example 6-6 REXX EXEC example for updating the library name in the TCDB

```
/* REXX */
/*************************************************************/
/* ALTERVOL                                                  */
/*                                                           */
/* Usage:  ALTERVOL  DSN(volserlist) LIB(libname)            */
/*                                                           */
/*         Before this EXEC is run, you must create the      */
/*         input data set "volserlist". The LISTCAT command  */
/*         can be used to generate the list of volumes       */
/*         to be altered to an output dataset.               */
/*                                                           */
/*         LISTCAT VOLUMEENTRIES(V*)                         */
/*                 LIBRARY(sourcelib)                        */
/*                 OUTFILE(ddname)                           */
/*                                                           */
/*         The list generated has the following format:      */
/*               VOLUME-ENTRY----Vvolser                     */
/*         For command specifics, refer to "Access Method    */
/*         Services for the Integrated Catalog Facility".    */
/*                                                           */
/*         For each volume in the "volserlist" specified,    */
/*         the library name in the volume record is updated  */
/*         to the library name specified on the invocation.  */
/*                                                           */
/*         ALTER Vvolser VOLUMEENTRY LIBRARYNAME(libname)    */
```

```
/************************************************************/
Arg parms
Dsn=''; Lib=''
If pos('DSN(',parms)>0 then do
   parse var parms front 'DSN(' dsn ')' back
   parms = front || back
   end
If pos('LIB(',parms)>0 then do
   parse var parms front 'LIB(' lib ')' back
   parms = front || back
   end
If dsn='' | lib='' then do
   'Usage:   ALTERVOL DSN(volserlist) LIB(libname) '
   exit 4
   end
/************************************************************/
/*  Get volume serials from source input dsn                */
/************************************************************/
Address TSO "FREE FI(INDD)"
Address TSO "ALLOCATE FI(INDD) DA("dsn") SHR"
Address TSO "EXECIO * DISKR INDD (STEM X."
Alter1 = "ALTER '"
Alter2 = "' VOLUMEENTRY LIBRARYNAME("lib")"
Volumes = 0
Do N=1 to X.0
   If Pos("VOLUME-ENTRY----",x.n)>0 then do
      Volumes = Volumes + 1
      Parse var x.n "VOLUME-ENTRY----" volser .
      Address TSO Alter1||volser||Alter2
   end
End
Say "Lines Read:      " format(x.0,9)
Say "Volumes Altered: " format(Volumes,9)
Address TSO "EXECIO * DISKR INDD (FINIS"
Address TSO "FREE FI(INDD)"
Exit 0
```

6. Drain the VTS Tape Volume Cache.
7. Extract the VTS database as explained in “Back up VTS data” on page 241.
8. Disconnect the VTS from its Tape Library (either IBM 3953/TS3500 or IBM 3494), and from its TS1120 or 3592 tape drives.
9. If the VTS connects to the existing 3953 or 3494 Library Manager, using 2 Gbit Fibre Channel switches, they must be replaced with the 4 Gbit switches at this time.
10. Complete the installation of the TS7700:
    a. Connect the TS7700 to the Library Manager.
    b. Connect the TS7700 to the TS1120 or 3592 drives.
    c. Teach the library.
11. Restore the VTS database to the TS7700, as explained in “Restore VTS data” on page 242.
12. Vary the TS7700 online from the TS7700 System Management Interface (SMIT) panel, selecting **3957-V06 Online/Offline Menus** → **Vary 3957-V06 Online**.

13. Now you are ready to test the new Single Cluster Grid:

a. Vary the defined channels online to the host.
b. Vary the logical devices online to the host.
c. Vary the Composite and Distributed Libraries online to the host.
d. Run test jobs to read and write from the TS7700 Single Cluster Grid.

Note: The IBM SSR will perform Step 1 and Steps 6 - 12 of the migration steps listed above. We show them for your information only.

New target tape library

Migration from a standalone B10/B20 VTS and tape library to a new TS7700 installed in a different IBM tape library includes moving the cartridges from the source library to the target library, as illustrated in Figure 6-7.

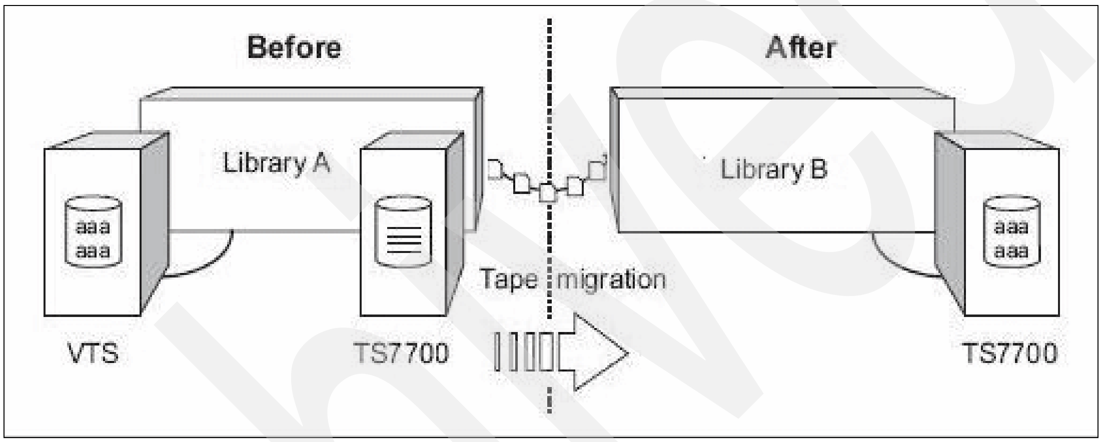

Figure 6-7 Standalone VTS to TS7700 in a different tape library with cartridge move

If the migration scenario occurs within the same physical TS3500 Tape Library, where one B10/B20 VTS is attached to one Logical Library with its own Library Manager (3953-L05 number 1), and the TS7700 is attached to another Logical Library with its own Library Manager (3953-L05 number 2), then no physical stacked cartridges that belong to the B10/B20 VTS will be moved, but the logical associations of the physical volumes from the VTS to the TS7700 will be transferred. You have to reassign the physical volume serial numbers that are dedicated to the B10/B20 VTS logical library to the TS7700 logical library using the TS3500 Specialist; see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789 for details. Transfer the following data from the source Library Manager (3953-L05 number 1) to the target Library Manager (3953-L05 number 2):

- 3953 Library Manager Administrative data, which includes the library constructs from the source library.
- 3953 Library Manager data, which includes the Volume Serial ranges for the media types.

The time required for the physical move of the cartridges from one library to another can vary depending on the number of cartridges and the distance they are transported. We recommend that you plan extra outage time for this activity.

If you are installing new drives together with the TS7700 in the new library, the TS7700 Virtualization Engine can be connected to the tape library, the Library Manager, and the drives, tested, and taught in advance to minimize the system outage window needed for the migration.

If you are reusing the existing VTS-attached drives in the new library, plan for extra outage to remove the drives from the existing library and install them in the new library. For the remainder of this scenario, we assume that you are not moving the existing VTS drives.

These are the steps required to accomplish the migration scenario described (see Figure 6-7):

1. Back up the Library Manager Administrative Data regarding the library constructs from the source library. You can perform this activity prior to the outage window.
2. Determine whether there is any conflicting information in the TCDB, the tape management system catalog, and the Library Manager database. This is only needed if you plan to change the library name as described in Step 4 on page 245. Example 6-1 on page 245 shows the sample job for using the RMM utility EDGUTIL for verification.
3. Stop all host activity.
   a. If there is another VTS or TS7700 system available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
   b. Complete or cancel all host jobs for the VTS.
   c. Vary off all device addresses associated with the library for all attached hosts.
   d. Vary the existing VTS offline to all hosts.
   e. Vary the existing channels offline.
4. Prepare the software changes. All of the following steps can be done concurrent with the hardware changes that follow:
   a. Defining the library names to SMS

      When you define the Single Cluster Grid TS7700, you need to define a Composite Library as well as one Distributed Library in SMS. We recommend that you reuse the existing library name of the VTS as the Composite Library name explained in Figure 6-6 on page 246.

   b. Changes to HCD channel definitions

      If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or might switch from ESCON to FICON with this process, these are planning considerations that should be completed at this time. You might define new FICON channels to ease the process.

   c. Changes to HCD LCU definitions

      Changes are needed because you change the Library-ID. If the VTS is a model with 64 or 128 logical units, you also need to define more LCUs, to enable the use of 256 logical units supported on a Single Cluster Grid. See Chapter 5, “Software implementation” on page 193 for more information.

      If you define the devices as offline in HCD and use any product for device sharing, you need to reflect the new addresses to that product.

   d. Missing interrupt values

      If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.5, “Set values for the Missing Interrupt Handler” on page 208.

   e. Completing the software definitions

      If you have redirected workload to other systems, during the migration, you might need to change the ACS routines back. If you are planning to use new SMS constructs and implement functions provided through outboard policy management, you should define them on the host now.

Activate the IODF and the SMS definitions and issue an OAM restart if not done after the SMS activation. For more information, see Chapter 5, “Software implementation” on page 193.

> **Note:** If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:
>
> ```
> CBR3006I Library library-name with Library ID library-ID unknown in I/O
> configuration.
> ```

5. Change volume entries in the TCDB

   If the Composite Library name is the same as the previous name, no changes are required, and you can continue with the next step. If you had to change the library name, refer to Step 5 on page 247 for more information about how to update the TCDB volume entries.
6. Drain the VTS Tape Volume Cache.
7. Extract the VTS database; see “Back up VTS data” on page 241 for details.
8. Back up IBM 3494 or 3953 Library Manager data. This includes the Volume Serial range for the media types.
9. Remove the physical cartridges that belong to the VTS being migrated from the source library.
10. Restore the VTS database in the TS7700 as explained in “Restore VTS data” on page 242.
11. Restore IBM 3494 or 3953 Library Manager data. This includes the Volume Serial range for the media types.
12. Insert the physical cartridges previously removed from the source library into the target library.
13. Initialize the Library Manager and perform an inventory of the library. This will allow the Library Manager to recognize the physically inserted cartridge.
14. Vary TS7700 online from the TS7700 System Management Interface (SMIT) panel, selecting **3957-V06 Online/Offline Menus** → **Vary 3957-V06 Online**.
15. Now you are ready to test the new Single Cluster Grid:
    a. Vary the defined channels online.
    b. Vary the logical devices online.
    c. Vary the Composite and Distributed Libraries online to the host.
    d. Run test jobs to read and write from the TS7700 Single Cluster Grid.

> **Note:** The IBM SSR will perform Steps 1, 6, 7, 8, 10, 11, 13, and 14 as part of the installation. We list them for your information.

6.2.2 Merge two VTSs to a Single Cluster Grid TS7700

In this section we describe the migration of two standalone VTSs to a TS7700 Single Cluster Grid.

With Release 1.4, you can migrate standalone VTS systems to a Single Cluster Grid TS7700 Virtualization Engine without the need to use host utilities to copy the data. The TS7700 can be attached to the same IBM 3494 or TS3500 Tape Library as the VTS systems (Figure 6-8 on page 253), or it can be attached to a different IBM 3494 or TS3500 Tape Library (Figure 6-9 on page 256).

If two standalone VTS systems are migrated to a Single Cluster Grid TS7700 Virtualization Engine, it is assumed that the physical volumes for one VTS system will remain in the current tape library to be used by the TS7700 Virtualization Engine and that physical volumes for the other VTS system will reside in a different tape library. The physical volumes belonging to the other VTS will be moved to the library where the TS7700 Virtualization Engine is attached. If both VTS systems share the same Library Manager, physical transfer of physical volumes does not occur. Instead, all the associations of the partitions to which the physical volumes belong, will be transferred.

Merge two standalone VTSs to one TS7700 in an existing Tape Library

If you are merging two standalone VTSs B10/B20 data into a Single Cluster Grid TS7700, and the merged TS7700 will be attached to the existing source Tape Library, the physical cartridges will not been moved. Use the procedure we describe in this section. Figure 6-8 on page 253 shows an overview of the process.

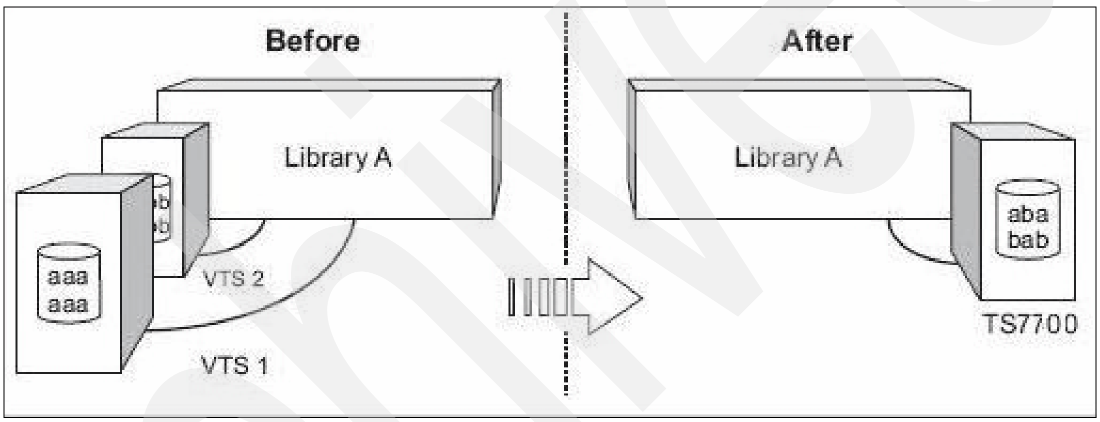

Figure 6-8 Two standalone VTS to a Single Cluster Grid TS7700 in the same Tape Library

These are the steps required to accomplish the migration scenario shown in Figure 6-8:

1. Logical volume ranges in a TS7700 must be unique. You should verify that the logical volume ranges for the two VTSs are unique before you start the merge procedure.
2. You should validate that there is no conflicting information between the TCDB, RMM and LM. This is an important step, because in this migration scenario the Library Names of volumes belonging to one of the two existing library names need to be changed. Only one of the existing library names can continue to exist after you merge two VTSs to one TS7700. An example of the job is shown in Example 6-7.

   Example 6-7 Verify conflicting information between RMM, LM and TCDB

   ```
   //EDGUTIL  EXEC PGM=EDGUTIL,PARM='VERIFY(ALL,VOLCAT)'
   //SYSPRINT DD SYSOUT=*
   //MASTER   DD DSN=your.database.name,DISP=SHR
   //VCINOUT  DD UNIT=3390,SPACE=(CYL,(900,500))
   ```

   The result of running such a job is that you will get information of all volumes with conflicting information that should be solved before the migration. For more information about this utility, refer to *z/OS V1R8.0 DFSMSrmm Implementation and Customization Guide,* SC26-7405. The job should be done days before the migration.

3. Stop all host activity.
   a. If there is another VTS or TS7700 system available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
   b. Complete or cancel all host jobs for the two VTSs.
   c. Vary off all device addresses associated with the library for all attached hosts.
   d. Vary the existing VTSs offline to all hosts.
   e. Vary the existing channels offline.
4. Prepare the software changes. Make sure that you apply the changes to all connected hosts. All the following steps could be done concurrent with the hardware changes that follow:
   a. Changes to SMS

      When you define a Single Cluster Grid you need to define a Composite Library as well as one Distributed Library. We recommend that you reuse the existing library name of one of the two existing VTSs as Composite Library name for the new TS7700. For the VTS that is removed you need to change all volume entries in the TCDB and in RMM to the remaining Composite Library name, before you can use the volumes. If the Composite Library name is not reused, you must change information for all volumes from both VTSs.

      Delete both of the old libraries and define the two new libraries in SMS through ISMF. Remember to write the new Library-IDs in the definitions. You must also remember to then relate all existing Storage Group definitions from the library name that is not reused to that remaining library name.

      Note: If you are reusing an existing library name, you only need to update its Library-ID in the library definition and do not need to delete it.

   b. Changes to HCD channel definitions

      If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or might switch from ESCON to FICON with this process—those are planning considerations that should be completed at this time. You might define new FICON channels to ease the process.
   c. Changes to LCU definitions

      Changes are needed because you change the Library-ID. If one or both VTSs are models with 64 logical units, you also need to define more LCUs, to enable the use of 256 logical units supported on a Single Cluster Grid. See Chapter 5, “Software implementation” on page 193 for more information.
   d. Device sharing

      If you define the devices as offline in HCD and use any product for device sharing, you need to reflect the new addresses to that product.
   e. Missing interrupt values

      If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.5, “Set values for the Missing Interrupt Handler” on page 208.

      Activate the IODF and SMS definitions and issue an OAM restart if not done after the SMS activation. For more information, see Chapter 5, “Software implementation” on page 193.

> **Note:** If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:
>
> `CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.`

5. Changes to volume entries in the TCDB

   The Composite Library name cannot remain the same for one of the old Library names. You must change all volume entries within the TCDB to relate to the new name by issuing a command to every single volume entry. You can obtain a list of all volumes in the TCDB by running a job as shown in Example 6-8.

   Example 6-8 JCL to list all entries in the TCDB

```
//STEP1   EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 LISTC VOLUMEENTRIES(V*) LIBRARY(vtsname)
```

   Update of all volume entries can be a time-consuming process. Because of the time used, you could change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible and after that take all the rest. See Example 6-9 for an example of the JCL for changing one scratch volume and one private volume.

   Example 6-9 JCL for changing the TCDB to a new TS7700

```
//*************************************************************
//**** Change TCDB for a scratch volume to a new TS7700      ****
//*************************************************************
//TCDBSCR EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLENTRY LIBRARYNAME(TSname) USEATTRIBUTE(SCRATCH)
//*************************************************************
//**** Change TCDB entry for a private volume to a new TS7700 ****
//**** Also change sgname to the one used (same as on the VTS)****
//*************************************************************
//TCDBPRIV EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLUMEENTRY LIBRARYNAME(TSname) -
       USEATTRIBUTE(PRIVATE) STORAGEGROUP(sgname)
```

   If you are using RMM, an RMM command for each volume is also needed. The command is shown in Example 6-10. The FORCE parameter is only used when needed and requires access to a specific Facility class in RACF named STGADMIN.EDG.FORCE. Verify that you have the needed authorization.

   Example 6-10 JCL for changing volumes in DFSMSrmm to a new TS7700

```
//PROCESS  EXEC PGM=IKJEFT01,DYNAMNBR=25,
//             TIME=100
//ISPLOG    DD DUMMY
//SYSPRINT  DD SYSOUT=*
//SYSTSPRT  DD SYSOUT=*
//SYSTSIN   DD *
 RMM CV volser LOCATION(TSname)
```

6. Drain the Tape Volume Cache of the first VTS and the second VTS partitions.

7. Extract the first VTS and the second VTS partitions database; see “Back up VTS data” on page 241 for details.
8. Disconnect the first VTS and the second VTS hardware from the Tape Library, either IBM 3953/TS3500 or IBM 3494, and the TS1120 or 3592 tape drives.
9. If the first VTS and the second VTS partitions connect to the existing 3953 or 3494 using 2 GB fiber switches, they must be replaced with the 4 GB switches at this time.
10. Install the Single Cluster Grid TS7700:
    a. Connect the TS7700 to the Tape Library, Library Manager.
    b. Connect the TS7700 to the TS1120 or 3592 drives.
    c. Teach the Library that now TS7700 is attached.
11. Restore the first VTS and the second VTS partitions database into the TS7700; see “Restore VTS data” on page 242 for details.
12. Vary TS7700 online from the TS7700 System Management Interface (SMIT) panel selecting **3957-V06 Online/Offline Menus** → **Vary 3957-V06 Online**.
13. Now you are ready to test the new Single Cluster Grid.
    a. Vary the defined channels online.
    b. Vary the logical devices online.
    c. Vary the Composite and Distributed Libraries online to the host.
    d. Run test jobs to read and write from the TS7700 Single Cluster Grid.

Note: The IBM SSR will perform Steps 7, 8, 10, 11, 12, and 13 as part of the installation. We list them for your information.

Two standalone VTSs to one TS7700 in a new Tape Library

If you are merging the data existing in two standalone B10/B20 VTSs into a Single Cluster Grid TS7700 Virtualization Engine, and the merged TS7700 will be attached to a new target Tape Library, and physical cartridges will be moved to the target library (*Move*), use the procedure we describe here. Figure 6-9 shows an overview for this migration.

When physical moving of the cartridges from one library to another occurs, the time required for this task might vary depending on the number of cartridges and the distance they are transported.

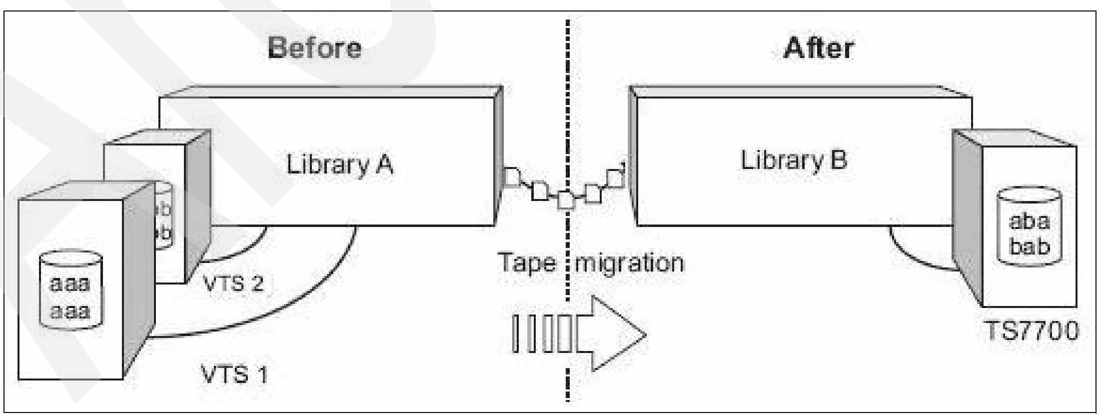

Figure 6-9 Two standalone VTSs to a Single Cluster Grid different Tape Library cartridge move

The TS7700 Virtualization Engine can be connected to the new Tape Library, Library Manager, and drives. This minimize the system outage window needed for the migration.

Here are the steps required to accomplish the migration scenario described above:

1. Back up the Library Manager administrative data including the library constructs from the source library. You can perform this activity prior to the outage window.
2. You should determine if there is any conflicting information between the TCDB, RMM, and the LM. This is an important step, because in this migration scenario volumes belonging to one of the two existing library names needs to be handled. Only one of the existing library names can continue to exist after you merge two VTSs to one TS7700. An example of the job is shown in Example 6-11.

   Example 6-11 Verify conflicting information between RMM, LM and TCDB

   ```
   //EDGUTIL  EXEC PGM=EDGUTIL,PARM='VERIFY(ALL,VOLCAT)'
   //SYSPRINT DD SYSOUT=*
   //MASTER   DD DSN=your.database.name,DISP=SHR
   //VCINOUT  DD UNIT=3390,SPACE=(CYL,(900,500))
   ```

   The result of running such a job is that you will get information of all volumes with conflicting information that should be solved before the migration. For more information about this utility, refer to *z/OS V1R8.0 DFSMSrmm Implementation and Customization Guide,* SC26-7405. The job should be done days before the migration.
3. Stop host activity
   a. If there is another VTS or TS7700 system available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
   b. Complete or cancel all host jobs for the two VTSs.
   c. Vary off all device addresses associated with the library for all attached hosts.
   d. Vary the existing VTSs offline to all hosts.
   e. Vary the existing channels offline.
4. Prepare the software changes. All the following steps could be done concurrent with all the hardware changes that follow.

   a. Changes to SMS

      When you define a Single Cluster Grid you need to define a Composite Library as well as one Distributed Library. We recommend that you reuse the existing library name of one of the two existing VTSs as Composite Library name for the new TS7700. For the VTS that is removed you need to change all volume entries in the TCDB and in RMM to the remaining Composite Library name, before you can use the volumes. If the Composite Library name is not reused, you must change information for all volumes from both VTSs.

      Delete both old libraries and define the two new libraries in SMS through ISMF. Remember to write the new Library-IDs in the definitions. You must also remember to then relate all existing Storage Group definitions from the library name that is not reused to that remaining library name. When you delete a library in the SCDS through ISMF, the change needs to be reflected in the TCDB as shown in Example 6-12.

      Example 6-12 Delete a Library in the TCDB

      ```
      //STEP1      EXEC   PGM=IDCAMS
      //SYSPRINT   DD     SYSOUT=*
      //SYSIN      DD     *
       DELETE (vtsname) -
       LIBRARYENTRY
      ```

      Note: If you are reusing one of the existing library names, you only need to update its Library-ID in the library definition and do not need to delete it.

b. Changes to HCD channel definitions

   If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or might switch from ESCON to FICON with this process—these are planning considerations that should be completed at this time. You might define new FICON channels to ease the process.

c. Changes to HCD LCU definitions

   Changes are needed because you change the Library-ID. If one or both VTSs models are with 64 logical units, you also need to define more LCUs, to enable the use of 256 logical units supported on a Single Cluster Grid. See Chapter 5, "Software implementation" on page 193 for more information.

d. If you define the devices as offline in HCD and use any product for device sharing, you need to reflect the new addresses to that product.

e. Missing interrupt values

   If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.5, "Set values for the Missing Interrupt Handler" on page 208.

f. Activate the IODF and SMS definitions and issue an OAM restart if not done after the SMS activation. For more information, see Chapter 5, "Software implementation" on page 193.

> **Note:** If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:
>
> ```
> CBR3006I Library library-name with Library ID library-ID unknown in I/O
> configuration.
> ```

5. Change volume entries in TCDB

   The Composite Library name cannot remain the same for one of the old Library names. You must change all volume entries within the TCDB to relate to the new name by issuing a command to every single volume entry. You can obtain a list of all volumes in the TCDB by running a job as shown in Example 6-13.

   Example 6-13 JCL to list all entries in the TCDB

```
//STEP1    EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 LISTC VOLUMEENTRIES(V*) LIBRARY(vtsname)
```

   Update of all volume entries can be a time-consuming process. Because of the time used, you could change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible and after that take all the rest. See Example 6-14 for an example of the JCL for changing one scratch volume and one private volume.

   Example 6-14 JCL for changing the TCDB to a new TS7700

```
//*****************************************************************
//**** Change TCDB for a scratch volume to a new TS7700       ****
//*****************************************************************
//TCDBSCR EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLENTRY LIBRARYNAME(TSname) USEATTRIBUTE(SCRATCH)
```

```
//*****************************************************************
//**** Change TCDB entry for a private volume to a new TS7700 ****
//**** Also change sgname to the one used (same as on the VTS)****
//*****************************************************************
//TCDBPRIV EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLUMEENTRY LIBRARYNAME(TSname) -
       USEATTRIBUTE(PRIVATE) STORAGEGROUP(sgname)
```

If you are using RMM, an RMM command for each volume is also needed. The command is shown in Example 6-15. The FORCE parameter is only used when needed and requires access to a specific Facility class in RACF named STGADMIN.EDG.FORCE. Verify that you have the needed authorization.

Example 6-15 JCL for changing volumes in DFSMSrmm to a new TS7700

```
//PROCESS  EXEC PGM=IKJEFT01,DYNAMNBR=25,
//             TIME=100
//ISPLOG    DD DUMMY
//SYSPRINT  DD SYSOUT=*
//SYSTSPRT  DD SYSOUT=*
//SYSTSIN   DD *
 RMM CV volser LOCATION(TSname)
```

6. Drain the Tape Volume Cache of the first VTS and the second VTS partitions.
7. Extract these first VTS and the second VTS partitions database; see “Back up VTS data” on page 241 for details.
8. Back up IBM 3494 or 3953 Library Manager data; this includes the volume serial range for the media types.
9. Remove the physical cartridges that belong to the first VTS and the second VTS partitions you are migrating to TS7700 from the source library.
10. Restore the first VTS and the second VTS partitions database into the TS7700; see “Restore VTS data” on page 242 for details.
11. Restore IBM 3494 or 3953 Library Manager data. This includes the volume serial range for the media types.
12. Insert the physical cartridges into the target library that belong to the first VTS and the second VTS partitions you are migrating to TS7700, previously removed from the source library.
13. Initialize the target Library Manager and perform an inventory of the library, which will allow the Library Manager to recognize the physically inserted cartridge.
14. Vary TS7700 online from the TS7700 System Management Interface (SMIT) panel, selecting **3957-V06 Online/Offline Menus** → **Vary 3957-V06 Online**.
15. Now you are ready to test the new Single Cluster Grid:
    a. Vary the defined channels online.
    b. Vary the logical devices online.
    c. Vary the Distributed and the Composite Library online to the hosts.
    d. Run test jobs to read and write from the TS7700 Single Cluster Grid.

Note: The IBM SSR will perform Steps 1, 6, 7, 8, 9, 10, 11, 13, and 14 as part of the installation. We list them for your information.

6.2.3 PtP VTS to Two-Cluster Grid

In this section we cover the migration scenario if you are migrating the data from two VTSs that are in a Peer-to-Peer configuration and restoring that data on two TS7700s that are in a grid configuration. We consider either that the new TS7700 Grid will be installed in the existing IBM 3494 or TS3500 Tape Libraries, in which case no physical cartridge movement will occur, or the TS7700 Grid will be installed in new target Tape Libraries, in which case physical cartridges will be moved from the source library to the target library. See Figure 6-11 on page 261 and Figure 6-12 on page 264

In this section we describe how the migrating data will be copied from one VTS system to the corresponding, empty TS7700 cluster in a grid, and the data from the other VTS system will then be copied to a second empty TS7700 cluster in the grid. Figure 6-10 illustrates the data migration process.

Note: Both installation teams must be on site at the same time and synchronized with the installation instruction steps.

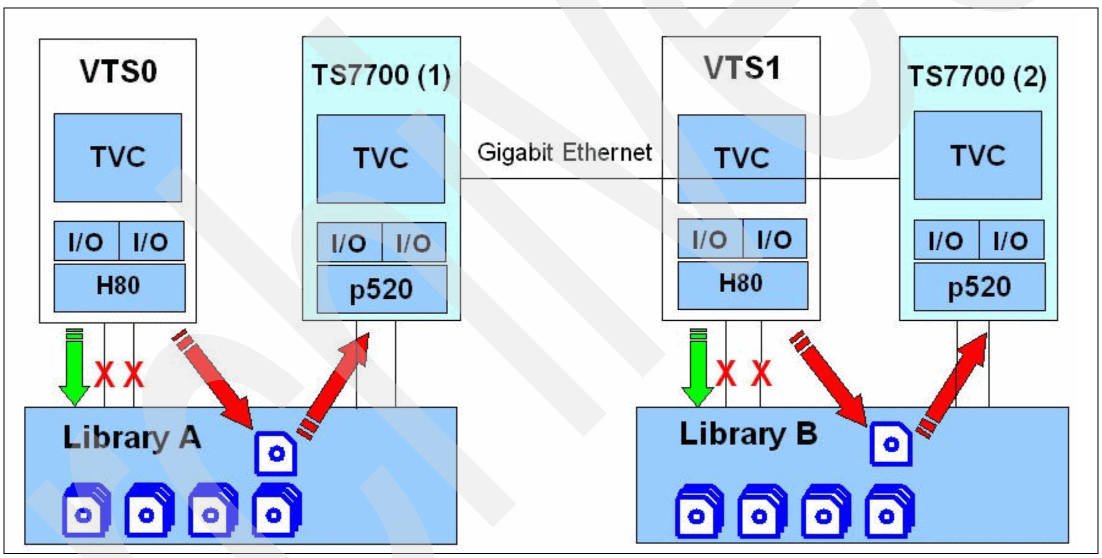

Figure 6-10 Data migration VTS PtP to TS7700 Two-Cluster Grid

Existing source Tape Libraries

The hardware migration from the Peer-to-Peer VTS B10 or B20 to the TS7700s in an IBM 3494 Tape Library or 3953 Library Controller connected to a TS3500 Tape Library, is performed in the steps we cover here. In this section we consider that the tape libraries remain the same and you reconnect the TS7700s in the existing tape libraries where the Peer-to-Peer VTSs were previously attached (see Figure 6-11).

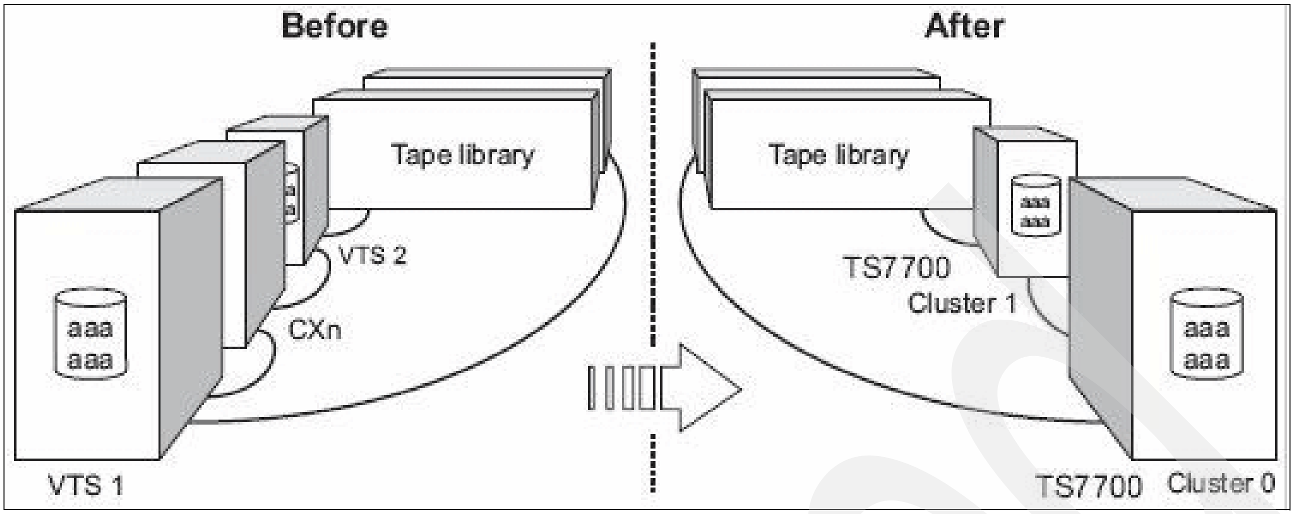

Figure 6-11 Migration VTS PtP to TS7700 Two-Cluster Grid same Tape Libraries

These are the steps required to accomplish the migration scenario described:

1. Start the installation of the two TS7700 hardware a few days prior to the outage window.
2. You should determine if there are is conflicting information between the TCDB, RMM and LM. This is only needed if you plan to change the library name as described in Step 4 on page 245. An example of the job is shown in Example 6-16.

   Example 6-16 Verify conflicting information between RMM, LM and TCDB

   ```
   //EDGUTIL  EXEC PGM=EDGUTIL,PARM='VERIFY(ALL,VOLCAT)'
   //SYSPRINT DD SYSOUT=*
   //MASTER   DD DSN=your.database.name,DISP=SHR
   //VCINOUT  DD UNIT=3390,SPACE=(CYL,(900,500))
   ```

   The result of running such a job is that you will get information of all volumes with conflicting information, which should be solved before the migration. For more information about this utility, refer to *z/OS V1R8.0 DFSMSrmm Implementation and Customization Guide,* SC26-7405. The job should be done days before the migration.

3. Stop host activity.
   a. If there is another VTS or TS7700 system available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
   b. Complete or cancel all host jobs for the PtP VTS.
   c. Vary off all device addresses associated with the library for all attached hosts.
   d. Vary the existing VTSs offline to all hosts.
   e. Vary the existing channels offline.
4. Prepare the software changes. The following steps could be done concurrent with all the hardware changes that follow.
   a. Changes to SMS.

      When you define a Two-Cluster Grid you need to define a Composite Library as well as two Distributed Libraries, the same as for PtP VTS. We highly recommend that you reuse the existing library name as Composite Library name. Then you only need to write the new Library-IDs in the existing library definitions. If you do not keep the existing library name, you must delete the old names and add the new names and then relate all existing Storage Group definitions to that new name. After that you must change all volume entries in the TCDB and in RMM to the new name before you can use the volumes.

b. Changes to HCD channel definitions.

   If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or might switch from ESCON to FICON with this process, These are planning considerations that should be completed at this time. You might define new FICON channels to ease the process.

c. Changes to LCU definitions.

   Changes are needed because you change the Library-ID of the Composite Library. If the PtP VTS model is with 64 or 128 logical units, you also need to define more LCUs to enable the use of 256 logical units supported on a Single Cluster Grid. See Chapter 5, “Software implementation” on page 193 for more information.

d. If you define the devices as offline in HCD and use any product for device sharing, you need to reflect the new addresses to that product.

e. Missing interrupt values.

   If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.5, “Set values for the Missing Interrupt Handler” on page 208.

f. Activate the IODF and SMS definitions and issue an OAM restart if not done after the SMS activation. For more information, see Chapter 5, “Software implementation” on page 193.

> **Note:** If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:
>
> `CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.`

5. Change volume entries in the TCDB.
   If the Composite Library is the same as the previous name, there are no changes and you can continue with Step 6 on page 249. If you change the name, you should change all volumes within the TCDB to relate to the new name by issuing a command to every single volume. You can get a list of all volumes in the TCDB by running a job as in Example 6-17.

   Example 6-17 JCL to list all entries in the TCDB

```
//STEP1   EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 LISTC VOLUMEENTRIES(V*) LIBRARY(vtsname)
```

   Update of all volume entries can be a time-consuming process. Because of the time used, you could change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible and after that take all the rest. See Example 6-18 for an example of the JCL for changing one scratch volume and one private volume.

   Example 6-18 JCL for changing the TCDB to a new TS7700

```
//*****************************************************************
//**** Change TCDB for a scratch volume to a new TS7700       ****
//*****************************************************************
//TCDBSCR EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLENTRY LIBRARYNAME(TSname) USEATTRIBUTE(SCRATCH)
//*****************************************************************
```

```
//**** Change TCDB entry for a private volume to a new TS7700 ****
//**** Also change sgname to the one used (same as on the VTS)****
//***************************************************************
//TCDBPRIV EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLUMEENTRY LIBRARYNAME(TSname) -
       USEATTRIBUTE(PRIVATE) STORAGEGROUP(sgname)
```

If you are using RMM, an RMM command for each volume is also needed. The command is shown in Example 6-19. The FORCE parameter is only used when needed and requires access to a specific Facility class in RACF named STGADMIN.EDG.FORCE. Verify that you have the needed authorization.

Example 6-19 JCL for changing volumes in DFSMSrmm to a new TS7700

```
//PROCESS  EXEC PGM=IKJEFT01,DYNAMNBR=25,
//             TIME=100
//ISPLOG    DD DUMMY
//SYSPRINT  DD SYSOUT=*
//SYSTSPRT  DD SYSOUT=*
//SYSTSIN   DD *
 RMM CV volser LOCATION(TSname)
```

6. Place the non-master PtP VTS in service preparation The team on the other site master VTS does not need to perform this function, it must wait for the non-master VTS to be in service. When the non-master VTS is in service, both teams start Step 7 synchronized with each other.
7. Drain the PtP VTSs Tape Volume Cache; this is done by both teams on the PtP sites.
8. Extract the PtP VTSs database; see “Back up VTS data” on page 241 for details. This step is done by both teams on the PtP sites.
9. Disconnect the VTSs from their Tape Library, either IBM 3953/TS3500 or IBM 3494, and TS1120 or 3592 tape drives; both teams perform this step.
10. If the PtP VTSs connect to the existing 3953 or 3494 using 2 GB fiber switches, they must be replaced with the 4 GB switches at this time; both teams must verify this.
11. Complete the installation of theTS7700s on both sites by both installation teams.
    a. Connect the TS7700s to the Tape Libraries and Library Managers.
    b. Connect the TS7700s to the TS1120s or 3592 drives.
    c. Teach the libraries that now TS7700s are attached.
12. Restore the PtP VTSs database into the TS7700s; see “Restore VTS data” on page 242 for details. Both installation teams perform this procedure on the PtP sites.
13. Vary TS7700s online from the TS7700s System Management Interface (SMIT) panel, selecting **3957-V06 Online/Offline Menus** → **Vary 3957-V06 Online**.
14. Both installation teams establish the WAN connection between the TS7700s to form the Two-Cluster Grid configuration.
15. Now you are ready to test the new Two-Cluster Grid.
    a. Vary the defined channels online.
    b. Vary the logical devices online.
    c. Vary the Distributed and Composite Libraries online to the hosts.
    d. Run test jobs to read and write from the Two-Cluster Grid TS7700.

Note: The IBM SSR will perform Step 1, and Steps 6 to 14 as part of the installation. We list them for your information.

New target tape libraries

In this section we describe the hardware migration from the Peer-to-Peer VTS to TS7700 Grids, and at the same time the movement of the physical cartridges from the existing library to the new source TS3500 libraries. The old libraries can be IBM 3494 Tape Libraries or 3953 Library Controllers connected to TS3500 Tape Libraries. The process is described in Figure 6-12 as well as in the following steps.

The time required for physically moving cartridges from one library to another can vary depending on the number of cartridges and the distance they are transported.

TS7700 Virtualization Engine should be connected to the new Tape Libraries, Library Managers, and drives, tested, and the libraries must be taught. This minimizes the system outage window needed for the migration.

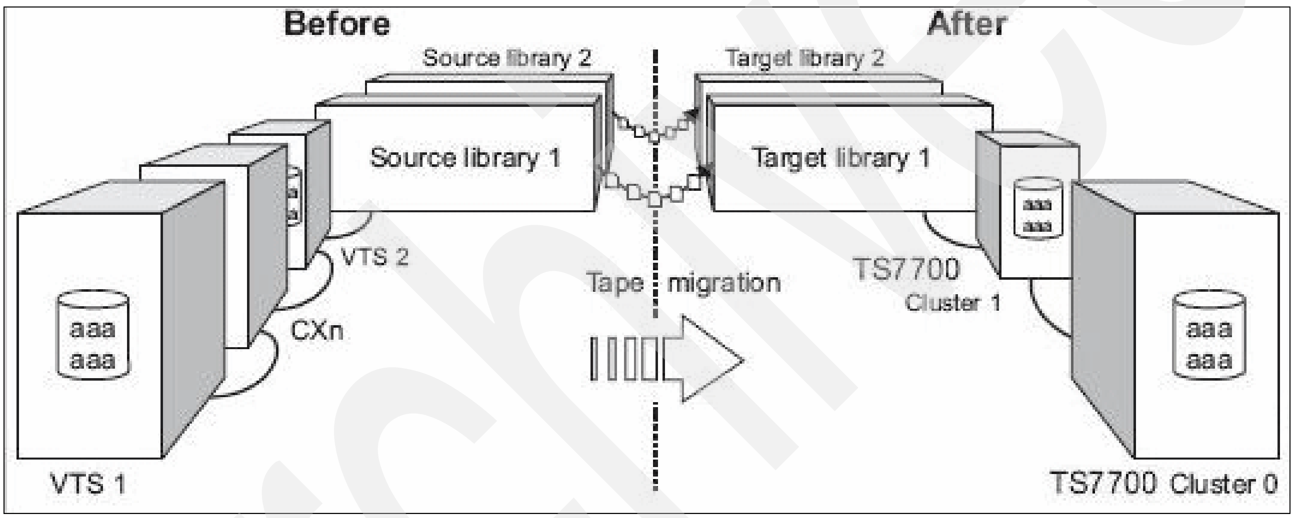

Figure 6-12 Migration of VTS PtP to TS7700 Two-Cluster Grid, different Tape Libraries

These are the steps required to accomplish the migration scenario described:

1. Back up the PtP Library Managers Administrative Data regarding the library constructs from the PtP source libraries. You can perform this activity prior to the outage window.
2. You should determine if there is any conflicting information between the TCDB, RMM and LM. This is only needed if you plan to change the library name as described in Step 4 on page 245. An example of the job is shown in Example 6-20.

   Example 6-20 Verify conflicting information between RMM, LM and TCDB

   ```
   //EDGUTIL  EXEC PGM=EDGUTIL,PARM='VERIFY(ALL,VOLCAT)'
   //SYSPRINT DD SYSOUT=*
   //MASTER   DD DSN=your.database.name,DISP=SHR
   //VCINOUT  DD UNIT=3390,SPACE=(CYL,(900,500))
   ```

   The result of running such a job is that you will get information of all volumes with conflicting information, which should be solved before the migration. For more information about this utility, refer to *z/OS V1R8.0 DFSMSrmm Implementation and Customization Guide,* SC26-7405. The job should be done days before the migration.

3. Stop host activity.
   a. If there is another VTS or TS7700 system available to the host during the migration, you can change the ACS constructs to direct allocation to that system.
   b. Complete or cancel all host jobs for the PtP VTS.
   c. Vary off all device addresses associated with the library for all attached hosts.
   d. Vary the existing VTSs offline to all hosts.
   e. Vary the existing channels offline.
4. Prepare the software changes. The following steps could be done concurrent with the hardware changes that follow.
   a. Changes to SMS.
      When you define a Two-Cluster Grid you need to define a Composite Library as well as two Distributed Libraries, the same as for PtP VTS. We highly recommend that you reuse the existing library name as Composite Library name. Then you only need to write the new Library-IDs in the existing library definitions.
      If you do not keep existing library names, you must delete the old names and add the new names and then relate all existing Storage Group definitions to those new names. After that you must change all volume entries in the TCDB and in RMM to the new name before you can use the volumes.
   b. Changes to HCD channel definitions.

      If you plan to reuse existing host channels and FICON adaptor connections, or plan to define new channels, or might switch from ESCON to FICON with this process—these are planning considerations that should be completed at this time. You might define new FICON channels to ease the process.
   c. Changes to LCU definitions.

      Changes are needed because you change the Library-ID of the Composite Library. If the PtP VTS is a model with 64 or 128 logical units, you also need to define more LCUs to enable the use of 256 logical units supported on a Single Cluster Grid. See Chapter 5, “Software implementation” on page 193 for more information.
   d. If you define the devices as offline in HCD and use any product for device sharing, you need to reflect the new addresses to that product.
   e. Missing interrupt values.

      If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.5, “Set values for the Missing Interrupt Handler” on page 208.
   f. Activate the IODF and the SMS definitions and issue an OAM restart if not done after the SMS activation. For more information, see Chapter 5, “Software implementation” on page 193.

      > **Note:** If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:
      >
      > `CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.`

   g. Change volume entries in TCDB.
      If the Composite Library is the same as the previous name, there are no changes and you can continue with Step 6 on page 249. If you change the name, you should change all volumes within the TCDB to relate to the new name by issuing a command to every single volume. You can get a list of all volumes in the TCDB by running a job as in Example 6-21.

Example 6-21 JCL to list all entries in the TCDB

```
//STEP1   EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 LISTC VOLUMEENTRIES(V*) LIBRARY(vtsname)
```

Update of all volume entries can be a time-consuming process. Because of the time used, you could change all existing scratch tapes first to prepare for non-specific mounts (write) from production as fast as possible and after that take all the rest. See Example 6-22 for an example of the JCL for changing one scratch volume and one private volume.

Example 6-22 JCL for changing the TCDB to a new TS7700

```
//****************************************************************
//**** Change TCDB for a scratch volume to a new TS7700       ****
//****************************************************************
//TCDBSCR EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLENTRY LIBRARYNAME(TSname) USEATTRIBUTE(SCRATCH)
//****************************************************************
//**** Change TCDB entry for a private volume to a new TS7700 ****
//**** Also change sgname to the one used (same as on the VTS)****
//****************************************************************
//TCDBPRIV EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD *
 ALTER Vvolser VOLUMEENTRY LIBRARYNAME(TSname) -
       USEATTRIBUTE(PRIVATE) STORAGEGROUP(sgname)
```

If you are using RMM, an RMM command for each volume is also needed. The command is shown in Example 6-23. The FORCE parameter is only used when needed and requires access to a specific Facility class in RACF named STGADMIN.EDG.FORCE. Verify that you have the needed authorization.

Example 6-23 JCL for changing volumes in DFSMSrmm to a new TS7700

```
//PROCESS  EXEC PGM=IKJEFT01,DYNAMNBR=25,
//             TIME=100
//ISPLOG    DD DUMMY
//SYSPRINT  DD SYSOUT=*
//SYSTSPRT  DD SYSOUT=*
//SYSTSIN   DD *
 RMM CV volser LOCATION(TSname)
```

5. Place the non-master PtP VTS in service preparation. The team on the other site at the master VTS does not need to perform this function, but must wait for the non-master VTS to be in service. When the non-master VTS is in service, both teams start with Step 7, synchronized with each other.
6. Drain the PtP VTSs Tape Volume Cache; this is done by both teams on the PtP sites.
7. Extract the PtP VTSs database; see “Back up VTS data” on page 241 for details. This step is done by both teams on the PtP sites.
8. Back up IBM 3494 or 3953 Library Managers data for both libraries of the PtP. This includes the volume serial range for the media types.

9. Remove the physical cartridges that belong to the PtP VTS you are migrating to TS7700s from the source libraries of the PtP VTS.
10. Restore the PtP VTSs database into the TS7700s; see “Restore VTS data” on page 242 for details. Both installation teams perform this procedure on the PtP sites.
11. Restore IBM 3494 or 3953 Library Managers data for both libraries of the PtP. This includes the volume serial range for the media types.
12. Insert the physical cartridges that belong to the PtP VTS you are migrating to TS7700s removed from the source libraries, into the new target libraries.
13. Initialize both Library Managers and perform an inventory of the libraries, which will allow the Library Managers to now recognize the physically inserted cartridges.
14. Vary TS7700s online from the TS7700 System Management Interface (SMIT) panel, selecting **3957-V06 Online/Offline Menus** → **Vary 3957-V06 Online**.
15. Both installation teams establish the WAN connection between the TS7700s to form the Two-Cluster Grid configuration.
16. Now you are ready to test the new Two-Cluster Grid.
    a. Vary the defined channels online.
    b. Vary the logical devices online.
    c. Vary the Distributed and Composite Libraries online to the hosts.
    d. Run test jobs to read and write from the Two-Cluster Grid TS7700.

> **Note:** The IBM SSR will perform Step 1, Steps 6 to 9, 11, 12, and 14 to 16 as part of the installation. We list them for your information.

6.2.4 Three-Cluster Grid

In this section, we cover the migration aspects needed to join an existing Two-Cluster Grid TS7700 with a Single Cluster Grid TS7700 to form a Three-Cluster Grid.

Planning for a Three-Cluster Grid

The TS7700 Grid configuration can be used to facilitate disaster recovery and high data availability. The TS7700 Grid configuration provides support for automatic replication of data and can be used in a variety of disaster recovery situations. If connected at the same location, a grid configuration might also help facilitate high availability of data.

To configure a Three-Cluster Grid for disaster recovery, you must plan for the following:

- Access from your local site’s hosts to the FICON channels on the TS7700 cluster located at the disaster recovery site(s). This might involve connections using Dense Wavelength-Division Multiplexing (DWDM) or channel extension equipment, depending on the distance separating the sites. If the local TS7700 cluster becomes unavailable, you would use this remote access to continue your operations using a remote TS7700 cluster.
- Because the virtual devices on a remote TS7700 cluster are connected to the host through channel extensions, there might be a difference in read or write performance as compared to the virtual devices on the local TS7700 cluster. If performance differences are a concern, you should consider only using the virtual device addresses in a remote TS7700 cluster when the local TS7700 Virtualization Engine is unavailable. If these

differences are an important consideration, in addition to the ownership takeover procedure, you would need to provide for operator procedures to vary the virtual devices in a remote TS7700 Virtualization Engine from online to offline.

- You might want to maintain separate copy consistency policies for disaster recovery data and data that requires high availability.

Note: To achieve Existing TS7700 cluster or TS7700 Grid upgrades, you must order FC 4015, Grid enablement against the TS7700 Server.

The Three-Cluster Grid is built upon an existing grid. It extends the configuration to three clusters combined into a single composite library. See Figure 6-13 for reference.

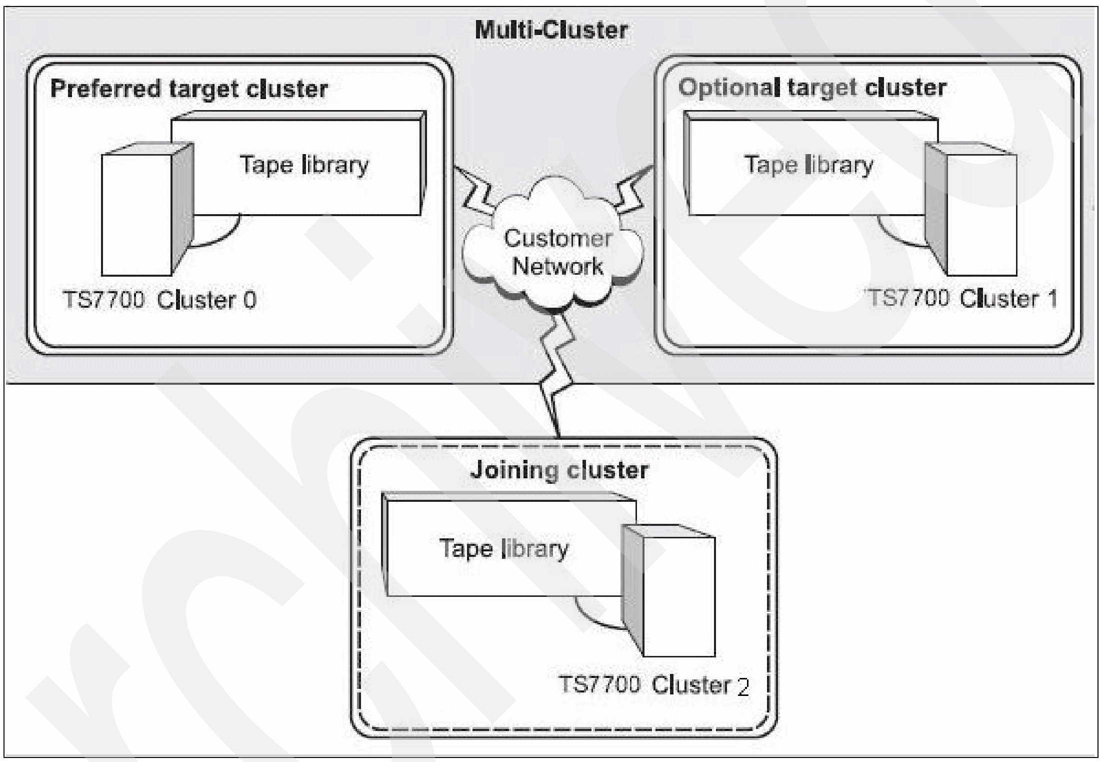

Figure 6-13 TS7700 Three-Cluster Grid configuration

Here we show the required steps for joining a TS7700 and a Two-Cluster in a Three-Cluster Grid configuration. First verify that the items in the following sections are pursued.

Identify the configuration

To identify the configuration, you need to fill out the pre-installation checklist that includes the following items:

- The IP addresses of the LAN configuration of Cluster 0, Cluster 1, and the Cluster that will join the existing clusters to form the Three-Cluster Grid.
- The clusters' identifiers (IDs). Check that they are unique.
- The Library Partition Composite identifiers.

Validate the prerequisites

To validate the prerequisites for the Three-Cluster Grid configuration, you need to verify the following items:

- Verify that the TS7700s microcode is 8.3.x.x or higher and that the associated Library Manager is at microcode level 535.xx or higher.
- Verify that no duplicated logical volumes exist between the three clusters' sites.
- Verify that FC4015 is available for the TS7700s clusters.
- Verify that the TS7700 that will be joining the existing grid has *no data*.
- Verify that each cluster in the final Multi Cluster configuration has a unique cluster identifier (ID).
- Verify that the customer network and Internet protocol addresses are available.

Three-Cluster Grid configurations

You can configure a TS7700 Grid to provide for both disaster recovery and high availability solutions.

The assumption is that two or three TS7700 clusters will reside in different locations, separated by a distance dictated by your company's requirements for disaster recovery. In a Three-Cluster Grid configuration, disaster recovery and high availability can also be achieved simultaneously by ensuring that two local, high availability clusters possess RUN volume copies and have shared access to the host, while the third and remote cluster possesses deferred volume copies for disaster recovery. During a single cluster outage, the Three-Cluster Grid solution maintains no single points of failure, which would prevent customers from accessing their data, assuming that copies exist on other clusters as defined in the copy consistency point. Refer to 4.6.1, "Define Grid copy mode control" on page 183 for more detail.

We consider these two configuration scenarios:

- Configuration 1:
  - Two of the three clusters will be contained within the same building or campus, as a High Availability pair, and both clusters will be connected to one or more production hosts. We consider that the Rewind Unload (RUN) consistency point has been configured between the Two-Cluster Grid.
  - The third cluster is remote outside the immediate region, and it will have backup host connectivity and devices varied offline. We consider that the third site that will join the Two-Cluster Grid to form the Three-Cluster Grid has copy consistency policy configured as deferred.
- Configuration 2:
  - Two independent production sites, where both clusters are connected to independent production host configurations. No Copy Consistency point has been configured between the production sites, and no copies done between the two clusters.
  - The third cluster is remote outside both regions of the production clusters, and it will have backup host connectivity with the devices varied offline. Copy consistency has been configured as deferred to the third site from both production clusters.

Implementing Configuration 1

In this section we show you the scenario from Two-Cluster Grid TS7700 to Three-Cluster Grid. As reference for this scenario, see Figure 6-14, which shows an overview of the configuration we are taking as example.

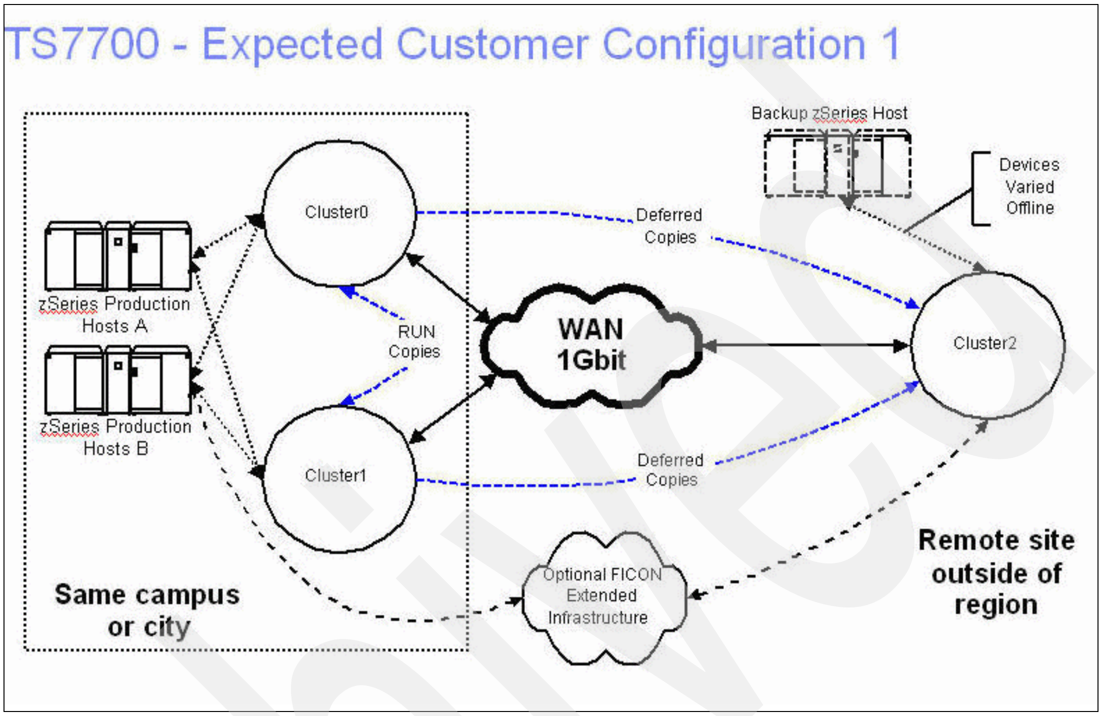

Figure 6-14 TS7700 Three-Cluster Grid Configuration 1

In the configuration shown in Figure 6-14, two clusters are in the same campus location or in the same city. The clusters can have one of these Copy Consistency points specified:

- Rewind Unload (RUN) Copy Consistency point

  If a data consistency point of RUN is specified, the data created on one TS7700 cluster is copied to the other TS7700 cluster as part of successful Rewind/Unload command processing, meaning that for completed jobs, a copy of the volume will exist on both TS7700 clusters. Access to data written by completed jobs (successful Rewind/Unload) prior to the failure is maintained through the other TS7700 cluster. Access to data of incomplete jobs that were in process at the time of the failure is not provided. We assume that the two existing clusters are using RUN Copy Consistency points specified in the Management Class storage construct, when volumes are written.

- Deferred Copy Consistency point

  If a data consistency point of Deferred is specified, the data created on one TS7700 cluster is copied to the other TS7700 cluster after successful Rewind/Unload command processing. Access to the data through the other TS7700 cluster is dependent on when the copy completes. Because there will be some delay in performing the copy, access might or might not be available when a failure occurs. This is the copy consistency policies we assume are assigned to the third cluster that joins the Two-Cluster Grid to form the Three-Cluster Grid configuration.

- No Copy Consistency point

  If a data consistency point of No Copy is specified, the data created on one TS7700 cluster is not copied to the other TS7700 cluster. If the TS7700 cluster that data was written to fails, the data for that logical volume is inaccessible until that TS7700 cluster's operation is restored. In this configuration, none of the clusters are set up for No Copy.

The new cluster that will join the Two-Cluster Grid to form the Three-Cluster Grid must already be installed. Every cluster in the system requires two network connections to a customer WAN for site-site operations, and the WAN connections between the three clusters in the Three-Cluster Grid must be completed. The grid network on the new cluster must be configured, containing the IP addresses of the three clusters.

The following are tasks done by the customer as well as the IBM System Service Representative (SSR).

1. Stop host activity on Cluster 0, which must go into Service Preparation mode.
   a. Vary the unit addresses offline.
   b. Complete or cancel all jobs for Cluster 0.
2. Set Cluster 0 in Service Preparation mode.
3. Extract the Library Manager database from Cluster 0 to be merged to the Library Manager on Cluster 2.
4. Changes to SMS.

   With a Three-Cluster Grid you need one Composite Library and three Distributed Libraries. You must now define the third Distributed Library in SMS. Make sure to enter the correct Library-ID delivered by the SSR.
5. Changes to HCD

   Define the new channels and the 256 units in HCD. We recommend that you define channels and units with OFFLINE=YES. Then the channels and units must be varied online manually, when and if they are to be used. See Chapter 5, “Software implementation” on page 193 for more information about HCD.
6. Missing Interrupt Handler values

   If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.5, “Set values for the Missing Interrupt Handler” on page 208.
7. Activate the IODF and the SMS definitions and issue an OAM restart if not done after the SMS activation. For more information, see Chapter 5, “Software implementation” on page 193. Cluster 2 is not ready yet, and will go online in a later step.

   > **Note:** If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:
   >
   > `CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.`

8. Configure the local grid network from the TS7700 System Management Interface (SMIT) panel.
9. Join Cluster 0 to Cluster 2 from SMIT:
   a. Merge Vital Product Data (VPD).
   b. Merge TS7700 databases data.
   c. Merge Library databases data.
10. Exit Cluster 0 from Service Preparation mode.
11. Vary devices from Cluster 0 online to all connected hosts.

12. Stop host activity on Cluster 1, which must go into Service Preparation mode.
    a. Vary the unit addresses offline.
    b. Complete or cancel all jobs for Cluster 1.
13. Set Cluster 1 in Service Preparation mode.
14. Extract the Cluster 1 Library Manager database to be merged to the third Cluster Library Manager.
15. Configure the local grid network from the TS7700 SMIT panel.
16. Join Cluster 1 to Cluster 2 from SMIT:
    a. Merge Cluster Vital Product Data (VPD).
    b. Merge TS7700 databases data.
    c. Merge Library databases data.
17. Exit Cluster 1 from Service Preparation mode.
18. Vary devices from Cluster 1 online to all connected hosts.
19. Now you are ready to validate the new Three-Cluster Grid.
    a. Vary the defined channels for Cluster 2 online to evaluate if FICON access to Cluster 2 is functioning.
    b. Vary logical devices for Cluster 2 online.
    c. Vary Cluster 2 online to the hosts.
    d. With the D SMS,LIB(libraryname),DETAIL commands, validate that the relation between Composite and Distributed Libraries is correct, as shown in Example 6-24 and Example 6-25.

Example 6-24 Display Composite Library

```
D SMS,LIB(COMPLIB),DETAIL
F OAM,D,LIB,COMPLIB,L=ST6T10-Z
CBR1110I OAM LIBRARY STATUS: 141
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY SCRTCH  ON OP
LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS   VOLS
COMPLIB   VCL  3957-V06  768  768  287      0      0 368298  Y  Y
----------------------------------------------------------------------
MEDIA       SCRATCH       SCRATCH       SCRATCH
TYPE          COUNT     THRESHOLD      CATEGORY
MEDIA1       170345             0          0001
MEDIA2       197953             0          0002
----------------------------------------------------------------------
DISTRIBUTED LIBRARIES:  DISTLIB0    DISTLIB1    DISTLIB2
----------------------------------------------------------------------
LIBRARY ID:  C0001
OPERATIONAL STATE:  AUTOMATED
ERROR CATEGORY SCRATCH COUNT:             33
CORRUPTED TOKEN VOLUME COUNT:              0
----------------------------------------------------------------------
LIBRARY SUPPORTS IMPORT/EXPORT.
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.
```

Example 6-25 Display Distributed Library

```
D SMS,LIB(DISTLIB1),DETAIL
F OAM,D,LIB,DISTLIB1,L=ST6T10-Z
CBR1110I OAM LIBRARY STATUS: 062
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY SCRTCH  ON OP
```

```
LIBRARY    TYP  TYPE       DRV  DRV  DRV  SLOTS  SLOTS   VOLS
DISTLIB1   VDL 3957-V06     0    0    0   1348    819      0  Y  Y
----------------------------------------------------------------------
COMPOSITE LIBRARY:      COMPLIB
----------------------------------------------------------------------
LIBRARY ID:  10001
OPERATIONAL STATE:  AUTOMATED
SCRATCH STACKED VOLUME COUNT:            222
PRIVATE STACKED VOLUME COUNT:            108
----------------------------------------------------------------------
LIBRARY SUPPORTS IMPORT/EXPORT.
LIBRARY SUPPORTS OUTBOARD POLICY MANAGEMENT.
CONVENIENCE I/O STATION INSTALLED.
CONVENIENCE I/O STATION IN OUTPUT MODE.
BULK INPUT/OUTPUT NOT CONFIGURED.
```

   e. Vary logical devices for Cluster 2 offline again, to be ready to test if original Two-Cluster Grid still works.

20. Run test jobs to read and write from the original Two-Cluster Grid.

21. Modify Copy Policies defined in the LM Management Class.

    The Copy Consistency points on all three clusters must be modified to support RUN copy between Cluster 0 and Cluster 1 as well as deferred copy from Cluster 0 and Cluster 1 to Cluster 2. The values must be updated in the Management Classes (from the Library Manager Interface). Make sure that the definitions will work when logical units are allocated from Cluster 2. Refer to 4.6.1, “Define Grid copy mode control” on page 183 for more information. Table 6-5 describes the setting needed for the scenario shown in Figure 6-14 on page 270.

Table 6-4 Copy Consistency point on Management Class -Three-Cluster Grid configuration one

Management Class	**Cluster 0**	**Cluster 1**	**Cluster2**
	From → To	**From → To**	**New Cluster**
MC Hosts A	RR → RRD	RR → RRD	DDD
MC Hosts B	RR → RRD	RR → RRD	DDD

22. Vary logical devices for Cluster 0 and Cluster 1 offline to be ready to validate the use of Cluster 2 as though there were a disaster, and set up Copy Consistency points that support the customer requirements; expected to be deferred copy mode?

23. Test write and thereafter read with Cluster 2 and validate the result.

24. Migration done. Return to normal production mode again.

25. It you want part or all of the existing logical volumes to have a logical copy on Cluster 2, this can be done in different ways. IBM has tools such as PRESTAGE to support you. The logical volumes must be read or referred in any way to redrive the new management policies that you define. The tools are available on the FTP site

    ftp://ftp.software.ibm.com/storage/tapetool/

> **Note:** The IBM SSR will perform Steps 2, 3, from 8 to 10, and from 13 to 17 as part of the installation. We list them for your information.

Implementing Configuration 2

In this configuration, two clusters are in independent campus, location, or city. We assume that the two existing clusters are running with Deferred Copy Consistency points. See Figure 6-15, which offers an overview of the configuration.

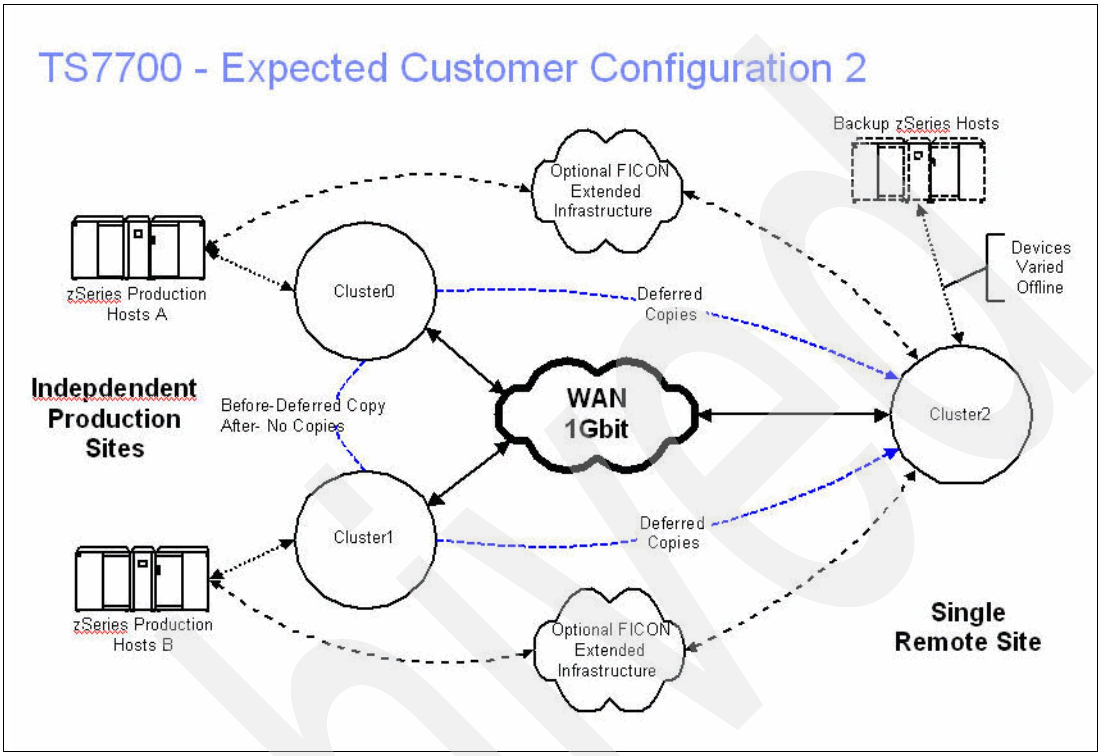

Figure 6-15 TS7700 Three-Cluster Grid Configuration 2

The new cluster that will join the Two-Cluster Grid to form the Three-Cluster Grid must be installed in advance. Every cluster in the system requires two network connections to a customer WAN for site-site operations, and the WAN connections between the three clusters in the Three-Cluster Grid must be completed. The grid network on the new cluster must be configured, containing the IP addresses of the three clusters.

The following are tasks done by the customer as well as the IBM System Service Representative (SSR):

1. Stop host activity on Cluster 0, which must go in Service Preparation mode. Note that the connected hosts are without access to tape from this moment.
   a. Vary the unit addresses offline.
   b. Complete or cancel all jobs for Cluster 0.
2. Set Cluster 0 in Service Preparation mode.
3. Extract the Library Manager database from Cluster 0 to be merged to the Library Manager on Cluster 2.
4. Changes to SMS.

   With a Three-Cluster Grid you need one Composite Library and three Distributed Libraries. You must now define the third Distributed Library in SMS. Make sure to enter the correct Library-ID delivered by the SSR. If the connected hosts are in separate SMS configurations, this step should be done twice.

5. Changes to HCD.

   Define the new channels and the 256 units in HCD. We recommend that you define channels and units with OFFLINE=YES. Then the channels and units must be varied online manually, when and if they are to be used. See Chapter 5, “Software implementation” on page 193 for more information about HCD.

6. Missing Interrupt Handler values.

   If you are defining specific address ranges in the IECIOSxx member in SYS1.PARMLIB, make sure that MIH values for the new devices are set. Proper values are described in 5.2.5, “Set values for the Missing Interrupt Handler” on page 208. If the connected hosts used different SYS1.PARMLIBs, this step must be repeated.

7. Activate the IODF and the SMS definitions and issue an OAM restart if not done after the SMS activation. For more information, see Chapter 5, “Software implementation” on page 193. Cluster 2 is not ready yet, and will go online in a later step. Again this step must be repeated to reach all hosts.

   > **Note** If the new SCDS is activated before the new library is ready, the host cannot communicate with the new library yet. Expect message CBR3006I to be issued:
   >
   > `CBR3006I Library library-name with Library ID library-ID unknown in I/O configuration.`

8. Configure the local grid network from the TS7700 System Management Interface (SMIT) panel.

9. Join Cluster 0 to Cluster 2 from SMIT:

   a. Merge Vital Product Data (VPD).
   b. Merge TS7700 databases data.
   c. Merge Library databases data.

10. Exit Cluster 0 from Service Preparation mode.

11. Vary devices from Cluster 0 online to all connected hosts.

12. Run test jobs to read and write from the original Cluster 0.

13. Stop host activity on Cluster 1, which must go into Service Preparation mode. Note that the connected hosts are without access to tape from this moment.

   a. Vary the unit addresses offline.
   b. Complete or cancel all jobs for Cluster 1.

14. Set Cluster 1 in Service Preparation mode.

15. Extract the Cluster 1 Library Manager database to be merged to the third Cluster Library Manager.

16. Configure the local grid network from the TS7700 SMIT panel.

17. Join Cluster 1 to Cluster 2 from SMIT:

   a. Merge Cluster Vital Product Data (VPD).
   b. Merge TS7700 databases data.
   c. Merge Library databases data.

18. Exit Cluster 1 from Service Preparation mode.

19. Vary devices from Cluster 1 online to all connected hosts.

20. Run test jobs to read and write from the original Cluster 1.

Note: The following steps apply to all connected hosts.

21. Now you are ready to validate the new Three-Cluster Grid.
    a. Vary the defined channels for Cluster 2 online to evaluate if FICON access to Cluster 2 is functioning.
    b. Vary logical devices for Cluster 2 online.
    c. Vary Cluster 2 online to the hosts.
    d. Evaluate with D SMS,LIB(libraryname),DETAIL commands that the relation between Composite and Distributed Libraries is correct, as shown in Example 6-24 on page 272 and Example 6-25 on page 272.
    e. Vary logical devices for Cluster 2 offline again, to be ready to test if the original Two-Cluster Grid still works.
22. Modify Copy Policies defined in the LM Management Class.

    The Copy Consistency points on all three Clusters must be modified to support deferred copy from Cluster 0 and 1 to Cluster 2. The values must be updated in the Management Classes (from the Library Manager Interface). Make sure that the definitions will work when logical units are allocated from Cluster 2. Refer to 4.6.1, “Define Grid copy mode control” on page 183 for more information. Table 6-5 describes the setting needed for the scenario shown in Figure 6-15 on page 274.

Table 6-5 Copy Consistency point on Management Class -Three-Cluster Grid Configuration 2

Management Class	Cluster 0	Cluster 1	Cluster2
	From → To	**From → To**	**New Cluster**
MC Hosts A	RD → RND	RD → RND	DND
MC Hosts B	RD → NRD	RD → NRD	NDD

23. Vary logical devices for Cluster 0 and Cluster 1 offline to be ready to validate the use of Cluster 2 as though there were a disaster, and set up Copy Consistency points that support the customer requirements; expected to be deferred copy mode?
24. Test write and thereafter read with Cluster 2 and validate the result.
25. Migration done. Return to normal production mode again.
26. It you want part or all of the existing logical volumes to have a logical copy on Cluster 2, this can be done in different ways. IBM has tools such as PRESTAGE to support you. The logical volumes must be read or referred in any way to redrive the new management policies that you define. The tools are available on the FTP site

    ftp://ftp.software.ibm.com/storage/tapetool/

Note: The IBM SSR will perform Steps 2, 3, from 8 to 10, and from 14 to 18 as part of the installation. We list them for your information.

6.3 Moving data in and out of the TS7700

To move data into the TS7700, it is not possible to simply move the cartridges out of an IBM 3494-Bxx VTS and insert them into a Single Cluster TS7700 and copy the control data sets. This migration approach is only supported for the scenarios described in 6.2, “TS7700 hardware migration scenarios” on page 243. In all other scenarios, migrating data into the

TS7700 requires that the TS7700 and the existing environment remain installed in parallel until the data has been migrated.

An existing VTS cannot be combined with a TS7700 to build a Peer-to-Peer (PtP) VTS or a Multi Cluster Grid configuration. To migrate your workload from an existing PtP VTS into an existing production TS7700 Grid, you must copy your data tape by tape from the PtP VTS into the TS7700 Grid.

Migration from native tape drives of any type to the TS7700 always requires host involvement to copy the data into the TS7700.

In this section, we discuss techniques for moving data in and out of the TS7700. You can start using the TS7700 by moving data into it. The best method depends on the application you want to manage with the TS7700. We describe two methods:

Phased method This method consists of starting to use TS7700 with new allocations. The migration of data takes longer, but it can be more controlled and flexible.

Quick method Use this method when you want to move existing data into the TS7700. It is called quick because it swiftly puts all data you want to move under TS7700 control.

Some hints on how to move data out of the TS7700 are provided in 6.3.6, “Moving data out of the TS7700” on page 283. However, the TS7700 is a closed-storage method, so you must be careful in selecting data to move into it. You do not want to store a large amount of data in the TS7700 that will need to be moved back out.

6.3.1 Moving data into the TS7700 - phased method

The data movement techniques outlined here depend more on changes in parameters, routines, or procedures than on overt data movement.

Select the data

If you select DFSMShsm-owned data, you can group your data as listed here:

- Migration data (DFSMShsm level 2)
- Backup copies (user data, CDS data, or both)
- Dump copies
- All of the above

You can select data based on data set name, by application, or by any other variable that you can use in the ACS routines. You can also select data based on type, such as SMF data, or DASD DUMP data.

Update the applicable parameters

If you select DFSMShsm-owned data, review the ARCCMDxx member according to the recommendations in 6.4, “Migration of DFSMShsm-managed data” on page 285 and update the following definitions:

- The Data Class ACS routines (if used)
- The Management Class ACS routines (if used)
- The Storage Class ACS routines (required)
- The Storage Group ACS routines (required)
- For BTLS, the unit parameter in the JCL

For DFSMSdss, update the following definitions:

- The Data Class ACS routines (if used)
- The Management Class ACS routines (if used)
- The Storage Class ACS routines (required)
- The Storage Group ACS routines (required)
- For BTLS, the unit parameter in the JCL

If you use database data, such as logs or image copy, direct new allocations into the TS7700 by updating the following definitions:

- The Data Class ACS routines (if used)
- The Management Class ACS routines (if used)
- The Storage Class ACS routines (required)
- The Storage Group ACS routines (required)

For data other than DFSMShsm and DFSMSdss, if you are using SMS tape, update the ACS routines to include the data you want to move. You decide what you filter for and how you write the ACS routines. You can also migrate based on the UNIT parameter in the JCL to reflect the applicable unit for the TS7700.

Update the tape management system

Even though you are not overtly copying data in this option, be sure to update the Tape Management System (TMS) catalog to reflect the changes that you expect. Check the retention rules and limits and update accordingly. If you change data set names when moving into TS7700 you must validate changes against retention rules in TMS. Refer to 6.5, “DFSMSrmm and other tape management systems” on page 293, for more information.

Watch the data move to the TS7700

Data movement using this option does not involve overt actions, such as COPY, RECYCLE, or DUMP. When you activate the ACS routines containing the code for the TS7700, all new data allocations for the data you have selected are written to the IBM TS7700. You simply verify that data is going where you expect it to go and add code to the ACS routines to manage more data as you see fit.

You can select data types that create large quantities of data, like SMF records or DASD DUMPS and you can also select data types that create many very small data sets. By observing how the TS7700 handles each type of data, you become familiar with the TS7700, its functions and capabilities.

6.3.2 Moving data into the TS7700 - quick method

The steps outlined in this section involve overt actions on your part to move data into the IBM TS7700. As with the techniques outlined in 6.3.1, “Moving data into the TS7700 - phased method” on page 277, you choose the data you want to move to the IBM TS7700.

Select the data to copy

The data you select influences all subsequent steps in this process. If you select DFSMShsm-owned data, the process for moving the data to the TS7700 is different from the process that you use for DFSMSdss data. You can select data based on the data's attributes, as when using expiration date as a basis for your choice. For example, you could select data that you keep for seven years. The choice is yours, because it is your data and your environment and you understand it better than anyone else. Probably the best method of selecting data to copy in the TS7700 is based on the data set name, application by application.

You need to be aware that some applications have knowledge of the VOLSER where the data is stored. There are special considerations for these applications. If you change the VOLSER that the data is on, the application will have no way of knowing where the data resides. For more information about this topic, see 6.3.4, “Considerations for static VOLSERs” on page 282.

An easy method is to gain information from the TMS database. Some reports can give you details on the data you actually have in the tape shop, helping to select the input volumes.

If you are using DFSMSrmm, you can easily acquire data from an RMM EXTRACT file, which is normally created as part of the regular housekeeping. Then with a REXX EXEC you extract the information needed, such as data set name, VOLSER and file sequence of the input volumes.

Copy data with a tape copy utility

If you are using SMS tape, update the ACS routines to include the data that you are copying to the TS7700.

After selecting data, the next step is to create a job stream that copies the data with the right utility from tape to tape. If you selected DFSMShsm-owned data, use DFSMShsm recycle to move the data to the TS7700. Use a COPYDUMP job to move DFSMSdss data to the TS7700.

The utility to use depends on the data selected. In most cases, it is sequential data that can be copied using the IEBGENER utility, DITTO/ESA. If you have DFSORT or a similar utility, ICEGENER and ICETOOL can probably give better performance.

You must use a specific utility when the input data is in a special format, for example, DFSMSdss dump data. DFSMSdss uses a 64 KB block size and only the proper DSS utility, like COPYDUMP, can copy with that blocksize. Also take care when copying multifile multivolume chains. You might want to separate these files because there is no penalty when they are in a TS7700.

Update TMS with the correct retention information

When the copy operation has been successful, update the tape management system catalog. This data must be updated on the output volume:

- File sequence number
- Creation date and time
- Last read and last write date
- Jobname

Optionally, you can also update:

- Stepname
- DDname
- Account number
- Device number

In RMM, this can easily be done with a CHANGEDATASET command with special authority to update O/C/EOV recorded fields.

Update the ICF catalog with the correct output volume

The next step is to uncatalog the input data sets (if they were cataloged) and recatalog the output data sets with the new volume information.

Release the input volume for SCRATCH processing

This final step needs to be done after you are sure the data has been correctly copied. You also need to verify that the retention and catalog information is correct.

Using this quick method sequence, you can copy every kind of tape data, including GDGs, without modifying the generation number.

In an RMM environment, you can use a REXX CLIST and RMM commands, listing data from the input volumes and then using the RMM REXX variables with the CD command to update the output. Afterward call IDCAMS to update the ICF catalog. When the operation completes and all errors have been corrected, use the RMM DELETEVOLUME command to release the input volumes.

Refer to the *z/OS DFSMSrmm Guide and Reference*, SC26-7404 for more information about RMM commands and REXX variables. If you are using a TMS other than RMM, refer to the appropriate product functions to obtain the same results.

The first approach when migrating data inside the TS7700 can be easier with products like DFSMShsm or Tivoli Storage Manager. If you are planning to put DFSMShsm or Tivoli Storage Manager data in the TS7700, see 6.4, “Migration of DFSMShsm-managed data” on page 285 and 6.6, “Tivoli Storage Manager” on page 294.

With DFSMShsm, you can change the ARCCMDxx tape device definitions addressing an esoteric name with TS7700 virtual drives (in a BTLS environment) or changing SMS ACS routines to direct DFSMShsm data in the TS7700. The DFSMShsm RECYCLE command can help speed the movement of the data.

A similar process can be used with Tivoli Storage Manager, changing the device class definitions for the selected data to put in the TS7700 and then invoking the space reclamation process.

If you are moving DB2 data into the TS7700, be sure that when copying the data, the DB2 catalog is also updated with the new volume information. You can use the DB2 MERGECOPY utility to speed up processing, using TS7700 virtual volumes as output.

6.3.3 Products to simplify the task

You might want to consider using a product designed to copy data from one medium to another. The first choice is the IBM enhancement for DFSMSrmm called Tape Copy Tool (see Table 6-6 on page 281). The Tape Copy Tool function of the internal IBM ADDONS package is designed to copy all types of MVS tape data sets from one or more volumes or volume sets to a new tape volume or tape volume set. Any tape media supported by DFSMSrmm is supported by this tool. The input tape media can be different than the output tape media. This tool should not be used to copy tape data sets owned by DFSMShsm (Hierarchical Storage Manager) or Tivoli Storage Manager or similar program products, where information of old VOLSER(s) is kept within the product itself and not reflected after a copy is made. This challenge typically applies to products where tapes are not cataloged in a ICF catalog.

The DFSMSrmm Copy Tool cannot be used when you have a Tape Management System other than DFSMSrmm. You need to choose another tape copy tool from Table 6-6 on page 281.

Be sure to consider these factors when evaluating a tape copy product:

- Interaction with your tape management system
- Degree of automation of the process
- Speed and efficiency of the copy operation
- Flexibility in using the product for other functions such as duplicate tape creation
- Ease of use
- Ability to create a pull list for any manual tape mounts
- Ability to handle multivolume data sets
- Ability to handle volume size changes whether from small to large or large to small
- Ability to review the list of data sets before submission
- Audit trail of data sets already copied
- Ability to handle failures during the copy operation such as input volume media failures
- Flexibility in being able to filter the data sets by wild cards or other criteria such as expiration or creation date

Table 6-6 lists some common tape copy products. You can choose one of these products or perhaps use your own utility for tape copy. Certainly you do not need any of these products, but a tape copy product will make your job easier if you have many tapes to move into the TS7700.

Table 6-6 Selection of tape copy tools

Product name	Vendor name	More information
Tape Copy Tool/ DFSMSrmm	IBM	Contact your IBM Representative for more information about this service offering.
Tape Optimizer	IBM	http://www.ibm.com/software/tivoli/products/tape-optimizer-zos
Beta55	Beta Systems Software AG	http://www.betasystems.com
CA-1/TLMS Copycat	Computer Associates International, Inc.	http://www.cai.com
Rocket Tape Optimizer for z/OS	Rocket Software, Inc.	http://www.rocketsoftware.com/portfolio/tapemedia/
Tape/Copy	OpenTech Systems, Inc.	http://www.opentechsystems.com/tape-copy.php
TelTape	Cartagena Software Ltd.	http://www.cartagena.com
Zela	Software Engineering of America	http://www.seasoft.com/zela.asp

In addition to using one of these products, consider using IBM Global Services Global Technology Services (GTS) to assist you in planning and moving the data into the TS7700. For more information about these services, see 3.8.6, “Implementation services” on page 126.

6.3.4 Considerations for static VOLSERs

Some applications have knowledge of the VOLSER of the volume where the data is stored. DFSMShsm is one such application. When moving data for these applications, you have two choices: You can utilize instructions from the application author to copy the data, or you can copy the data to a volume with the same VOLSER. For assistance with DFSMShsm tapes, see 6.3.2, “Moving data into the TS7700 - quick method” on page 278.

The preferred method for moving this type of data is to utilize instructions from the application author. If, however, you must copy the data to a volume with the same VOLSER, take these points into consideration:

- The source and target media might not be the exact same size.
- You cannot mount two volumes with the same VOLSER at the same time.
- If the source tape is a system-managed tape, you cannot have two volumes with the same VOLSER.

This is not the preferred method for moving data to the TS7700. This method applies only if you have to maintain the data-set-to-VOLSER relationship, and it has limitations and weaknesses.

- Copy the non-TS7700 tape volumes to DASD or other tape volumes.
- If the source volume is 3494-resident, eject the cartridge from the 3494, using the LIBRARY EJECT command or the ISMF EJECT line operator command from the ISMF panel.
- Delete the ejected volume from the tape management system.
- Define the VOLSER range, including the once-duplicated number, to the TS7700.
- Update ACS routines so that the data is directed to the TS7700. For BTLS, update the UNIT parameter in the JCL.
- Create a job to copy the data currently on DASD to the TS7700.
- Run the copy job.
- Update the tape management system records and any other catalog structures.

6.3.5 Combining methods to move data into the TS7700

You will most likely want to use a combination of the phased and quick methods for moving data into the TS7700. One approach is to classify your data as static or dynamic.

Static data is information that is going to be around for a long time. This data can only be moved into the TS7700 with the quick method. You have to decide how much of this data is to be moved into the TS7700. One way to make this decision is to examine expiration dates. You can then set a future time when all volumes, or a subset, are copied into the TS7700. There might be no reason to copy volumes that are going to expire in two months. By letting these volumes go to scratch status, you will save yourself some work.

Dynamic data is of a temporary nature. Full volume backups and log tapes are one example. These volumes typically have a very short expiration period. You can move this type of data with the phased method. There is no reason to copy these volumes if they are going to expire soon.

6.3.6 Moving data out of the TS7700

There are many reasons why you would want to move data out of the TS7700. The most common would be for disaster recovery or data interchange. You can move data out of the TS7700 in three ways.

The *first* method of moving data out of the TS7700 is to use a host-based tool to copy the data from the TS7700 to the target. This method is described next.

With this method the data is reprocessed by the host and copied to another medium. This method is described in 6.3.1, “Moving data into the TS7700 - phased method” on page 277. The only difference is that you need to address the TS7700 as input and the non-TS7700 drives as output.

The *second* method is to use the Copy Export function. This method is described next.

Copy Export provides a new function that allows a copy of selected logical volumes written to the TS7700 to be removed and taken offsite for disaster recovery purposes. The benefits of volume stacking, which places many logical volumes on a physical volume, are retained with this function. In addition, because the data being exported is a copy of the logical volumes, the logical volumes data remains accessible by the production host systems. We describe this method in detail in 5.3.2, “Implementing Copy Export” on page 212, and 9.4, “Disaster recovery using Copy Export” on page 504.

The Copy Export function for standalone configurations was introduced with TS7700 Virtualization Engine code level 8.3.x.x and Library Manager code level 534.x. For grid configurations, the Copy Export function was introduced with TS7700 Virtualization Engine code level 8.4.x.x. Although there are no host software updates required to support the Copy Export function, there are other functions supported in the TS7700 that do require a later level of host software for support. One of those, host console request, does require z/OS support, and that is provided at z/OS V1R6 and later. Refer to OAM APAR OA20065 and device services APARs OA20066, OA20067, and OA20313.

> **Note:** It is expected that customers will execute the Copy Export operation on a periodic basis, possibly even more than once a day. Because the purpose is to get a copy of the data offsite for disaster recovery purposes, performing it soon after the data is created minimizes the time for the recovery point objective.

The *third* method is to copy the data with the DFSMShsm ABARS function. This method is described next.

ABARS is the command-driven DFSMShsm function that backs up a user-defined group (called an aggregate group) of data sets (usually for recovery purposes) at another computer site or at the same site. ABARS can be used to back up and recover both SMS- and non-SMS-managed data, on DASD and on tape. Using the DFSMShsm ABARS function, group the data you want to move outside the TS7700, then start addressing other tape drives outside the IBM TS7700, or use the Copy Export function. In this way, you obtain an exportable copy of the data that can be put in an off-site location.

We now give an overview of the process.

Create a selection data set

Before you can run an aggregate backup, you have to create one or more selection data sets. The selection data set lists the names of the data sets to be processed during aggregate backup.

You can identify the data set names in a single selection data set, or you can divide the names among as many as five selection data sets. You can specify six types of data set lists in a selection data set. The type you specify determines which data sets are backed up and how they are recovered.

An INCLUDE data set list is a list of data sets to be copied by aggregate backup to a tape data file where they can be transported to the recovery site and recovered by aggregate recovery. The list can contain fully qualified data set names or partially qualified names with place holders. DFSMShsm expands the list to fully qualified data set names.

Using a selection data set with the names of the data sets you want to export from the IBM TS7700, you obtain a list of files on logical volumes that the ABARS function copies to non-TS7700 drives.

You can also use the Copy Export function to move the ABAR tapes to a Data Recovery site, out of the library.

Define an aggregate group

You have to define an aggregate group and related Management Class to specify exactly which data sets are to be backed up.

Define the aggregate group and Management Class used for aggregate backup to DFSMS through ISMF panels.

The aggregate group lists the selection data set names, instruction data set name and additional control information used by aggregate backup in determining which data sets are to be backed up.

Execute the ABACKUP VERIFY command

You have two options with the ABACKUP command. You can choose to verify the contents of the aggregate backup without actually backing up any data sets, which is the same as performing a test run of aggregate backup.

An example of the ABACKUP command is:

```
HSEND ABACKUP agname VERIFY UNIT(non_TS7700_unit) PROCESSONLY(USERTAPE)
```

With the PROCESSONLY(USERTAPE) keyword, only tape data sets are processed. In this way you can be sure that only the input data from TS7700 logical volumes is used.

Execute the ABACKUP EXECUTE command

When you are ready, you can perform the actual backup, using this command:

```
HSEND ABACKUP agname EXECUTE UNIT(non_TS7700_unit) PROCESSONLY(USERTAPE)
```

When you issue the ABACKUP command with the execute option, the following tape files are created for later use as input to aggregate recovery:

- Data file: Contains copies of the data sets that have been backed up.
- Control file: Contains control information needed by aggregate recovery to verify or recover the application's data sets.
- Instruction/activity log file: Contains the instruction data set, which is optional.

At the end of this process, you obtain an exportable copy of the TS7700 data, which can be used for disaster recovery and stored off-site using other physical tapes.

You should consider the use of the Copy Export function, which allows you to move a copy of the original logical volumes to an offsite location without reading the tape data twice. The Copy Export function operates on another Physical Volume Pool in the library and creates the copy in the background without any process required on the host. But Copy Export requires an empty TS7700 on your disaster site. For more information of Copy Export, refer to 5.3.2, “Implementing Copy Export” on page 212, and for more information about Copy Export Recovery refer to 9.4, “Disaster recovery using Copy Export” on page 504.

For more information about using the DFSMShsm ABARS function, refer to *DFSMShsm Storage Administration Guide*, SC35-0421.

6.4 Migration of DFSMShsm-managed data

The z/OS DFSMS hierarchical storage manager (DFSMShsm) has its own functions that allow full utilization of the capacity of tape cartridges. So at first sight, using the IBM TS7700 in a DFSMShsm environment does not bring any special advantage from a capacity point of view, because you create virtual volumes, which, after copied, fill a stacked volume completely. You can achieve the same results by writing directly to a physical tape drive, leaving the TS7700 storage for other applications that cannot use the full cartridge capacity and are therefore better candidates for TS7700 management.

Even though DFSMShsm is an application capable of using the full cartridge capacity, there are a number of reasons why you might want to consider using the TS7700 instead of native physical drives for DFSMShsm data. For example, when writing ML2 data onto a cartridge with an uncompressed capacity of 300 GB, chances are higher that a recall request needs exactly this cartridge that is currently being written to by a space management task. This incident is known as *recall takeaway*.

The effects of recall takeaway can be a real disadvantage when writing ML2 data onto native, high capacity cartridges, because the space management task must set aside its output tape to make it available to the recall task. While the partially-filled output tape remains eligible for subsequent selection the next time that space management runs, it is possible to accumulate a number of partial tapes beyond DFSMShsm's needs if recall takeaway activity occurs frequently. Excess partial tapes created by recall takeaway activity result in poor utilization of native cartridges. In addition, because recall takeaway activity does not cause the set-aside tape to be marked full, it is not automatically eligible for recycle, despite its poor utilization.

High capacity cartridges are more likely to experience not only frequent recall takeaway activity, but also frequent *piggy-back recall* activity in which recalls for multiple data sets on a single tape are received while the tape is mounted. While piggy-back recalls have a positive effect by reducing the number of mounts required to perform a given number of recalls, one must also keep in mind that multiple recalls from the same tape must be performed serially by the same recall task. Were those same data sets to reside on different tapes, the recalls could potentially be performed in parallel, given a sufficient number of recall tasks. In addition, the persistence of the virtual tape in the Tape Volume Cache after it has been demounted allows DFSMShsm to perform ML2 recalls from the disk cache for a period of time without requiring that a physical tape be mounted.

There are also other reasons for directing DFSMShsm data into a TS7700. The amount of native drives limits the number of DFSMShsm tasks that can run concurrently. With the large number of up to 256 virtual drives in a Single Cluster Grid configuration or 512 virtual drives in a Two-Cluster Grid configuration, you can dedicate a larger number of virtual drives to each DFSMShsm function and allow for higher throughput during your limited backup and space management window.

When increasing the number of DFSMShsm tasks to take advantage of the large number of virtual drives in a TS7700, consider adding more DFSMShsm auxiliary tasks (MASH), rather than simply increasing the number of functional tasks within the existing DFSMShsm started tasks. While each DFSMShsm started task can support up to 15 AUTOBACKUP tasks, for example, greater throughput will be achieved by requesting five AUTOBACKUP tasks on each of three DFSMShsm started tasks.

Other reasons for using the TS7700 with DFSMShsm are the greatly reduced run times of DFSMShsm operations that process the entire volume, such as AUDIT MEDIACONTROLS and TAPECOPY.

DFSMShsm can benefit from TS7700’s high throughput and from its large tape volume cache size allowing for long periods of peak throughput.

DFSMShsm data is well suited for the TS7700, given appropriate tailoring of those parameters that can affect DFSMShsm performance. The subsequent sections describe this tailoring in more detail.

For more details, refer to *DFSMShsm Storage Administration Guide,* SC26-7402.

6.4.1 Volume and data set sizes

The size of user data sets is very important when you choose between a TS7700 or native drives like 3592. DFSMShsm Migration, Backup, and Recycle use only Single File Format to write to tape cartridges.

z/OS supported data set sizes

Different data set sizes are supported for disk and tape data sets, based on the data set organization and the number of volumes a single data set can span:

- DASD data sets are limited to 59 volumes, except for PDS and PDSE data sets, which are limited to one volume.
- A data set on a VIO simulated device is limited to 65,535 tracks and to one volume.
- Tape data sets are limited to 255 volumes.

Table 6-7 lists the maximum data set sizes supported in z/OS environments.

Table 6-7 Maximum supported data set sizes

Storage medium	Max. volume size	Max. number of volumes	Max. data set size
DASD: DS8000®	65,520 CYL = 54 GB	59	3.18 TB
Tape: IBM TS1120	700 GB x 2.5 Compression	255	446.25 TB
Tape: TS7700	4,000 MB x 2.5 Compression	255	2.55 TB

DFSMShsm-supported data set sizes

Single file format, as used by DFSMShsm, reduces I/O and system serialization because only one label is required for each connected set (as opposed to multiple file format tapes that require a label for each data set). The standard-label tape data set that is associated with the connected set can span up to the allocation limit of 255 tapes. This standard-label tape data set is called the DFSMShsm tape data set. Each user data set is written in 16K logical blocks to the DFSMShsm tape data set.

> **Important:** A single DFSMShsm user data set can span up to 40 tapes. This is the limit for Migration, Backup, and Recycle.

After DFSMShsm writes a user data set to tape, it checks the volume count for the DFSMShsm tape data set. If the volume count is greater than 215, the DFSMShsm tape data set is closed, and the currently mounted tape is marked full and is de-allocated. The number 215 is used so that a data set spanning 40 tapes fits within the 255-volume allocation limit. DFSMShsm selects another tape, and then starts a different DFSMShsm tape data set.

Data set spanning can be reduced using the SETSYS TAPESPANSIZE command.

DFSMShsm and large logical volumes

The TS7700 supports logical volume sizes of 400, 800, 1000, 2000, and 4000 MB. With a maximum of 40 volumes supported and assuming a compression ratio of 2.5:1, the maximum user data set size for 800 MB volumes is:

800 MB x 2.5 x 40 = 80 GB

Let us assume that you have a very large data set of 300 GB. Such a data set does not fit on 40 volumes of 800 MB each, but it would fit on 4000 MB large virtual volumes, as shown in the following example:

4000 MB x 2.5 x 40 = 400 GB

Any single user data set larger than 400 GB is a candidate for native TS1120 tape drives; assuming a compression rate of 2.5:1, they might not fit onto the supported number 40 volumes. In this case, you should consider using native TS1120 tape drives rather than TS7700.

IDCAMS DCOLLECT BACKUPDATA and MIGRATEDATA can be used to determine the maximum size of backed-up and migrated data sets, respectively, in DFSMShsm's inventory.

Important: DFSMShsm can have more than one address space on one LPAR (MASH = Multi Address Space Support) or have a different setup on different LPARs in your PARMLIB member ARCCMDxx and separated using ONLYIF statements. One DFSMShsm address space can have a MIGUNIT(3590-1) and the other address space a MIGUNIT(TS7700). The same is true for BUUNIT. The DFSMShsm instance which has the 3592 as Migration or Backup Unit can run Space Management or Auto Backup only for that Storage Group (SG) where all your large data sets like z/FS reside. The other DFSMShsm instance would migrate and back up all the smaller data sets onto TS7700. You can use a command such as F DFSMS2,BACKDS or F DFHSM2,BACKVOL(SG2) to issue the command to the second address space of DFSMShsm:

- SETSYS TAPEMIGRATION(ML2TAPE(TAPE(unittype))
- SETSYS RECYCLEOUTPUT(MIGRATION(unittype))
- SETSYS BACKUP(TAPE(unittype))
- SETSYS RECYCLEOUTPUT(BACKUP(unittype))

Migration to a different logical volume size

To make sure that DFSMShsm starts using larger data sets, you must mark full any empty or partially filled tapes written using the previous logical volume size. To identify these tapes, issue the following DFSMShsm command:

```
LIST TTOC SELECT(NOTFULL)
```

Each tape identified as being empty or partially filled must be marked full using the following DFSMShsm command:

```
DELVOL volser MIGRATION(MARKFULL)
```

Or:

```
DELVOL volser BACKUP(MARKFULL)
```

As DFSMShsm migrates data and creates backup copies, it prefers to add to an existing migration/backup volume. As the volume nears full, it handles spanning of data sets as described in “Tape spanning” on page 290. If a data set spans across DFSMShsm volumes, it becomes a “connected set” in DFSMShsm terms.

A key point, however, is that whether the data set spans or not, DFSMShsm uses FEOV processing to get the next volume mounted. So, the system believes that the volume is part of a multi-volume set regardless of whether DFSMShsm identifies it as a connected set. Because of the EOV processing, the newly mounted DFSMShsm volume will use the same Data Class and other SMS constructs as the previous volume.

With the DFSMShsm SETSYS PARTIALTAPE MARKFULL option, DFSMShsm will mark the last output tape full, even though it has not reached its physical capacity. By marking the last volume full, the next time processing starts, DFSMShsm will use a new volume, starting a new multi-volume set and allowing for the use of a new Data Class and other SMS constructs. If the volume is not marked full, the existing multi-volume set will continue to grow and to use the old constructs.

We strongly recommend that you use the SETSYS PARTIALTAPE MARKFULL option as it reduces the occasions in which DFSMShsm will append to a partial tape, which results not only in the need to mount a physical tape, but also in the invalidation of the existing virtual tape, which will eventually need to be reclaimed by the TS7700.

This discussion is relevant to outboard policy management and the implementation of different logical volume sizes. If all volumes have been marked full, you can simply update your ACS routines to assign a new Data Class and other SMS constructs and, from then on, each new migration or backup volume will use the new size.

6.4.2 TS7700 implementation considerations

This section summarizes some DFSMShsm implementation considerations with regard to the TS7700.

Mount wait time

If you direct DFSMShsm data into the TS7700, we recommend that you modify your DFSMShsm mount wait timer to be 12 minutes. This allows for possibly needed extra time on specific mounts for the TS7700 to stage the data back into cache.

Logical volume size

Consider using large logical volumes such as 4000 MB for backup and smaller logical volumes for migration. If you have a high recall rate from ML2, you might not even want to use the entire capacity of a MEDIA1 or MEDIA2 virtual volume. Installations in which recalls from ML2 are rare and installations in which extremely large data sets are migrated that could result in the 40-volume limit being reached should use the maximum capacity of the virtual volume. You need to select another SMS DATACLAS for backup than for migration through your ACS routine.

Refer to Table 6-8 on page 291 when tailoring the ARCCMDxx SETSYS parameters. The PERCENTFULL value will depend upon the size of the logical volume.

DFSMShsm recognizes only that the Data Class specifies either Media1 or Media2; it is unaware of large logical volume sizes (1000 MB, 2000 MB and 4000 MB) specified in the Data Class. For that reason, the SETSYS TAPEUTILIZATION command must be used to inform DFSMShsm that the virtual tapes have a capacity other than 400 MB or 800 MB. You must tune the PERCENTFULL value for your requirements.

Other applications might have a similar existing TAPECAPACITY-type specification or a PERCENTFULL-type specification enabling applications to write beyond the default volume sizes for MEDIA1 (cartridge system tape) and MEDIA2 (enhanced capacity cartridge system tape).

In the case with OAM's Object Tape Support, the TAPECAPACITY parameter in the SETOAM statement of the CBROAMxx PARMLIB member is used to specify the larger logical volumes sizes. For a 1000 MB logical volume, the capacity specified should be 1 000 000 KB (1000 X 1000). Then, for a 2000 MB logical volume, the capacity specified should be 2 000 000 KB (2000 x 1000), and for a 4000 MB logical volume, the capacity specified should be 4 000 000 KB (4000 x 1000). For additional information about outboard policy management, refer to *z/OS DFSMS Object Access Method Planning, Installation and Storage Administration Guide for Tape Libraries*, SC35-0427.

Multi-system considerations

If multiple TS7700 Virtualization Engines are eligible for a request, consideration should also be given so that the same logical volume size is used for the request across all libraries. When displaying the volumes through your tape management system, the tape management system might continue to display the volume capacity based on the default volume size for the media type with the volume usage (or a similar parameter) showing how much data has actually been written to the volume reflecting its larger capacity.

Scratch volumes

The default volume size is overridden at the library through the Library Manager Data Class policy specification and is assigned or re-assigned when the volume is mounted for a scratch mount. Using a global scratch pool, you benefit from a fast mount time by using the fast-ready attribute for the scratch category as explained in 4.3.5, “Define Fast Ready categories” on page 154. Consider using the following definitions to benefit from the fast scratch mount times:

```
SETSYS SELECTVOLUME(SCRATCH)
SETSYS TAPEDELETION(SCRATCHTAPE)
SETSYS PARTIALTAPE(MARKFULL)
```

The first of these commands requests DFSMShsm to use volumes from the common scratch pool, and the second one defines that DFSMShsm returns tapes to the common scratch pool. The last one defines that an DFSMShsm task will mark the last tape it used in a cycle to be full, thus avoiding a specific mount during the next cycle. The MARKFULL parameter does not mean a waste of space using TS7700, because the stacked volume contains only the written data of each logical volume copied and the same applies to the TVC.

Tape spanning

You can use the optional TAPESPANSIZE parameter of the SETSYS command to reduce the spanning of data sets across migration or backup tape volumes, for example:

```
SETSYS TAPESPANSIZE(4000)
```

The value in parentheses represents the maximum number of megabytes of tape (ML2 or backup) that DFSMShsm can leave unused while it tries to eliminate spanning of data sets. To state this differently, this value is the minimum size of a data set that is allowed to span tape volumes. Data sets whose size is less than the value do not normally span volumes. Only those data sets whose size is greater than or equal to the specified value are allowed to span volumes.

This parameter offers a trade-off: to reduce the occurrences of a user data set spanning tapes in exchange for writing less data to a given tape volume than its capacity would otherwise allow. The amount of unused media can vary from 0 to nnnn physical megabytes, but roughly averages 50% of the customer's median data set size. For example, if you specify 4000 MB and your median data set size is 2 MB, then on average only 1 MB of media is unused per cartridge.

Installations that currently experience an excessive number of spanning data sets might want to consider specifying a larger value in the SETSYS TAPESPANSIZE command.

Using a high value reduces tape spanning. In a TS7700 this reduces the number of virtual volumes that need to be recalled to satisfy DFSMShsm recall/recover requests. You can be generous with the value, as no space is wasted. For example, TAPESPANSIZE of 4000 would mean that any data set with less than 4,000 MB that will not fit on the remaining space of a virtual volume will be started on a fresh new virtual volume.

6.4.3 DFSMShsm task-related considerations

To better understand the use of DFSMShsm with TS7700, we now summarize the different DFSMShsm functions that use tapes and analyze the benefit of tape virtualization for these functions.

Backups of DFSMShsm control data sets

The backup of DFSMShsm control data sets (CDSs) can easily be done in a TS7700, exploiting the benefit of using virtual volumes instead of physical volumes, which might otherwise be underutilized.

Volume dumps

When using TS7700 as output for the DFSMShsm AUTODUMP function, do not specify the DEFINE DUMPCLASS(dclass STACK(nn)) or BACKVOL SG(sgname)|VOLUMES(volser) DUMP(dclass STACK(10)) parameters. These parameters were introduced to force DFSMShsm to use the capacity of native physical cartridges. If used with TS7700, they cause unnecessary multivolume files and reduce the level of parallelism possible when the dump copies are restored. We therefore recommend using the default value, which is NOSTACK.

Migrate/recall (DFSMShsm Migration Level 2)

When using a TS7700 as DFSMShsm Level 2 you need to consider the number of simultaneous recall processes. Consider how many recall tasks are started at the same time and compare that number with the number of physical drives assigned to your TS7700.

For example, if your installation often has more than 10 tape recall tasks at one time, you will probably need twelve TS1120 back-end drives to satisfy this throughput request, because all migrated data sets might have been removed from the TVC and need to be recalled from tape.

Backup and recovery

Unlike the DFSMShsm RECALL operation, RECOVERY usually has a lower frequency in an DFSMShsm environment. Therefore, using TS7700 for DFSMShsm backup and recovery functions benefits you without impacting DFSMShsm performance. However, carefully review your DFSMShsm performance requirements before moving DFSMShsm BACKUP to the TS7700.

Tapecopy

The DFSMShsm TAPECOPY function requires that original and target tape volumes are of the same media type and use the same recording technology. Using a TS7700 as target for the TAPECOPY operation can cause problems in DFSMShsm, because TS7700 virtual volumes have different volume sizes, even though they are defined as CST or ECCST.

Use Table 6-8 to tailor your TAPECOPY environment.

Table 6-8 TAPECOPY utilization

ORIGINAL volume unit name	ALTERNATE volume unit name	Percent full that should be defined (assuming 2:1 compression)
TS7700 (CST) - 400 MB	3490E (CST)	106%
TS7700 (ECCST) - 800 MB	3490E (ECCST)	107%
3490E (CST) - 400 MB	TS7700 CST - 400 MB	45%
3490E (ECCST) - 800 MB	TS7700 (ECCST)- 800 MB	45%
TS7700 (CST) - 400 MB	TS7700 (CST) - 400 MB	106%
TS7700 (CST) - 1 GB	TS7700 (CST) - 1 GB	270%

ORIGINAL volume unit name	ALTERNATE volume unit name	Percent full that should be defined (assuming 2:1 compression)
TS7700 (CST) - 2 GB	TS7700 (CST) - 2 GB	543%
TS7700 (CST) - 4 GB	TS7700 (CST) - 4 GB	1090%
TS7700 (ECCST) - 800 MB	TS7700 (ECCST) - 800 MB	107%
TS7700 (ECCST) - 1 GB	TS7700 (ECCST) - 1 GB	135%
TS7700 (ECCST) - 2 GB	TS7700 (ECCST) - 2 GB	271%
TS7700 (ECCST) - 4 GB	TS7700 (ECCST) - 4 GB	545%

For example, if you are planning to put DFSMShsm alternate copies into an IBM TS7700, a tape capacity of 45% might not be enough for the input non-TS7700 ECCST cartridges. TAPECOPY will fail if the (virtual) output cartridge encounters EOV before the input volume has been copied completely.

Alternatively, using TS7700 logical volumes as original and 3490E native as TAPECOPY target might cause EOV at the alternate volume because of the higher IBMLZ1 compression seen on the virtual drive compared to the IDRC compression on the native drive.

For special situations where copying from standard to enhanced capacity media is needed, the following patch command can be used:

```
PATCH .MCVT.+4F3 BITS(.....1..)
```

DUPLEX TAPE

For duplexed migration both output tapes have to be of the exact same size and unit type. We recommend to use a Multi Cluster Grid and let the hardware do the duplex rather than the DFSMShsm software function. This also helps to more easily manage the disaster side. You can use GDPS and switch to the remote DASD side and the tape VOLSER itself has not to be changed. No TAPEREPL or SETSYS DISASTERMODE commands are needed.

When using the TS1120 Tape Drive for duplexed migration output, performance is degraded because of the back-to-back SYNCDEV operations done for the original and the alternate tapes. APAR OA09928 provides a patch allowing syncs on the alternate tape to be disabled. The performance improvement varies with data set size, with the greatest improvements seen for the smaller data sets. Performance improvements can be quite substantial.

RECYCLE

The DFSMShsm RECYCLE function reduces the number of logical volumes inside the TS7700, but when started can cause bottlenecks in the TS7700 recall process. If you have a TS7700 with four physical drives, use a maximum of two concurrent DFSMShsm RECYCLE tasks. If you have an IBM TS77000 with six physical drives, use no more than five concurrent DFSMShsm RECYCLE tasks.

Select the RECYCLEPERCENT carefully, bearing in mind that:

- You will free up logical volumes residing on a stacked volume with hundreds of other logical volumes.
- The space occupied by the logical volume will be freed up only if and when the logical volume is used (overwritten) again, unless you are using Expired Volume Management.
- To RECYCLE, the TS7700 has to load the input volumes into the TVC.

We recommend to use a RECYCLEPERCENT value depending on the logical volume size, for example:

```
 5      for 1000, 2000, or 4000 MB volumes
10      for 400 or 800 MB volumes
```

Using RECYCLE SELECT (INCLUDE(RANGE(nnnnn:mmmmm))) or RECYCLE SELECT (EXCLUDE(RANGE(nnnnn:mmmmm))) for RECYCLE input can be helpful while selecting and migrating data to and from a TS7700. Its immediate purpose is to enable you to set up volume ranges for different media types and different emulation types, like TS7700 logical volumes and 3490-emulated cartridges. There are no special data set names for RECYCLEOUTPUT, although you must code your ACS routines to route RECYCLEOUTPUT to the library, using the &UNIT variable.

Refer to *DFSMShsm Primer*, SG24-5272, for more information about implementing DFSMShsm.

6.5 DFSMSrmm and other tape management systems

No changes are required to any tape management system to support basic TS7700. You need only review the retention and movement criteria for the data in the TS7700.

When you direct allocations inside the TS7700, the Vital Record Specifications (VRSs) or vault rules should tell the tape management system that the data set will never be moved outside the library.

Stacked volumes cannot be used by the host; they are managed exclusively by the TS7700. Do not allow any host to either implicitly or explicitly address these stacked volumes. To indicate that the stacked VOLSER range is reserved and cannot be used by any host system, define the VOLSERs of the stacked volumes to RMM.

Use the following PARMLIB parameter, assuming that VT is the prefix of your stacked TS7700 cartridges:

REJECT ANYUSE(VT*)

This will cause RMM to deny any attempt to read or write those volumes on native drives.

You do not need to explicitly define the virtual volumes to RMM. During entry processing, the active RMM automatically records information about each volume in its control data set. RMM uses the defaults you specified in ISMF for the library entry values if there is no existing RMM entry for an inserted volume; set the default entry status to SCRATCH.

When adding 1,000,000 virtual volumes, the size of the RMM CDS and the amount of secondary space available must be checked. RMM uses 1 MB for every 1,000 volumes defined in its CDS; an additional 1,000,000 volumes would need 1000 MB of space. However, we do not recommend that you add all volumes initially. Refer to 4.5.8, “Insert logical volumes using the TS7700 Management Interface” on page 182.

To increase the size of the RMM CDS, you have to stop RMM activities, back up the CDS, then reallocate a new CDS with a larger size and restore the CDS from the backup copy. To calculate the correct size of the RMM CDS, refer to the *z/OS DFSMSrmm Guide and Reference*, SC26-7404.

Other tape management systems, such as BrightStor, CA-1 Tape Management Copycat Utility (BrightStor CA-1 Copycat), and BrightStor CA-Dynam/TLMS Tape Management

Copycat Utility (BrightStor CA-Dynam/TLMS Copycat) need to reformat their database to add more volumes.

Additionally, some tape management systems do not allow the specification of tape volumes with alphanumeric characters or require user modifications to do so. Refer to the proper product documentation for this operation.

In both RMM and the other tape management systems, the virtual volumes do not have to be initialized. The first time a VOLSER is used, TS7700 marks the virtual volume with VOL1, HDR1 and a tape mark, as though it had been done by EDGINERS or IEHINITT.

6.6 Tivoli Storage Manager

Customers who use Tivoli Storage Manager with TS1120 or 3590 tape drives might be less satisfied with performance of Tivoli Storage Manager and TS7700 compared to the connection to native drives. In general, Tivoli Storage Manager operations will perform better with native drives than with a TS7700, and you should understand the trade-off between performance, economy, and the TS7700 functionality.

When using a Multi Cluster Grid you can get two or more copies instantly without any use of the Tivoli Storage Manager COPYSTORAGEPOOL function and CPU-cycles related to do the function. You can even move a copy of the logical volumes offsite with the Copy Export function.

6.6.1 Recommendations for TS7700 usage

Tivoli Storage Manager, like DFSMShsm, can automatically fill a native 3590 or 3592 cartridge. It can use the tape up to the EOV, independent of the media type. If you plan to store Tivoli Storage Manager data into the TS7700, consider the following suggestions for placing data on your TS7700.

Use TS7700 for Tivoli Storage Manager archiving for archiving and backup of large files or databases for which you do not have a high performance requirement during backup and restore. TS7700 is ideal for Tivoli Storage Manager archive or long term storage because archive data is not frequently retrieved. Archives and restores for large files should see less impact from the staging overhead. Small files, such as individual files on file servers, can see performance impacts from the TS7700 staging. (If a volume is not in cache, the entire volume must be staged before any restore can be done.)

Set Tivoli Storage Manager reclamation off by setting the reclamation threshold to 100 percent. Tivoli Storage Manager, like DFSMShsm, has a reclamation function to consolidate valid data from tapes with a low valid data percentage onto scratch tapes so that tapes can be freed up for reuse. Tivoli Storage Manager Reclamation with TS7700 can be slower because all volumes have to be staged to the cache. You should periodically set Tivoli Storage Manager reclamation on, by setting the threshold to a lower value to regain the use of TS7700 volumes with a small amount of valid data that will not expire for a longer period of time. Tivoli Storage Manager reclamation should be scheduled for off-peak hours.

Use collocation to reduce the number of TS7700 volumes required for a full restore. Tivoli Storage Manager has a collocation function to group Tivoli Storage Manager client data onto a minimum set of tapes to provide a faster restore and to provide separation of client data onto different physical tapes. Collocation with TS7700 does not minimize the physical tapes used but minimizes the number of logical volumes used. Collocation with TS7700 can improve restore time for large amounts of data. TS7700 will not ensure physical tape separation when collocation is used because different logical volumes can reside on the same physical tape.

Use TS7700 for Tivoli Storage Manager database backups that are to be used for recovery from local media and use TS7700 at a recovery site or native drives for backups that are to be used for recovery from off-site media. Tivoli Storage Manager requires a different tape for every backup of the Tivoli Storage Manager database, therefore a large number of logical volumes with less data is created. Using the TS7700 you do not have to worry about the unused capacity of logical volumes.

Use TS7700 for backups of primary pools, noting that similar considerations apply to copy storage pools. If only one copy pool is used for local backups, then that storage pool should not be in the TS7700, because there can be no guarantee that data in the copy storage pools are on different physical volumes. If storage pools for local and off-site backups are used, the copy storage pools for local backups can be in the TS7700. The copy storage pools for off-site backups should use native drives or a TS7700 at the recovery site.

Use TS7700 in server-to-server configurations for multiple Tivoli Storage Manager server implementations. If you are using Tivoli Storage Manager server-to-server configuration, the data from your remote Tivoli Storage Manager servers are stored as virtual volumes, which appear as sequential media volumes on the source server and which are actually stored as archive files on a target server. These are ideal candidates for a TS7700.

6.6.2 Recommendations for native drives

Use native drives for data that will be used for frequent individual file restores or have a requirement for high performance for backup and restore without any delays because of staging activity. Tivoli Storage Manager uses EXPORT to move data from one Tivoli Storage Manager server to another. This requires that both servers have compatible devices for the EXPORT media. Native drives should be used for Tivoli Storage Manager EXPORT unless you have the advanced function IMPORT/EXPORT at both TS7700s.

6.6.3 Tivoli Storage Manager parameter settings

The settings for the following parameters can affect the performance of Tivoli Storage Manager with TS7700.

MAXSCRATCH (storage pool definition) - As for DFSMShsm, Tivoli Storage Manager should use a scratch pool for tapes, because you do not have to predefine tapes to Tivoli Storage Manager and you could benefit from the faster TS7700 scratch mounts.

MOUNTLimit (device class definition) - With one TS7700, you have up to 256 virtual drives available. The number of drives available for Tivoli Storage Manager use can probably be increased, taking into account TS7700 performance. Set MOUNTLimit high enough so that the number of available drives does not limit the performance of Tivoli Storage Manager tape operations.

MOUNTRetention (device class definition) - When storing data in the TS7700, you can set this parameter to zero, because you have a greater chance of finding the virtual volume still in the TVC when Tivoli Storage Manager needs it. This avoids the need to keep the virtual volume mounted and frees a virtual drive for other users.

MAXCAPacity (device class definition) - With this parameter you can tailor the size of the data written in a virtual volume. Having smaller virtual volumes can speed up recall processing.

Using the full capacity of the virtual volume can limit the number of volumes used by Tivoli Storage Manager.

BAckup DB (database back up): Use SCRATCH=YES to use tapes from the TMS scratch pool and benefit from the faster TS7700 scratch mounts.

For details on setting up Tivoli Storage Manager, refer to the *Tivoli Storage Manager Administrators Guide* at:

http://www.tivoli.com/support/public/Prodman/public_manuals/td/TD_PROD_LIST.html

6.7 DFSMSdss

DFSMSdss full volume dumps are a good use of the IBM TS7700. You have to plan carefully, however, and make sure you can achieve the required throughput: A DFSMSdss full volume physical dump can easily provide a data transfer rate of 10 MB/s and higher for a single job. However, with today's TS7700 throughput capabilities of up to more than 400 MB/s, the TS7700 throughput capabilities are unlikely to be a limiting factor. In the past, the data rate was often limited by the bandwidth of the DASD subsystem as the weakest part in the chain.

With TS7700, you fill the stacked cartridge completely without changing JCL, using multiple virtual volumes. TS7700 then moves the virtual volumes created onto a stacked volume.

The only problem you might experience when using TS7700 for the DSS volume dumps is related to the size of the virtual volumes. If a single dump does not fit onto five logical volumes, you can either use an SMS DATACLAS specification, Volume Count nn, to enable more than five volumes. A better way is if you have TS7700 release 7.4 installed and can choose a 4000 MB logical volume through your SMS DATACLAS. This prevents unneeded multi-volume files.

Using the COMPRESS keyword of the DUMP command, you obtain a software compression of the data at the host level. As data is compressed at the TS7700 before being written into the tape volume cache, host compression is not required, unless channel utilization is high already.

6.7.1 Full volume dumps

The only problem you might experience when using TS7700 for the DSS volume dumps is related to the size of the virtual volumes. If a single dump does not fit onto five logical volumes, you can either use an SMS DATACLAS specification, Volume Count nn, to enable more than five volumes. A better way is to choose a 4000 MB logical volume through your SMS DATACLAS. This prevents unneeded multi-volume files.

Using the COMPRESS keyword of the DUMP command, you obtain a software compression of the data at the host level. As data is compressed at the TS7700 before being written into the tape volume cache, host compression is not required, unless channel utilization is high already.

6.7.2 Standalone Services

Standalone Services of DFSMSdss provides a standalone restore function that enables you to restore vital system packs without needing to rely on a System z environment.

Standalone Services supports the 3494 and the Virtual Tape Server. It enables you to restore from native as well as virtual tape volumes in a TS7700. Standalone Services lets you specify the input volumes on the RESTORE command and sends the necessary mount requests to the Library Manager.

You can initial program load (IPL) the Standalone Services core image from a virtual tape device and use it to restore dump data sets from virtual tape volumes.

Standalone Services is provided as a replacement to the previous DFDSS V2.5 and DFSMS/MVS V1 standalone functions. The installation procedure for Standalone Services retains, rather than replaces, the existing standalone restore program so you do not have to immediately change your recovery procedures. We recommend that you implement the procedures as soon as you can and start using the enhanced Standalone Services.

To use Standalone Services, create a standalone core image suitable for IPL, using the new BUILDSA command of DFSMSdss. Create a new virtual tape as non-labeled and then put the standalone program on it.

Note: The BUILDSA command does not write over the label. A tape that the TS7700 labeled initially *cannot* be changed to unlabeled. It is not possible to alter an LVOL from labeled to unlabeled. This criteria is valid for all other standalone programs as well.

Use the following steps to IPL the Standalone Services program from a virtual device and to restore a dump data set from virtual volumes. Refer to “Standalone Mount Logical Volume” on page 340 on how to use Library Manager menus to set a device in standalone mode.

- Ensure that the virtual devices you will be using are offline to other host systems. Tape drives to be used for standalone operations must remain offline to other systems.
- Set in standalone mode the virtual device from which you will load the Standalone Services program by using the Setup Standalone Device window on the Library Manager console. Select Mount a single volume and specify the virtual volume that contains the core image for IPL.
- Load the Standalone Services program from the device you just set in standalone mode. As part of this process, select the operator console and specify the input device for entering Standalone Services commands.
- When the IPL is complete, enter the Standalone Services RESTORE command from the specified input device. Example 6-26 shows a group of statements for use.

  Example 6-26 RESTORE command

  ```
  RESTORE FROMDEV(TAPE) FROMADDR(0A40) TOADDR(0900) -
     NOVERIFY TAPEVOL((L00001),(L00002))
  ```

  L00001 and L00002 are virtual volumes that contain the dump data set to be restored, 0A40 is the virtual device used for reading source volumes L00001 and L00002, and 0900 is the device address of the DASD target volume to be restored.

  Standalone Services requests the Library Manager to mount the source volumes in the order in which they are specified on the TAPEVOL parameter. It automatically unloads each volume, then requests the Library Manager to demount it and to mount the next volume.
- When the restore is complete, unload and demount the IPL volume from the virtual device by using the Library Manager’s Setup Standalone Device window.
- Take the IPL device out of standalone mode by using the Library Manager’s Reset Standalone Device window.

Standalone Services issues the necessary mount and demount orders to the library. If you are using another standalone restore program which does not support the mounting of library resident volumes, you would have to set the source device in standalone mode and manually instruct the Library Manager to mount the volumes using the Setup Standalone Device window.

For details on how to use Standalone Services, refer to *z/OS DFSMSdss Storage Administration Reference,* SC35-0424.

6.8 Object Access Method

Tape cartridges provide a low-cost storage medium for storing primary or backup copies of Object Access Method (OAM) objects.

Allowing objects to be stored on tape volumes in conjunction with DASD and optical media provides flexibility and more efficiency within the storage management facility.

OAM stores objects on a TS7700 as in a normal TS3500 Tape Library, with up to 256 virtual drives and many virtual volumes available.

Consider using the TAPEPERCENTFULL parameter with object tape data, because the retrieve time of an OAM object is important. The recall time for smaller logical volumes can be reduced considerably.

There are also functional enhancements associated with the TS7700 Virtualization Engine, one of which includes support for larger logical volume sizes. Application configuration-related changes might be needed to fully utilize the larger logical volume sizes defined at the library. For example, with OAM's object tape support, the existing SETOAM TAPECAPACITY keyword can be used to fill the larger logical volumes. However, before using this support with OAM's object tape support, install existing field APAR OA08963.

Virtual volumes in a TS7700 can contain primary or backup copies of OAM objects, addressing either OBJECT or OBJECT BACKUP Storage Groups. An advisable solution is to address TS7700 with the OBJECT Storage Group and other non-TS7700 tape devices with the OBJECT BACKUP Storage Group.

A virtual volume can contain multiple OAM objects, separated by a buffer space. To optimize the use of TS7700 storing OAM object data, consider the following suggestions:

- Review the MOUNTWAITTIME parameter when using TS7700 to store OAM object tape data. The default (5 minutes) should probably be increased. Twelve minutes is a better number in case you have to recall a logical volume to read object data and there are other recall requests queued at the time. The TS7700 might need to stage the data back in to cache and this accounts for the extra mount time.
- Review the MAXTAPERETRIEVETASKS and MAXTAPESTORETASKS parameters when using TS7700, because you have more virtual tape drives available.
- There are also other parameters, such as DEMOUNTWAITTIME, TAPEPERCENTFULL and TAPEFULLTHRESHOLD, that might need to be reviewed when using TS7700 to store OAM data.

6.9 Database backups

Using a TS7700 Virtualization Engine as output confers several advantages on database backups. In this section we provide a detailed description of these benefits for database products, such as DB2.

6.9.1 DB2 data

DB2 uses tapes for two purposes: for storing archive logs and for storing image copies. Either one can be created in multiple copies, to be stored both on-site for local recovery and off-site for disaster recovery purposes. To use DB2 tape data with TS7700, we recommend the approaches described here.

Archive logs

DB2 keeps track of database changes in its active log. The active log uses up to 31 DASD data sets (up to 62 with dual logging) in this way: When a data set becomes full, DB2 switches to the next one and automatically offloads the full active log to an archive log.

Archive logs are sequential data sets that are allocated on either DASD or tape. When archiving to tape, a scratch tape volume is requested each time.

Archive logs contain unique information necessary for DB2 data recovery. Therefore, to ensure DB2 recovery, customers usually make backups of archive logs. You can use general backup facilities or DB2's dual archive logging function.

When creating a dual copy of the archive log, usually one is local and the other is for disaster recovery. The local copy can be written to DASD, then moved to tape, using Tape Mount Management (TMM). The other copy can be written directly to tape and then moved to an off-site location.

With TS7700, you can write the local archive log directly inside the TS7700. Avoiding the use of TMM saves DASD space, saves DFSMShsm CPU cycles and simplifies the process. The disaster-recovery copy must be created on non-TS7700 tape drives, so that it can be moved off-site.

The size of an archive log data set varies from 150 MB to 1 GB. The size of a virtual volume on a TS7700 can be up to 4,000 MB, so be sure that your archive log can fit in only one virtual volume. This is because it is suggested to use a single volume when unloading an archive log to tape. The size of a virtual volume on a TS7700 can be up to 12,000 MB, assuming a 3:1 compression ratio.

Tailoring the size and number of active log DASD data sets allows you to obtain an archive log on tape whose size does not exceed the virtual volume size.

Limiting data set size might increase the frequency of offload operations and reduce the amount of active log data on DASD. However, this should not be a problem because TS7700 does not require manual operation and archive logs will stay in the TVC for some time and be available for fast recovery.

One form of DB2 recovery is backward recovery, typically done after a processing failure, where DB2 backs out uncommitted changes to resources. When doing so, DB2 processes log records in reverse order, from the latest back toward the oldest.

If the application being recovered has a large data set and makes only a few commits, you probably need to read the old archive logs that are on tape. When archive logs are on tape, DB2 uses read-backward channel commands to read the log records. Read-backward is a slow operation on tape cartridges processed on real IBM 3480 (if IDRC enabled) and IBM 3490 tape drives. On a TS7700 it is only about 20% slower than a normal I/O because data is retrieved from the TVC, so the tape drive characteristics are replaced by the random access disk characteristics. Another benefit TS7700 can provide to DB2 operations is the availability of up to 256 (single cluster) or 512 virtual drives (Multi Cluster Grid configuration), because DB2 often needs a large number of drives concurrently to perform recovery or backup functions.

Image copies

Image copies are backup copies of table spaces in a DB2 database. DB2 can create both full and incremental image copies. A full image copy contains an image of the whole table space at the time the copy was taken. An incremental image copy contains only those pages of a table space that have changed since the last full image copy was taken. Incremental image copies are typically taken daily, whereas full image copies are typically taken weekly.

DB2 provides the option for multiple image copies. You can create up to four identical image copies of a table space, one pair for local recovery use and one pair for off-site storage.

The size of the table spaces to be copied varies from a few megabytes to several gigabytes. The TS7700 solution is best for small and medium sized table spaces, because you need a higher bandwidth for large table spaces.

When a database is recovered from image copies, a full image copy and the subsequent incremental image copies need to be allocated at the same time. This can potentially tie up many tape drives and, in smaller installations, can prevent other work from being run. With one TS7700, with its 256 virtual drives, this is not an issue.

The large number of tape drives is important also for creating DB2 image copies. Having more drives available allows you to run multiple copies concurrently and use the MERGECOPY DB2 utility without impact. An advisable solution is to run a full image copy of the DB2 databases once a week outside the TS7700 and run the incremental image copies daily using TS7700. (The smaller incremental copy fits better with the TS7700 volume sizes.)

6.9.2 CICS and IMS

Like DB2, both CICS® and IMS use tapes to store logs and image copies of databases.

CICS is only a data communication product, whereas IMS has both the data communication and the database function (IMS-DL/1). CICS uses the same DL/1 database function to store its data.

CICS journals and IMS logs

CICS keeps track of database changes in its journal data sets. IMS keeps track of database changes in its online log data sets. After these data sets become full, both CICS and IMS offload the logs to tape.

CICS and IMS logs are sequential data sets. When offloading these logs to tape, you must request a scratch volume every time.

The logs contain information necessary to recover databases and usually those logs are offloaded, as with DB2, in two copies, one local and one remote. You can write one local copy and then create the second for disaster recovery purposes later, or you can create the two copies in the same job stream.

With TS7700, you can create the local copy directly on TS7700 virtual volumes, then copy those volumes to non-TS7700 tape drives, or to a remote TS7700.

Having a local copy of the logs written inside the TS7700 allows you faster recovery, because the data will stay in the tape volume cache for some time.

When recovering a database, you can complete backout operations in significantly less time with the TS7700, because, when reading logs from tape, IMS uses the slow read backward operation (100 KB/s) on real tape drives. With the TS7700, the same operation is much faster, because the data is read from TVC. Lab measurements do not see much difference between read forward and read backward in a TS7700; both perform much better than on physical drives. The reason is not just that the data is in the tape volume cache, but the TS7700 code also fully buffers the records in the reverse order that they are on the volume when in read backwards mode.

Another benefit TS7700 gives to recovery operations is the availability of up to 256 virtual drives per cluster, allowing you to mount several logs concurrently and therefore to back out the database to be recovered faster.

The IMS change accumulation utility is used to accumulate changes to a group of databases from several IMS logs. This implies the use of many input logs that will be merged into an output accumulation log. Using the TS7700, you can use more tape drives for this function.

Image copies

Image copies are backup copies of the IMS databases. IMS can create only full image copies. To create an image copy of a database, use a batch utility, copying one or more databases to tape.

With the TS7700 you do not have to stack multiple small image copies to fill a tape cartridge. Using one virtual volume per database does not waste space, because the TS7700 then groups these copies into a stacked volume.

IMS, unlike DB2, has a batch function that works with databases through an IMS batch region. If you do not use logs when running an IMS batch region, then in order to recover the database, you must use an image copy taken before running the batch job. Otherwise, you can use logs and checkpoints, which allows you to restart from a consistent database image taken during the batch execution processing. Using TS7700, you can access these image copies and logs at a higher speed.

The TS7700 volume stacking function is the best solution for every database backup, because it is transparent to the application and does not require any JCL procedure change.

6.9.3 Batch data

Other applications that write to tape and benefit from using the TS7700 include:

- VSAM REPRO
- IEBGENER / IEBCOPY / ICETOOL
- DSS data set COPY or DUMP
- DFSMSrmm Tape Copy Tool
- Any other tape copy utility

The amount of data from these applications can be huge if your environment does not use TMM or if you do not have DFSMShsm installed. All such data benefit from using the TS7700 for output.

With TS7700, the application can write one file per volume, using only part of the volume capacity and TS7700 takes care of completely filling the stacked cartridge for you, without JCL changes.

The only thing you must remember is that, if you need to move the data off-site, you must address a device outside the local TS7700, or use other techniques to copy TS7700 data on other movable tapes, as described in 6.3.6, “Moving data out of the TS7700” on page 283.

7

Operation

In this chapter, we describe operational considerations and usage guidelines unique to the IBM TS7700 Virtualization Engine. For general guidance on how to operate the IBM TS3500/3953 Tape Library, refer to:

- *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.
- *z/OS Object Access Method Planning, Installation and* Storage Administration Guide for Tape Libraries, SC35-0427

We provide information about how to operate with the IBM TS7700 Virtualization Engine, in the following main topics:

- User interfaces
- Operational states and modes
- z/OS with system-managed tape
- Tape management
- Expired Volume Management
- Ejecting logical volumes
- Querying the IBM 3953 Library Manager database
- Inventory
- Standalone support
- Additional basic operations
- Messages and displays
- Recovery scenarios

7.1 User interfaces

For successful operation of your TS7700, it is important that you understand its concepts and its components. In this chapter, we combine the components and functions of the TS7700 in two groups: the logical view and the physical view. Each component and each function belongs to only one view.

The logical view is also the host view. From the host allocation point of view, there is only one library to deal with: the Composite Library. A Composite Library has up to 768 virtual addresses for tape mounts. The logical view includes virtual volumes and tape drives.

The host is only aware of the existence of the physical libraries because they are defined through Interactive Storage Management Facility (ISMF) in a z/OS environment. The term *distributed library* is used to denote the physical libraries and TS7700 components that are part of one cluster of the Multi Cluster Grid configuration. The physical view is the hardware view that deals with the hardware components of a Single Cluster or the two clusters of a Multi Cluster Grid configuration, with TS3500 Tape Libraries and TS1120 Tape Drives, with the 3953 and 3494 Library Managers and stacked volumes.

The operator interfaces for providing information about the TS7700 are:

- The Library Manager operator panels deal with information in the Library Manager database. This is part of the physical view. There is limited informations available on the TS7700 in the 3953 and 3494 Library Managers. You can display them directly on the Library Manager console, or you can use the Distributed Console Access Facility (DCAF) for remote operations of the Library Manager console. For more information, see *IBM System Storage 3953 Library Manager Model L05 Operator Guide*, GA32-0448.
- OAM commands are available from the host operator console for information regarding the TS7700 in a standalone and grid environment. This information represents the host view of the components within the TS7700. Other z/OS commands can be used against the virtual addresses. These commands are not aware that the 3490E addresses are part of a TS7700 configuration.
- ETL Specialist functions are available through a Web-based user interface:
  - You can access the different Web interfaces with the Microsoft Internet Explorer Version 6.0 or a fully compatible alternative browser with JavaScript and Java enabled.
  - The following Specialists are available for tape library management:
    - Two Tape Library Specialists (TS3500 Tape Library Specialist and the ETL Specialist). The TS3500 Library Specialist allows for management (configuration and status) of the TS3500 Library. The ETL Specialist is mostly related to Library Manager configuration and monitoring. Panels display information related to the physical view.
    - One Tape Library Specialist that is called *Management Interface* (MI). The TS7700 MI is used to perform TS7700 specific actions, such as inserting logical volumes. The MI operates on the logical view of the TS7700.

Some actions can be done only on the Library Manager operator panel. In this case, we explain the usage of this function on the operator panels. For all other functions we explain the usage with the corresponding Web-based user interfaces. For further information about using the operator panels, refer to the following documentation:

- *IBM System Storage 3953 Library Manager Model L05 Operator Guide,* GA32-0448
- *IBM TotalStorage Library Operator Manual,* GA32-0280
- *IBM System Storage TS3500 Tape Library Operator Guide,* GA32-0560

In this chapter, we focus on the interfaces related to the operation of the TS7700 Virtualization Engine. For detailed information about the TS3500-related operational tasks, refer to the respective operator's guide or to *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

7.1.1 IBM System Storage TS3500 Tape Library Specialist System z

The IBM System Storage TS3500 Tape Library Specialist (TS3500 ETL specialist) interface lets you perform many library functions from the Web.

Figure 7-1 shows the TS3500 Tape Library Specialist Welcome page. You can now choose to view either the physical or the logical library view. Also, you see the status of ALMS. In our example, ALMS is enabled, as always when an IBM 3953 Library Manager is installed.

Figure 7-1 TS3500 Specialist Welcome panel

Figure 7-2 shows a flowchart of the functions that are available depending on the configuration of your TS3500 Tape Library.

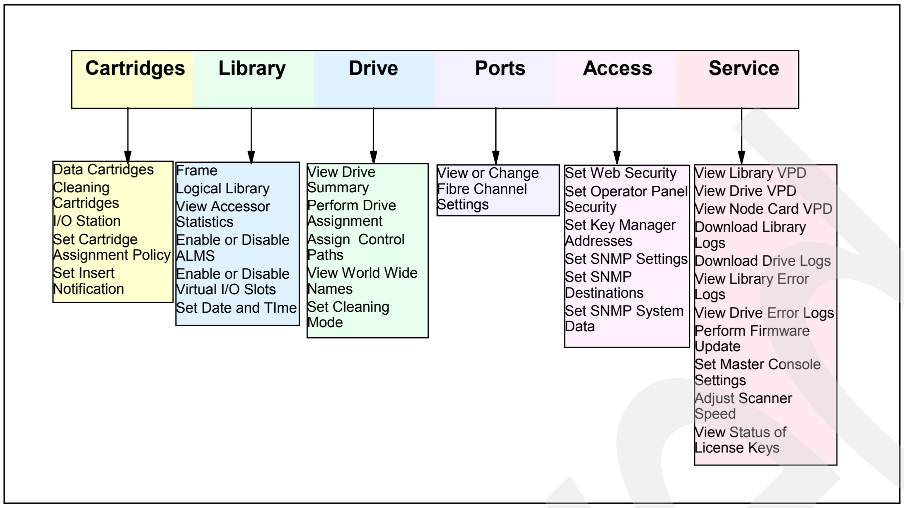

Figure 7-2 TS3500 Specialist functions

7.1.2 IBM Library Manager consoles

The IBM 3953-L05 Library Manager and IBM 3494 Library Manager process all System z host requests and control functions for their associated logical library. The user interface, or operator console, enables you to obtain information about the operation of the Library Manager and attached components. It also instructs the Library Manager to perform specific tasks through the use of the console (display and keyboard with its pointing device).

The user interface recognizes three types of user or authorization levels; general operator, system administrator, and IBM Systems Service Representative (SSR). You can use password protection for the level of authorization. You can optionally password protect some functions of Library Manager, for example, service menus, shutdown, and request inventory upload.

Library Manager Status window

Figure 7-3 and Figure 7-4 on page 308 show the Library Manager Status drop-down menu at the operator console in the IBM 3494 Tape Library and in the IBM 3953 Enhanced Library Controller (ELC). From here you can navigate as usual to the System Summary window, which shows the current Library Manager Mode.

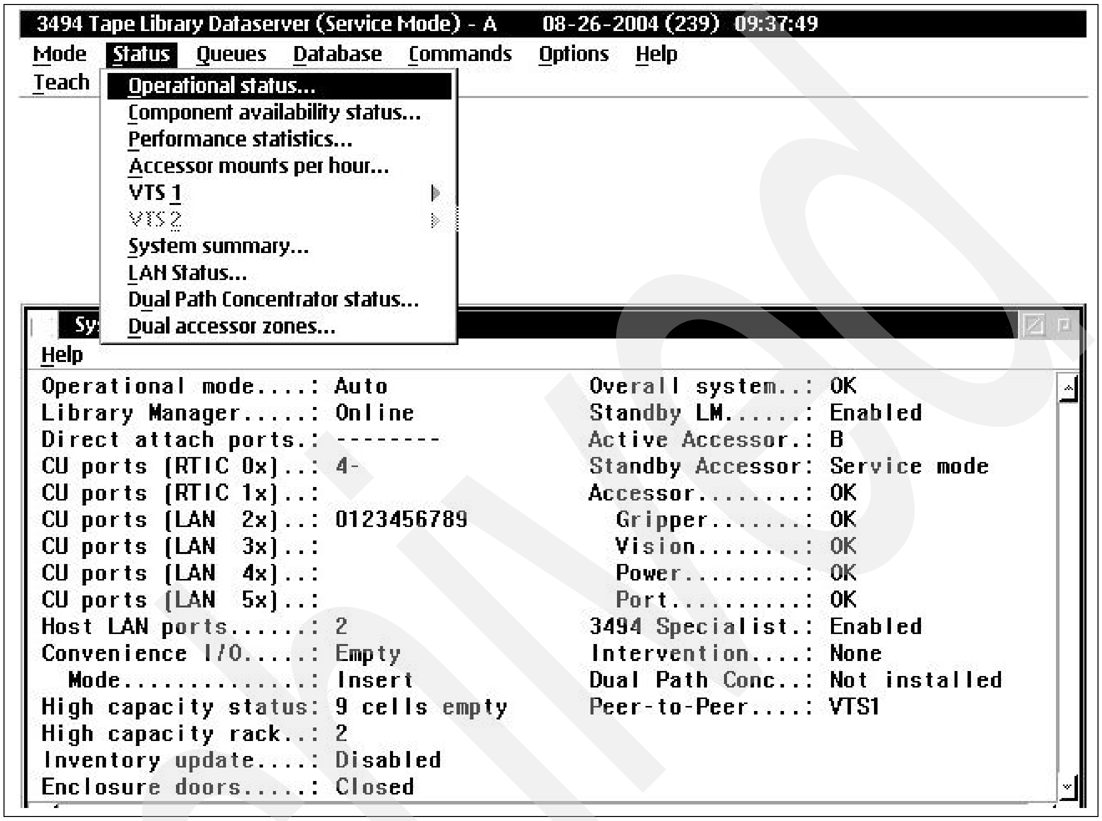

Figure 7-3 Example of IBM 3494 LM Status menu

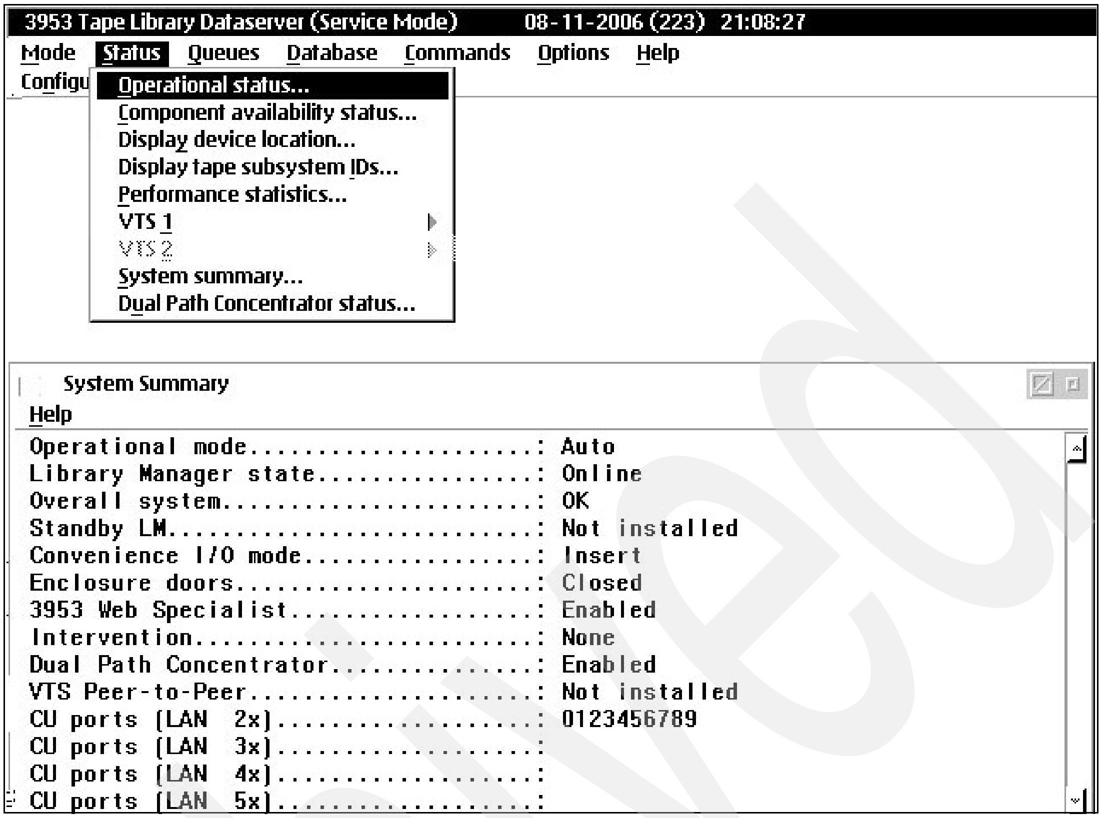

Figure 7-4 Example of IBM 3953 LM Status menu

The System Summary window shows the Library Manager Operational Mode. Note that the Mode is simply a reflection of the state of the physical tape library. For example, the Library Manager will enter Auto mode when the physical tape library enters its Ready state. Otherwise, the Library Manager will remain in Pause mode.

The Operational Mode field shows the actual state of the TS3500 Tape Library, which can be Auto, Auto Pending, Manual, Pause, or Pause Pending:

- **Auto** is displayed when the IBM TS3500 robotics is ready to perform movement. Auto is indicated in this example, showing normal operation. Auto will be shown if the IBM TS3500 is READY.
- **Pause** indicates that the tape library has stopped performing automated operations or that the Library Manager can no longer communicate with the tape library. Pause is displayed when an IBM TS3500 library door is open or if the IBM TS3500 library is not ready for some unknown reason.
- **Auto Pending** is displayed when the tape library has become ready and the Library Manager is uploading the volume inventory.

The Library Manager state is shown in the row following the Operational Status of the tape library (see Figure 7-4). The following states are possible:

- **Online**: The Library Manager and the attached logical libraries are available and accept workload from the attached System z hosts.
- **Offline**: The Library Manager and the attached logical library are not available and do not accept workload from the attached System z hosts.
- **Pending Offline**: Operator switched the library to offline mode. Some active processes are still running, and the offline state is not reached yet.

- **Manual**: Only necessary if a problem in the TS3500 Tape Library prevents automatic operations. The LM then displays the pending operations and gives you detailed instructions, which mounts, demounts, or ejects are requested. When in Manual mode, you follow the instructions on the Library Manager display and confirm as necessary when you complete the instructions. Mounts and dismounts for logical cartridges in a TS7700 are continued by the LM. Only if a recall needs a physical volume, this will be shown as action for the operator.

> **Important:** If you want to use Manual mode operations on the 3953 Library Manager, you must first put the TS3500 Tape Library in Pause mode.

Figure 7-5 shows the Library Manager Mode drop-down menu at the operator console in the IBM 3953; in the IBM 3494 Tape Library you have the same display. The Mode menu displays the current operating state of the Library Manager by showing bullets next to appropriate line items. When you select a new state, the bullets move to the new line items, in this case showing the operating state as Online.

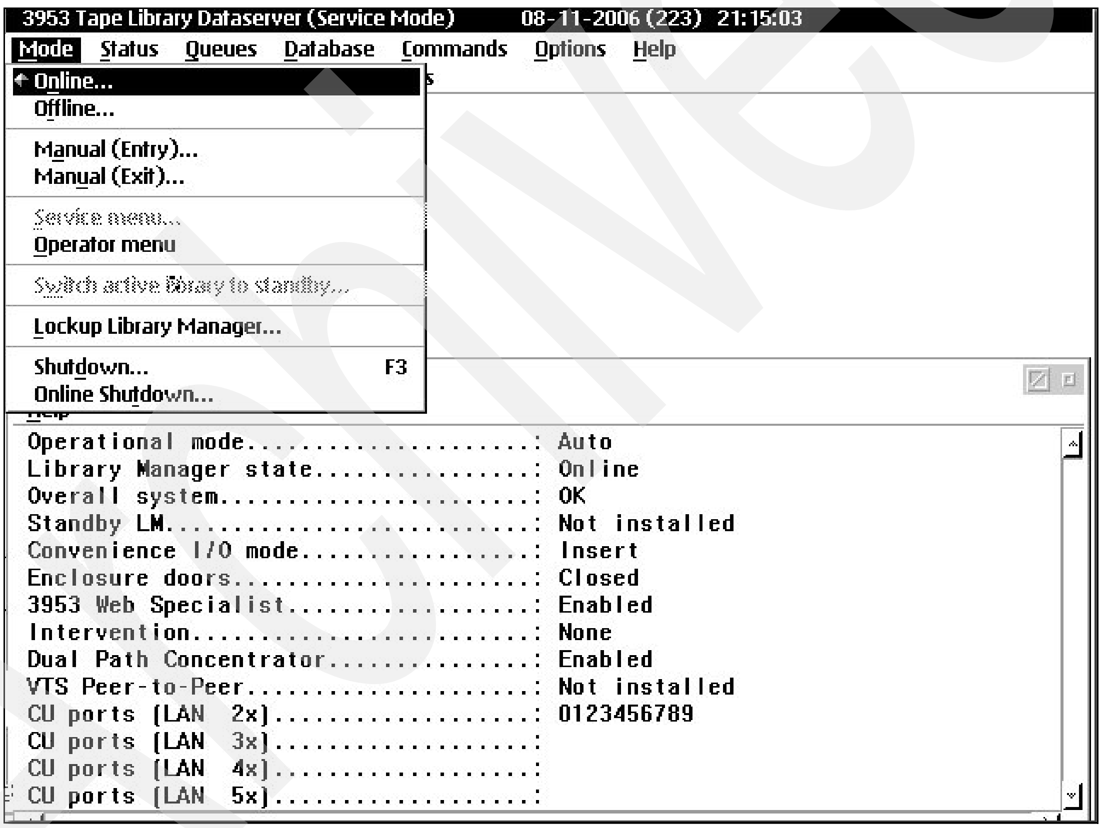

Figure 7-5 Example of IBM 3953 LM Mode menu

You can use this Mode menu to change the operating state of the Library Manager—that is, between Online and Offline, to enter and exit Manual mode of operation or the Service menu, or to shut down the Library Manager.

> **Note:** You cannot use this menu to change the operational mode in a TS3500 between Auto and Pause. These are TS3500 Tape Library functions that affect the whole physical library and not just the System z logical library. Use the operator panel on the IBM TS3500 itself to Pause the library.

7.1.3 The IBM System Storage Enterprise Automated Tape Library Specialist

The IBM TotalStorage Enterprise Automated Tape Library (ETL) Specialist is the Web-based interface of the IBM 3953 Manager and IBM3494 Tape Library. Figure 7-6 shows the Home Page of the ETL Specialist.

If you are not receiving Tape Library Specialist messages, turn off all pop-up blocking software when using the Specialist or make sure that the software allows pop-ups from the URL of the Specialist. This type of software blocks important messages from the Specialist, so it is important that it is turned off or disabled while using the Specialist.

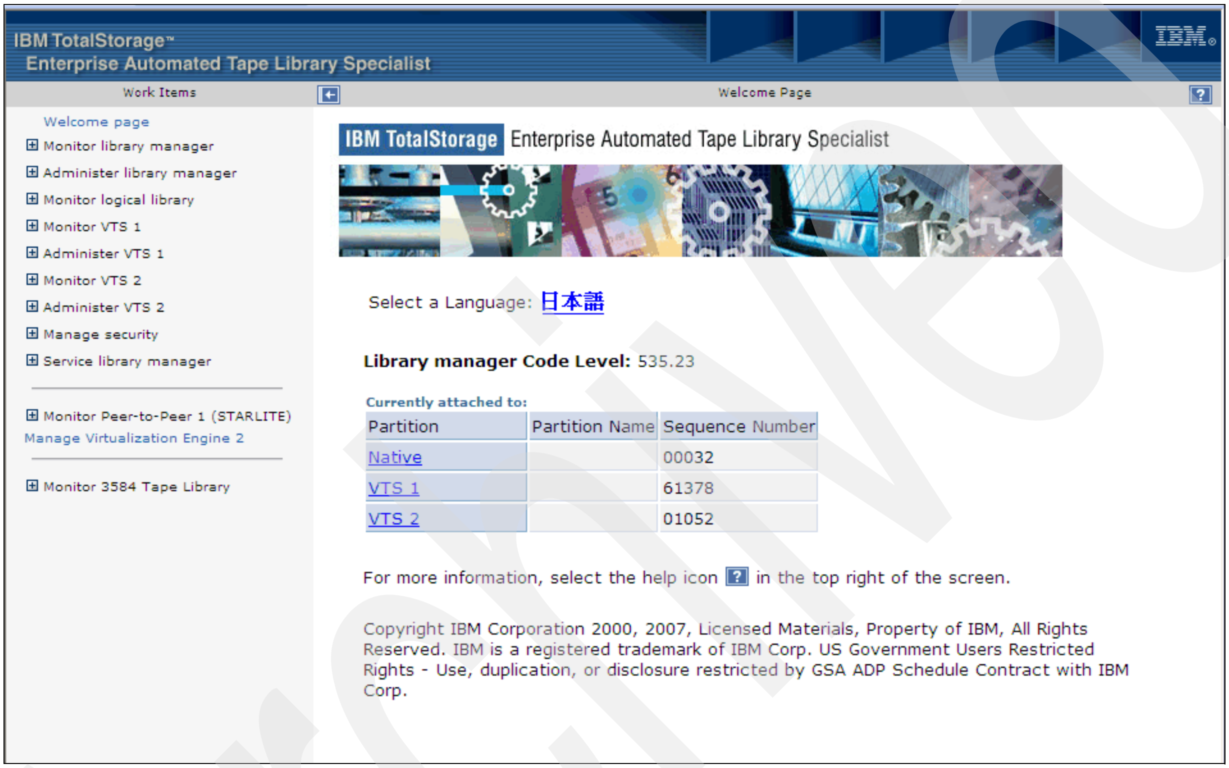

Figure 7-6 IBM TotalStorage Enterprise Automated Tape Library Specialist

> **Note:** The ETL Specialist panels show VTS for any virtual tape solution managed by this Library Manager. This VTS can be a TS7700 Virtualization Engine, or an IBM 3494 Virtual Tape Server Model B10 or B20.

Using the Specialist, you can access information such as current Library Manager status and limited TS7700 statistics from your Web browser by connecting to the Web server on the Library Manager PC, and also perform numerous administrative functions. The Web server serves HTML pages to a remote Web browser over your LAN connection. The Specialist allows multiple active server connections at the same time (service and several user connections).

The Specialist provides display functions and most of the control functions available from the Library Manager console. The following are available on the Specialist:

- A home page
- A set of Library Manager pages
- A set of Library Partition pages
- A set of TS7700 pages
- A set of Security pages

You must enable the Specialist panels at the Commands drop-down menu on the Library Manager console. When enabled, the ETL Specialist provides a set of monitoring and control functions for up to two VTS or TS7700 systems and for the native System z-attached TS1120 Tape Drives, as shown in Figure 7-7.

The left side shows the monitoring tasks; the right side shows the control functions you can perform through the ETL Specialist.

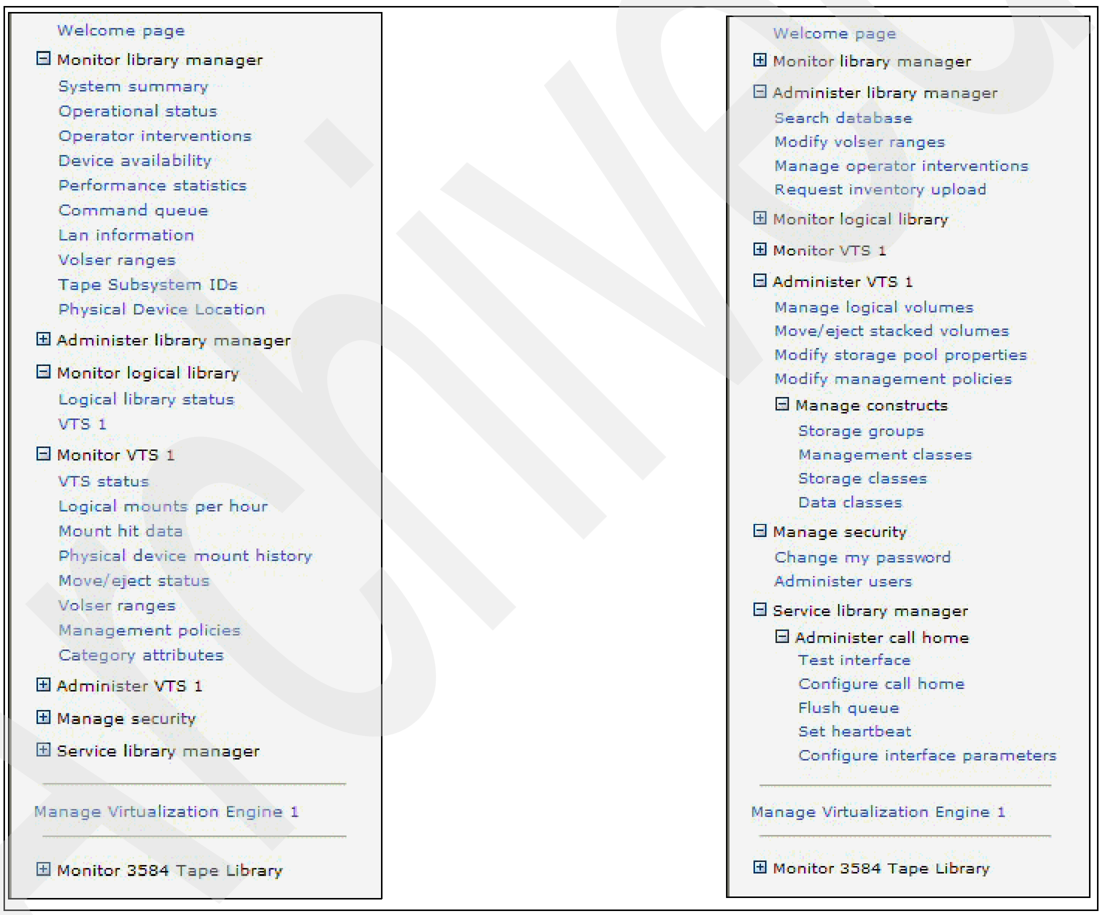

Figure 7-7 Work items from the ETL Specialist Welcome page

7.1.4 Simple Network Management Protocol (SNMP) monitoring

The IBM 3494 and the TS3500 Tape Library with the 3953 Library Manager provide a standard TCP/IP protocol called Simple Network Management Protocol (SNMP) to send alerts (called *SNMP traps*) over a TCP/IP LAN network to one or more SNMP monitoring stations. These monitoring stations, along with other user-supplied software, can alert operations staff to possible problems or operator interventions that occur at the TS3500 Tape Library, at 3953 Library Manager, or at IBM3494 Tape Library.

To activate this SNMP traffic, you must enable it in the IBM 3494 Library Manager panel or in the TS3500 Tape Library with the 3953. See Figure 7-8 to enable an SNMP trap on the 3494 or 3953 Library Manager; the panel does not differ from the 3953 to 3494 Library Manager.

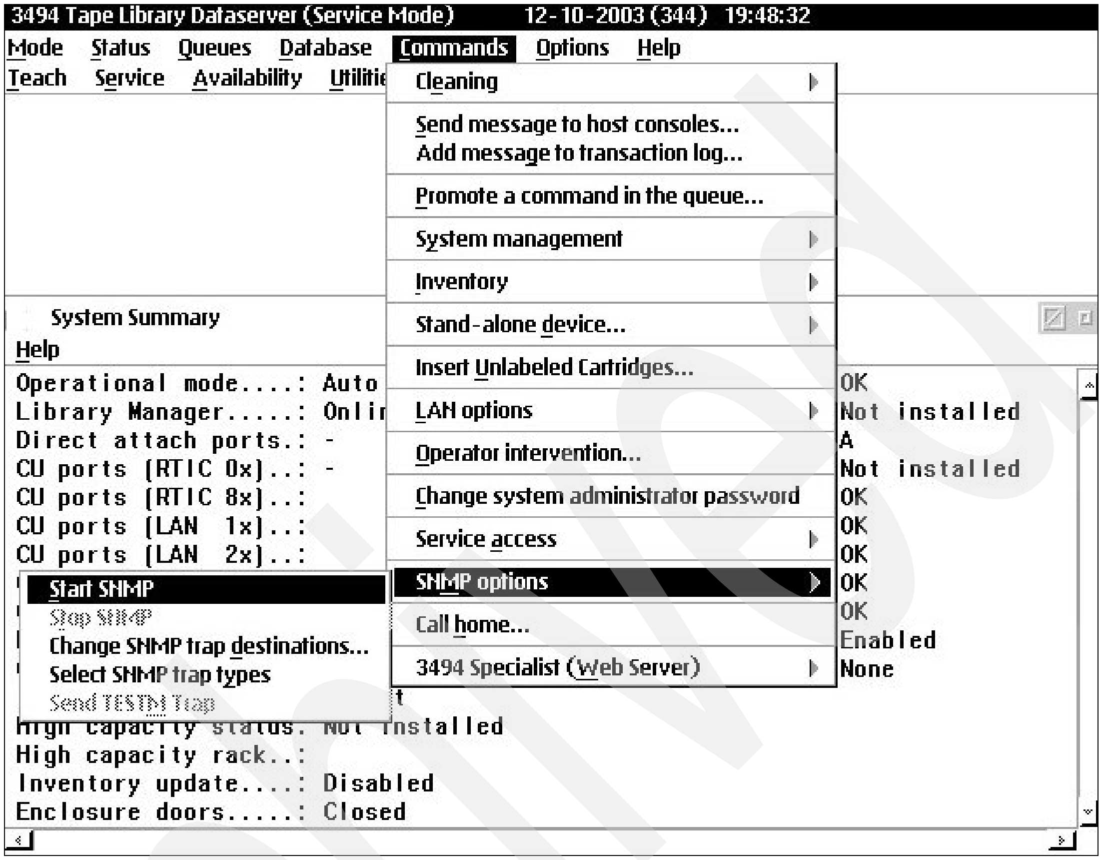

Figure 7-8 3494 SNMP panel

For further information how to select and activate the SNMP function, refer to:

- *IBM System Storage Tape Library Guide for Open Systems*, SG24-5946
- *IBM TotalStorage Automated Tape Library (3494) Operator Guide*, GA32-0449
- *IBM System Storage 3953 Library Manager Model L05 Operator Guide*, GA32-0558
- *IBM System Storage TS3500 Tape Library Operator Guide*, GA32-0560

7.1.5 Call home

The IBM TotalStorage 3494 Library family, and IBM System Storage TS3500 Tape Library include several external interfaces not directly associated with data paths. Rather, these interfaces are associated with library control, service, and status information. They support customer interaction or feedback, as well as attachment to IBM remote support infrastructure for product service and support.

The call home function generates a service alert automatically when a problem occurs with one of the following components:

- IBM 3953 Model L05 Library Manager
- IBM 3494 Library Manager and Robotics
- IBM TS3500 Tape Library
- IBM 3592 Model J70 and TS1120 Model C06 controller
- IBM TS7700 Virtualization Engine

Error information is transmitted to the IBM System Storage TS3000 System Console for service, and then to the IBM Support Center for problem evaluation; the IBM Support Center can dispatch an IBM SSR to the installation. Call home can send the service alert to a pager service to notify multiple people, including the operator. The IBM SSR can deactivate the function through service menus if you require that.

7.1.6 Electronic Customer Care

Here we give an overview of the communication path between a customer's purchased subsystems and IBM support. The method that has been used for call home since the introduction of the TotalStorage System Controller, or TSSC, previously known as the *TotalStorage Master Console*, or TSMC, was related to TSSC microcode V4.1.14 and before. Later a new method was introduced with the release of the new microcode V4.2.0.

The TSSC console uses analog phone lines and a modem to connect to IBM Remote Technical Assistance Information Network, better known as RETAIN®. The code running in RETAIN then decides what to do with the information. A Problem Management Record (PMR) will be opened in the case of a problem. After the PMR is created, RETAIN will do auto consult with IBM Knowledge Based Systems (RKBS) in order to add in any already known information about that type of problem. Finally, in the case where RETAIN detects that the call home is a data package, the data is forwarded to dumpers which move the data to an IBM Internal Server called DFS Cell. From there it is pulled into the IBM Tape Call Home Database.

In Figure 7-9 on page 314 we describe Electronic Customer Care call home. We define three security zones:

- *Red Zone* is defined as the customer data center. This zone is where the IBM Storage subsystem and the TSSC reside.
- *Yellow Zone* is defined as the open Internet. This is open to all outside communication.
- *Blue Zone* is defined as IBM Support. This sits inside the IBM intranet and is only accessible to IBM-authenticated users.

In Electronic Customer Care, the TSSC will use either a modem connection or a broadband connection to connect to an electronic customer care gateway located in the Yellow Zone. This is an IBM-managed server used as a relay for IBM support information. This gateway then forwards the information to the Inter Enterprise Process Directory, or IEPD. IEPD is located within IBM support in the Blue Zone. IEPD then checks the problem report that has been submitted and monitors what system the report has come from. The system's credentials are then checked to make sure a valid support contract exists. If the credentials pass for this system, it will consult Technical Services Knowledge Base System (TSKBS) for known information to add to the problem report. The problem report will then be used to open a Problem Management Record (PMR) in IBM RETAIN.

In the event of a data call home, the data is sent from the same TSSC connection to an IBM-managed server located in the Yellow Zone known as Testcase. Dumpers monitor this server for new information. When they see this new information, they move the data package to DFS space where it gets pulled into the RMSS Call Home Database.

Initial ECC handshaking communication utilizes HTTP communication. After initial handshaking, all information is sent using secure HTTPS communication. Because all traffic is outbound in nature and only uses these two ports, it should work through most customers' firewall without additional firewall rules being added by the customer.

Problem Reporting communication is then sent to IEPD, which consults Technical Services Knowledge Base System TSKBS for the systems listed, and opens a PMR in RETAIN. Of course all of these details are for knowledge purposes only. After initial setup, these details are utilized in the background, without knowledge by the user.

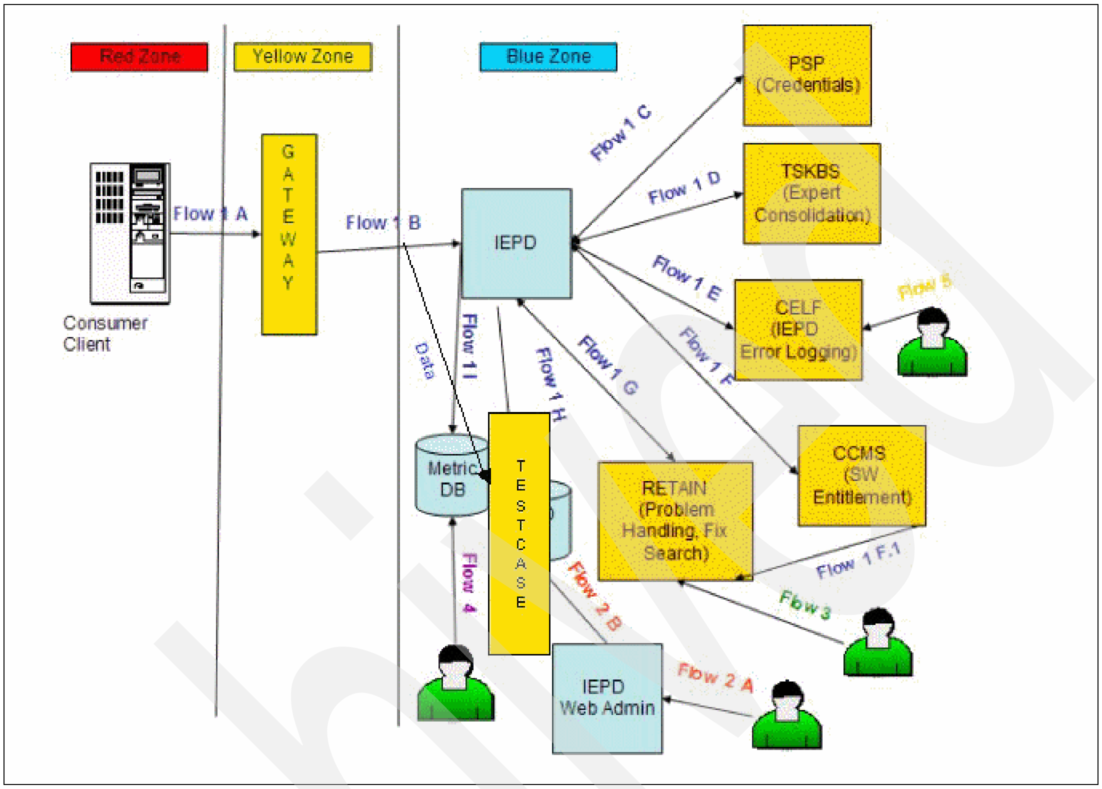

Figure 7-9 Electronic Customer Care

All in-bound connection by IBM service personnel will still be through dial-in modem connection. Nothing pertaining to this process has changed.

The outbound communication associated with ECC call home can be through a customer-supplied Ethernet connection, a modem, or both in the form of failover. The local subnet LAN connection between the TSSC and the attached subsystems remains the same. It is still isolated without any outside access.

ECC adds another Ethernet connection to the TSSC, bringing the total number to three. These connections are labeled External Ethernet Connection, which is the ECC Interface; Grid Ethernet Connection, which is used for the TS7740 Autonomic Ownership Takeover Manager (AOTM); and the Internal Ethernet Connection, used for the local attached subsystems subnet. All of these connections are set up using the Console Configuration Utility User Interface located on the TSSC.

7.2 Virtualization Engine TS7740 Management Interface

The IBM TS7740 Virtualization Engine Management Interface (MI) is the primary interface to monitor and administer the TS7700 Virtualization Engine.

7.2.1 Connecting to the MI

To connect to the IBM TS7700 Virtualization Engine management interface:

1. The TS7700 must first be installed and configured.
2. In the address bar of a supported Web browser, enter `http://` followed by the virtual IP assigned during installation. Press the Enter key on your keyboard or Go button on your Web browser. The virtual IP is one of three customer IP addresses given during installation; see 3.2.3, “TCP/IP configuration considerations” on page 93.

   If you are using your own Name Server, where you can associate a name with the virtual IP address, you can use the name instead of the hardcoded address for reaching the MI.
3. The login page for the management interface will load, as shown in Figure 7-10. The default login name is *admin* and the default password is *admin*.

Figure 7-10 TS7700 Management Interface login

After entering your password, you see the first Web page presented by the MI, the Virtualization Engine Grid Summary shown in Figure 7-11.

Figure 7-11 MI Virtualization Engine Grid Summary

Figure 7-11 shows a graphical summary of the TS7700 grid. Note that this page does not imply that you have a Multi Cluster Grid. Each TS7740 is considered a grid. In our example, you see a Three-Cluster Grid configuration.

In the center of the picture is the grid, which is composed of one or more clusters shown connected to the grid. The cluster that you are currently connected to is highlighted by a blue background with a light blue border. The clusters are also shown connected to the hosts.

The connection status between the grid and cluster along with the cluster and host is indicated by an icon tied to a state that can be determined from the Legend on the page. To get more details on the cluster, click the picture of the cluster to bring up the Cluster Summary page. You can go back to the Virtualization Engine Grid Summary page by clicking the picture of the grid.

For pages applicable to the entire grid, there is a link to the Grid Summary at the top of the page.

Customer Network Interfaces

The TS7700 Virtualization Engine provides two Ethernet connections to the customer's network for access to the TS7700 Management Interface. The following user security policy settings are available:

- **Disable account expiration**: User accounts will not expire.
- **Enable account expiration**: Allow accounts to expire after a set number of days from when their passwords are set. The amount of days can be set in the "Password expires after" text box.

- **Disable account lockout**: User accounts will not be locked out when a wrong password is entered.
- **Enable account lockout**: The account will be allowed a certain number of failed logon attempts before locking the account out of logging into the management interface. The number of failed attempts is entered in the “Number of successive incorrect password attempts allowed” text box.

Note: To unlock an account, an administrator can modify the user account and change the password from the Manage Users page.

The 3494 Enterprise Automated Tape Library and the 3593-F05 use a Library Manager to control library operations. The Library Manager can optionally provide network connectivity on an external customer-specified network. The 3494 Library Manager connectivity can be used for establishing connections to open-systems hosts through a LAN or Serial host connection. Library Managers provide the Automated Tape Library Specialist, also called Web Specialist. The Web Specialist is a Web-based monitoring and management interface to the Library Managers. To enable this functionality in an IBM 3494 Library Manager, either 3494 Feature Code 5219 (for Token-Ring connections) or Feature Code 5220 (for Ethernet connections) must be properly installed and configured in the Library Manager.

The 3953-L05 comes standard with an Ethernet connection for the Specialist. In environments where the tape configuration is separated from the LAN-attached hosts or Web clients by a firewall, these are the only ports that must be opened on the firewall, all others might be closed; see Table 7-1 for reference.

Table 7-1 Customer Network Interface firewall

Function	Port	Direction (From Library)	Protocol
Library Operations	3494	Bidirectional	TCP/IP
ETL Specialist	80	Inbound	TCP/IP
SNMP Traps	161/162	Bidirectional	UDP/IP
Encryption Key Manager	1443	Outbound	SSL
Encryption Key Manager	3801	Outbound	TCP/IP

For more information about the customer network interface, refer to:

- *IBM System Storage TS3500 Tape Library Operator Guide,* GA32-0560
- *IBM System Storage TS3500 Tape Library Introduction and Planning Guide,* GA32-0559
- *IBM System Storage TS1120 Tape Drive and Controller Operator Guide,* GA32-0556
- *IBM System Storage TS1120 Tape Drive and Controller Introduction and Planning Guide,* GA32-0555
- *IBM System Storage 3953 L05 Library Manager Operator Guide,* GA32-0558
- *IBM System Storage 3953 Tape System Introduction and Planning Guide,* GA32-0557
- *IBM TotalStorage 3494 Tape Library Introduction and Planning Guide,* GA32-0448
- *IBM Encryption Key Manager component for the Java platform Introduction, Planning, and User's Guide,* GA76-0418
- *IBM Virtualization Engine TS7700 Series Introduction and Planning Guide,* GA32-0567

7.2.2 Using the Management Interface

Before we explain in detail the tasks that you can perform from the IBM TS7700 Virtualization Engine Management Interface, we describe some common page and table elements.

Standard page elements

The following elements might be present on a page and allow you to use the management interface more effectively:

- Refresh button

  Clicking **Refresh** at the top of the page refreshes the page's data. The date and time of the page's last refresh are displayed next to Last refresh.

- Current cluster being accessed

  For cluster-specific pages, the current cluster being accessed is displayed at the top of the page. Clicking the graphic or hyperlink for the cluster opens the Cluster Summary page. If more than one cluster is forming a grid with another, there will be an option to change the active cluster by selecting a new cluster in the drop-down of the current cluster being accessed.

Standard table elements

Tables in the management interface contain a number of features that allow for a variety of filtering and sorting options. These options are selectable through the Select Action drop-down and through buttons above the table header.

Edit Sort	Selecting this option allows you to select a primary, secondary and tertiary sorting category and whether the sort will be ascending or descending.
Clear All Sorts	Removes any current table sorting.
Show/Hide Filter Row	Shows or hides the Filter option available under a table header. Selecting **Filter** displays the option to filter for that table column by entering a Condition and Text.
Clear All Filters	Clears any current filtering in place.

Note: In order to support Japanese input, a Japanese front-end processor needs to be installed on the computer where a Web browser is accessing the Management Interface.

In the following sections, 7.2.3, "Health and Monitoring" on page 318 through 7.2.9, "Service & Troubleshooting" on page 364, we explain the MI panels that enable you to control and monitor the TS7700.

For details about the Performance and Statistics selections, refer to 8.4.1, "Performance and statistics" on page 431.

7.2.3 Health and Monitoring

This section of the TS7700 Management Interface gives you the choice to query the current Cluster status as well acquire logical and physical volume details. You can access the following selections if you navigate within the TS7700 Management Interface to the Health and Monitoring section by opening with a mouse click.

We include the following topics in this section:

- Cluster Summary
- Logical Volume Details
- Physical Stacked Volume Details
- Physical Drives
- Physical Media Inventory
- Tape Volume Cache

Cluster Summary

Use this page (Figure 7-12) for a graphical summary of the IBM Virtualization Engine TS7700 cluster. It shows a graphical summary of the TS7700 cluster selected from the Grid Summary page. The graphical description shows the cluster, its key components such as the vNode, hNode, Tape Library and drives, and their connectivity.

The health of the system is checked and updated automatically at times determined by the TS7700. In order to populate the summary with updated health status, you can click **Scan Cluster**. This operation can take some time. Its status can be tracked on the Cluster Operation History page as described later in “Cluster Operation History” on page 330.

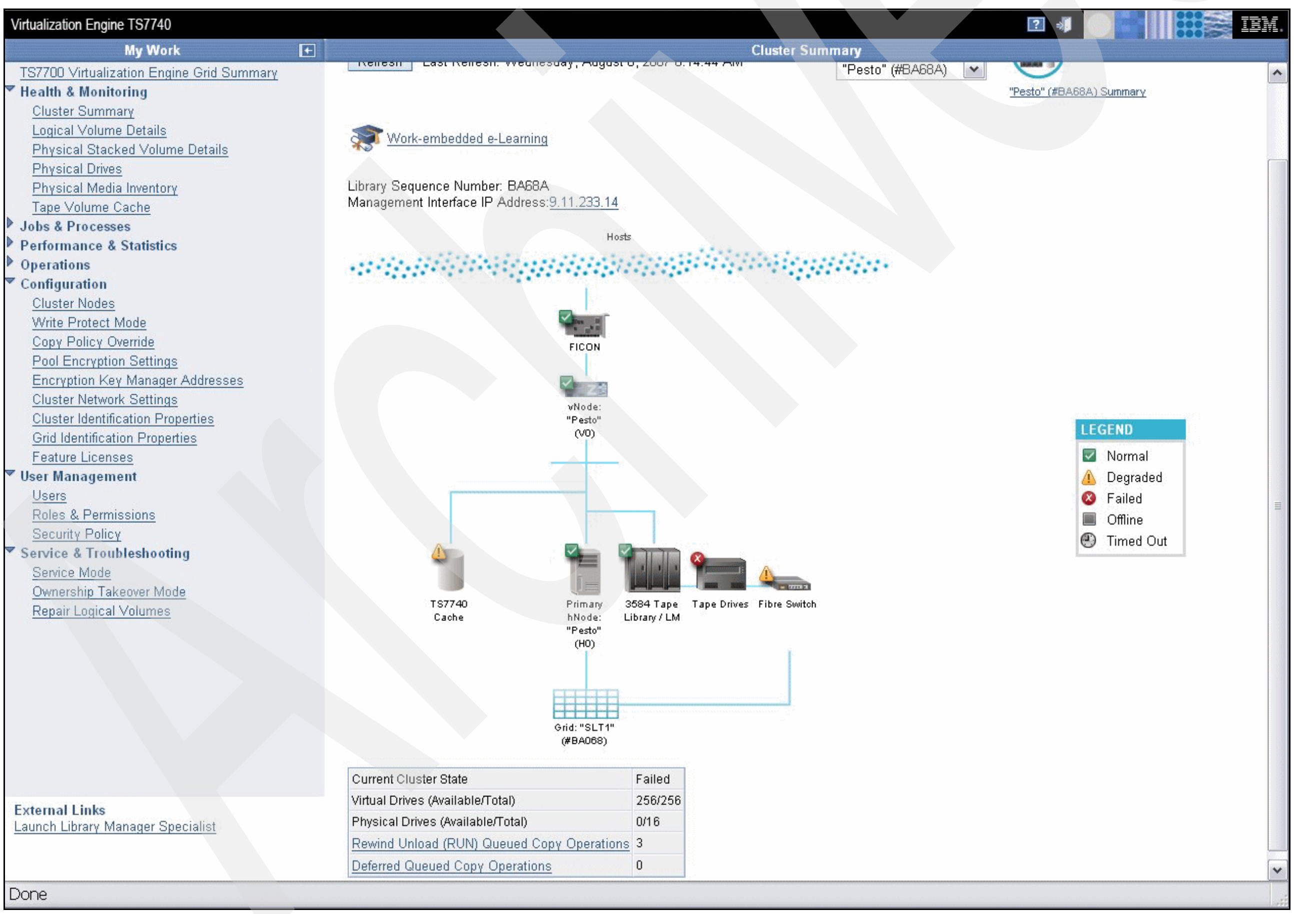

Figure 7-12 TS7700 Management Interface Cluster Summary

The connected components have colored borders along with pictures to indicate their status, which is indicated by Legend on the page. You can click failed or degraded components, except nodes, to obtain information about how to correct this state. To get more details on either the vNode or hNode, click their picture, and the Cluster Nodes page will be displayed.

The components of the cluster are also connected with either a solid or dashed line. A dashed line occurs if a component is offline.

At the top of the page, the Library Sequence Number and Management Interface IP Address are displayed. Clicking the IP address displays the Cluster Network Settings page where network settings for the cluster can be viewed or altered; see “Cluster Network Settings” on page 352.

The following components are displayed in the Cluster Summary:

Host adapter	Adapter connecting the host to the TS7700 cluster
vNode	vNode on the TS7740 Virtualization Engine
Fibre switch	Fibre switch connecting the tape library and the TS7700 cluster
TS7740 cache	Disk cache located on the TS7700 cluster
Primary hNode	Primary hNode located on the TS7700 cluster
Tape Library	Tape library connected to the TS7700 cluster
Tape Drives	Physical drives located in the tape library

Components can be in a normal, failed, degraded, or offline state as indicated in the component's graphic. A degraded state indicates that a component is working but one of its redundant parts has stopped functioning.

The following information can be found at the bottom of the page (Figure 7-12 on page 319) below the Cluster diagram.

- **Current Cluster State** shows the general state the cluster is in. Possible values are:
  - Normal
  - Degraded
  - Failed
  - Offline
  - Service
  - Service Prep
  - Timed Out
- **Virtual drives (Available/Total)** shows the number of virtual drives available for this TS7740 Virtualization Engine and the total amount for the TS7740.
- **Physical drives (Available/Total)** shows the available and total physical drives for this cluster.
- **RUN Queued Copy Operations** shows the number of rewind-unload copies queued. Clicking the available hyperlink will display the Logical Volume Incoming Copy Queue page. See Figure 7-13 on page 321.
- **Deferred Queued Copy** shows the number of deferred copies queued. Clicking the available hyperlink will display the Logical Volume Incoming Copy Queue page. See Figure 7-13.

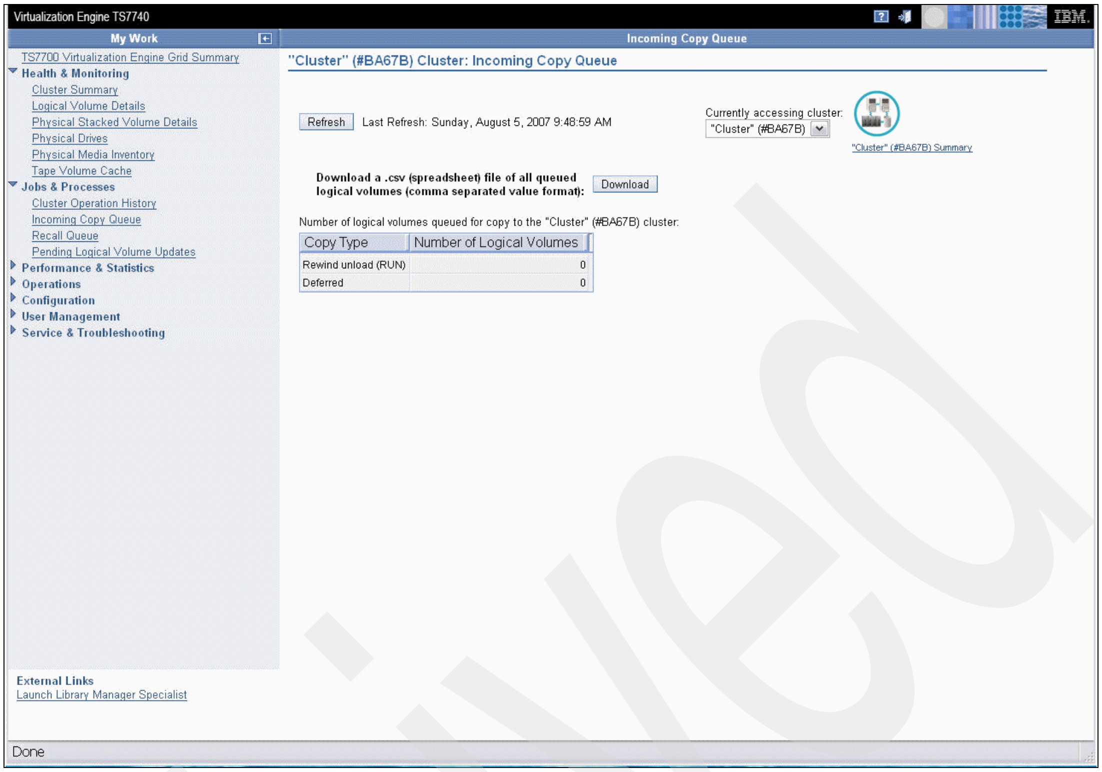

Figure 7-13 Incoming Copy Queue

Logical Volume Details

Use this page (Figure 7-14 on page 322) for getting detailed information about a logical volume in the IBM Virtualization Engine TS7700 Grid.

Whenever you need to know specific details about a logical volume, such as its copy status, the size, or the stacked volumes where it resides, this page allows you to enter the logical volume serial number and obtain the desired information, as detailed in Figure 7-14 on page 322.

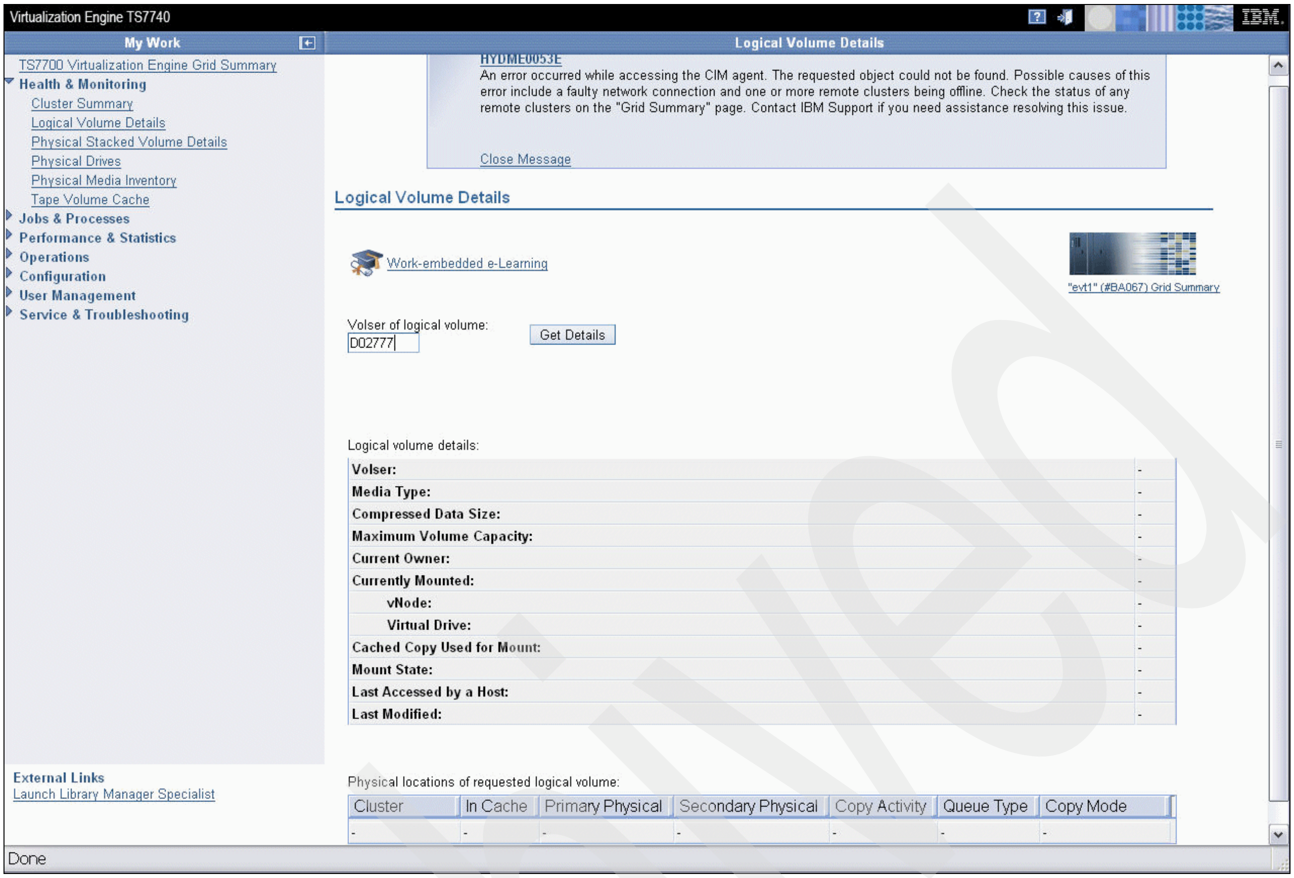

Figure 7-14 TS7700 Management Interface Logical Volume Details

To obtain details about a logical volume, enter the logical volume's VOLSER in the available text field, then select **Get Details**. The VOLSER must be 6 characters in length.

The following information is displayed when details for a logical volume are retrieved:

- **VOLSER**: The 6 character VOLSER identifier of the logical volume
- **Media Type**: The Media type of the logical volume. Possible values are:
  - Cartridge System Tape (400 MB)
  - Enhanced Capacity Cartridge System Tape (800 MB). The use of 1,6 Gb, 2 Gb and 4 Gb is also possible based on defining ECCST tapes. This must then be set up by using Dataclass (DC) in SMS.
- **Compressed Data Size**: The compressed file size of a logical volume expressed in bytes (B), kilobytes (KB), megabytes (MB), or gigabytes (GB)
- **Maximum Volume Capacity**: The configured size of the logical volume based upon media type or the volume override setting, expressed in bytes (B), kilobytes (KB), megabytes (MB), or gigabytes (GB)
- **Current owner**: The cluster that currently owns this logical volume
- **Currently Mounted**: Whether or not the logical volume is currently mounted on the owning cluster
- **vNode**: The Node the logical volume is mounted on
- **Virtual Drive**: The ID of the virtual drive the logical volume is mounted on. Note that this ID is not the device address used by the host.

- **Cache Copy Used for Mount**: The cluster name of the cache chosen for I/O operations for mount, based on consistency policy, volume validity, residency, performance, and cluster mode.
- **Mount State**: The current mount state of the logical volume. Possible values are:

  Mounted The volume is mounted.
  Mount Pending A mount request has been received and is in progress.
  Recall Queued A mount request has been received and requires a recall.
  Recalling A mount request has been received and a recall from physical tape is in progress.

- **Last Accessed by a Host**: The date and time the logical volume was last accessed by a host. Date and time are obtained from the current owner of the logical volume.
- **Last Modify**: The date and time the logical volume was last modified by a host. Date and time are obtained from the current owner of the logical volume.

The following information is available for all physical locations of the logical volume:

- **Cluster**: The cluster location of the logical volume copy.
- **In Cache**: Whether the logical volume is in cache for this cluster.
- **Primary Physical**: Physical volume that contains the specified logical volume. Click the VOLSER hyperlink to open the Physical Stacked Volume Details page for this physical volume. A value of “-” means that no secondary physical copy is to be made.
- **Secondary Physical**: Secondary physical volume that contains the specified logical volume. Click the VOLSER hyperlink to open the Physical Stacked Volume Details page for this physical volume. A value of “-” means that no secondary physical copy is to be made.
- **Copy Activity**: Status information about the copy activity of the logical volume copy. Possible values are:

  Complete A consistent copy exists at this location.
  In Progress A copy is required and currently in progress.
  Required A copy is required at this location but has not started or completed.
  Not Required A copy is not required at this location.
  Reconcile Pending updates exist against this location’s volume. The copy activity is updated after the pending updates are resolved.

- **Queue Type**: The type of queue as reported by the cluster. Possible values are:

  RUN The copy will occur during rewind-unload and before the rewind-unload operation completes at the host.
  Deferred The copy will occur some time after the rewind-unload operation at the host.
  Immediate Deferred A RUN copy that has been moved to the deferred state due to copy time-outs or TS7700 Grid states.

- **Copy Mode**: The copy behavior of the logical volume copy. Possible values are:

  RUN The copy will occur during rewind-unload and before the rewind-unload operation completes at the host.
  Deferred The copy will occur some time after the rewind-unload operation at the host.
  No Copy No copy will be made.
  Exist A copy exists at this location although No Copy is intended. A consistent copy existed at this location at the time the logical volume was mounted. After the volume is modified, the Copy Mode of this location changes to No Copy.

Physical Stacked Volume Details

Use this page (Figure 7-15) for getting detailed information about a physical stacked volume in the IBM Virtualization Engine TS7700.

To obtain details about a physical stacked volume, enter the volume's VOLSER identifier in the available text field, then select **Get Details**. The VOLSER identifier must be 6 characters in length.

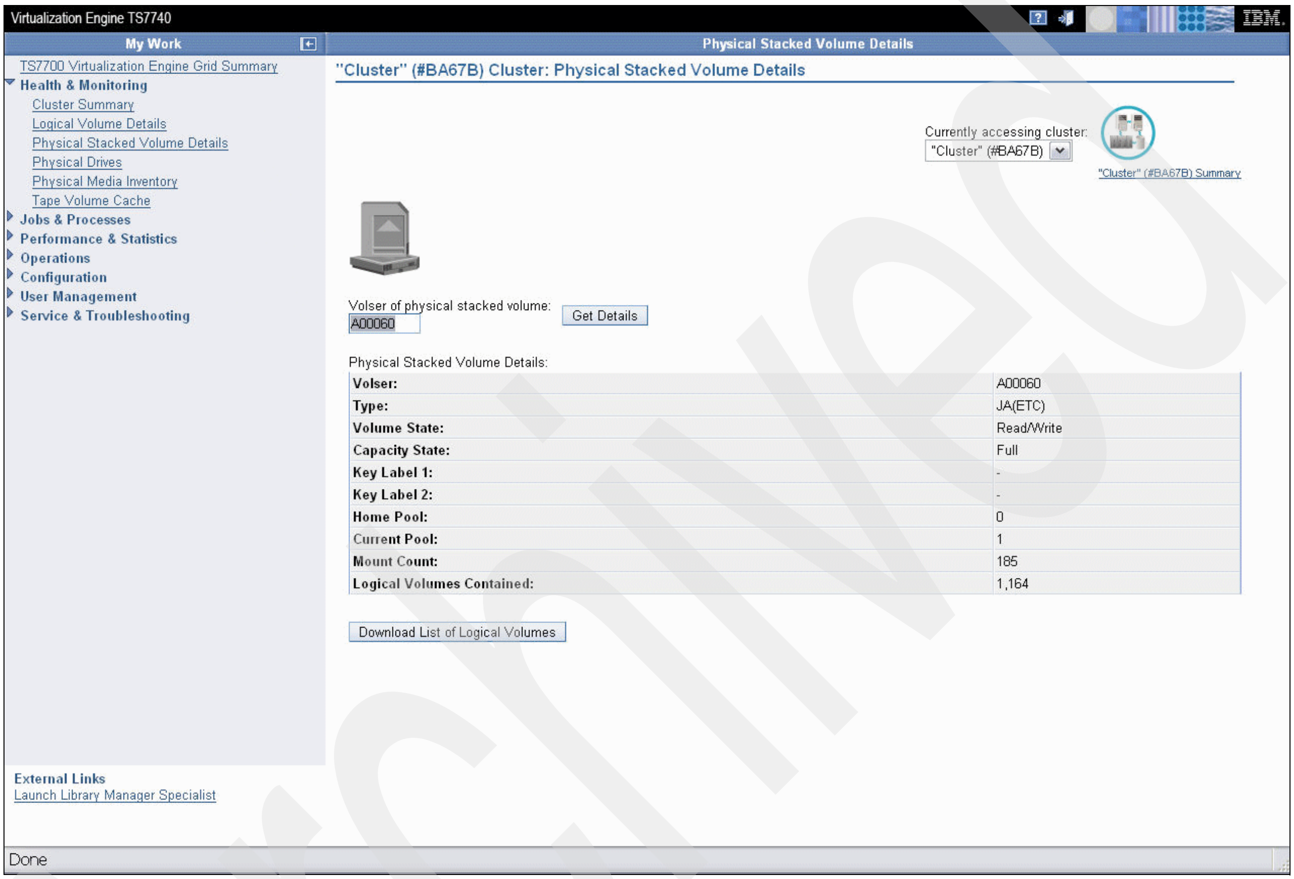

Figure 7-15 TS7700 Management Interface Physical Volume Details

The following information is displayed when details for a physical stacked volume are retrieved:

- **VOLSER**: The 6 character VOLSER identifier of the physical stacked volume
- **Type**: The media type of the physical stacked volume. Possible values are:

JA (ETC)	Enterprise Tape Cartridge
JB (EEETC)	Enterprise Tape Cartridge Extended Data
JJ (EETC)	Enterprise Economy Tape Cartridge

- **Volume state**: Possible values are:

Read-Only	The volume is in a read only state.
Read/Write	The volume is in a read/write state.
Unavailable	The volume is in use by another task or is in a pending eject state.
Destroyed	The volume is damaged and unusable for mounting.

- **Capacity state**: Possible values are empty, filling, and full.

- **Key Label 1**: This field identifies the encryption key label that is associated with a physical volume when it is used as a scratch volume for the pool. Key Label 1 if the encryption of the physical volume is enabled.
- **Key Label 2**: This field identifies the encryption key label that is associated with a physical volume when it is used as a scratch volume for the pool. Key Label 2 if the encryption of the physical volume is enabled.
- **Home Pool**: This field identifies the pool number the physical volume was assigned to when it was inserted into the library, or the pool it was moved to through the Library Manager Move/Eject Stacked Volumes function.
- **Current Pool**: This field identifies the current storage pool the physical volume resides in.
- **Mount Count**: The number of times the physical volume has been mounted since being inserted into the library.
- **Logical Volumes Contained**: The number of logical volumes contained on this physical stacked volume.

To obtain details about the specific logical volumes on the physical stacked volume, click **Download List of Logical Volumes** under the table. A dialog box for the download appears and you can download the provided information as a CSV File (Figure 7-16).

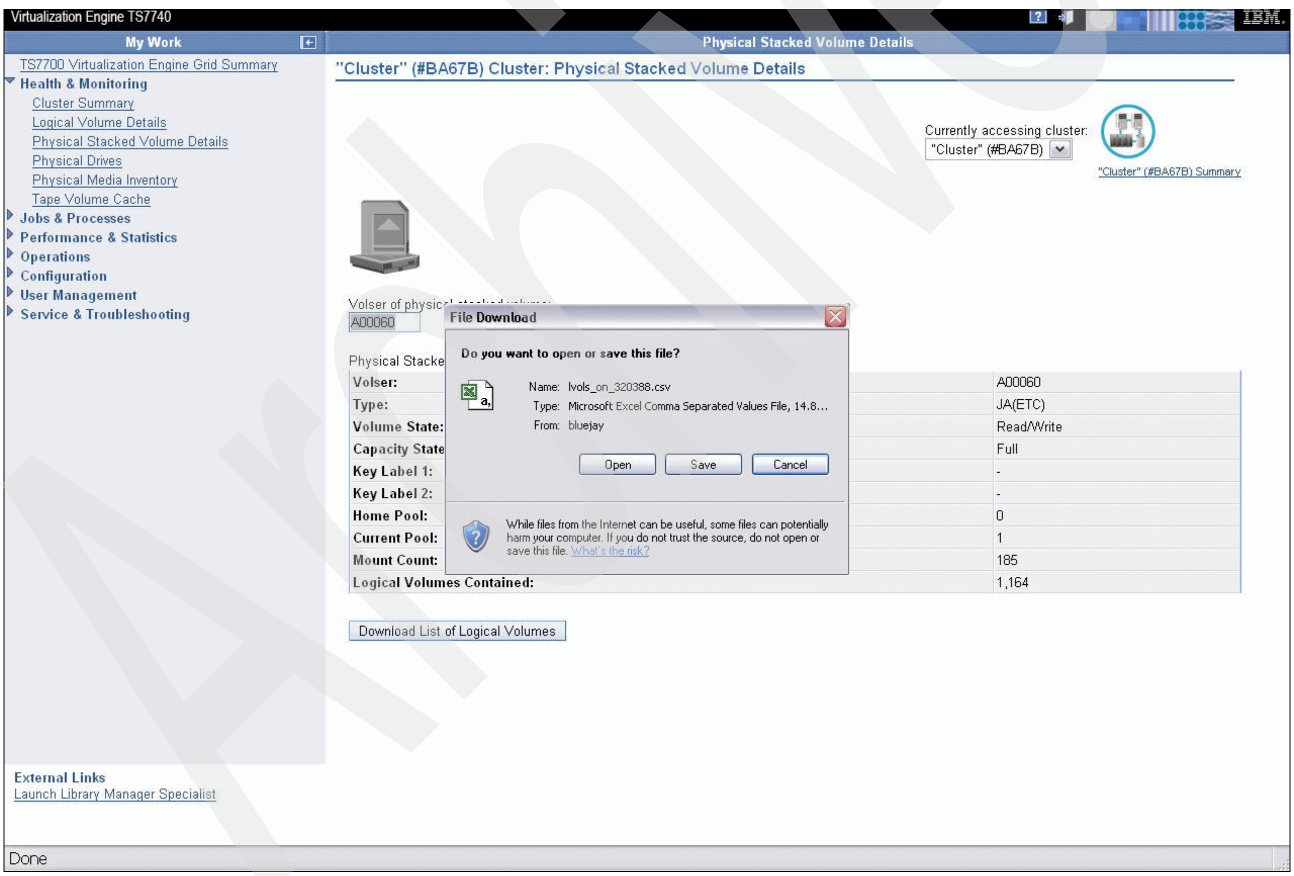

Figure 7-16 TS7700 Management Interface Download of Physical Volume details

Note: In case there is no download dialog box presented after selecting the download of a CSV file, use the InfoCenter in the TS7700 Management Interface to find a procedure for Windows® XP SP2 users for configuration information of Microsoft Internet Explorer.

Physical Drives

Use this page to view the state of physical drives accessible to the IBM Virtualization Engine TS7700 cluster.

Whenever you need to know specific details about a physical drive, such as its serial number, the drive type, the drive format, or the stacked volumes mounted on it, this page allows you to see the desired information, as shown in Figure 7-17.

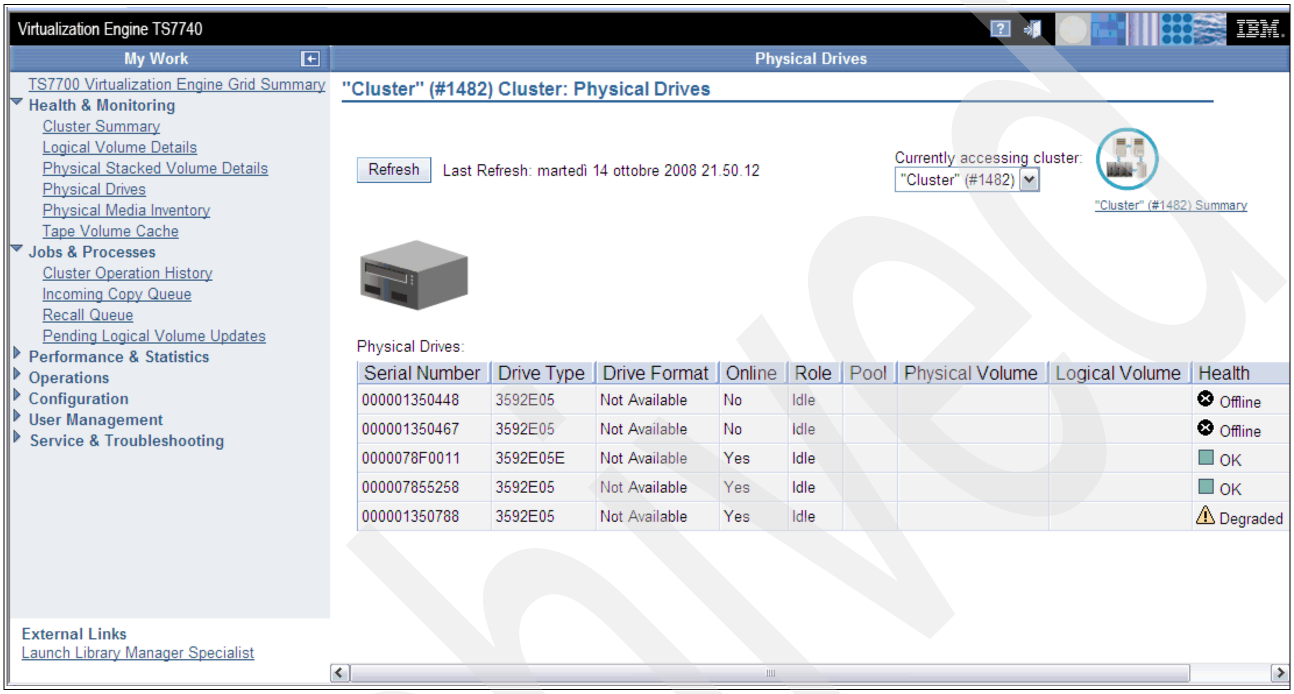

Figure 7-17 TS7700 Management Interface Physical Drives

The following information is displayed when details for physical drives are retrieved:

- **Serial Number**: The serial number that identifies the physical drive.
- **Drive Type**: The machine type and model number of the drive. Possible values are:

3592J1A	3592 drives J1A model type.
3592E05	3592 drives E05 model type.
3592E05E	The drive is operating in an E05E encryption capable format.
Not Available	The format is unable to be determined because there is no physical media in the drive.

- **Drive Format**: The characters that identify the drive format.
- **Online**: The current state of the drive, possible values are:

YES	The drive is online to the TS7740.
NO	The drive is offline to the TS7740.

- **Role**: The actual role of the drive at the time of the display. Possible values are:

Idle	The drive is currently not in use.
Recall	The drive is being used to recall a logical volume from a physical volume to the tape volume cache.
Reclaim Source	The drive has a physical stacked volume mounted that is the source for the reclamation task.
Reclaim Target	The drive has a physical stacked volume mounted that is the target for the reclamation task.

Secure erase	The drive is being used to erase expired volumes from the physical volume securely and permanently.
Migration	The drive is being used to copy a logical volume from the tape volume cache to a physical volume.

- **Pool**: The pool name of the physical volume mounted by the drive. This field is blank if the drive is idle.
- **Physical Volume**: The 6 character volume serial identifier of the physical stacked volume mounted on the drive at display time.
- **Logical Volume**: The 6 character volume serial identifier of the logical volume that resides on the physical stacked volume mounted on the drive at display time.
- **Health**: The health of the physical drive. This value is obtained automatically at times determined by the TS7700. Possible values are:

OK	The drive is fully functioning.
WARNING	The drive is functioning but reporting errors. Action should be taken to correct the errors.
DEGRADED	The drive is functioning but at lesser redundancy and performance.
FAILURE	The drive is not functioning and immediate action should be taken to correct it.
OFFLINE/TIMEOUT	The drive is out of service or could not be reached within a certain time frame.

Physical Media Inventory

Use this page (Figure 7-18) to obtain detailed information about a physical media inventory in the IBM Virtualization Engine TS7700.

Whenever you need to know specific details about a physical media, such as its pool number, media type pool related, or the quantity of physical media in full status, this page allows you to see all that desired information, as shown in Figure 7-18.

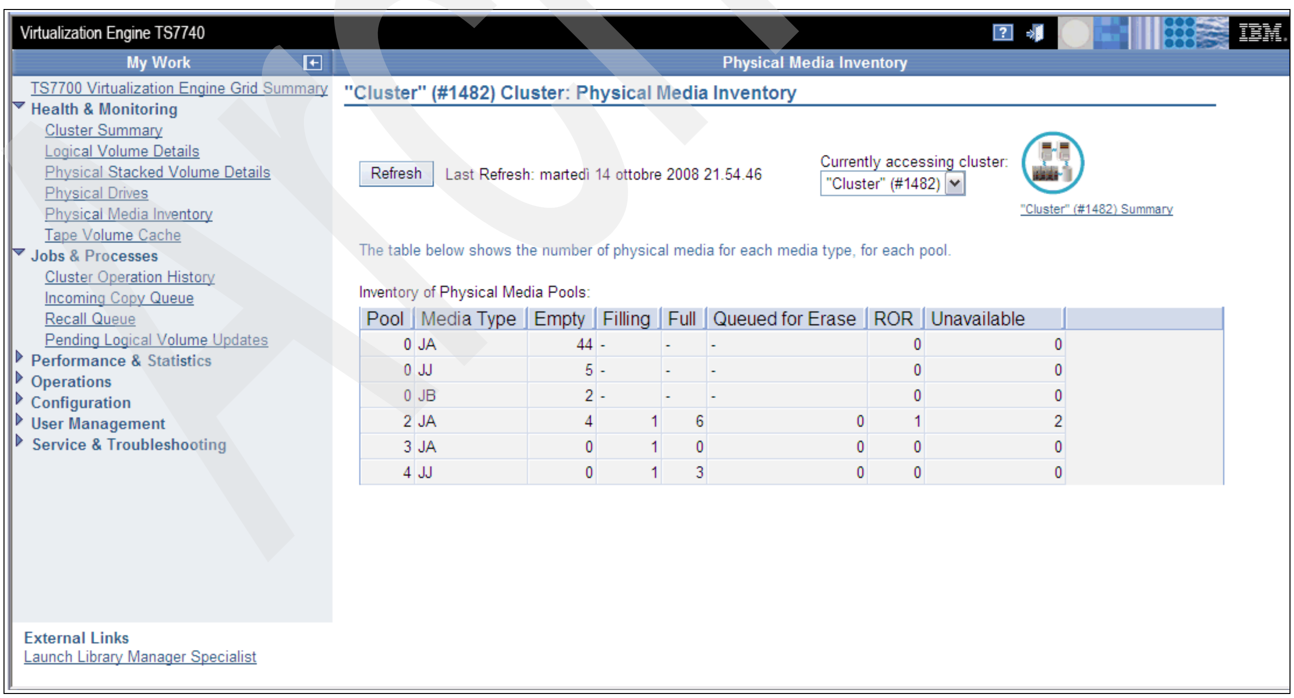

Figure 7-18 TS7700 Management Interface Physical Media Inventory

The following information is displayed when details for physical media inventory are retrieved:

- **Pool**: The number that identifies the pool for the specific media type.
- **Media Type**: The media type defined for the pool. A storage pool can have multiple media types and each media type will be displayed separately. Possible values are:

JA (ETC)	Enterprise Tape Cartridge (ETC)
JB (ETCL)	Enterprise Extended-Length Tape Cartridge (ETCL)
JJ (EETC)	Enterprise Economy Tape Cartridge (EETC)

- **Empty**: The count of physical volumes that are empty for the pool.
- **Filling**: The count of physical volumes that are filling for the pool. This field is blank for Pool 0.
- **Full**: The count of physical volumes that are full for the pool. This field is blank for Pool 0.
- **Queued for Erase**: The count of physical volumes that are reclaimed but need to be erased before they can become empty. This field is blank for Pool 0.
- **ROR (Read Only Recovery)**: The count of physical volumes that are in the Read Only Recovery state. That does not imply that there is an error on the physical stacked volume. The ROR is activated for the reclamation process.
- **Unavailable**: The count of physical volumes that are in the unavailable or destroyed state.

Tape Volume Cache

Use this page (Figure 7-19) for obtaining information about the tape volume cache in the IBM Virtualization Engine TS7700.

Whenever you need to know specific details about a tape volume cache, such as its installed size, premigration size, migration or copy throttling time, this page allows you to see the desired information, as shown in Figure 7-19.

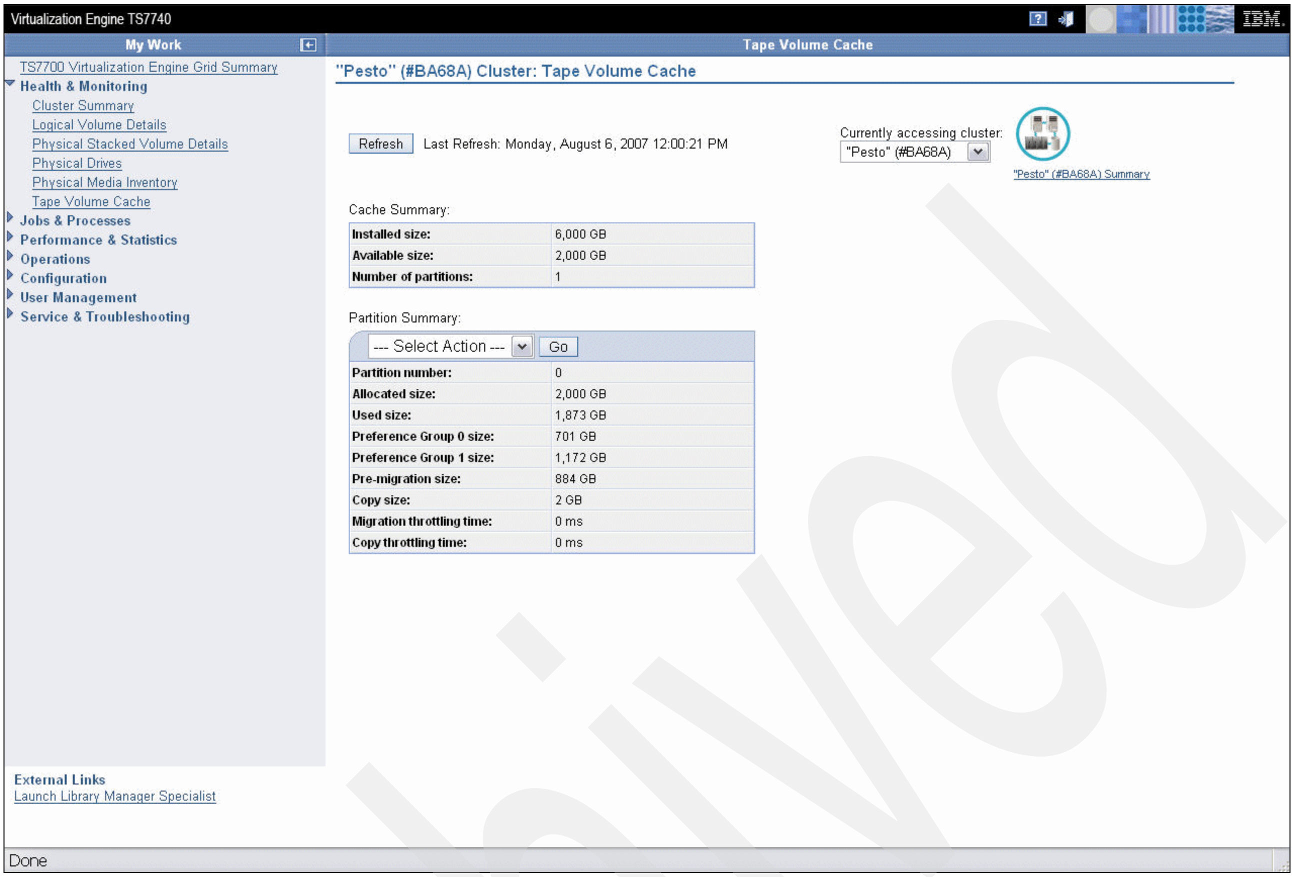

Figure 7-19 TS7700 Management Interface Tape Volume Cache

The following information is displayed for the tape volume cache and for each partition in the cache. To view information for an individual partition, select the partition from the Select Partition drop-down.

- **Cache summary**: The table contains status information about the summary of the cache of the TS7740. The following information can be obtained:

  Installed size The cache size that is physically installed in gigabytes (GB).

  Available size The cache size that is logically enabled in gigabytes (GB).

  Number of Partitions The number of partitions in the cache. The maximum number of partitions is eight. Figure 7-20 on page 330 shows only one partition configured.

- **Partition Summary**: The table contains status information related to the cache partition. In order to view details on the desired cache partition, open the Select Action drop-down list, select the partition, and click **GO**.

  Partition number The partition number within the cache. Information displayed in this table will be for this partition. In Figure 7-19 you see that partition zero is selected.

  Allocated size The amount of cache that has been allocated for this partition.

  Used size Amount of cache that is currently in use by the partition.

  Preference Group 0 size Size of preference group zero. Volumes in this group have preference to be removed from cache over other volumes.

  Preference Group 1 size Size of preference group one. Volumes in this group have preference to be retained in cache over other volumes.

Pre-migration size Amount of data that needs to be copied to a physical volume.

Copy size Amount of data that needs to be copied to another cluster.

Migration throttling time Indicates whether migration throttling is active for the partition. If throttling is active, a value in milliseconds is displayed. If throttling is not active, zero is displayed.

Copy throttling time Indicates whether copy throttling is active for the partition. If throttling is active, a value in milliseconds is displayed. If it is not active, zero is displayed.

7.2.4 Jobs & Processes

This section of the TS7700 Management Interface gives you access to some panels to query the cluster operation history, as well monitoring the incoming copy queues and view pending logical volume updates if the cluster is part of a TS7700 Multi Cluster Grid environment. You can access each selection if you navigate within the TS7700 Management Interface to the Jobs & Processes pull-down section and open it with a mouse click.

In this section, we discuss the following topics:

- Cluster Operation History
- Incoming Copy Queue
- Recall Queue
- Pending Logical Volume Updates

Cluster Operation History

Use this page for viewing information about currently running and completed tasks on a TS7700 Virtualization Engine cluster, as shown in Figure 7-20.

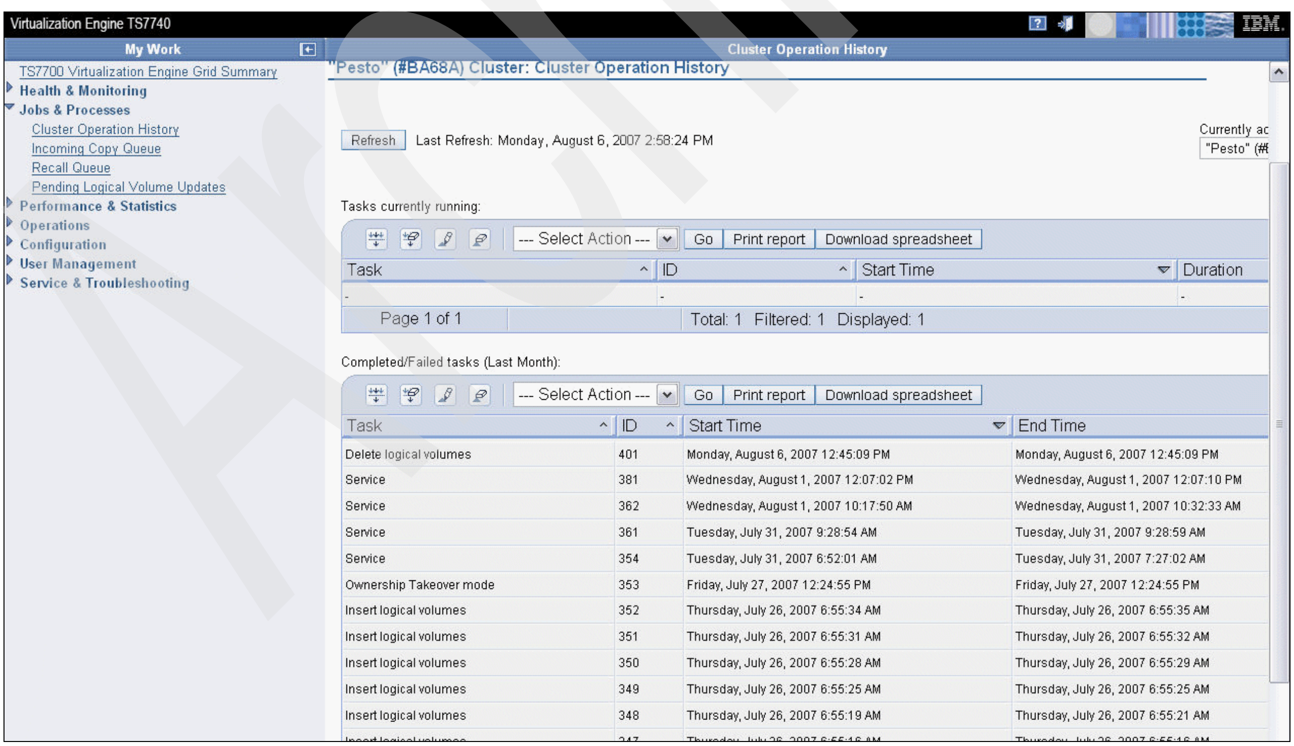

Figure 7-20 TS7700 Management Interface Cluster Operation Interface

The following information is available for currently running and completed tasks and is presented in two tables. The first table displays information for the currently running tasks on the TS7700 cluster. The second table displays information about completed tasks. You can print the table data by clicking **Print Report** next to the Select Action menu. A comma-separated value (.csv) file of the table data can be downloaded by clicking **Download spreadsheet**.

- **Task**: Name of the task
- **ID**: ID number of the task
- **Start Time**: The timestamp for the start of the task
- **End Time**: The timestamp for when the task ends. This is only available for completed tasks.
- **Duration (h:m:s)**: In the case of a completed task, this field displays the time the task took to complete. For currently running tasks, this displays the time the task has taken up to this point.
- **Initiating user**: The user login that initiated the task
- **Current step**: Additional contextual information about the task's current progress
- **Status**: A link to additional information about the status of the task. This is only available for running tasks.
- **Result**: The result of the task. This information is only available for completed tasks. A failed task has a link to additional information about the failure.

Incoming Copy Queue

Use this page (Figure 7-21) to view the logical volume Incoming Copy Queue for an IBM Virtualization Engine TS7700 cluster. This view can only be selected if the cluster is part of a Multi Cluster Grid configuration.

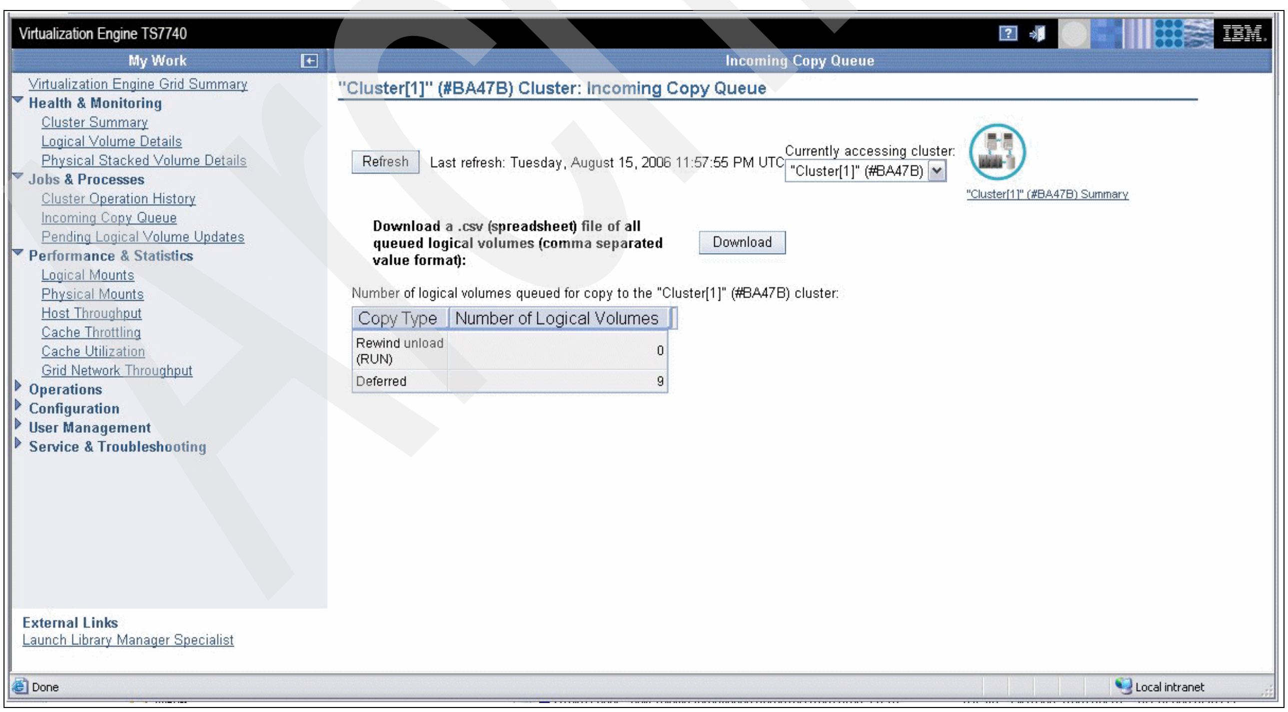

Figure 7-21 TS7700 Management Interface Incoming Copy Queue

The following information is available about the logical volume copy queue:

- **Copy Type**: Possible values are:

  RUN The copy occurs during rewind-unload and before the rewind-unload operation completes at the host.

  Deferred The copy occurs some time after the rewind-unload operation at the host.

- **Number of Logical Volumes**: The total number of queued logical volume copies of the indicated copy type

The following information about individual queued logical volume copies with Ownership Takeover status will be available in a downloadable comma-separated value format(.csv) report that can be obtained by clicking **Download**:

- **VOLSER**: VOLSER of the logical volume being copied.
- **Queue Type**: The type of queue as reported by the cluster. Possible values are:

  RUN The copy occurs during rewind-unload and before the rewind-unload operation completes at the host.

  Deferred The copy occurs some time after the rewind-unload operation at the host.

  Immediate Deferred A RUN copy that is moved to the deferred state due to copy time outs or TS7700 Grid states.

- **Copy Activity**: Status information about the copy activity of the logical volume copy. Possible values are:

  Complete A consistent copy exists at this location.
  In Progress A copy is required and currently in progress.
  Required A copy is required at this location but has not started or completed.
  Not Required A copy is not required at this location.
  Reconcile Pending updates exist against this location's volume. The copy activity will be updated after the pending updates get resolved.

- **Age in Queue**: Elapsed time that the copy has been in the queue
- **Volume Size**: Bytes that will be written during the logical volume copy

Recall Queue

Use this page (Figure 7-22) to view the logical volume Recall Queue and to promote logical volumes to the front of the queue in the IBM Virtualization Engine TS7700.

A recall of a logical volume retrieves the logical volume from physical tape and places it in the cache. A queue is used to process these requests.

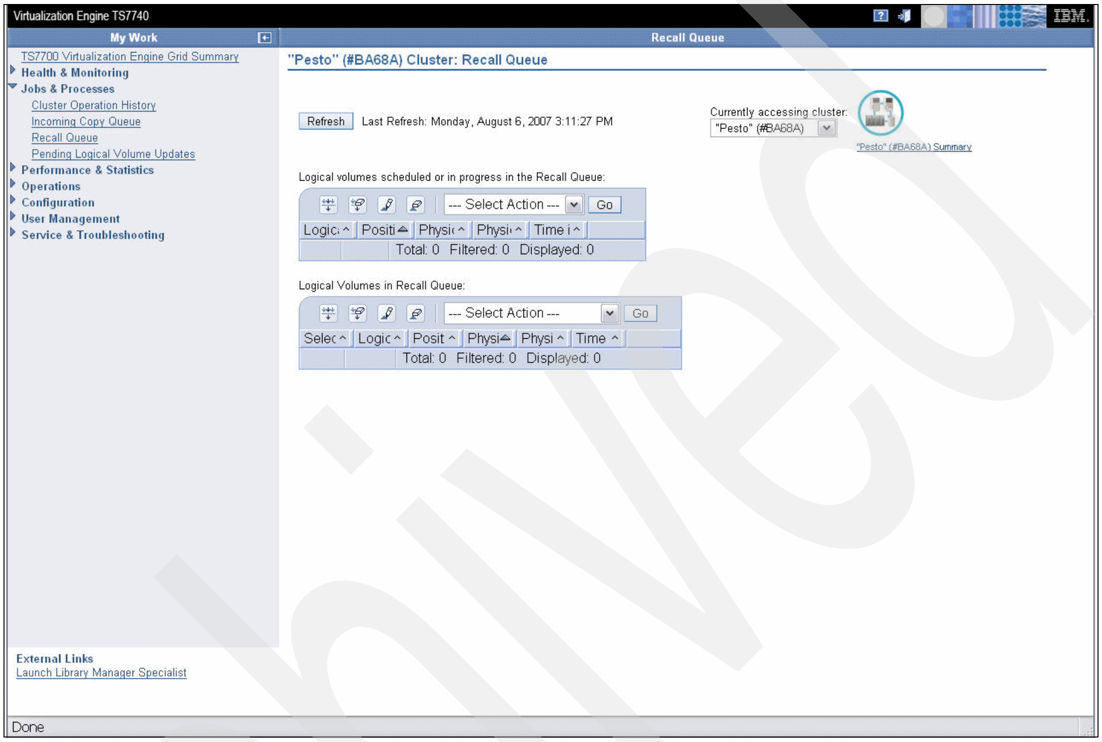

Figure 7-22 TS7700 Management Interface Recall Queue

The following information is displayed in two tables:

- One table shows logical volumes scheduled for recall or where a recall is already in progress.
- The other table displays the queue of volumes waiting to be scheduled for a recall.

Under the Logical Volumes in the Recall Queue table, logical volumes can be moved to the top of the queue by selecting the logical volume in the table, selecting **Move To Head Of Queue** from the Select Action drop-down, and clicking **Go**:

- **Logical volumes scheduled or in progress in the Recall Queue**

  Logical Volume The logical volume to be recalled.

  Position If the logical volume has not been scheduled for recall, the position in the queue will be shown. If it has been scheduled for recall, the following values are possible:
  In Progress:
  The recall for the logical volume is in progress.
  Scheduled:
  The volume is scheduled to be recalled. If optimization is enabled, the TS7700 will schedule recalls to be processed from the same physical volume.

  Physical Volume 1 The physical volume the logical volume resides on.

Physical Volume 2	A second physical volume the logical volume resides on if the logical volume spans a physical volume. Spanning will only occur with migrated data.
Time in Queue	Time in seconds the logical volume has been in the queue.

- **Logical Volumes in Recall Queue**

Select	Select the Logical Volume in the Table.
Logical Volume	The logical volume to be recalled.
Position	If the logical volume has not been scheduled for recall, the position in the queue will be shown. If it has been scheduled for recall, the following values are possible: **In Progress**: The recall for the logical volume is in progress. **Scheduled**: The volume is scheduled to be recalled. If optimization is enabled, the TS7700 will schedule recalls to be processed from the same physical volume.
Physical Volume 1	The physical volume the logical volume resides on.
Physical Volume 2	A second physical volume the logical volume resides on if the logical volume spans a physical volume. Spanning will only occur with migrated data.Pending logical volume updates
Time in Queue	Time, in seconds, the logical volume has been in the queue.

Pending Logical Volume Updates

Use this page (Figure 7-23) to view the pending logical volume updates for the IBM Virtualization Engine TS7700 Grid.

This view can only be selected if the cluster is part of a Multi Cluster Grid configuration.

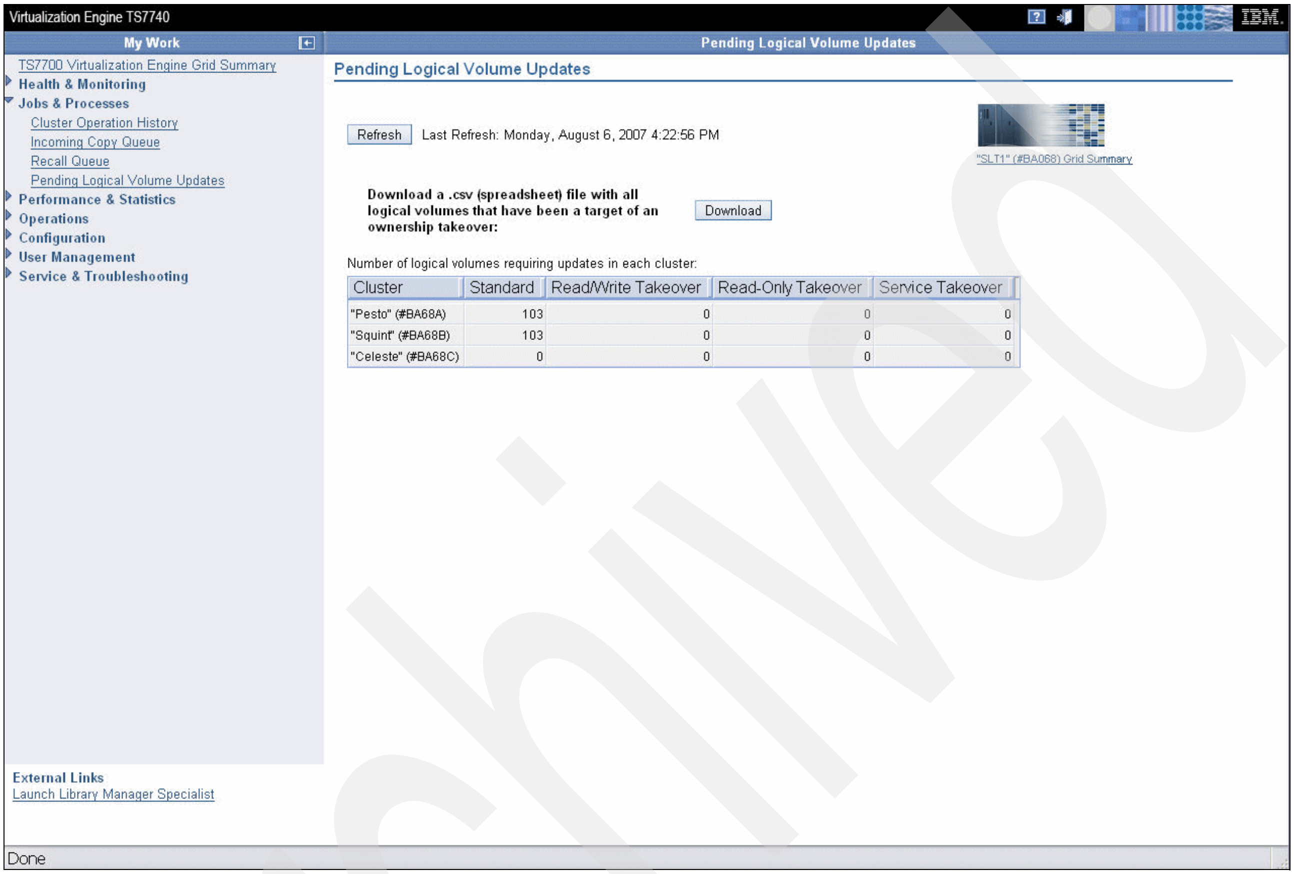

Figure 7-23 TS7700 Management Interface Pending Logical Volume Updates

The following information is available about the pending logical volume updates for each cluster:

- **Cluster**: The cluster name
- **Standard**: Total number of logical volumes that require an update but were not taken over while the cluster was down. These logical volumes were not owned by this cluster.
- **Read/Write Takeover**: Total number of logical volumes that require updates as a result of a read/write ownership takeover while the cluster was down.
- **Read-Only Takeover**: Total number of logical volumes that require updates as a result of a read-only ownership takeover while the cluster was down.
- **Service Takeover**: Total number of logical volumes that require updates as a result of a service ownership takeover while the cluster was in service mode.

The following information about individual pending logical volume updates is available in a downloadable CSV file that can be obtained by clicking **Download a .csv (spreadsheet) file with all logical volumes that have been a target of an ownership takeover**.

- **Logical Volume**: The logical volume's VOLSER number.
- **Original Owner**: The name of the cluster the logical volume update is for.
- **Takeover Cluster**: The cluster that first took over the volume from the original owner.
- **Current owner**: The current owner of the logical volume.
- **Type of Takeover**: Possible values are:

  Read-Only Ownership Takeover The logical volume requires an update due to a read-only ownership takeover.

  Read/Write Ownership Takeover The logical volume requires an update due to a read/write ownership takeover.

  Service The logical volume requires an update due to a service takeover.

- **Time of Ownership Takeover**: Date and time that the logical volume was first taken over from the original owner.

7.2.5 Performance & Statistics

This section cover graphical informations related to performance and statistics for the IBM TS7700 Virtualization Engine. The graphical view are from the last 15 minutes activities. Figure 7-24 shows the Logical Mounts statistics.

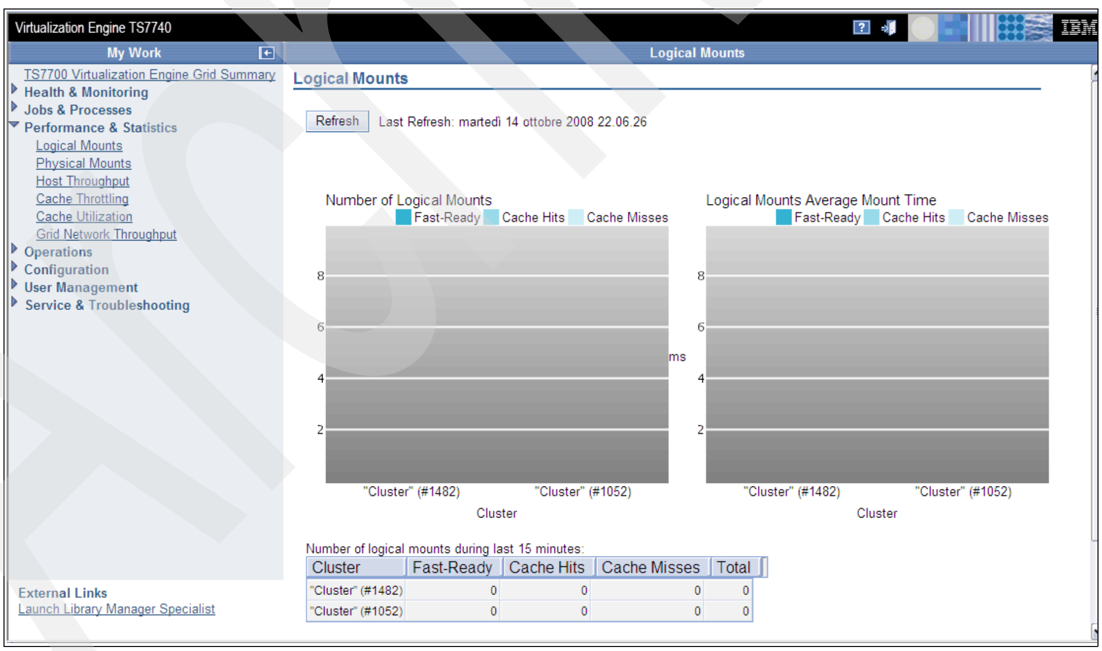

Figure 7-24 Logical Mounts statistics

Note: For deeper informations refer to Chapter 8, “Performance and monitoring” on page 409.

7.2.6 Operations

In this section we present information that is related to performing various operations in the TS7700 Virtualization Engine, such as actions relating to inserting and deleting logical volumes and standalone mount/demount of logical volumes.

In this section, we discuss the following topics:

- Insert Logical Volumes
- Delete Logical Volumes
- Standalone Mount Logical Volume

Insert Logical Volumes

Use this page (Figure 7-25) for inserting a quantity or range of logical volumes in the IBM Virtualization Engine TS7700 Grid.

Logical volumes inserted on an individual cluster will be available to the entire grid. During cartridge entry processing, even if the library is online and operational to a given host, at least one device needs to be online (or have been online) to that host for the library to be able to send the cartridge entry attention interrupt to that host. If the library is online and operational, yet there are no online devices to a given host, that host will not receive the attention interrupt from the library unless a device had previously been VARYed online. To get around this, ensure that at least one device is online (or had been online) to each host or use the LIBRARY RESET,CBRUXENT command to initiate cartridge entry processing from the host. This is especially important if you only have one host attached to the library that owns the volumes being entered. In general, after you have entered volumes into the library, if you do not see the expected CBR36xxI cartridge entry messages being issued, you can use the LIBRARY RESET,CBRUXENT command from MVS to initiate cartridge entry processing. The LIBRARY RESET,CBRUXENT command causes the host to ask for any volumes in the insert category.

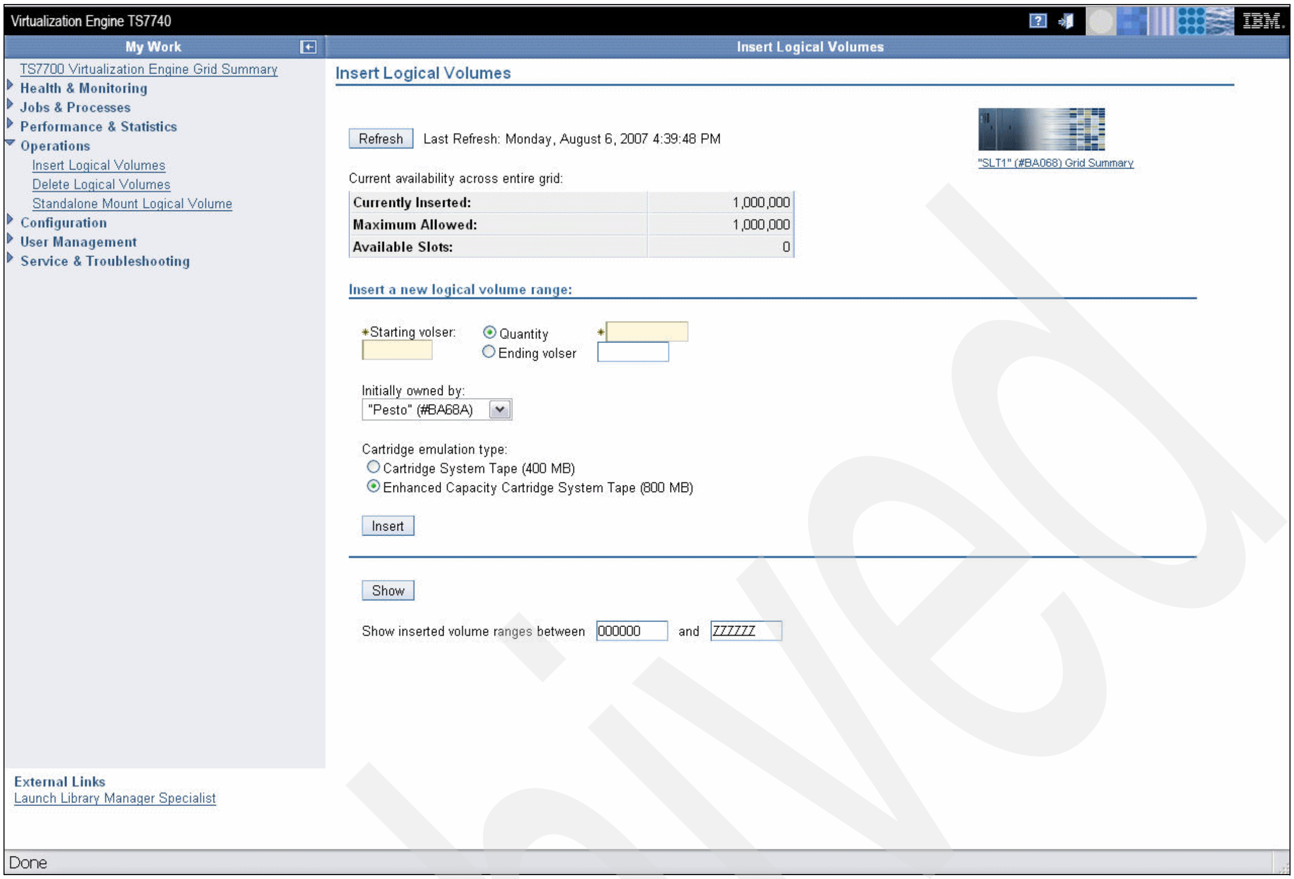

Figure 7-25 TS7700 MI Insert Logical Volumes

The table shown at the top of Figure 7-25 shows current information about the number of logical volumes in the TS7700.

Currently Inserted The total number of logical volumes inserted into the TS7700

Maximum Allowed The maximum number of logical volumes that can be inserted

Available Slots The available slots remaining for logical volumes to be inserted, obtained by subtracting the Currently Inserted logical volumes from the Maximum Allowed.

The following information about the page is displayed and is used for inserting logical volumes:

- **Starting VOLSER**: The first logical volume to be inserted. The range for inserting logical volumes begins with this VOLSER identifier.
- **Quantity**: Select this option to insert a set amount of logical volumes starting with the Staring VOLSER. The adjoining text field is where the quantity of logical volumes to be inserted is entered.
- **Ending VOLSER**: Select this option to insert a range of logical volumes. The adjoining text field is where the ending VOLSER identifier is entered.
- **Initially owned by**: Used to select the initial owning cluster of the logical volume.
- **Cartridge emulation type**: Media type of the logical volume. Possible values are:
  - Cartridge System Tape (400 MB)
  - Enhanced Capacity Cartridge System Tape (800 MB)

- **Insert**: Click this to insert the logical volumes with the selected options.
- **Show**: Click this to display an inserted volume range between a specific range.

> **Restriction:** Up to 10 000 logical volumes can be inserted at one time. This applies to both inserting a range of logical volumes and inserting a quantity of logical volumes. Attempting to insert amounts over ten thousand will return an error.

Delete Logical Volumes

Use this page (Figure 7-26) to delete unused logical volumes from the IBM Virtualization Engine TS7700.

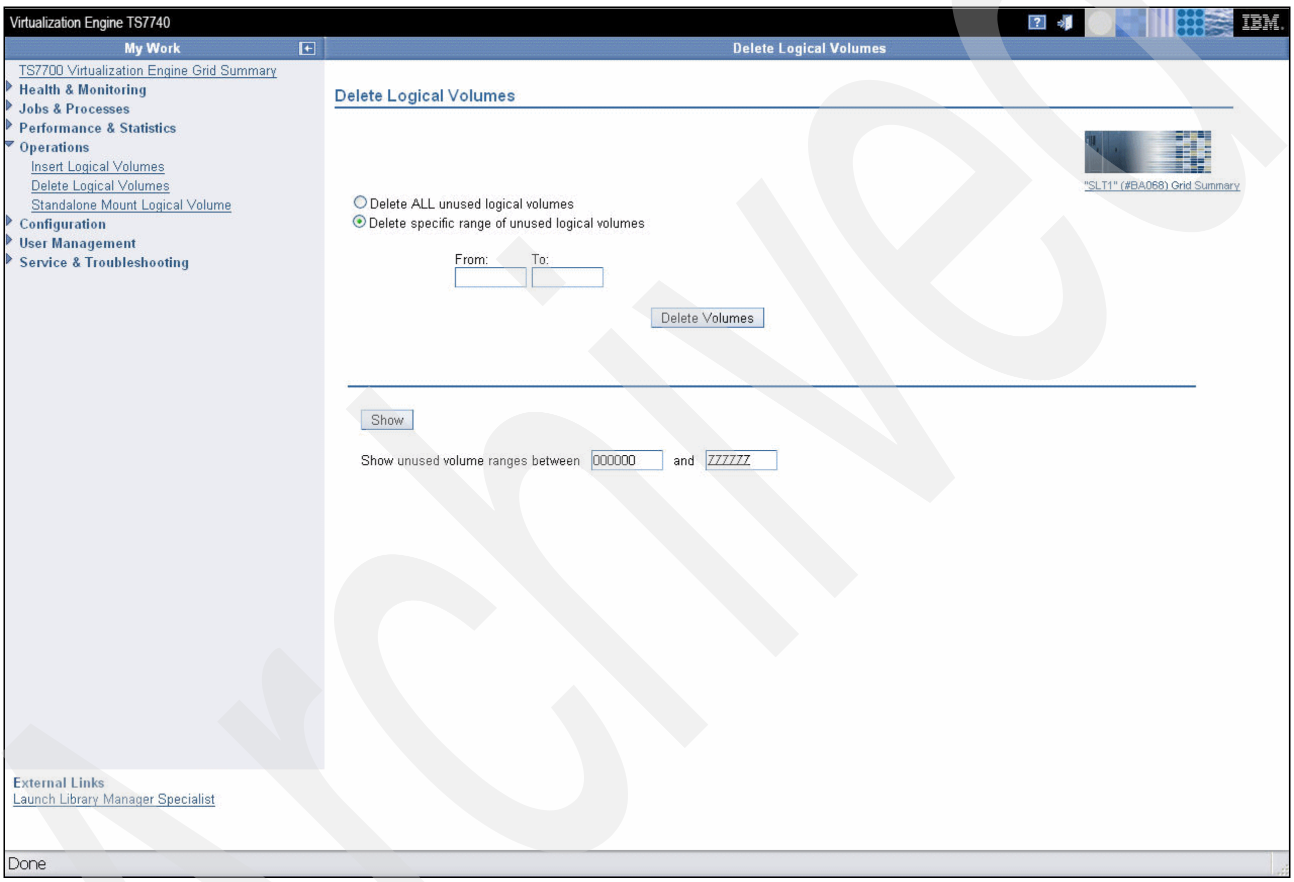

Figure 7-26 TS7700 MI Delete Logical Volumes

To delete unused logical volumes, select one of the options described below and click **Delete Volumes**. A confirmation page will be displayed. Click **Yes** to delete or **No** to cancel. You can use the following fields to specify which volumes to delete:

- **Delete ALL unused logical volumes**: Deletes all logical volumes still in the insert category. A logical volume in the insert category has not been mounted by a host.
- **Delete specific range of unused logical volumes**: All unused logical volumes in the entered VOLSER range will be deleted.

  From The start of the VOLSER range to be deleted if **Delete specific range of unused logical volumes** is selected.

  To The end of the VOLSER range to be deleted if **Delete specific range of unused logical volumes** is selected.
- **Show**: Click this to display an unused volume range between a specific range.

Standalone Mount Logical Volume

Use this page (see Figure 7-27) to perform standalone mounts/demounts of logical volumes on an IBM TS7700 cluster.

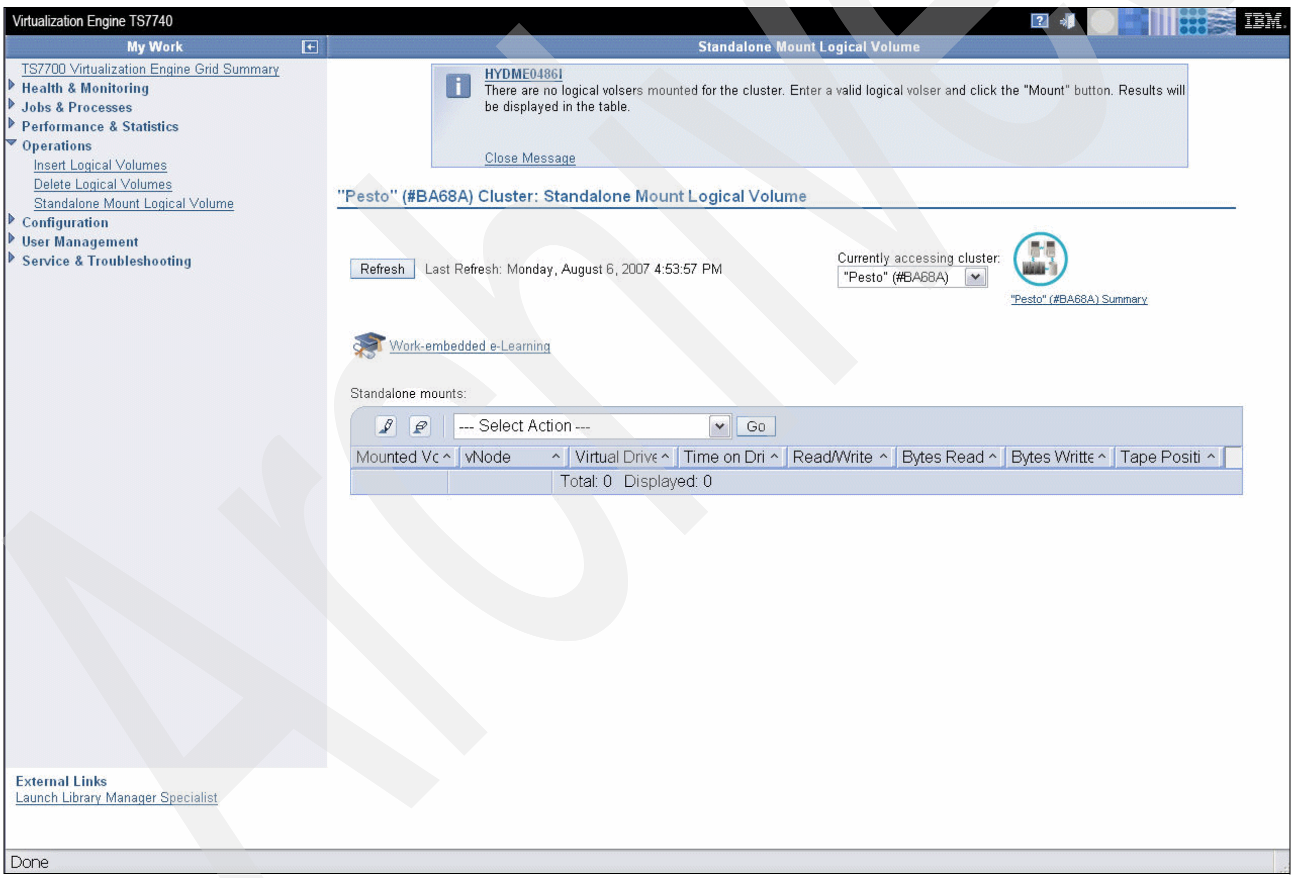

Figure 7-27 TS7700 MI Standalone Mount Logical Volume

This function enables you to mount a logical volume on a selected virtual drive. If you have an operating system image written to a logical volume, performing a standalone mount of a logical volume to a virtual drive can allow you to boot from that mounted logical volume.

You can use these columns to select an entry for a table action:

Mounted Volume The VOLSER of the logical volume mounted.

vNode The name of the vNode that contains the virtual drive the logical volume is mounted on.

Virtual Drive Address The address of the virtual drive that has the logical volume mounted.

Time on Drive (h:m:s) Elapsed time since the logical volume was first mounted on the virtual drive.

Read/Write state Possible values are Read, Write, and Idle

Bytes read Amount of data read from the mounted logical volume.

Bytes written Amount of data written to the mounted logical volume.

Tape position Logical position of the mounted logical volume that the drive is currently reading.

To perform a standalone mount, proceed as follows:

1. From the Standalone Mount Logical Volume page (see Figure 7-27 on page 340), select **Mount New Logical Volume** from the Select Action drop-down list and click **Go**. A form for mounting a new logical volume is presented; see Figure 7-28.
2. Enter the VOLSER of the logical volume to be mounted in the Volser* text field.
3. Select a vNode from the vNode drop-down list. The list of Available Virtual Drives will be populated based on the vNode selected. When the page is loaded, a vNode and virtual drive will be preselected for the user. These values do not have to be changed.

   Note: The virtual drives presented are those that are not in use at the current time. The system will continue to operate in the background. Select a virtual device address which is not currently allocated to a running system.

4. Select a virtual drive to mount the logical volume from the Available Virtual Drives drop-down list.
5. Click **Mount** to mount the logical volume on the selected virtual drive or **Cancel** to abort the mount and return to the previous page.

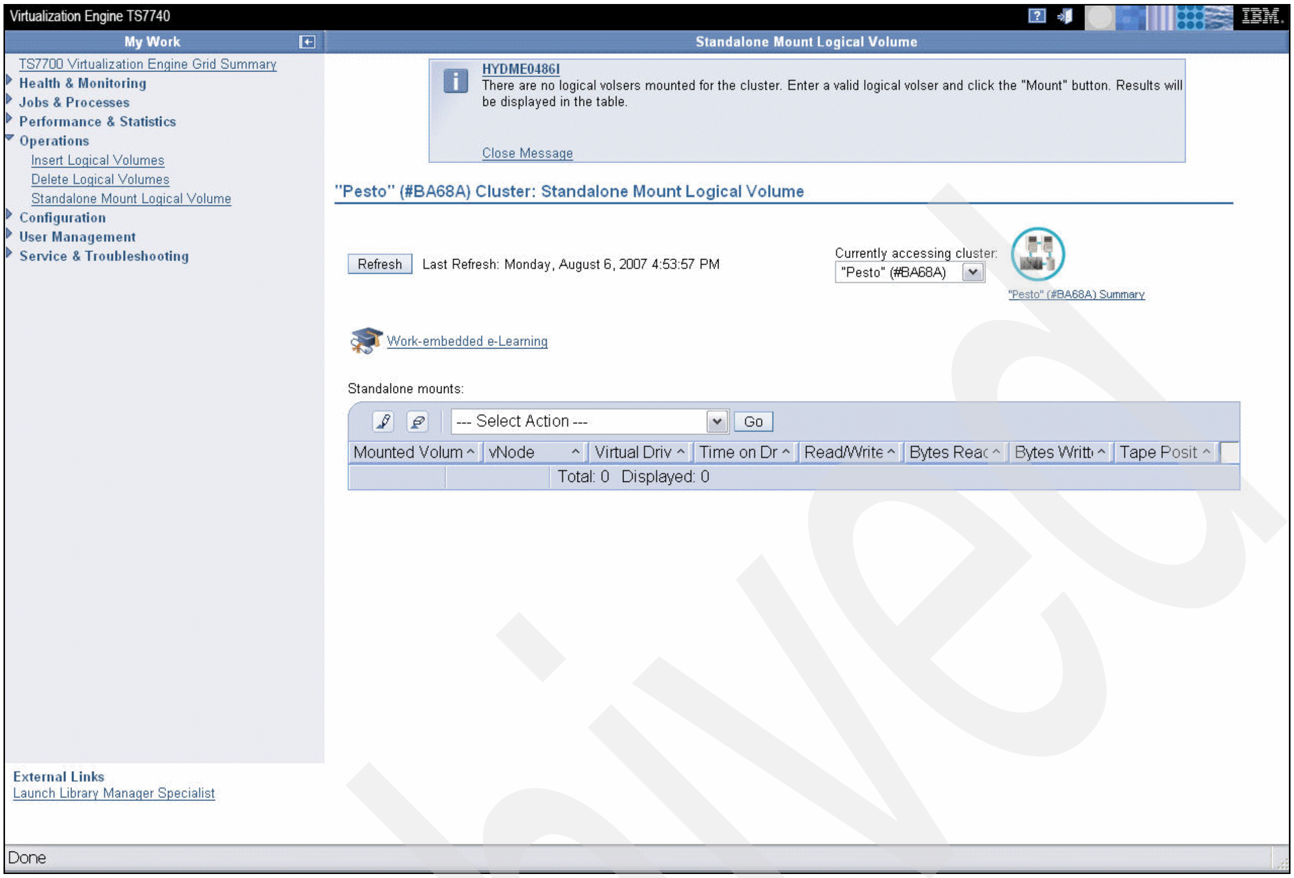

Figure 7-28 TS7700 MI Standalone Mount New Logical Volume

> **Note:** For a standalone mount of a TS7700 logical volume, use the panels shown in Figure 7-27 on page 340 and Figure 7-28 on page 342. A standalone mount request from the IBM 3953 Library Manager will not work for a TS7700 logical volume.

7.2.7 Configuration

The functions included in the Configuration selection of the TS7700 MI are related to configuring the IBM TS7700.

In this section, we include the following topics:

- Cluster Nodes
- Write Protect Mode
- Copy Policy Override
- Pool Encryption Settings
- Encryption Key Manager Addresses
- Cluster Network Settings
- Cluster Identification Properties
- Grid Identification Properties
- Feature Licenses

Cluster Nodes

You can use this page for viewing information and performing actions related to Cluster Nodes on the IBM TS7700 Virtualization Engine. The example provided in Figure 7-29 is taken from a TS7700 cluster that is part of a Multi Cluster Grid configuration.

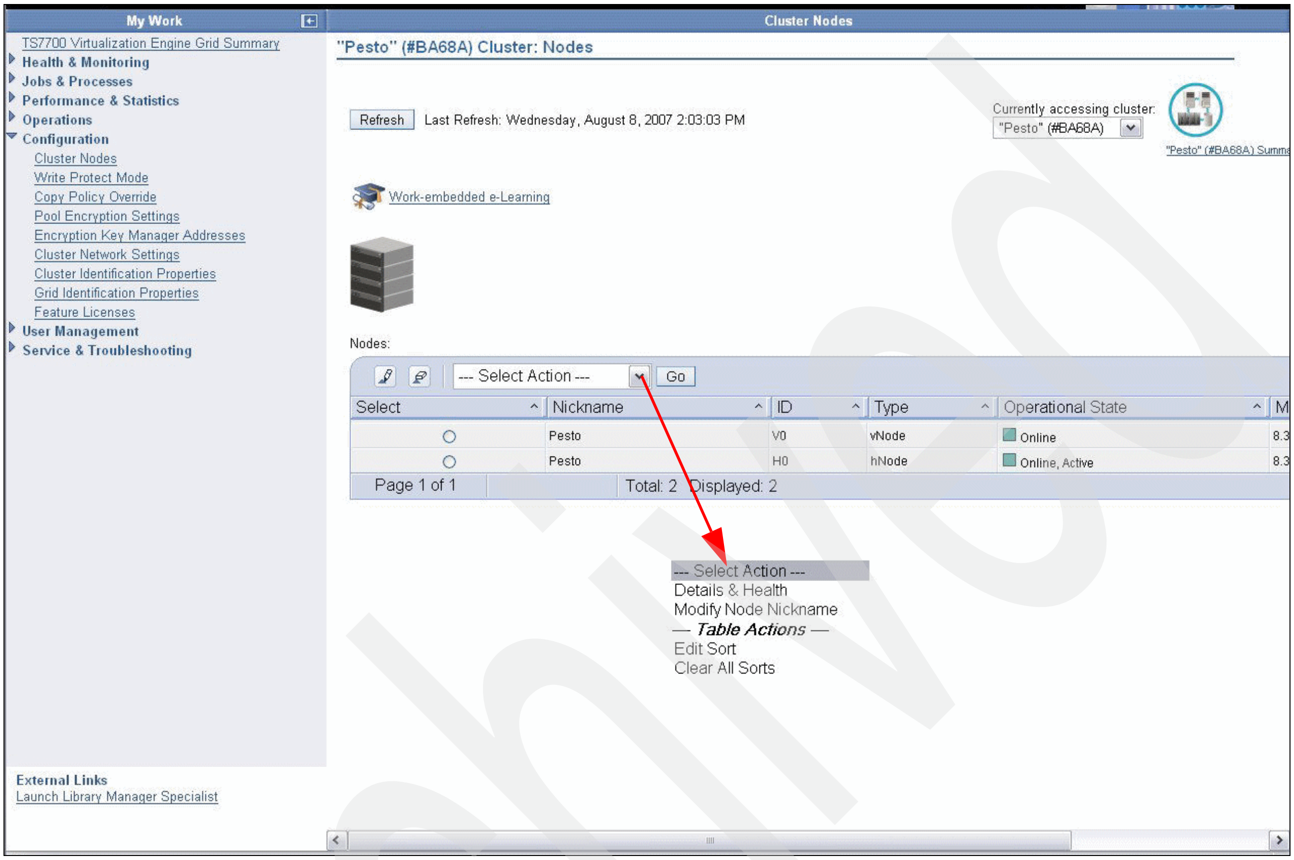

Figure 7-29 TS7700 Management Interface Cluster Node

The columns in the table show:

Nickname Nickname of the node followed by the library sequence number.

ID Node ID.

Type Type of the node. Possible values are vNode or hNode.

Operational State Possible values are Online, Coming Online, Going Offline, and Offline.

Microcode level Microcode level of the node.

There are some actions available for cluster nodes, accessed through the Select Action drop-down list:

Details & Health View details about the selected node. Details shown include general node information, node state and health information, and networking settings for the selected node. See Figure 7-30 as example.

Modify nickname Modify the nickname of the selected node. See Figure 7-31 on page 344 and Figure 7-32 on page 345 for an example.

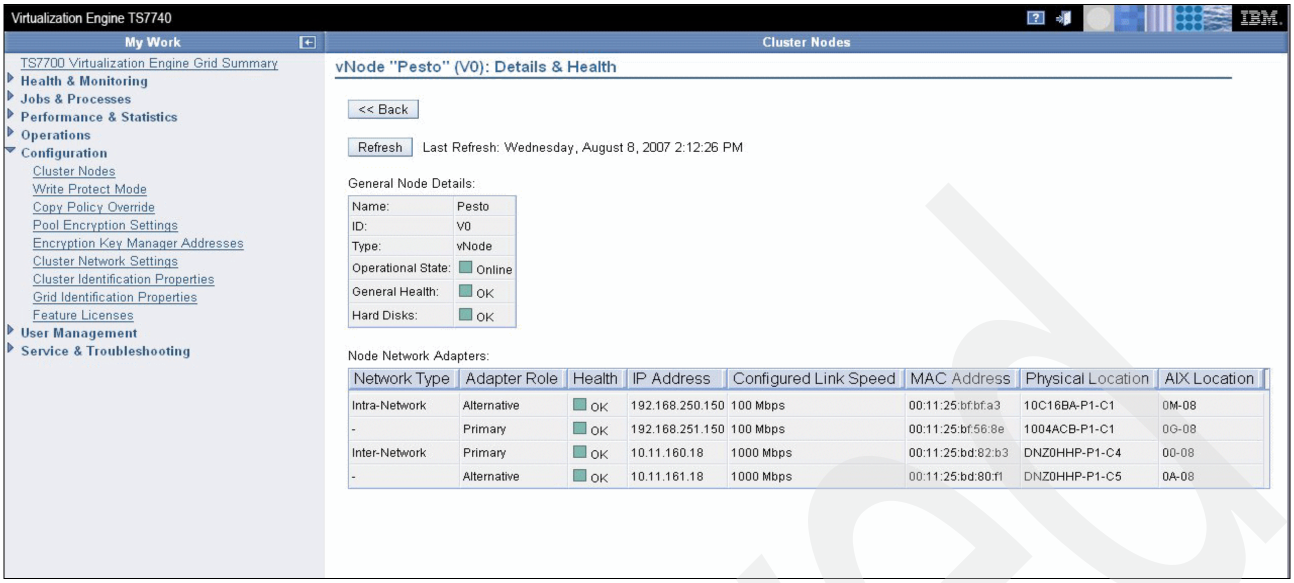

Figure 7-30 TS7700 Management Interface Cluster Node, Details & Health

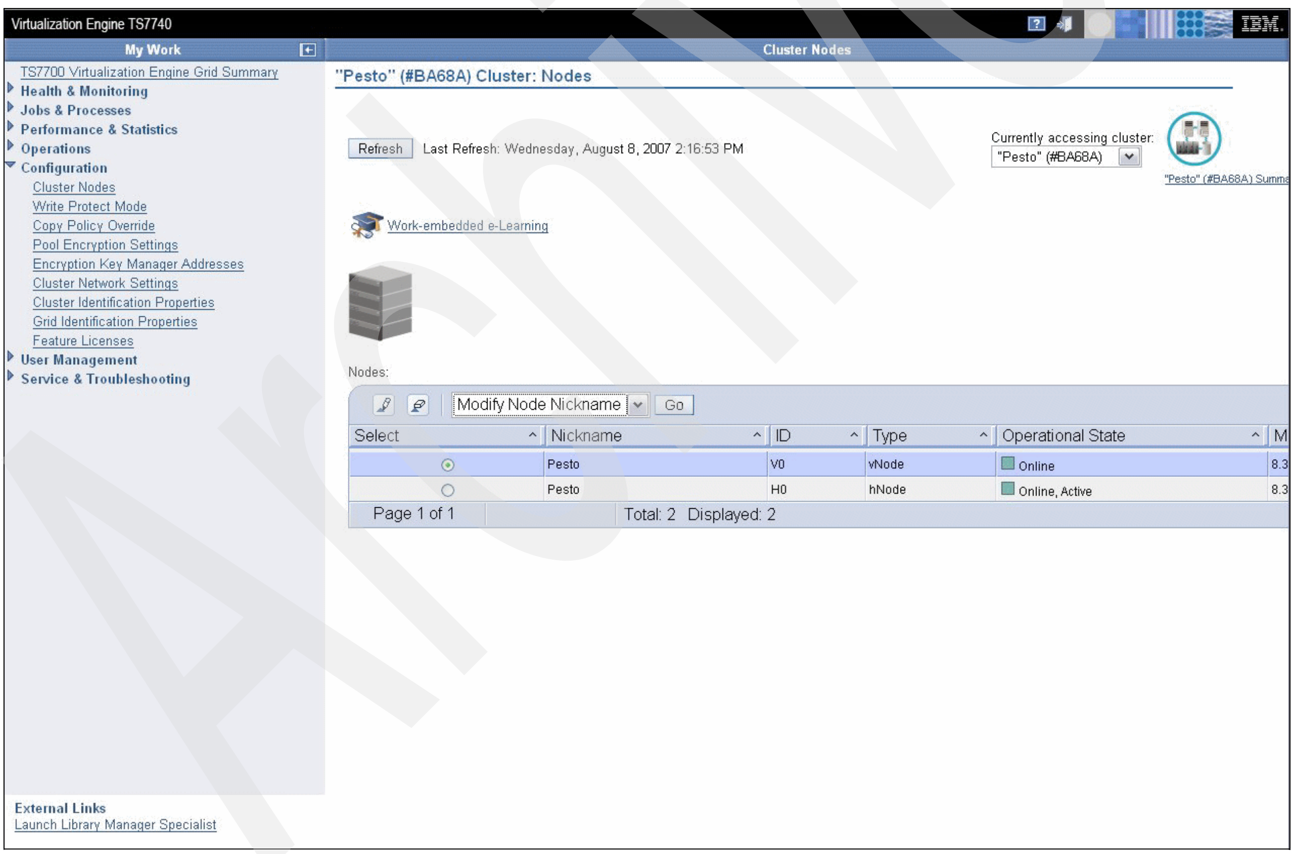

Figure 7-31 TS7700 Management Interface Cluster Node, Modify Node Nickname Step 1

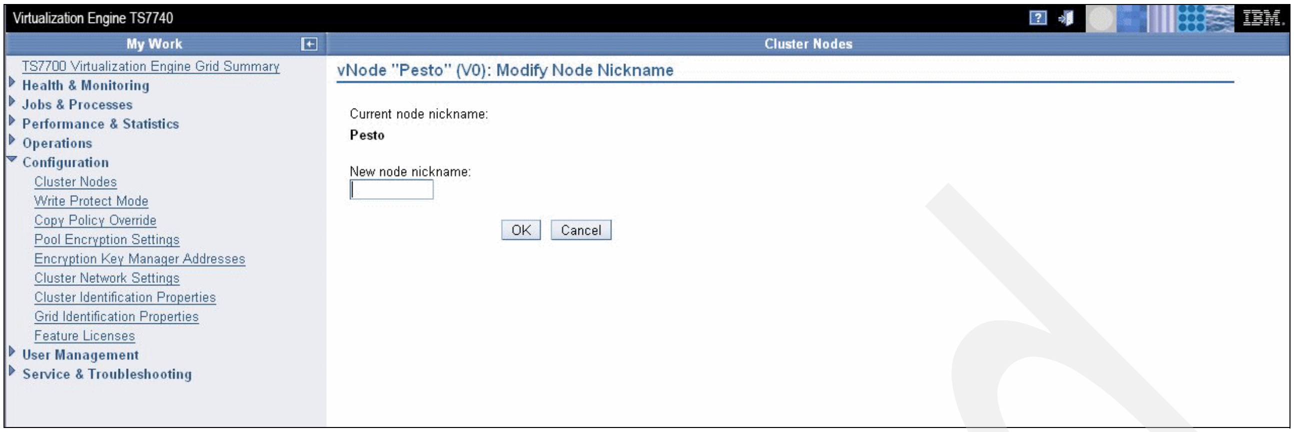

Figure 7-32 TS7700 Management Interface Cluster Node, Modify Node Nickname Step 2

Write Protect Mode

Use the panel shown in Figure 7-33 for enabling or disabling Write Protect Mode in the TS7700 cluster.

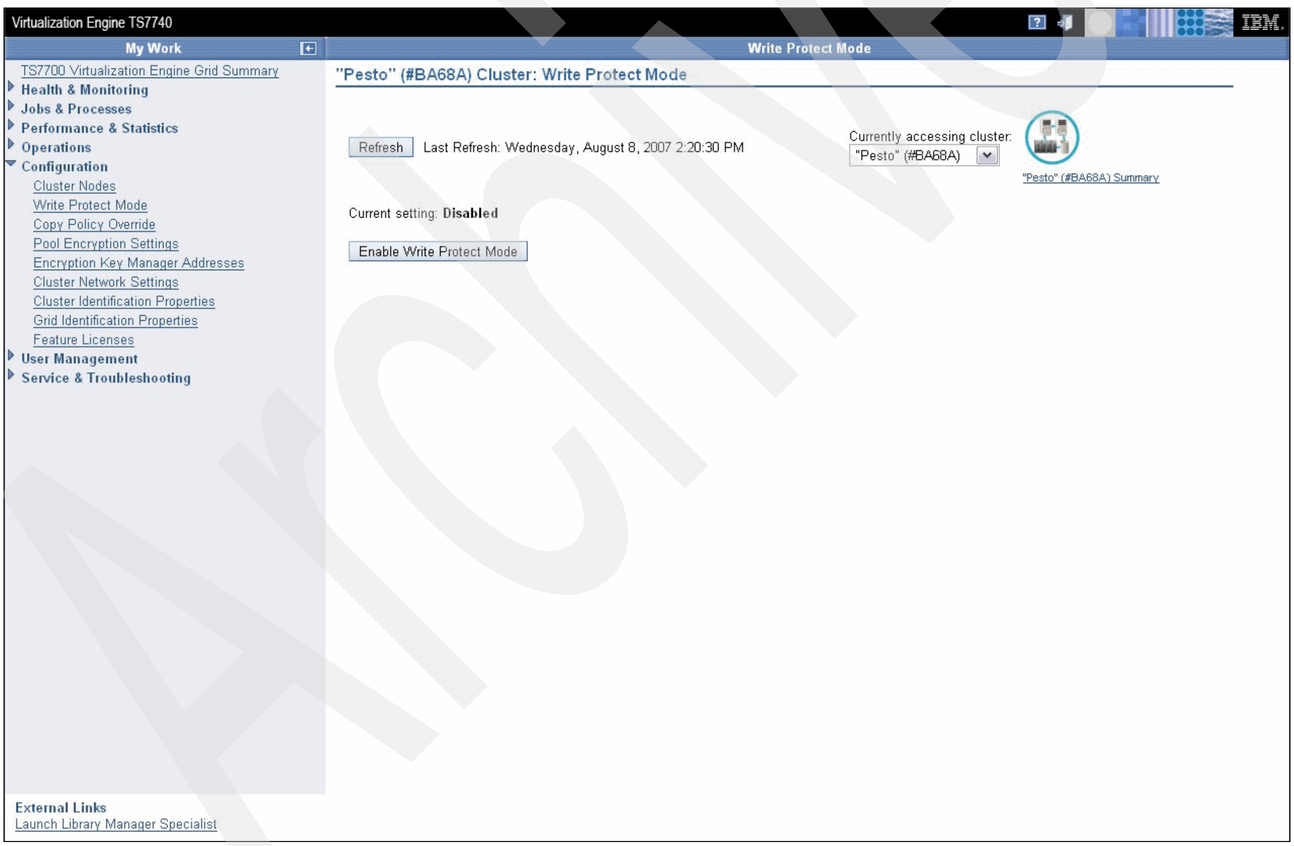

Figure 7-33 TS7700 Management Interface Write Protect Mode

A cluster can be placed in Write Protect Mode if it can communicate with all other clusters in the grid. While in this mode, any attempt by a host attached to the cluster to modify a volume or its attributes will fail. Access to a volume while in this mode does not change the cluster ownership of the volume. This mode is provided for testing customer disaster recovery processes.

The current Write Protect Mode setting for the cluster is displayed and a button for enabling and disabling Write Protect Mode is available.

- **Enable Write Protect Mode**: Only available if Write Protect Mode is currently disabled. Choosing this enables Write Protect Mode.
- **Disable Write Protect Mode**: Only available if Write Protect Mode is currently enabled. Choosing this disables Write Protect Mode.

Copy Policy Override

Use the panel shown in Figure 7-34 to override Copy Policy in the TS7700 cluster.

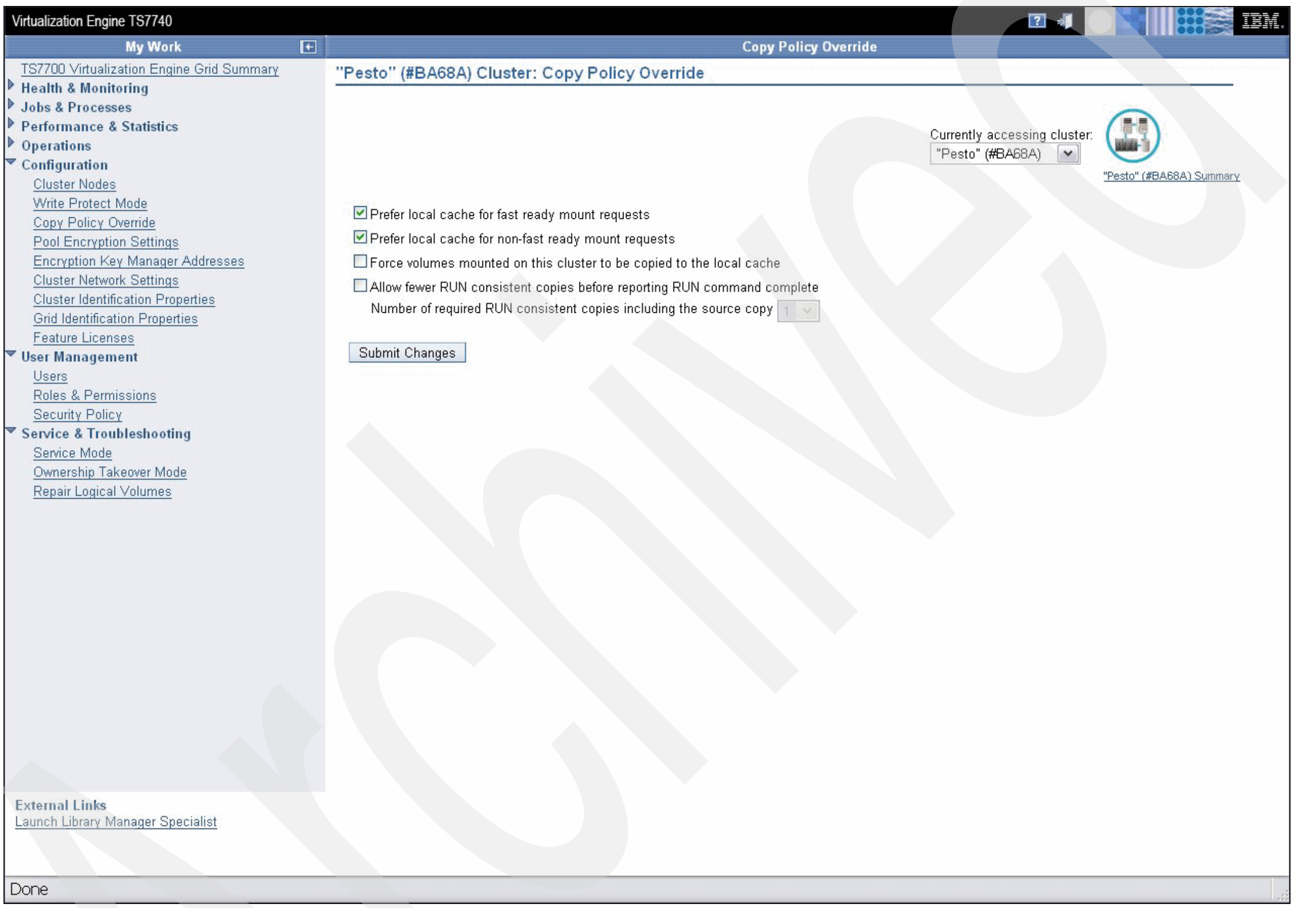

Figure 7-34 TS7700 Management Interface Copy Policy Override

The following settings allow a user to specify cluster overrides for certain I/O and copy operations. These settings override default TS7700 behavior and can be different for every cluster in a grid.

To change a cluster policy override setting for the cluster, text boxes are available to enable or disable policies. Check mark the desired item to change policy. The available options are described here:

- **Prefer local cache for fast ready mount requests**

  A fast ready mount will select a local copy as long as a copy is available and a cluster copy consistency point is not specified as No Copy in the management class for the mount. The cluster is not required to have a valid copy of the Prefer local cache for non-fast ready mount requests.

- **Prefer local cache for non-fast ready mount requests**

  This override will cause the local cluster to satisfy the mount request as long as the cluster is available and the cluster has a valid copy of the data, even if that data is only resident on physical tape. If the local cluster does not have a valid copy of the data, then default cluster selection criteria applies.

- **Force volumes mounted on this cluster to be copied to the local cache**

  For a non-fast ready mount, this override causes a copy to be performed on the local cluster as part of the mount processing. For a fast ready mount, this setting has the effect of overriding the specified management class with a copy consistency point of *Rewind/Unload* for the cluster. This does not change the definition of the management class, but serves to influence the replication policy.

- **Allow fewer RUN consistent copies before reporting RUN command complete**

  If selected, the value entered at `Number of required RUN consistent copies including the source copy` will be used to determine the number of copies to override before the Rewind/Unload operation reports as complete. If this option is not selected, the management class definitions are to be used explicitly. Thus, the number of RUN copies can be from one to the number of clusters in the grid.

Pool Encryption Settings

Use these panels (from Figure 7-35 to Figure 7-39 on page 350) for viewing and modifying storage pool encryption settings on the IBM Virtualization Engine TS7700.

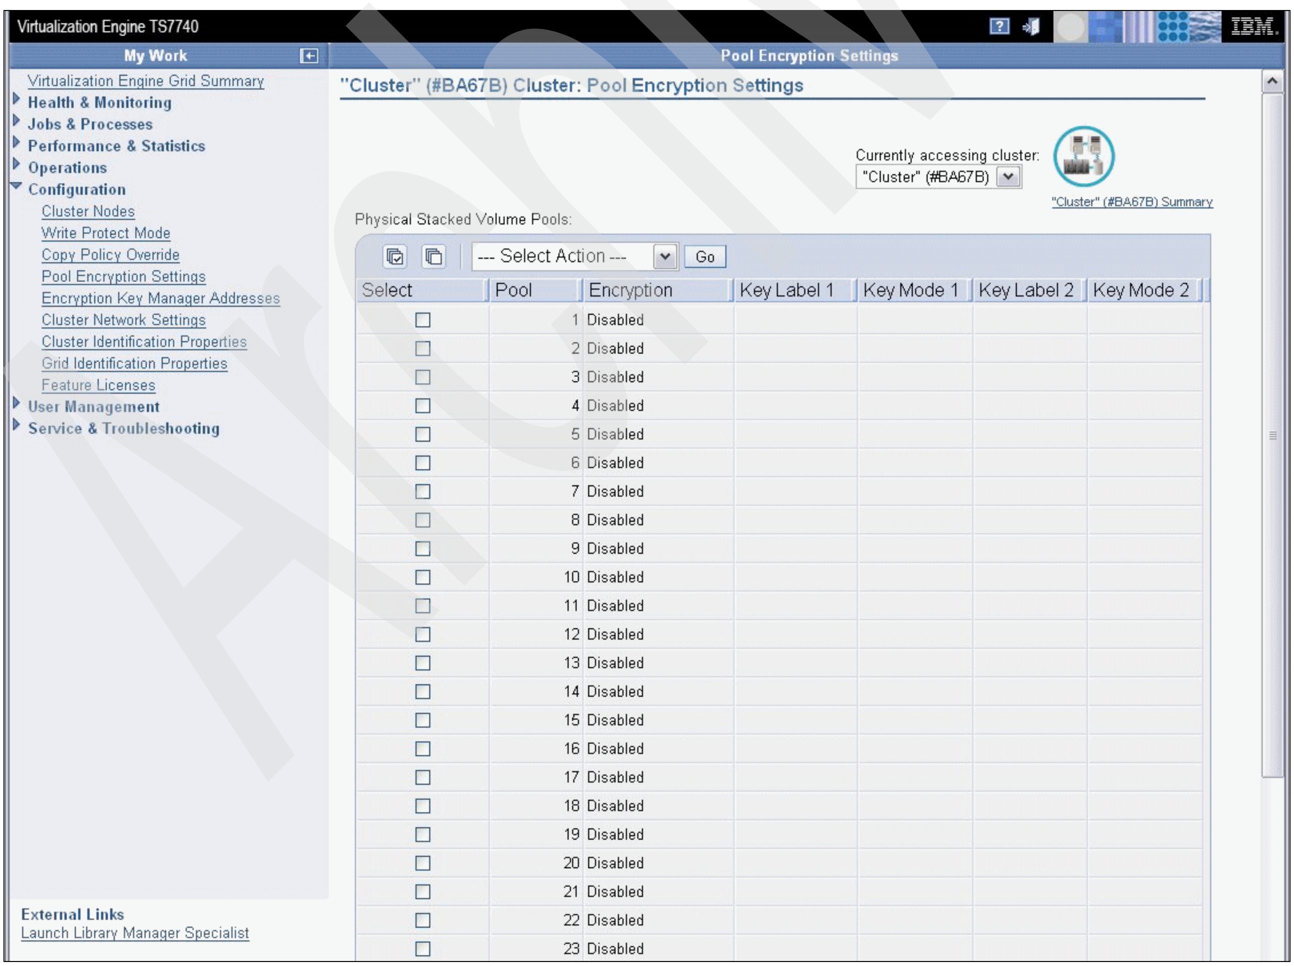

Figure 7-35 Pool Encryption Setting Step 1

The TS7700 Virtualization Engine allows encryption by storage pool. A storage pool contains physical stacked volumes. Within a pool some of the volumes are scratch and some are private. The private volumes contain logical volumes with active data. Every physical stacked volume is assigned to a storage pool.

More information about your current storage pool settings can be obtained from the Library Manager.

A storage pool encryption setting can be modified by selecting one or more storage pools and clicking **Modify Encryption**.

- **Pool**: The number of the storage pool. The range of storage pools is 1 to 32.
- **Encryption**: Enabled or disabled for the pool.
- **Key Label 1/2**: An alias for the Key Encrypting Key (KEK) used by the encryption key manager. A user can select up to two key labels. The maximum length of each label can be 64 characters. Key Label 2 will be enabled only if Key Label 1 is enabled. A previously selected label can be used by selecting one from the previously selected key labels drop-down.
- **Key Mode 1/2**: The method by which an Encryption Key Manager (EKM) identifies the public/private keys that were used to encrypt it.
- **Clear Label**: The key label points to an externally encoded data key (EEDK). It is possible that the same key can be in one keystore with one label and be imported into another keystore with a different label. This can cause problems across sites, such as a TS7700 Grid with a main site and a disaster recovery site. In the case that there are inconsistent label names, the disaster recovery site would be unable to decrypt.
- **Hash Label**: A numerical value is calculated from the public key, independent of any key label that might be assigned to the key. The hash is valid across multiple EKM/keystores even if different key labels are assigned. This mode avoids the inconsistency problems that occur with the Clear Label mode.

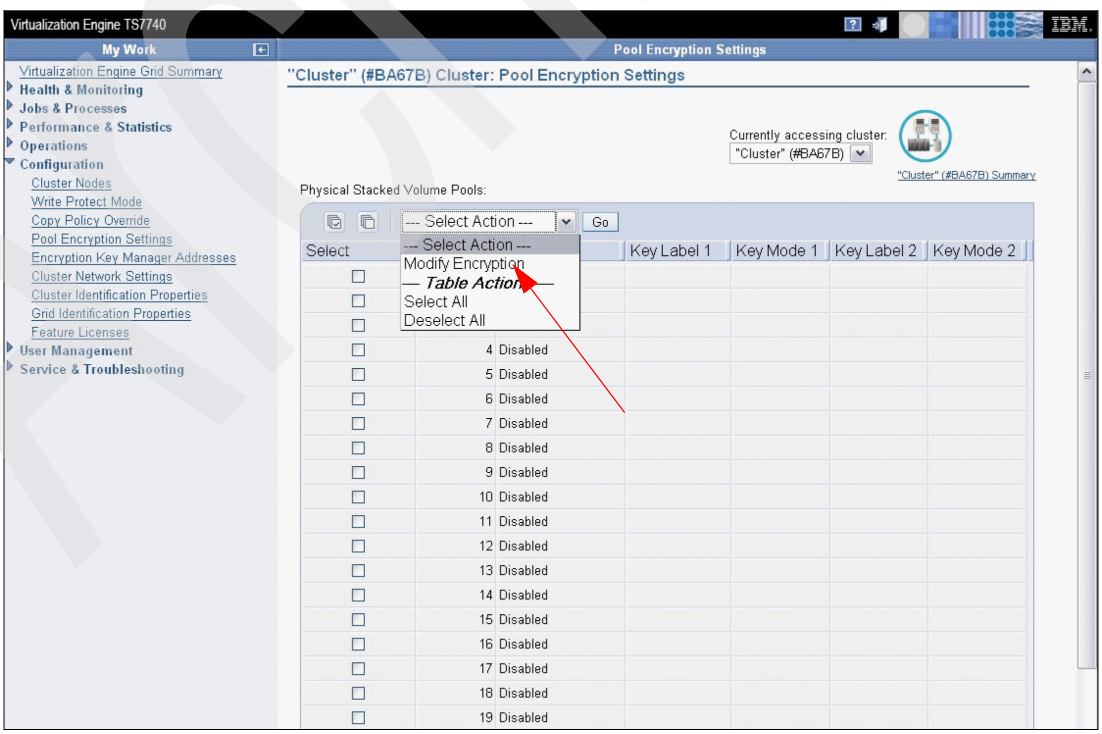

Figure 7-36 Pool Encryption Setting Step 2

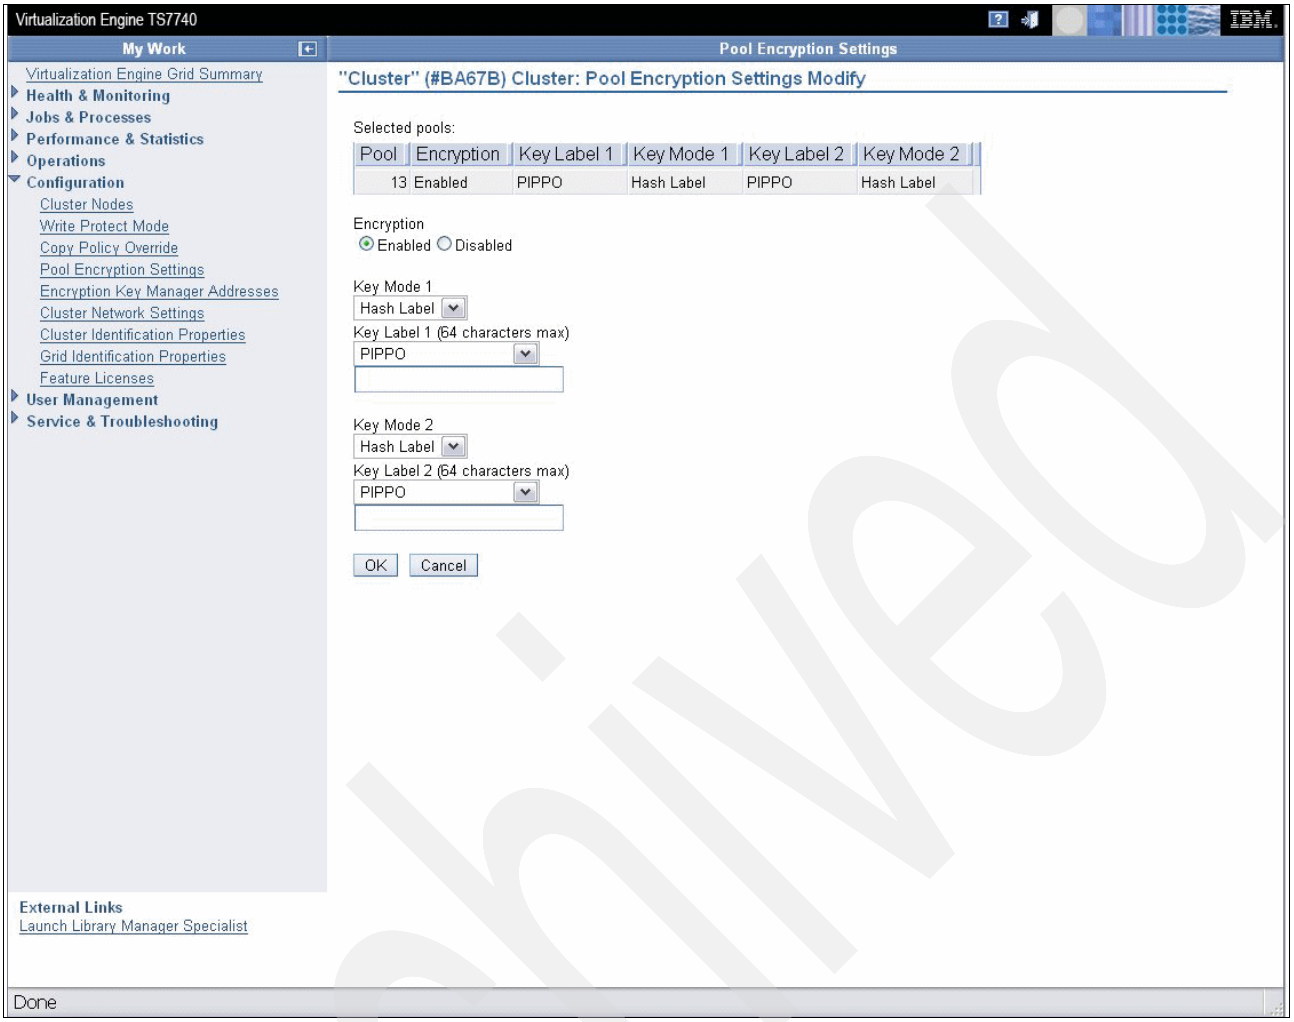

Figure 7-37 Pool Encryption Setting Step 3

You will see drives encryption capable available for encryption setting, if in the column named Encryption `enabled` is displayed for the related pool. If you attempt to modify the encryption setting from the Select Action drop-down list for a pool that shows `disable`, you receive the message shown in Figure 7-38.

Figure 7-38 TS7700 Management Interface Encryption Message

Figure 7-39 shows the successful change of the pool encryption settings.

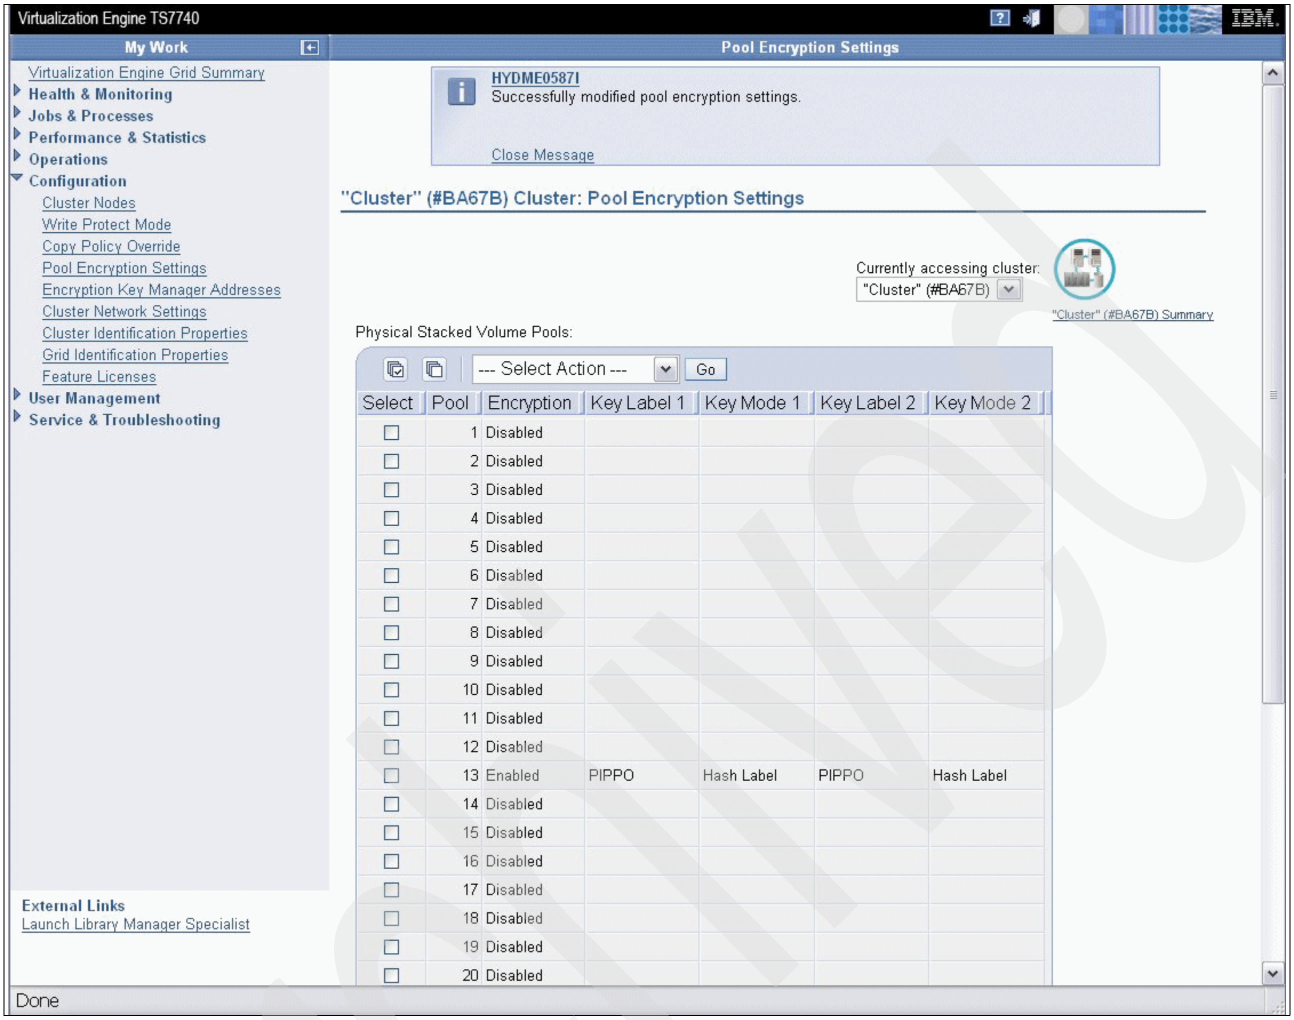

Figure 7-39 Pool Encryption Setting Step 4

Encryption Key Manager Addresses

Use this page (Figure 7-40) for setting the Encryption Key Manager Addresses in the IBM Virtualization Engine TS7700.

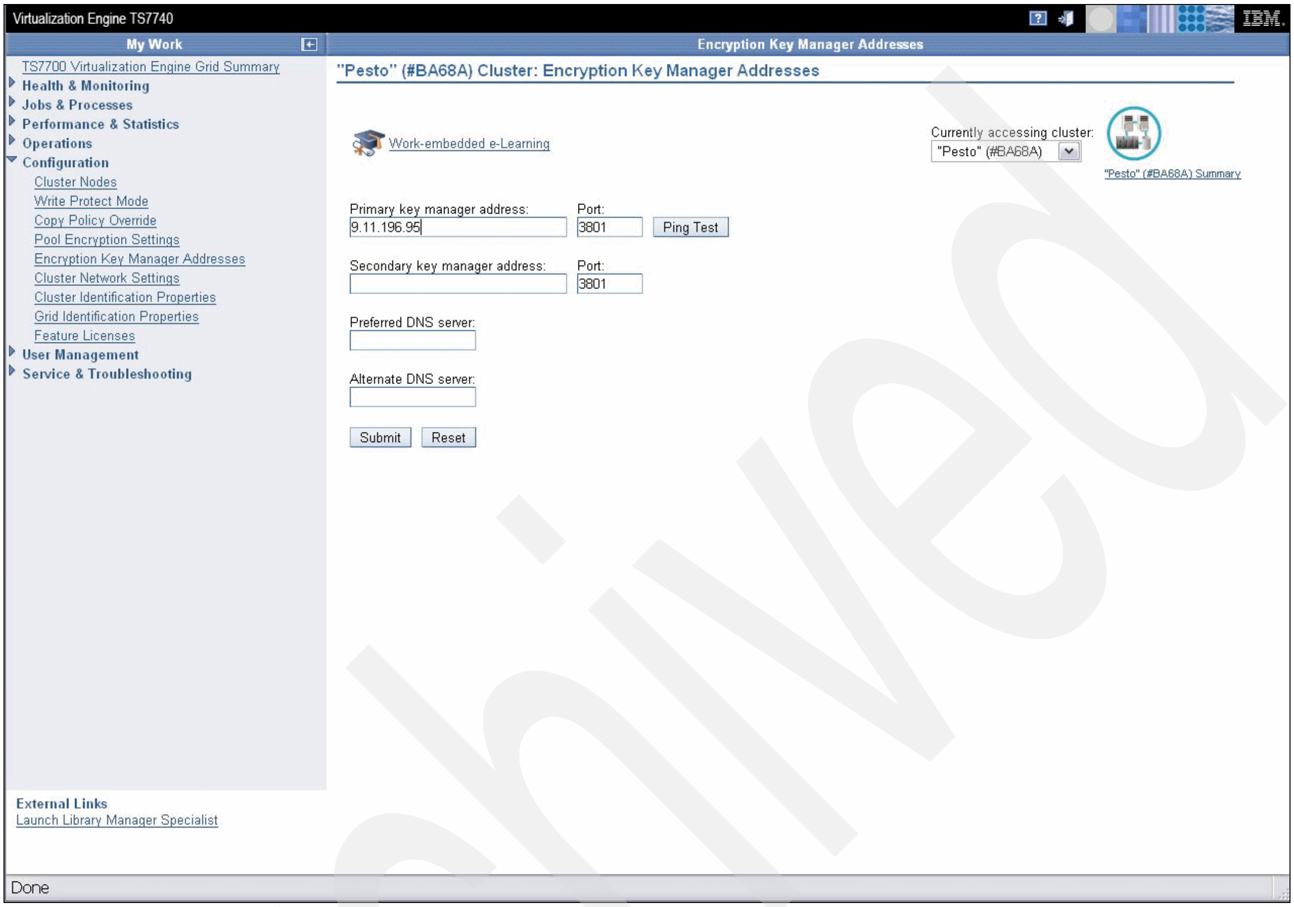

Figure 7-40 TS7700 Management Interface Encryption Key Manager Addresses

The Encryption Key Manager (EKM) assists encryption-enabled tape drives in generating, protecting, storing, and maintaining encryption keys that are used to encrypt information being written to, and decrypt information being read from, tape media (tape and cartridge formats).

The following settings are used to configure the TS7700 connection to an Encryption Key Manager.

- **Primary key manager address**: It is the key manager server name or IP address that is primarily used.
- **Secondary key manager address**: It is the key manager server name or IP address that is used when the primary key manager is unavailable.
- **Port**: The port number for the key manager address. Default setting is 3801.
- **Preferred DNS server**: The Domain Name Server (DNS) that is primarily used. DNS addresses are only needed if you specify a symbolic domain name for one of the key manager addresses rather than a numeric IP address. If you need to specify a DNS, be sure to specify both a primary and an alternate so you do not lose access to your EKM due to one of the DNS servers being down or inaccessible.

- **Alternate DNS server**: The Domain Name Server that is used in case the preferred DNS server is unavailable. If a Preferred DNS server is specified, it is recommended to specify an alternate DNS as well.

Cluster Network Settings

Use the panel shown in Figure 7-41 for viewing and altering Cluster Network Settings for the TS7700.

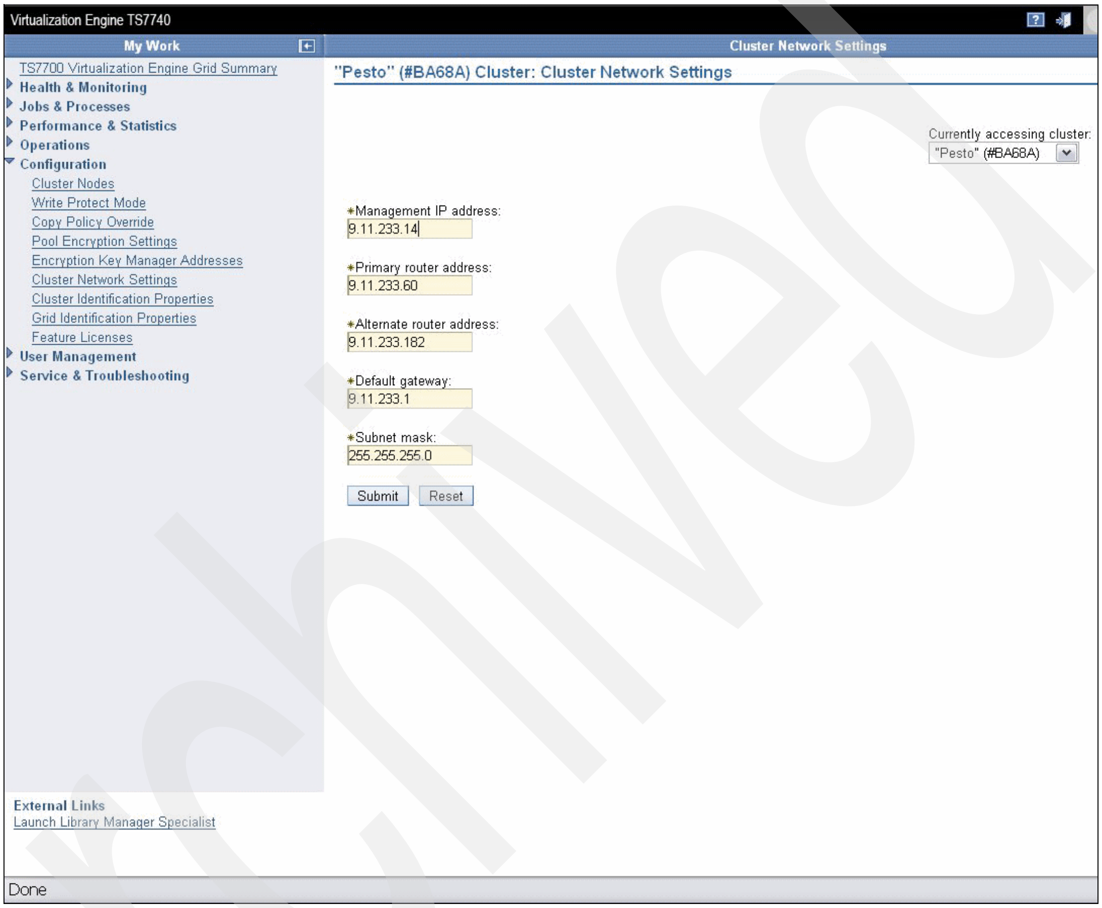

Figure 7-41 TS7700 MI Cluster Network Settings

To change cluster network settings, edit the available fields and click **Submit Changes**. The fields on the panel in Figure 7-41 contain the following:

Management IP address The IP address used for connecting to the Management Interface

Primary router address The IP address of the primary router contained in the TS7700

Alternate router address The IP address of the secondary router contained in the TS7700

Default gateway The gateway IP address of the network the cluster resides on

Subnet mask The subnet mask of the network the cluster resides on

> **Note:** Changing the Management IP address can result in loss of connection to the Management Interface.

Cluster Identification Properties

Use this page (Figure 7-42) for viewing and altering Cluster Identification Properties for the IBM Virtualization Engine TS7700.

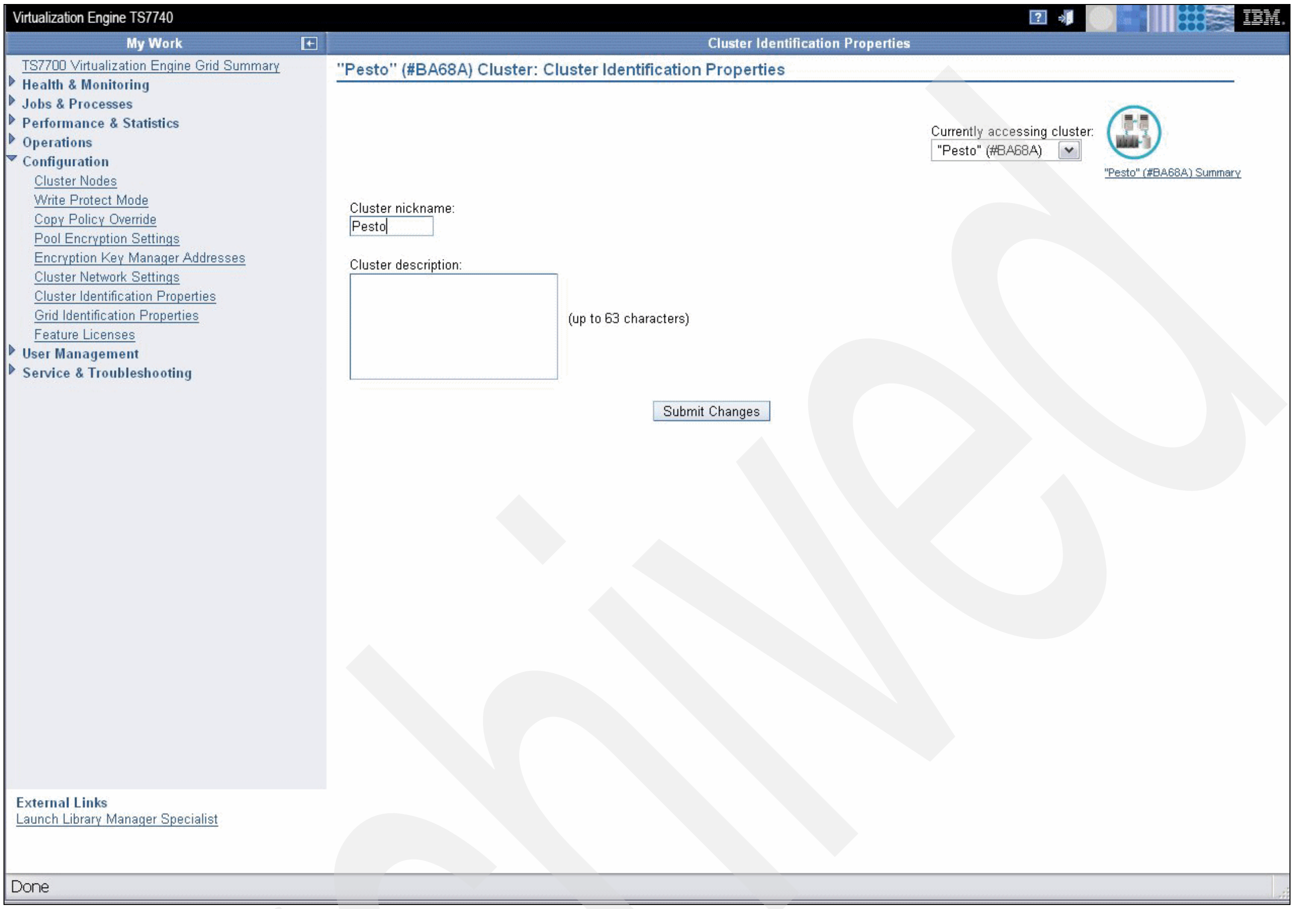

Figure 7-42 TS7700 Management Interface Cluster Identification Properties

The following information related to cluster identification is displayed in Figure 7-42.

To change cluster identification properties, edit the available fields and click **Submit Changes**.

- **Cluster nickname**: The cluster nickname must be one to eight characters in length composed of alphanumeric characters with no spaces. The characters @, ., -, and + are also allowed.
- **Cluster description**: A short description of the cluster of up to 63 characters.

Grid Identification Properties

Use this page (Figure 7-43) for viewing and altering identification properties for the IBM Virtualization Engine TS7700 Grid.

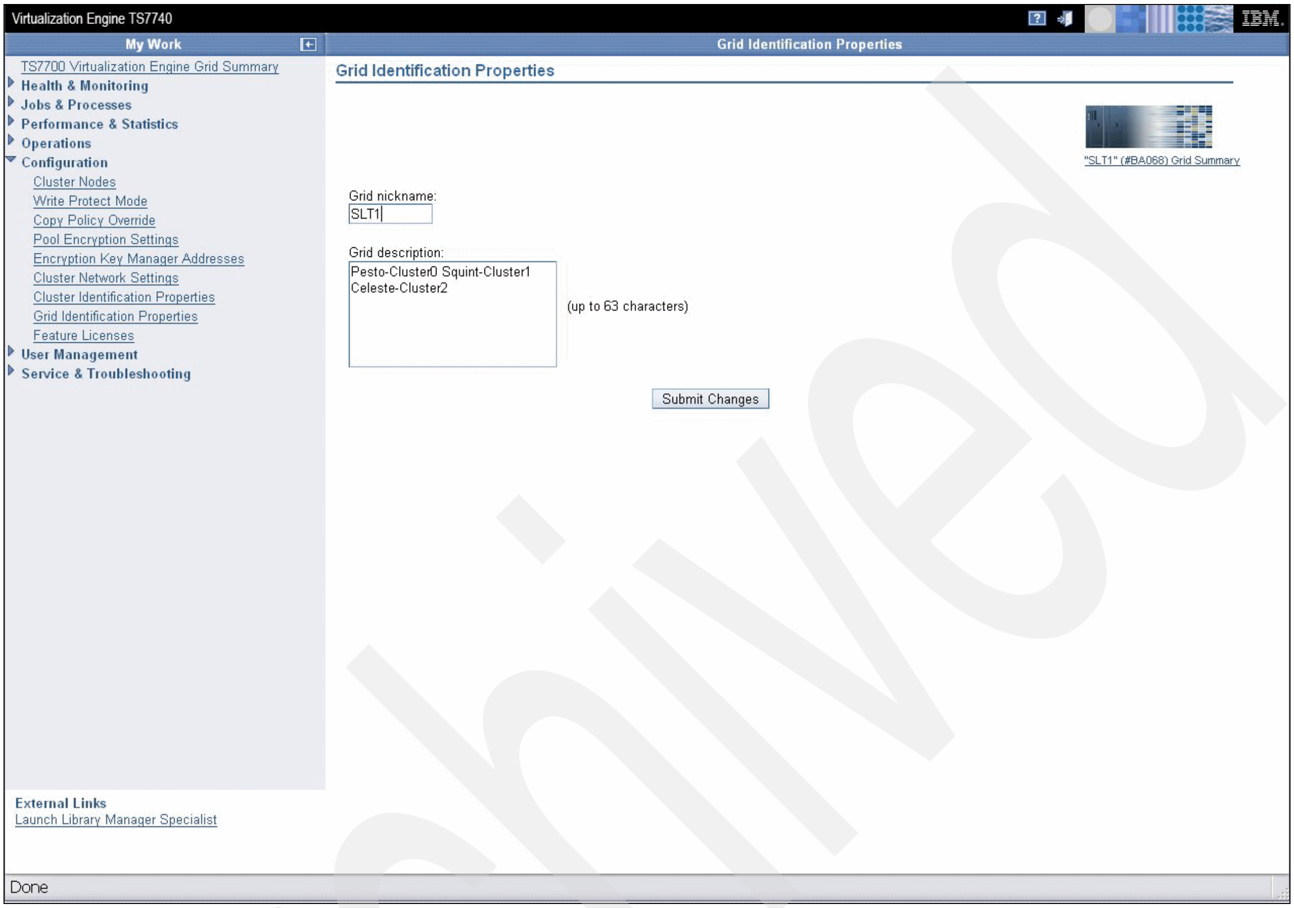

Figure 7-43 TS7700 Management Interface Cluster Grid Identification Properties

The following information related to grid identification is displayed in Figure 7-43.

To change grid identification properties, edit the available fields and click **Submit Changes**.

- **Grid nickname**: The grid nickname must be one to eight characters in length composed of alphanumeric characters with no spaces. The characters @, ., -, and + are also allowed.
- **Grid description**: A short description of the cluster of up to 63 characters.

Feature Licenses

Use the panel shown in Figure 7-44 for activating and removing feature licenses in the IBM TS7700.

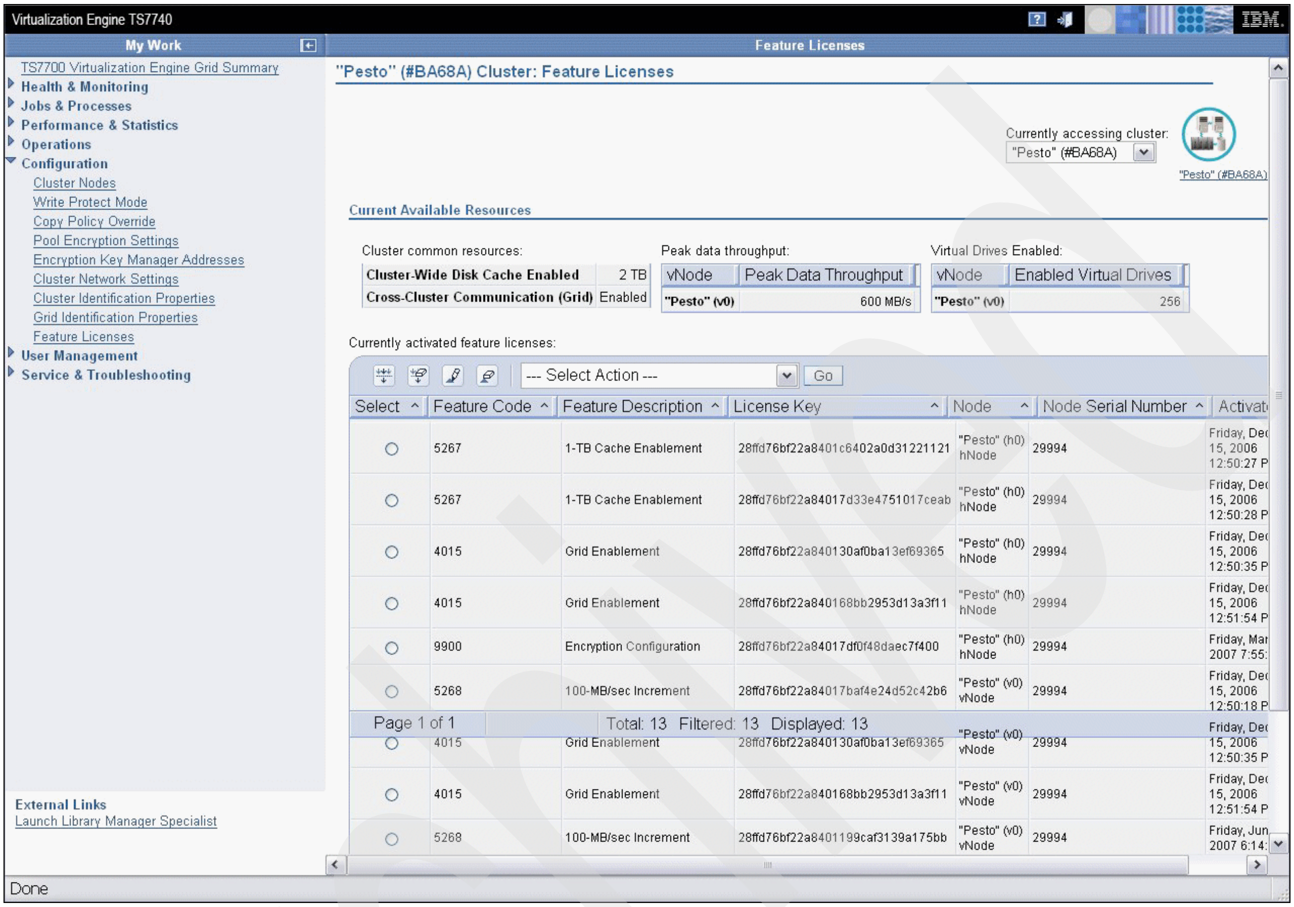

Figure 7-44 TS7700 MI feature licenses

The following information is displayed in the panel shown in Figure 7-44:

- Under “Cluster common resources”:
  - **Cluster Wide Disk Cache Enabled**: The amount of disk cache enabled for the entire cluster.
  - **Cross-Cluster Communication** (Global Mirror): Options are Enabled or Disabled. Enabling this allows multiple clusters to form a grid.
- Under “Peak data throughput”:
  - **vNode**: Name of the vNode
  - **Peak data throughput**: The upper limit of the data transfer speed between the vNode and the host
- Under “Virtual Drives Enabled”:
  - **vNode**: The name of the vNode
  - **Enabled Virtual Drives**: The amount of virtual drives enabled for the specified vNode

  Note that there are currently no license features for virtual drives.

Table Fields shown in Figure 7-44 are:

Feature Code	The feature code number of the installed feature
Feature Description	A description of the feature installed by the feature license
License Key	The 32 character license key for the feature
Node	The node the feature is installed on
Node Serial Number	The serial number of the node the feature is installed on
Activated	Date and time the feature license was activated

Activate or remove a new feature license

Use the panel shown in Figure 7-45 to activate or remove a feature license.

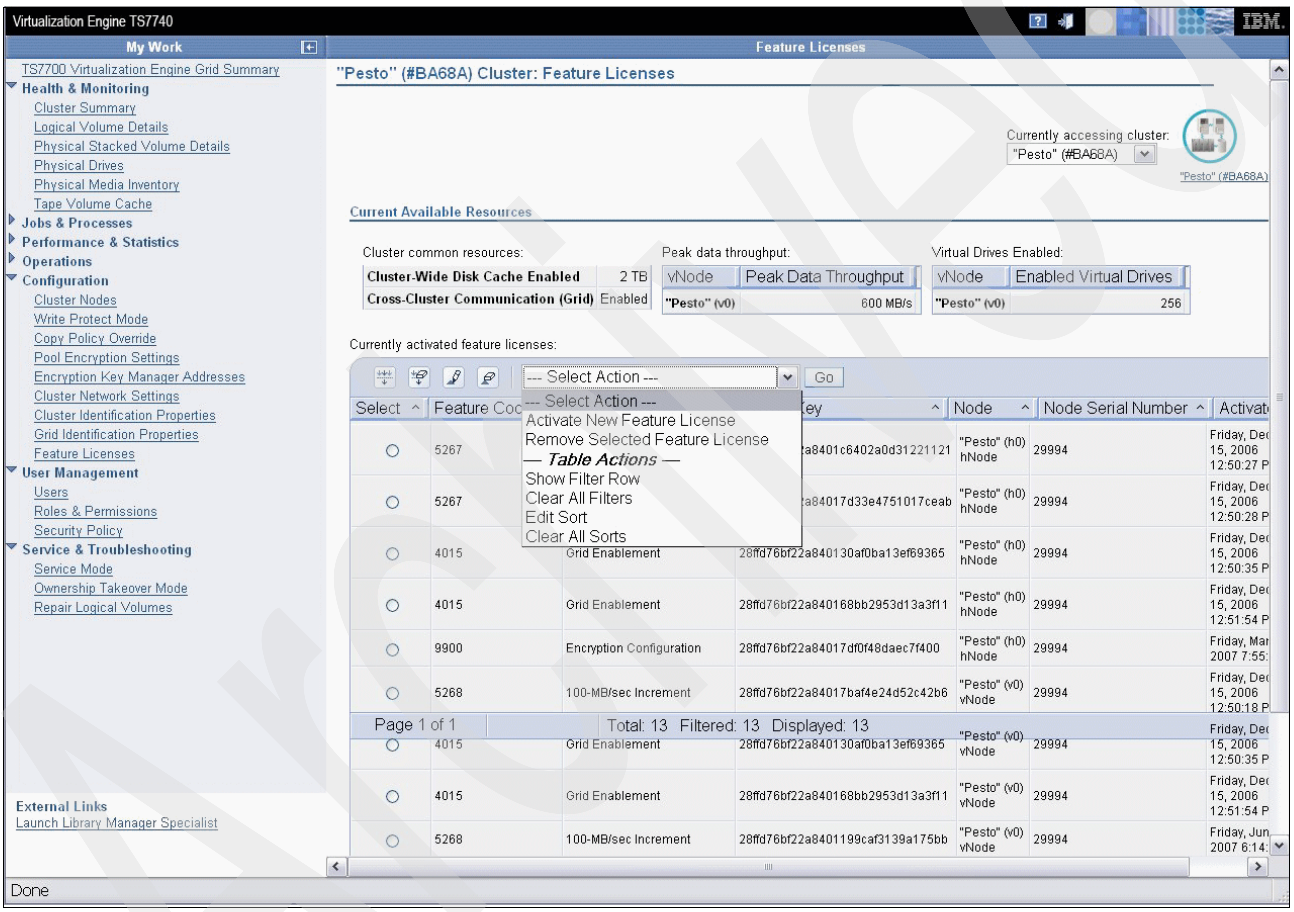

Figure 7-45 TS7700 MI Activate New Feature License

To activate a new feature license, follow these steps:

1. Select **Activate New Feature License** from the Select Action drop-down and click **Go**.
2. The Activate New Feature License page will be displayed; see Figure 7-46 on page 357. Enter the 32 character feature license in the available text field and click **Activate** to continue or **Cancel** to cancel the activation.
3. The Confirm Feature Activation page will display details of the feature license. To activate the feature license click **Yes**, to cancel activation click **No**.

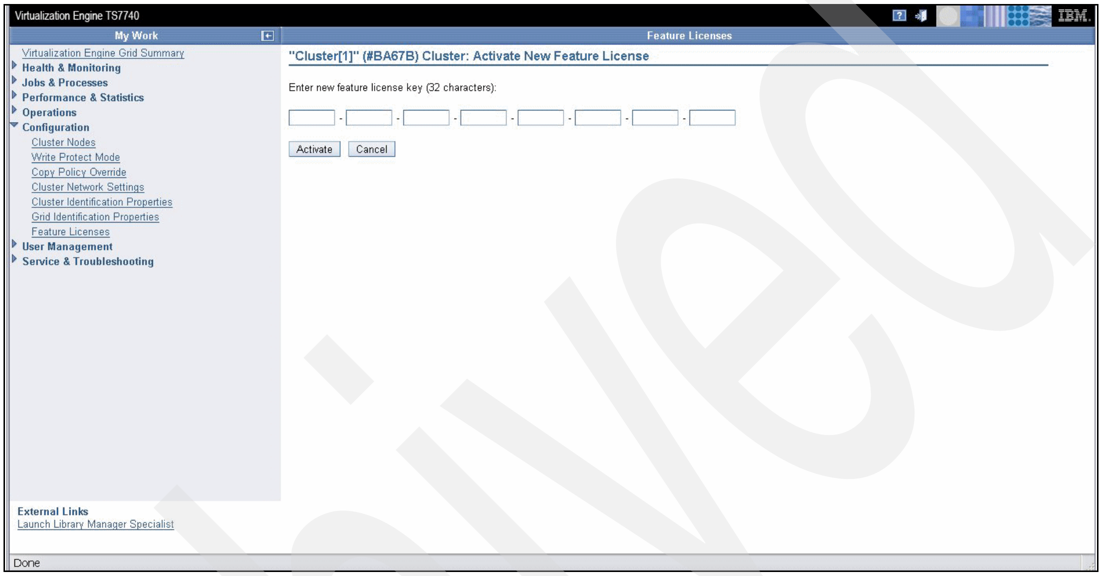

Figure 7-46 TS7700 MI 32 character feature license

A feature license can also be removed by selecting a feature license in the table, selecting **Remove Selected Feature License** from the Select Action drop-down, and clicking **Go**.

> **Note:** You cannot remove Cache Enablement Features (FC5267) while the TS7700 is operating, because there are currently no options to add those features again.

7.2.8 User Management

The topics in this section present information that is related to managing user access in the IBM TS7700 Virtualization Engine.

In this section, we include the following topics:

- Users
- Roles & Permissions
- User roles
- Security Policy

Users

Use the panel shown in Figure 7-47 to add, modify, or remove users in the IBM TS7700 Virtualization Engine.

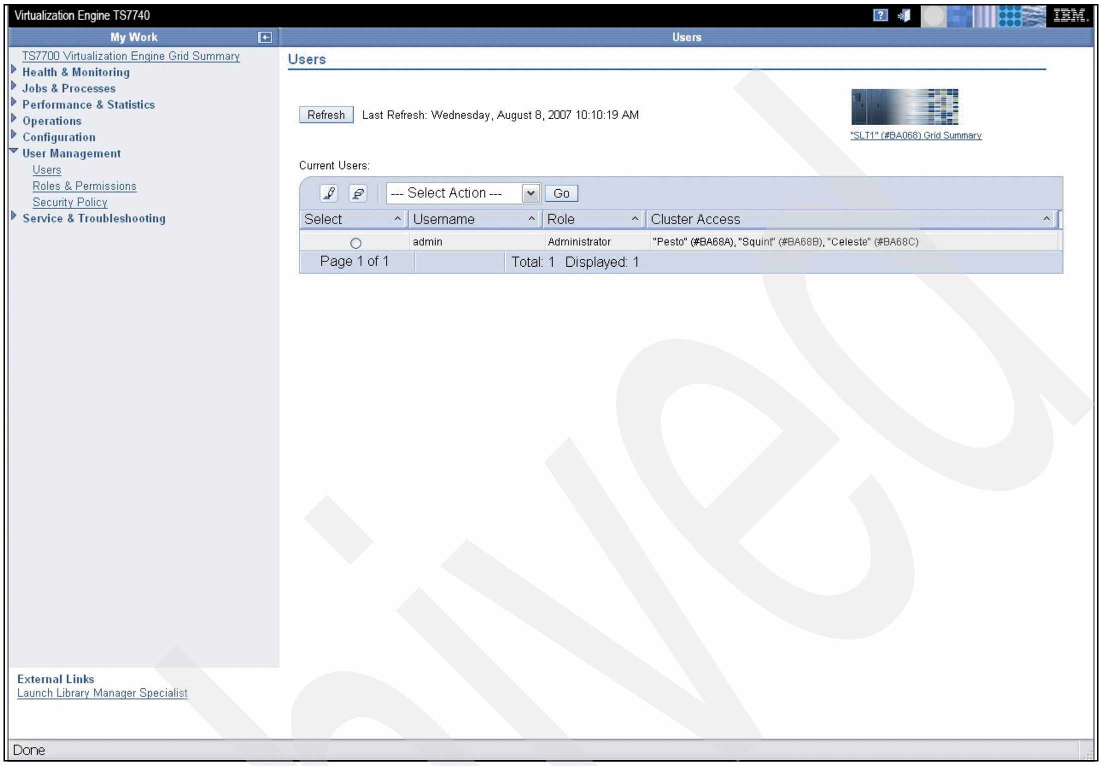

Figure 7-47 TS7700 Management Interface Users

Manage users - add new user

To add a new user, select **Add** from the Select Action drop-down list and click **Go**. Then, follow these steps:

1. Enter a Username. The user name must be one to 16 characters in length composed of alphanumeric characters with no spaces. The characters **@**, **.**, and _ are also allowed.
2. Enter a Password. This password must be six to 16 characters in length, composed of alphanumeric characters with no spaces. The first or last character of the password cannot be a number.
3. Reenter the same password in the Confirm password field.
4. Enter the Email address of the user.
5. Select the **Role of user** from the drop-down menu. A predefined or custom role can be selected. Custom roles will display (`custom`) next to their name. A custom role will only be shown if it has been configured and given a unique name. For details on the predefined and custom role permissions and to define custom roles, click **View permissions of default roles and configure custom roles** to be taken to the Roles & Permissions page.
6. Select the clusters the user should have access to. More than one can be selected.
7. Click **OK** to create the user or click **Cancel** to abort creating a user.

You are then taken back to the Users page. If a user was created, it will now be available in the table.

Manage users - modify user

Use this page to modify a user in the IBM TS7700. This page is also where you can change user passwords.

To modify a user from the Users page:

1. Select a user from the table.
2. Select **Modify** from the Select Action drop-down list and click **Go**.
3. Edit the user's information.
4. Click **OK** to modify the user or click **Cancel** to abort modifying a user.

You are then taken back to the Users page. If a user was modified, updated information is displayed in the table.

Manage users - remove user

To remove a user from the IBM TS7700, from the Users page:

1. Select a user from the table.
2. Select **Remove** from the Select Action drop-down list and click **Go**.
3. A confirmation panel will be presented. Select **Yes** to remove the user or **No** to cancel the action.

The user will be removed from the system.

Roles & Permissions

Use the panel shown in Figure 7-48 to modify permissions for custom user roles and to view the predefined user roles in the IBM Virtualization Engine TS7700 Grid.

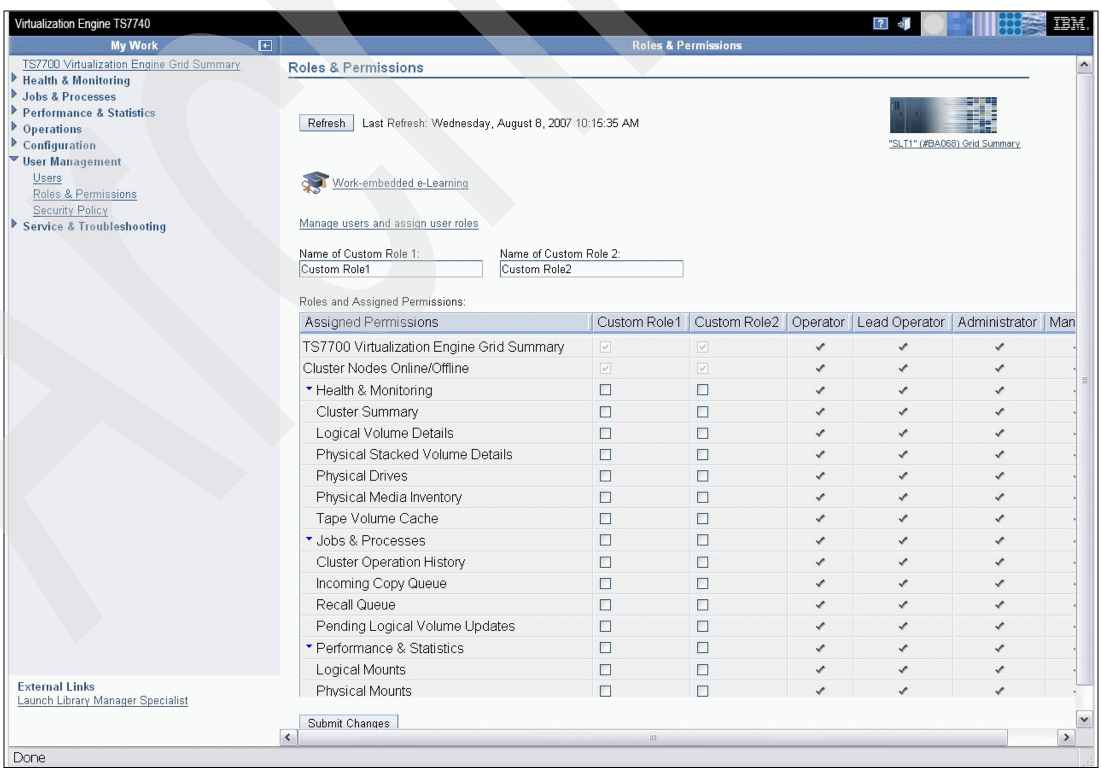

Figure 7-48 TS7700 MI Roles & Permissions

The user-defined roles can be renamed at the top of the page. To edit them, change their name in the text boxes and click **Submit Changes** at the bottom of the page. Clicking the **Manage users and assign user roles** link opens the Users page.

User roles

Users of the TS7700 Virtualization Engine can be assigned one or more roles associated with performance of certain system functions.

Table 7-2 shows the available tasks and how they map to the predefined user roles. The table also allows for assigning tasks to the user-defined roles. Check marks indicate whether a particular task is assigned to a role. Tasks for custom roles are editable while predefined roles are not.

Admin

The administrator has the highest level of authority, including the authority to add or remove user accounts. The administrator has access to all service functions and TS7700 resources.

Lead Operator

The lead operator has almost all of the same permissions as the administrator but cannot change network configuration, feature licenses, user accounts, or custom roles.

Manager

The manager has access to monitoring information and performance data and functions, but is restricted from changing most settings, including those for logical volume management, network configuration, feature licenses, user accounts or custom roles.

Operator

The operator has access to monitoring information and can insert logical volumes but is restricted from changing settings for performance, network configuration, feature licenses, user accounts or custom roles. The operator is also restricted from deleting logical volumes.

Custom roles

The administrator can name and define two custom roles by selecting the individual tasks permitted to each custom role. In Table 7-2 the question mark (?) symbol in the custom role columns indicates that the capability listed in that row can be defined for the custom role.

Table 7-2 User role capabilities

Capability	Custom role 1	Custom role 2	Operator	Lead operator	Manager	Administrator
TS7700 VE Grid summary	?	?	X	X	X	X
Cluster Nodes Online/Offline	?	?	X	X	X	X
Health & Monitoring						
Cluster summary	?	?	X	X	X	X
Logical volume details	?	?	X	X	X	X
Physical Stacked Volume Details	?	?	X	X	X	X
Physical Drives	?	?	X	X	X	X
Physical Media Inventory	?	?	X	X	X	X
Tape Volume Cache	?	?	X	X	X	X

Capability	**Custom role 1**	**Custom role 2**	**Operator**	**Lead operator**	**Manager**	**Administrator**
Jobs & Processes						
Cluster Operation History	?	?	X	X	X	X
Incoming Copy Queue	?	?	X	X	X	X
Recall Queue	?	?	X	X	X	X
Pending Logical Volume Updates	?	?	X	X	X	X
Performance & Statistics						
Grid network path utilization	?	?	X	X	X	X
Logical Mounts	?	?	X	X	X	X
Physical Mounts	?	?	X	X	X	X
Host Throughput	?	?	X	X	X	X
Cache Throttling	?	?	X	X	X	X
Cache Utilization	?	?	X	X	X	X
Grid Network Throughput	?	?	X	X	X	X
Operations						
Insert Logical Volumes	?	?		X		X
Delete Logical Volumes	?	?		X		X
Standalone Mount Logical Volume	?	?				X
Perform Standalone mount/Demount	?	?	X	X	X	X
Configuration						
Cluster Nodes	?	?	X	X	X	X
Cluster Nodes Details	?	?	X	X	X	X
Cluster Nodes Modify Node Nickname	?	?				X
Write protect mode	?	?	X	X	X	X
Set write protect mode	?	?				X
Copy Policy Override	?	?	X	X	X	X
Modify Copy Policy Override	?	?				X
Pool Encryption Settings	?	?	X	X	X	X
Modify Pool Encryption Settings	?	?				X
Encryption Key Manager Addresses	?	?	X	X	X	X
Modify Encryption Key Manager Addresses	?	?				X

Capability	Custom role 1	Custom role 2	Operator	Lead operator	Manager	Administrator
Cluster Network Settings	?	?	X	X	X	X
Modify Cluster Network Settings	?	?				X
Cluster Identification Properties	?	?	X	X	X	
Modify Cluster Identification Properties	?	?				X
Grid Identification Properties	?	?	X	X	X	X
Modify Grid Identification Properties	?	?				?
Feature Licenses	?	?	X	X	X	X
Feature License Modify	?	?				?
User Management						
Users	?	?	X	X	X	X
Add User	?	?	X	X	X	X
Modify User	?	?			X	X
Remove User	?	?			X	X
Roles & Permissions	?	?	X	X	X	X
Modify Custom Roles	?	?			X	X
Security Policy	?	?	X	X	X	X
Modify Security Policy	?	?			X	X
Service & Troubleshooting						
Service Mode	?	?	X	X	X	X
Set Service Mode	?	?				X
Ownership Takeover Mode	?	?	X	X	X	X
Ownership Takeover Mode Modify	?	?				X
Repair Logical Volumes	?	?	X	X	X	X
Repair Logical Volumes Confirm	?	?		X		X
Damaged Logical Volumes Details	?	?	X	X	X	X
Copy Export Recovery	?	?	X	X	X	X
Start a Copy Export Recovery or a Clean Up	?	?		X		X
Copy Export Recovery State	?	?	X	X	X	X
Launch Library Manager Specialist	?	?	X	X	X	X

Security Policy

Use this page (Figure 7-49) to view and change user security policy settings for the IBM Virtualization Engine TS7700 Grid.

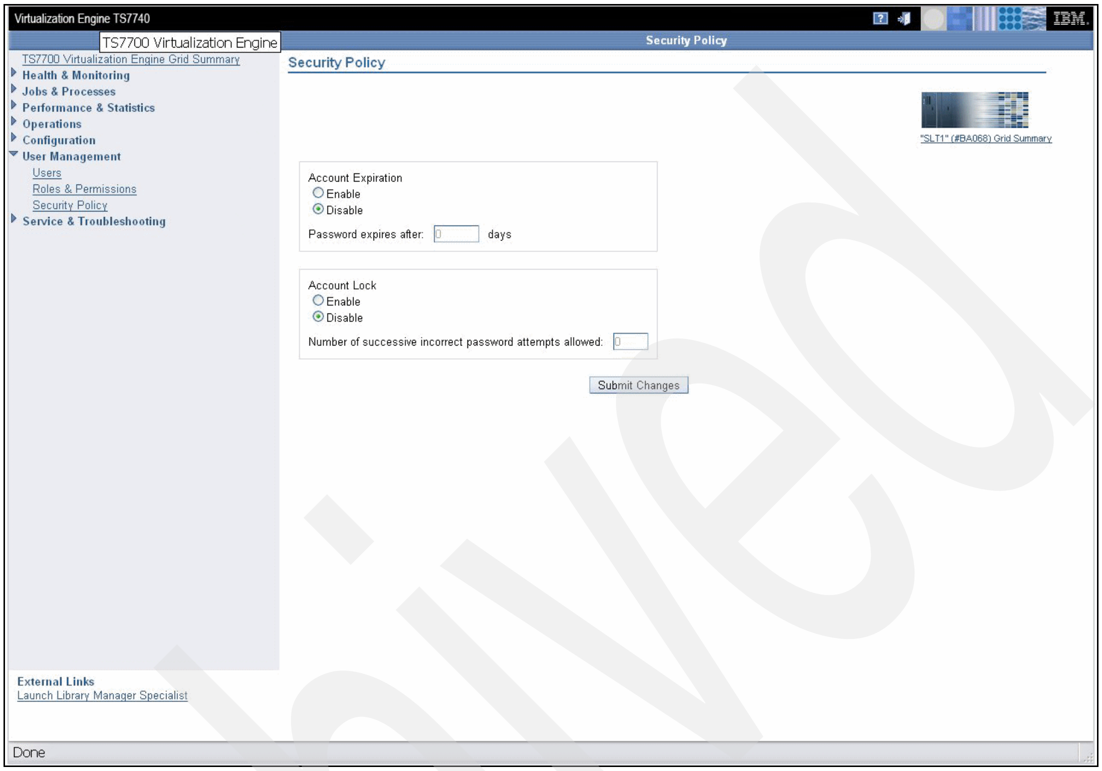

Figure 7-49 TS7700 Management Interface Security Policy

The following user security policy settings are available. To change any setting, edit the available field and click **Submit Changes**.

- Under Account Expiration:

  Enable Allow accounts to expire after a set number of days from account creation. The amount of days can be set in the **Password expires** after text box.

  Disable User accounts will not expire.

- Under Account Lock:

  Enable The account will be allowed a certain amount of failed login attempts for a user before locking the account out of logging into the management interface. The number of failed attempts is entered in the “Number of successive incorrect password attempts allowed text box”.

  Disable User accounts will not be locked out due to inputting the wrong password.

Note: To unlock an account, an administrator can modify the user account and change the password from the Users page.

7.2.9 Service & Troubleshooting

The topics in this section present information that is related to performing service operations and troubleshooting problems in the IBM Virtualization Engine TS7700.

In this section, we discuss the following topics:

- Service Mode
- Ownership Takeover Mode
- Repair Logical Volumes

Service Mode

From the panel shown in Figure 7-50, you can perform limited service operations in the IBM TS7700 Virtualization Engine.

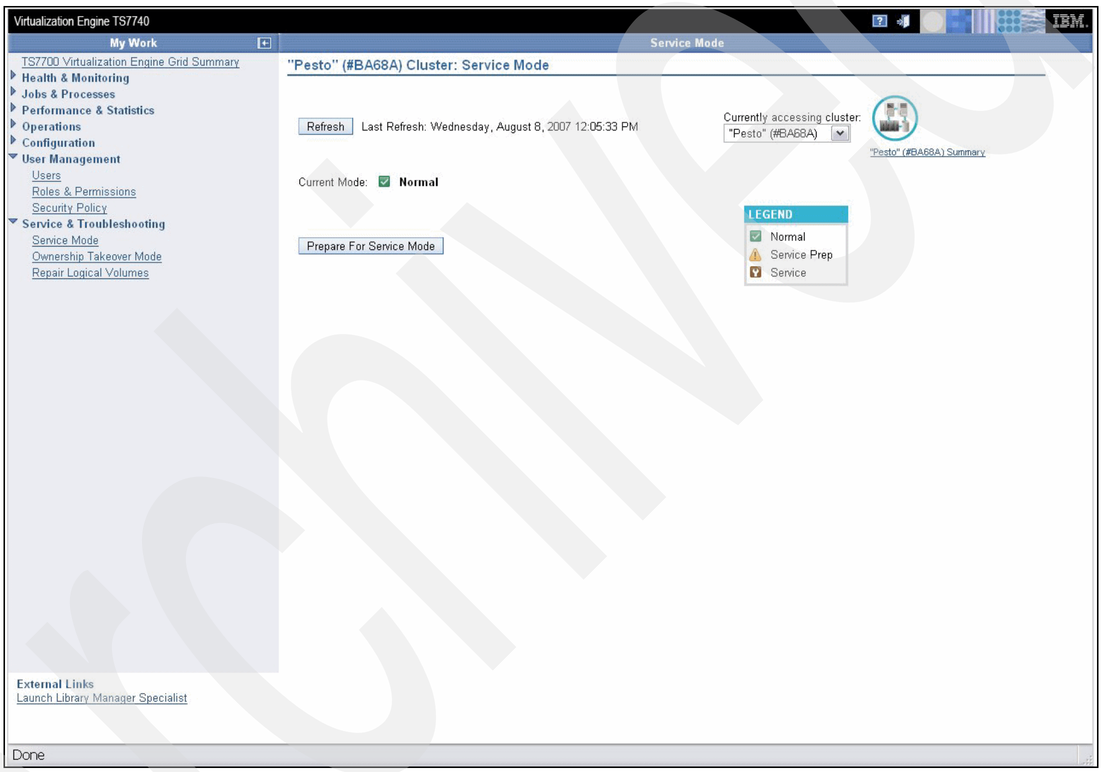

Figure 7-50 TS7700 MI Cluster Service Mode

From the panel shown in Figure 7-50, you can set the TS7700 Cluster in Service Mode or take it back into Normal Mode. The Current Mode field can show:

Normal The cluster is in a normal operation state and can be brought into Service Mode.

Service The cluster is in Service Mode.

> **Note:** Service Mode is an Administrator function. When Service Mode is enabled, all logical devices go offline.

Depending on what mode the cluster is in, a different action is presented on the button below where the Current Mode is displayed.

Ownership Takeover Mode

Use this page (Figure 7-51) to enable or disable Ownership Takeover Mode for a failed cluster in the IBM Virtualization Engine TS7700.

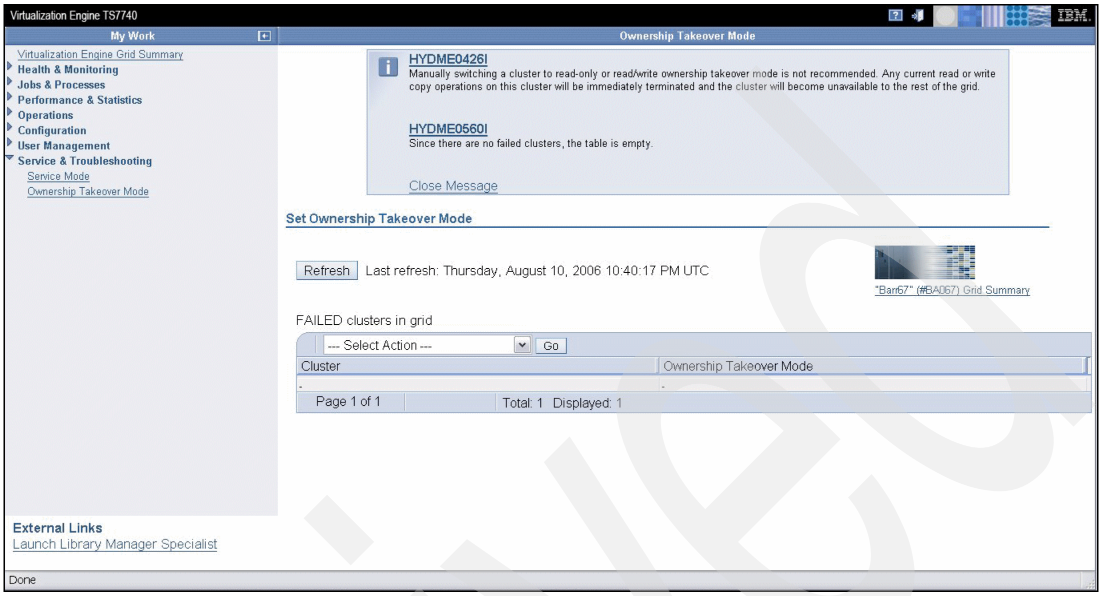

Figure 7-51 TS7700 MI Ownership Takeover Mode

The following Ownership Takeover Mode information is available for failed clusters:

- **Cluster**: The failed cluster's name.
- **Ownership Takeover Mode** (Figure 7-52 on page 366):
  - **No Ownership Takeover**: The failed cluster is not in any ownership takeover mode. For a cluster that is in a failed state, enabling ownership takeover mode allows other clusters in the grid to obtain ownership of logical volumes owned by the failed cluster. Normally, ownership is transferred from one cluster to another through communications between the clusters. When a cluster fails or the communication path between clusters fails, the normal means of transferring ownership is not available. Enabling a read/write or read-only takeover mode should not be done if only the communication path between clusters has failed. A mode should only be enabled for a cluster that is no longer operational. Integrity of logical volumes in the grid can be compromised if a takeover mode is enabled for a cluster that was not actually in a failed state.
  - **Read/Write Ownership Takeover**: Read/write ownership takeover mode allows other non-failed clusters in the grid to perform read and write operations as well as change to private or scratch status on logical volumes owned by the failed cluster. Ownership of the logical volume can only be obtained if there is a consistent copy of the logical volume available on the grid or if the logical volume is in a scratch category. If there was no cluster failure, there is the possibility that two sites could be writing data to the same logical volume.
  - **Read-only Ownership Takeover**: Read-only ownership takeover mode allows other non-failed clusters in the grid to perform read operations on the logical volumes owned by the failed cluster. The status of the volume, private/scratch cannot be changed in this mode. Scratch mounts are failed to a cluster in this mode. If there was no cluster failure, there is the possibility a logical volume accessed by another cluster in Read-only takeover mode contains older data than the one on the owning cluster. This

can occur if the logical volume was modified on the owning cluster while the communication path between the clusters was down.

If the Autonomic Ownership Takeover Manager (AOTM) is installed and configured, it will perform an additional check to determine if the unavailable cluster is in a failed state when an ownership takeover is initiated. If AOTM determines that the cluster is not in a failed state, the ownership takeover attempt will fail.

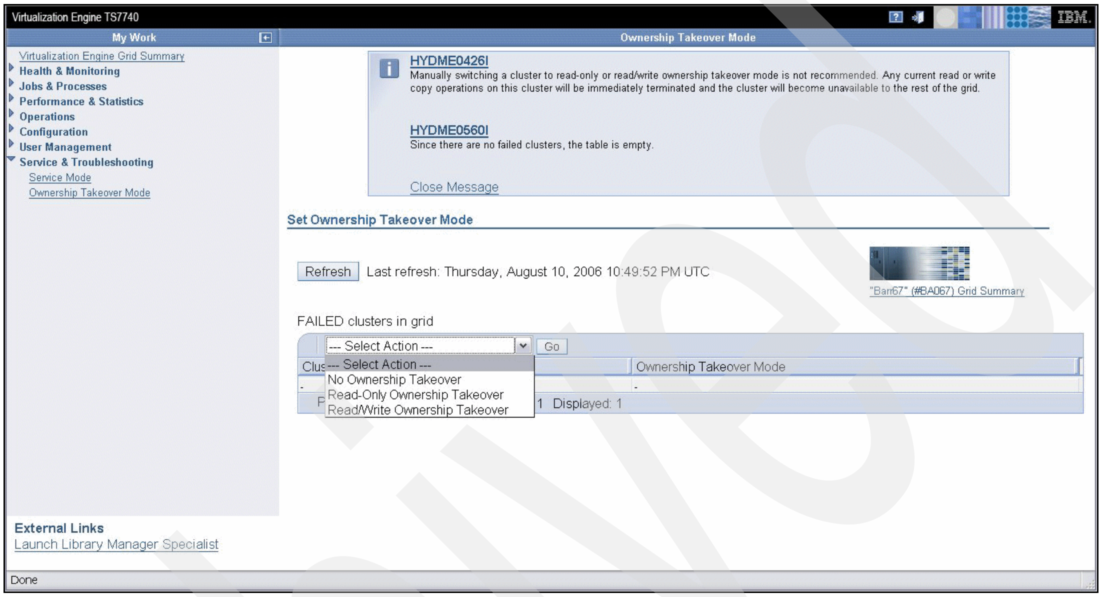

Figure 7-52 TS7700 MI Ownership Select Action

Repair Logical Volumes

Use this panel (Figure 7-53) to repair logical volumes in the damaged category for the IBM Virtualization Engine TS7700 Grid.

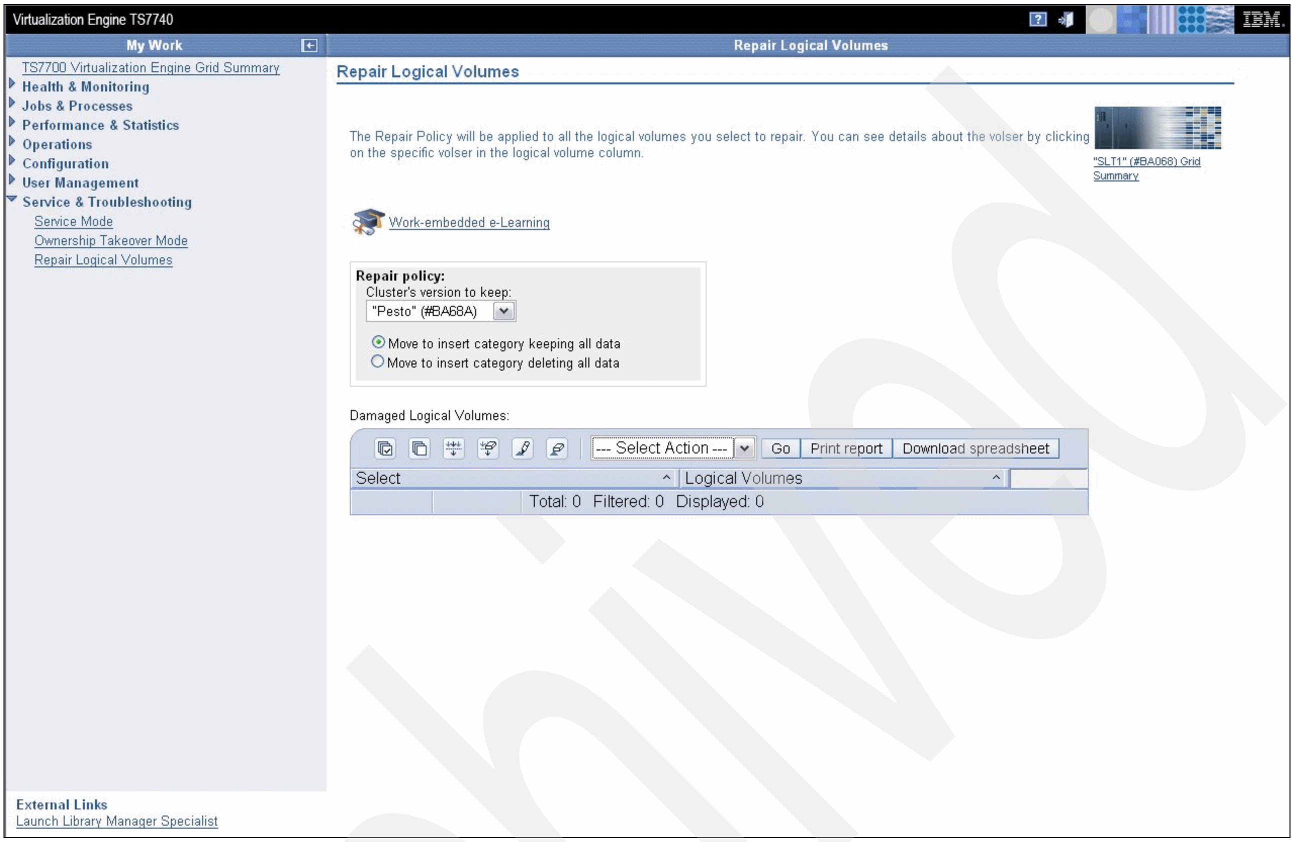

Figure 7-53 TS7700 Management Interface Repair Logical Volumes

- **Repair Policy**: The Repair Policy will be applied to all the logical volumes you select to repair. You can see details about the volser by clicking the specific volser in the Logical Volume column. At the top of the panel shown in Figure 7-53 you see the repair policy that will be applied to all damaged logical volumes selected. Table fields are:
  - **Cluster's version to keep**: The selected cluster will obtain ownership of the logical volume when the repair is complete. The cluster's version of the logical volume will be the basis for repair if the “Move to insert category keeping all data” option is selected.
  - **Move to insert category keeping all data**: This option should be used if the data on the logical volume is intact and still relevant. If there has been data loss, then this option should not be used. If the cluster chosen in the repair policy has no data for the logical volume to be repaired, choosing this option will be the same as choosing “Move to insert category deleting all data.”
  - **Move to insert category deleting all data**: The repaired logical volumes will be moved to the insert category and all data will be erased. This option should be used if the volume has been returned to scratch or if data loss has rendered the volume obsolete. If the volume has been returned to scratch, then the data on the volume is no longer needed. If data loss has occurred on the volume, then data integrity issues could occur if the data on the volume is not erased.

- **Damaged Logical Volumes**: The following information about the damaged logical volumes is presented in a table and should assist in selecting a cluster for a repair action if data is to be retained. The information is presented for each cluster, as follows:
  - **Cluster**: The cluster name.
  - **Last Modified**: Date and time the logical volume was last modified.
  - **Last Mounted**: Date and time the logical volume was last mounted.
  - **Data Exists**: Possible values are:
    - **Exists**: Data exists in the local cache or was migrated to physical tape.
    - **Does not exist**: Data for this logical volume does not exist on this cluster.
  - **Size/MBs**: Size of data on the logical volume in megabytes.
  - **Category**: Category attribute of the logical volume. This is used to denote a grouping a logical volumes.
  - **Media Type**: Media type of the logical volume.
  - **Ownership Takeover Time**: Date the last time an ownership takeover occurred for this logical volume.
  - **Data Level**: Every time the logical volume is written to, this value increases. Inserted logical volumes start with a data level of 100. This is a secondary factor to Insert Level when choosing a cluster for the repair policy.
  - **Insert Level**: When a group of logical volumes is inserted, they are assigned an Insert Level. Later inserts are given a higher level. This is the most important factor when choosing a cluster in the repair policy if data is going to be retained. A higher value means higher consistency for the logical volume's data.
  - **Data Consistent**: If the cluster's logical volume copy's data level is considered the latest data level, then data is consistent.

7.2.10 Drive cleaning

Each drive in the IBM 3494 and in TS3500 Tape Libraries needs cleaning. The 3592 drives and TS1120 drives request the cleaning by themselves if a cleaning is needed. To insert a cleaning cartridge in a 3494 is simple: you just insert the cleaning cartridge into the Convenience I/O Station and the Library Manager takes care of inserting it in the library and Library Manager database. This section details this procedure.

The cleaning in the IBM TS3500 does not involve the Library Manager at all. It is the responsibility of the TS3500 Tape Library itself. However, you must enable the automatic cleaning, and provide the necessary cleaning cartridges.

> **Important:** ALMS as a requirement for System z attachment; it is always installed in a TS7700 z/OS environment, has automatic cleaning enabled, and cannot be disabled.

Inserting cleaning cartridges into an IBM 3494

The tape drives need to be cleaned on a regular basis, and cleaning requires special cleaning cartridges. To ensure that the operator can identify a cleaning cartridge, external labels are placed on these cartridges by the manufacturer. However, the 3590 tape drive requires a 3590 cleaning cartridge, and the 3592 or TS1120 tape drives require the 3592 cleaning cartridges. Within a library cleaning cartridges are identified by a particular set of volume serial masks, typically CLN***. The 3592 cleaning cartridges must be further identified by media type label (JA after the VOLSER). The cleaner volumes in the tape library also need an external label. From the Library Manager console you must define one or more masks that identify which VOLSERs identify cleaning volumes. For major information, refer to *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

After a cleaner mask has been defined, you just insert the cleaning cartridge into the Convenience Input Output Station (CIO) of the 3494 Tape Library, and the Library Manager will take care of inserting it into the library and Library Manager database.

Removing cleaning cartridges from an IBM 3494

You can eject a cleaning cartridge by using a Library Manager command. The cleaning cartridge is automatically ejected if the end-of-life time (depending on the clean mode) is reached. Ejecting a cleaning cartridge is, therefore, typically unnecessary. If it is necessary for any reason, refer to *IBM TotalStorage Automated Tape Library (3494) Operator Guide*, GA32-0449.

Inserting cleaning cartridges into a TS3500

The process to insert cleaning cartridges varies depending on the setup of the IBM TS3500 Tape Library.

You can use the IBM TS3500 Specialist or the library's operator panel to insert a cleaning cartridge.

To insert a cleaning cartridge using the IBM TS3500 Specialist, perform the following steps:

1. Open the door of the I/O station and insert the cartridge so that the bar code label faces the interior of the library and the write-protect switch is on the right.
2. Close the door of the I/O station.
3. Type the Ethernet IP address on the URL line of the browser and press Enter. The Welcome Page displays.
4. Select **Manage Cartridges** → **I/O Station**. The I/O Station panel displays.

Follow the instructions on the panel.

Removing cleaning cartridges from a TS3500

In this section we describe how to remove a cleaning cartridge by using the TS3500 Tape Library Specialist. You can also use the Operator Panel. Refer to *IBM System Storage TS3500 Tape Library Operator Guide*, GA32-0560 for more information.

Perform the following steps:

1. Type the Ethernet IP address on the URL line of the browser and press Enter. The Welcome Page displays.
2. Select **Manage Cartridges** → **Cleaning Cartridges**. The Cleaning Cartridges panel displays.
3. Follow the instructions on the panel to specify the cleaning cartridge and remove it to the I/O station.
4. Look at the Activity panel on the operator panel to determine whether the I/O station that you want to use is locked or unlocked. If the station is locked, use your application software to unlock it.
5. Open the door of the I/O station and remove the cleaning cartridge.
6. Close the door of the I/O station.

Determine the Cleaning Cartridge Usage in the IBM 3494

To determine the cleaning cartridge usage in the 3494 Tape Library, you can select from the IBM 3494 Library Manager console (see Figure 7-54 for reference) the following items:

- Database
- Search database
- Volser, Category, Device, and so forth

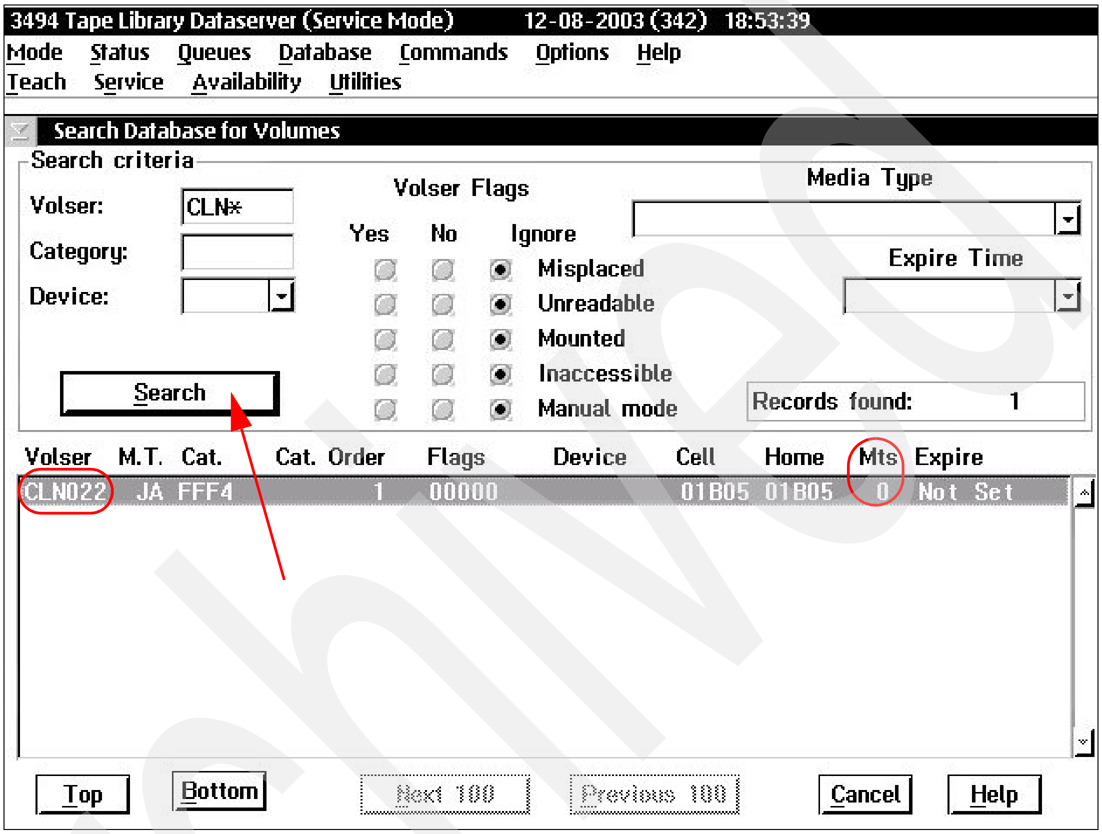

Figure 7-54 IBM 3494 cleaning usage

Whenever you need to know the cleaning cartridge usage, the panel shown in Figure 7-54 allows you to enter CLN* in the Volser field, and click **Search**. A list of cleaning cartridges will be shown, with the respective Volser. The Mts column identifies the Mount usage for the specific cleaning cartridge.

In the example shown in Figure 7-54, you can see that the mount usage for the cleaning cartridge that has Volser CLN022 is actually zero. For more information, refer to *IBM TotalStorage Automated Tape Library (3494) Operator Guide,* GA32-0449.

Determine the Cleaning Cartridge Usage in the IBM TS3500

You can determine the usage of the cleaning cartridge in the same panel that is used for the removal of the cleaning cartridges; see Figure 7-55.

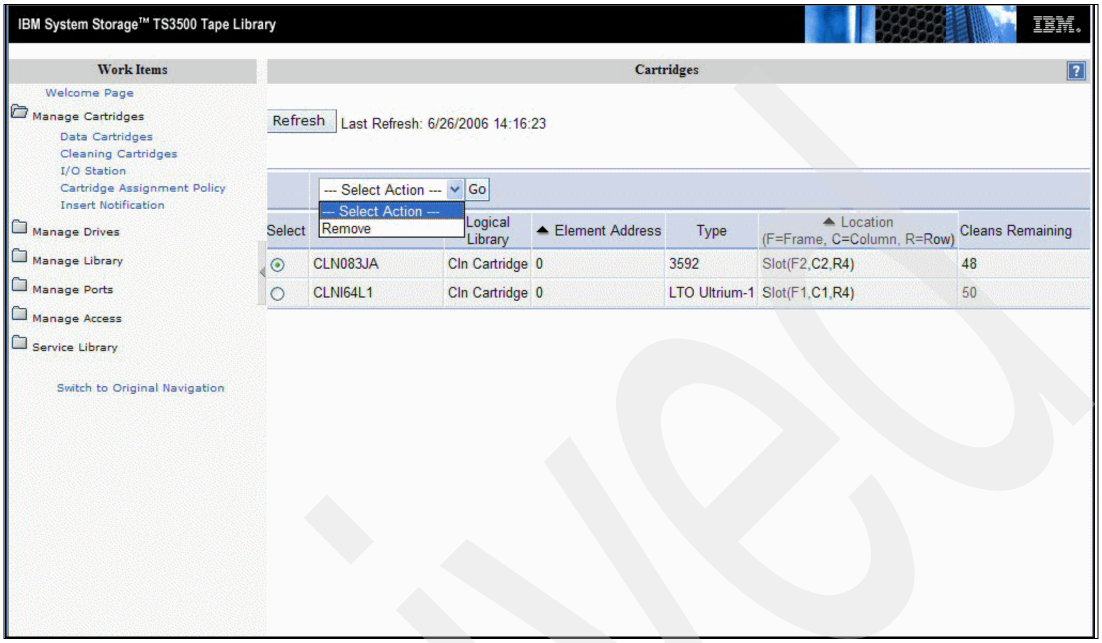

Figure 7-55 TS3500 cleaning cartridge usage

7.3 System-managed tape

This section describes the commands to operate a tape library in a z/OS and system-managed tape environment. It is not intended to replace the full description of operational procedures in the product documentation. It is a quick reference for the needed DFSMS and MVS commands.

7.3.1 DFSMS operator commands

Some of the commands contain “libname” as a variable. In this case, the SMS-defined library name is required. Depending on whether you refer to a TS7700 composite, TS7700 distributed, or your native drives partition, the output will be slightly different for some of these commands. For more information about DFSMS commands, refer to *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427.

Information from the 3494 or TS3500 Tape Library itself is contained in some of the outputs. However, you cannot switch the modes of the 3494 or 3593 Library Manager or the TS3500 Tape Library with z/OS commands.

> **Note:** DFSMS and MVS commands apply only to SMS-defined libraries. The library name defined during the definition of a library in ISMF is required for “libname” in the DFSMS commands.

The following DFSMS operator commands support the tape library:

- `LIBRARY EJECT,volser{,PURGE|KEEP|LOCATION}{,BULK}`

  This command is used to request the ejection of a volume from a tape library. The variations available in this command are:

  - Eject to the convenience I/O station for physical volumes (no additional specification). Delete from the tape library for logical volumes.
  - Eject to the bulk output station (BULK or B) for physical volumes. Delete from the tape library for logical volumes.
  - Remove the volume record from the TCDB (PURGE or P).
  - Keep the volume record in the TCDB and update it to indicate that the cartridge has been ejected (KEEP or K). If the record contains information in the SHELF location field, it is not changed. If the SHELF location field is empty, the operator must enter information about the new location as a reply to WTOR. The reply can be up to 32 characters long.
  - Keep the volume record in the TCDB and update it, including updating the SHELF location even if there is some information in this field (LOCATION or L). The operator has to enter the new information as a reply to WTOR.

  If none of the variations (PURGE, KEEP, or LOCATION) are indicated in the command, a default decides whether the record is kept or purged. This default can be set separately for each library through the ISMF Library Definition panel.

  This command is available for the operator to eject single cartridges. Mass ejection of cartridges is usually performed through program interfaces such as ISMF, a tape management system, or batch job.

- `LIBRARY SETCL, device-number, media-type`

  This command allows the setting of the media type of the scratch volume that is to be loaded into the ICL of the specified tape drive. You must issue the command on the system on which the drive is online. The other hosts are notified when the drive is varied online on the system. This does not apply to logical volumes.

  If the media assignment by this command is different from the current assignment, the ICL is emptied, and the proper cartridges are loaded.

- `VARY SMS,LIBRARY(libname),OFFLINE`

  This command acts on the SMS library, which is referred by libname. That is, it stops tape library actions and gradually makes all of the tape units within this logical library unavailable. The units are varied offline “for library reasons”, which means that they are not accessible because the whole SMS-defined library is offline.

  This simple form is a single-system form. The status of the library remains unaffected in other MVS systems.

  > **Note:** This command does not change the status of the Library Manager or the TS3500 Tape Library itself. It only applies to the SMS-defined logical libraries.

- `VARY SMS,LIBRARY(libname),ONLINE`

  This command is required to bring the SMS-defined library back to operation after it has been offline.

  The logical library does not necessarily go offline as a result of an error in some component of the physical library. Therefore, some message explanations for error situations request the operator to first vary the library offline and then back online. This usually clears all error indications and returns the library back to operation. Of course, this is only the MVS part of error recovery. You must clear the hardware, software, or operational error within the physical library and the 3953 LM itself before you bring the library back to work with MVS.

- `VARY SMS,LIBRARY(libname,sysname,...),ON/OFF and VARY SMS,LIBRARY(libname,ALL),ON/OFF`

  These extended form of the VARY command can affect more than one system. The first form affects one or more named MVS systems. The second form performs the VARY action on all systems within the SMSplex.

  The VARY SMS command allows the short forms ON and OFF as abbreviations for ONLINE and OFFLINE, respectively.

- `DISPLAY SMS,OAM`

  This command gives a single line of information about all tape libraries (if present), their tape units, storage cells, and scratch cartridges.

  This is the view of the single system where the command was executed. The number of unallocated, online drives is given under the heading AVL DRV (available drives).

  If both optical libraries and tape libraries are defined in the SMS configuration, two multiline WTOs are displayed. The first multiline display produced by the library control system (LCS) is the display of optical library information. The second multiline display contains tape library information.

- `DISPLAY SMS,LIBRARY(libname|ALL),STATUS`

  The library status display shows the SMS view of either one SMS-defined library or all SMS-defined libraries. The result contains one line of information for each library. This is a multihost view, which basically indicates whether the SMS-defined library is online, offline, or pending offline.

  STATUS is the default parameter.

- `DISPLAY SMS,LIBRARY(ALL),DETAIL`

  The DETAIL display, although a single-system view, gives slightly more information. The display is similar to the result of DISPLAY SMS,OAM, but each library gets its own line of information.

- `DISPLAY SMS,LIBRARY(libname),DETAIL`

  This command provides details about the status of a single library. It is the only command that displays the library state (auto, pause, or manual mode). Reasons for the mode and indications of inoperative parts of the library are given on additional status lines. Examples of special situations are:

  – Safety enclosure interlock open
  – Vision system not operational
  – Convenience output station full
  – Out of cleaner volumes

- `DISPLAY SMS,STORGRP(grpname|ALL)`

  There are no new parameters in the Storage Group display command because the optical library request formats are adequate here.

  This display command is a general form of a request and gives the total SMS multihost view of the situation. The result is a display of the status of either all Storage Groups (DASD, optical, and tape) or a single Storage Group. There is no format to display one category only.

- `DISPLAY SMS,STORGRP(grpname|ALL),DETAIL`

  The DETAIL display is not much more detailed than the general display. Only the library names of this Storage Group are indicated. This display is, in fact, more restricted than the general display. It gives the view of only one system, the view of its OAM, as the header line indicates.

  The LISTVOL parameter of DISPLAY SMS,STORGRP is not used for tape Storage Groups. Although you can view a volume list through ISMF, a similar listing on the console is too long to be meaningful.

- `DISPLAY SMS,VOLUME(volser)`

  This command displays all information that is stored about the volume in the TCDB (the VOLCAT) and some nonpermanent state information, such as "volume mounted on library-resident drive".

7.3.2 Library LMPOLICY command

Use the LIBRARY LMPOLICY command to assign or change a volume's policy names outboard at the library. You can use this command only for private, library-resident volumes that reside in a library that supports outboard policy management.

The processing for the LIBRARY LMPOLICY command invokes the LCS external services FUNC=CUA function. Any errors that the CUA interface returns can also be returned for the LIBRARY LMPOLICY command. If the change use attribute installation exit (CBRUXCUA) is enabled, the CUA function calls the installation exit. This can override the policy names that you set using the LIBRARY LMPOLICY command.

The results of this command are specified in the text section of message CBR1086I. To verify the policy name settings and to see whether the CBRUXCUA installation exit changed the policy names you set, display the status of the volume.

The syntax of the LIBRARY LMPOLICY command to assign or change volume policy names is as shown in Example 7-1.

Example 7-1 LIBRARY LMPOLICY command syntax

```
LIBRARY|LI LMPOLICY|LP , volser ,SG= storage group name |*RESET*
                               ,SC= storage class name |*RESET*
                               ,MC= management class name |*RESET*
                               ,DC= data class name |*RESET*
```

The following parameters are required:

- LMPOLICY | LP

  Specifies a request to set one or more of a private volume's policy names outboard in the library in which the volume resides. The library must support outboard policy management.

- Volser

  Volser specifies the volume serial number of a private volume which resides in a library with outboard policy management support.

- You must specify *at least one of the following optional parameters*. These parameters can be specified in any order.
  - `SG={storage group name | *RESET*}`

    Specifies a construct name for the SG parameter. If the request is successful, the construct name becomes the Storage Group for the volume in the TCDB and the Storage Group policy name in the library. If you specify the *RESET* keyword, you are requesting that OAM set the volume's Storage Group name to blanks in the TCDB, and to the default Storage Group policy in the library, which is also blanks.
  - `SC={storage class name | *RESET*}`

    Specifies a construct name for the SC parameter. If the request is successful, the construct name becomes the Storage Class policy name for the volume in the library. If you specify the *RESET* keyword, you are requesting that OAM set the volume's Storage Class name to the default Storage Class policy in the library, which is blanks.
  - `MC={management class name | *RESET*}`

    Specifies a construct name for the MC parameter. If the request is successful, the construct name becomes the Management Class policy name for the volume in the library. If you specify the *RESET* keyword, you are requesting that OAM set the volume's Management Class name to the default Management Class policy in the library (blanks).
  - `DC={data class name | *RESET*}`

    Specifies a construct name for the DC parameter. If the request is successful, the construct name becomes the Data Class policy name for the volume in the library. If you specify the *RESET* keyword, you are requesting that OAM set the volume's Data Class name to the default Data Class policy in the library, which is blanks.

The values you specify for the SG, SC, MC, and DC policy names must meet the Storage Management Subsystem (SMS) naming convention standards:

- Alphanumeric and national characters only
- Name must begin with an alphabetic or national character ($*@#%)
- No leading or embedded blanks
- Eight characters or less

7.3.3 Host Console Request

The Library Request command provides a way for an administrator or an operator to determine status and obtain information of resources of the TS7700 as well as prioritizing some of the active work. It will also enable the customer to perform problem determination and act proactively for instance on delays of single important jobs waiting for recalls from backend drives.

The Library Request was introduced with the TS7700 Virtualization Engine code level 8.3.0.106. z/OS support is also required. Refer to OAM APAR OA20065 and device services APARs OA20066, OA20067 and OA20313. The Library Request command for Host Console Request is supported in z/OS V1R6 and above.

A detailed description of the Host Console Request commands and responses is available from White Paper WP101091 - *IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide Version 1.0* which is available at the following Techdocs Web site listed. Search for *TS7700*.

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs

Command syntax for the Host Console Request command

The Host Console Request is also referred to as the Library Request command. The syntax of the command is shown in Example 7-2.

Example 7-2 Host Console Request command syntax

```
LIBRARY|LI REQUEST|REQ,library_name,
,keyword1|,keyword2|,keyword3|,keyword4|
,L=a|name |name-a
```

The following parameters are required:

`REQUEST | REQ` Specifies a request to obtain information from the TS7700 Virtualization Engine or to perform an outboard operation.

`library_name` Specifies the library name associated with the TS7700 to which the request should be directed. The library name specified can be a composite or a distributed library and which library is applicable depends on the other keywords specified.

`keyword1` Specifies which operation is to be performed at the TS7700.

The following optional parameters are dependent on the first keyword specified. Based on the first keyword specified, zero or more of the additional keywords might be appropriate.

`keyword2` Specifies additional information in support of the operation specified with the first keyword.

`keyword3` Specifies additional information in support of the operation specified with the first keyword.

`keyword4` Specifies additional information in support of the operation specified with the first keyword. Keyword4 is prepared for future use.

`L={a | name | name-a}` Specifies where to display the results of the inquiry: the display area (`L=a`), the console name (`L=name`), or both the console name and the display area (`L=name-a`). The name parameter can be an alphanumeric character string.

> **Note:** The keywords specified must be from one to eight characters in length and can consist of alphanumeric (A-Z and 0-9), the national character set ($@#), and mask and wildcard-type characters (*%). The only checking done by the host is to verify that the specified keywords conform to the supported character set. The validity of the keywords themselves and the keywords in combination with each other is verified when the command is received at the library. Any errors that the library encounters are reported back to the host and displayed as part of the command output for CBR1280I.

Overview of the Host Console Request commands

Table 7-3 lists all the command keywords and a short description of each of them.

Table 7-3 Overview of Host Console Request commands

Keyword1	Keyword2	Keyword3	Description	Comp. Library	Dist. Library
CACHE			Requests information about the current state of the cache and the data managed within it associated with the specific distributed library.	N/A	Y
COPYEXP	*volser*	RECLAIM	Requests that the specified physical volume that has been exported previously in a Copy Export operation, be made eligible for priority reclaim.	N/A	Y
COPYEXP	*volser*	DELETE	Requests that the specified physical volume that has been exported previously in a Copy Export operation, be deleted from the TS7700 database. The volume must be empty.	N/A	Y
GRIDCNTL	COPY	DISABLE	Requests that copy operations for the specified distributed library be disabled. Copies that are in progress are allowed to complete, but no new copies using the specified distributed library as the source or target are initiated.	N/A	Y
GRIDCNTL	COPY	ENABLE	Requests that copy operations for the specified distributed library be enabled. Copy operations can again use the specified distributed library as the source or target for copies.	N/A	
LVOL	volser		Requests information about a specific logical volume.	Y	N/A
PDRIVE			Requests information about the physical drives and their current usage associated with the specified distributed library.	N/A	Y
POOLCNT	00-32		Requests information about the media types and counts, associated with a specified distributed library, for volume pools beginning with the value in keyword2.	N/A	Y
PVOL	volser		Requests information about a specific physical volume.	N/A	Y
RECALLQ	volser		Requests the content of the recall queue starting with the specified logical volume. Keyword2 could be blank.	N/A	Y

Keyword1	Keyword2	Keyword3	Description	Comp. Library	Dist. Library
RECALLQ	volser	PROMOTE	Requests that the specified logical volume be promoted to the top of the recall queue.	N/A	Y
STATUS	GRID		Requests information about the copy, reconcile and ownership takeover status of the libraries in a grid configuration	Y	N/A
STATUS	GRIDLINK		Requests information about the status and performance of the links between the TS7700s in the Grid configuration	N/A	Y

Examples of the Host Console Request commands

Let us go through some examples of the commands and the responses retrieved. Many of these commands could be subject for automation based on your own automation products. You could create your own actions to be taken by periodically issuing the Library Request commands and react on responses automated and without operator interference. This could be used for proactive handling.

CACHE keyword

Example 7-3 shows the command with the CACHE keyword. The response is:

- Number of GB installed and enabled
- One line per partition for future partition support
- Allocated GB and used GB in the cache and how they are split between PG0 and PG1
- Number of GB that must be copied to another cluster
- Premigration and copy throttling values in milliseconds

Example 7-3 Library request command: CACHE

```
LI REQ,BARR68A,CACHE
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,CACHE.
CBR1280I LIBRARY BARR68A REQUEST. 149
KEYWORDS: CACHE
----------------------------------------------------------------------
TAPE VOLUME CACHE STATE V1
 INSTALLED/ENABLED GBS   6000/  2000
 PARTITION   ALLOC   USED    PG0    PG1  PMIGR   COPY    PMT   CPYT
      0       2000   1750     47   1703      0      0      0      0
      1          0      0      0      0      0      0      0      0
      2          0      0      0      0      0      0      0      0
      3          0      0      0      0      0      0      0      0
      4          0      0      0      0      0      0      0      0
      5          0      0      0      0      0      0      0      0
      6          0      0      0      0      0      0      0      0
      7          0      0      0      0      0      0      0      0
```

GRIDCNTL keyword

Response from GRIDCNTL DISABLE as in Example 7-4 shows that copies have been stopped.

Example 7-4 Library request command: GRIDCNTL DISABLE

```
LI REQ,BARR68,GRIDCNTL,DISABLE
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,GRIDCNTL,DISABLE
CBR1280I LIBRARY BARR68 REQUEST. 509
KEYWORDS: GRIDCNTL,DISABLE
------------------------------------------------------------
GRID COPY CAPABILITIES V1
 DISABLED FOR SOURCE AND TARGET COPIES
```

Response from GRIDCNTL ENABLE as in Example 7-5 shows that copies have been restarted.

Example 7-5 Library request command GRIDCNTL ENABLE

```
LI REQ,BARR68,GRIDCNTL,ENABLE
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,GRIDCNTL,ENABLE
CBR1280I LIBRARY BARR68 REQUEST. 509
KEYWORDS: GRIDCNTL,ENABLE
------------------------------------------------------------
GRID COPY CAPABILITIES V1
 ENABLED FOR SOURCE AND TARGET COPIES
```

LVOL keyword

Response from LVOL as in Example 7-6 shows detailed information of the logical volume. The response is:

- Shows if a logical volume is ECCST or CST and the size of the volume.
- Number of copies and VOLSER of physical volumes where the logical volume resides.
- Copy policy used.

Example 7-6 Library request command: LVOL

```
LI REQ,BARR68,LVOL,693023
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,LVOL,693023.
CBR1280I LIBRARY BARR68 REQUEST. 509
KEYWORDS: LVOL,693023
------------------------------------------------------------
LOGICAL VOLUME INFORMATION V1
 LOGICAL VOLUME:              693023
 MEDIA TYPE:                  CST
 COMPRESSED SIZE (MB):        392
 MAXIMUM VOLUME CAPACITY (MB): 0
 CURRENT OWNER:               Pesto
 MOUNTED LIBRARY:
 MOUNTED VNODE:
 MOUNTED DEVICE:
 TVC LIBRARY:                 Pesto
 MOUNT STATE:
 CACHE PREFERENCE:            PG1
CATEGORY:                    300F
LAST MOUNTED (UTC):          2007-08-05 07:59:14
```

```
LAST MODIFIED (UTC):         2007-08-05 07:19:25
LAST MODIFIED VNODE:         00
LAST MODIFIED DEVICE:        1D
TOTAL REQUIRED COPIES:       3
KNOWN CONSISTENT COPIES:     3
IMMEDIATE-DEFERRED:          N
DELETE EXPIRED:              N
RECONCILIATION REQUIRED:     N
---------------------------------------------------------------------
 LIBRARY  RQ CACHE PRI PVOL  SEC PVOL  COPY ST  COPY Q  COPY CP
   Pesto   N     N   JA6321    ------     CMPT       -      RUN
  Squint   N     N   JA4238    ------     CMPT       -      RUN
 Celeste   N     N   JB0450    ------     CMPT       -      RUN
```

PVOL keyword

Response from PVOL as in Example 7-7 shows detailed information of the logical volume. The response is:

- Media type, drive mode, format of the volume and Volume State (read-write).
- Capacity in MB, valid data in percent, and number of logical volumes.
- Shows whether the physical volume is exported and whether it is encrypted.

Example 7-7 Library request command: PVOL

```
LI REQ,BARR68A,PVOL,JA5313
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,PVOL,JA5313.
CBR1280I LIBRARY BARR68A REQUEST. 225
KEYWORDS: PVOL,JA5313
----------------------------------------------------------------
PHYSICAL VOLUME INFORMATION V1
 PHYSICAL VOLUME:  JA5313
 MEDIA TYPE:       JA
 DRIVE MODE:       E05
 FORMAT:           TS7700
 VOLUME STATE:     READ-WRITE
 CAPACITY STATE:   FILLING
 CURRENT POOL:     2
 MBYTES WRITTEN:   140718
 % ACTIVE DATA:    100.0
 LAST INSERTED:    2007-03-16 20:06:06
 WHEN EXPORTED:    N/A
 MOUNTS:           2
LOGICAL VOLUMES:  374
ENCRYPTED:        N
```

POOLCNT keyword

Response from POOLCNT as in Example 7-8 shows detailed information about each physical volume pool. The response is:

- Shows detailed info from each pool.
- Details about media type, volumes that are empty, filling or full.
- Volumes eligible for erase.
- Volumes that are in Read-only state, unavailable, or in Copy Export state.

Example 7-8 Library request command: PVOL

```
LI REQ,BARR68A,POOLCNT
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,POOLCNT.
 CBRXLCS SCRATCH PROCESSING FAILED FOR: V66654  RC = 0004 RSN = 0004
CBR1280I LIBRARY BARR68A REQUEST. 441
KEYWORDS: POOLCNT
----------------------------------------------------------------------
PHYSICAL MEDIA COUNTS V1
 POOL MEDIA  EMPTY  FILLING   FULL  ERASE    ROR  UNAVAIL  CXPT
    0    JA   1618
    1    JA      3        1    576      0      0        0     0
    2    JA      2        8    712      0      7        0     0
```

RECALLQ keyword

Response from RECALLQ on Distributed Library as in Example 7-9 shows detailed information about the logical volume recall queue. The response is:

- Recall is in progress for volume L00121 and L99356.
- Volume Y30458 and L54019 have recall scheduled, meaning a RECALLQ,*volser*,PROMOTE has been issued.
- Volume L67304 is in position 1 for recall and has been in the recall queue for 135 seconds.
- Volume T09365 spans from physical volume AD5901 to P00167.

Example 7-9 Library request command: RECALLQ

```
LI REQ,BARR68A,RECALLQ
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68A,RECALLQ.
CBR1280I LIBRARY BARR68A REQUEST. 820
KEYWORDS: RECALLQ
---------------------------------------------------------
RECALL QUEUE V1
POS    LVOL    PVOL1   PVOL2           TIME
 IP  L00121   AB0456                    175
 IP  L99356   AA0350                    201
 SC  Y30458   AB0456                    148
 SC  L54019   AA0350                    145
  1  L67304   AC0101                    135
  2  T09356   AD5901   P00167           102
```

STATUS keyword

Response from STATUS,GRID on the Composite Library as in Example 7-11 shows detailed information about the Multi Cluster Grid. The response is:

- Seen from the Composite Library, Squint is in service mode and there is a queue of data that needs to be copied. Squint is in “service ownership takeover” mode, meaning that the other two TS7700s must do recalls even though a logical volume resides in cache in Squint.
- Seen from the Distributed Library View, Squint is unknown, and Celeste has an unavailable link.

Example 7-10 Library request command: STATUS,GRID

```
LI REQ,BARR68,STATUS,GRID
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,STATUS,GRID.
CBR1280I LIBRARY BARR68 REQUEST. 373
KEYWORDS: STATUS,GRID
------------------------------------------------------------
GRID STATUS V1
 COMPOSITE LIBRARY VIEW
                 IMMED-DEFERRED  OWNERSHIP-T/O  RECONCILE
 LIBRARY  STATE    NUM        MB  MODE      NUM        NUM
   Pesto     ON      0         0   -          0          8
  Squint    SVC     12      8713  SOT        24         12
 Celeste     ON      0         0   -          0          0
------------------------------------------------------------
 DISTRIBUTED LIBRARY VIEW
                RUN-COPY-QUEUE  DEF-COPY-QUEUE  LIBK STATE
 LIBRARY  STATE    NUM        MB    NUM       MB    01234567
   Pesto     ON     40     15691      0        0    -AA
  Squint     UN      -         -      -        -    ---
 Celeste     ON     61     23929      0        0    AU-
```

STATUS GRIDLINK keyword

Response from STATUS,GRIDLINK on Composite Library as in Example 7-11 shows detailed information about the links between clusters. The response is:

- Detailed information about link throughput and their status.
- Latency seen by each cluster. This value is gathered every five minutes.
- Amount of data sent by this cluster.
- Percent of packets retransmitted; this means the amount of sent operations failed. In general any value below 1,5% is acceptable.

Example 7-11 Library request command: STATUS,GRIDLINK

```
LI REQ,BARR68,STATUS,GRIDLINK
CBR1020I PROCESSING LIBRARY COMMAND: REQ,BARR68,STATUS,GRIDLINK.
CBR1280I LIBRARY BARR68 REQUEST. 373
KEYWORDS: STATUS,GRIDLINK
GRIDLINK STATUS V1
 CAPTURE TIMESTAMP: 2008-02-08 12:45:32
 LINK VIEW
  LINK NUM    CFG   NEG    READ  WRITE  TOTAL ERR  LINK STATE
                           MB/S   MB/S   MB/S        01234567
         0   1000  1000    87.2  102.4  189.6   0    -AA
         1   1000  1000    74.9  104.6  179.5   0    -AA
```

```
       2      0     0     0.0    0.0    0.0   0
       3      0     0     0.0    0.0    0.0   0
----------------------------------------------------------------------
 LINK PATH LATENCY VIEW
  LIBRARY         LINK 0         LINK 1         LINK 2         LINK 3
           LATENCY IN MSEC
   TS001B               6             7              0              0
   TS001C              19            20              0              0
----------------------------------------------------------------------
CLUSTER VIEW
  DATA PACKETS SENT:              103948956
  DATA PACKETS RETRANSMITTED:     496782
  PERCENT RETRANSMITTED:          0.4778
----------------------------------------------------------------------
 LOCAL LINK IP ADDRESS
  LINK 0 IP ADDR:  9.11.200.60
  LINK 1 IP ADDR:  9.11.200.61
  LINK 2 IP ADDR:
  LINK 3 IP ADDR:
```

7.3.4 MVS commands

The following commands are described in detail in the *z/OS MVS System Commands*, SA22-7627:

- DS QT,devnum,1,RDC

  This command displays identification, status, and diagnostic information about tape devices. You can use the command to display the LIBRARY-ID and the LIBPORT-ID that are stored for a device in an IBM TS3500.

  Example 7-12 shows the sample output of a DS QT system command.

Example 7-12 Sample output of a DS QT system command

```
DS QT,1699,1,RDC
IEE459I 12.30.05 DEVSERV QTAPE 970
UNIT DTYPE  DSTATUS CUTYPE  DEVTYPE  CU-SERIAL  DEV-SERIAL ACL LIBID
1699 3490L  ON-NRD  3490A20 3490B40  0177-10619 0177-10619 I   10007
  READ DEVICE CHARACTERISTIC
34902034904000000 1FF8808000000000 0000000000000000 0000000000000000
0100070100000000 4281000000000000 0000000000000000 0000000000000000
------
    | --
    | |------->           4. Byte = LIBPORT-ID
    |
    |---------------->  1.-3. Byte = LIBRARY-ID (omit first half byte)
LIBRARY-ID=10007
LIBPORT-ID=01
```

- `DS QT,devnum,MED,nnn`

  This command displays information about the device type, media type, and the cartridge volume serial number. *devnum* is the device address in hexadecimal. *nnn* is the number of devices to query.

  Example 7-13 shows the sample output of a DS QT system command.

  Example 7-13 Sample output of a DS QT system command

```
IEE459I 11.32.31 DEVSERV QMEDIUM 608
UNIT RDTYPE EDTYPE EXVLSR INVLSR RMEDIA EMEDIA
0940 3590-E 3590-1 003700         3
 UNIT, the device address
RDTYPE, the real device type (physical)
EDTYPE, emulated device type
EXVLSR, external volume label
INVLSR, internal volume label
RMEDIA, real media type
EMEDIA, emulated media type
```

- `VARY unit,ONLINE/OFFLINE`

  The VARY unit command in itself is no different from what it used to be. However, new situations are seen when the affected unit is attached to a library.

  When the library is offline, the tape units cannot be used. This is internally indicated in a new status, `offline for library reasons`, which is separate from the normal `unit offline` status. A unit can be offline for both library and single-unit reasons.

  A unit that is offline for library reasons only cannot be taken online with VARY unit,ONLINE. Only VARY SMS,LIBRARY(...),ONLINE helps.

  You can bring a unit online that was individually varied offline and was offline for library reasons by varying it online individually and varying its library online. The order of these activities is not important, but both are necessary.

  Currently no display directly gives the reason why the unit is offline, nor is there a display that gives the name of the library to which this unit belongs.

- `DISPLAY U`

  The DISPLAY U command displays the status of the requested unit. If the unit is part of a tape library (either manual or automated), device type 348X is replaced by 348L. An IBM 3490E is shown as 349L, and a 3590 or 3592 as 359L.

  For a manual tape library, this might create a situation where it is no longer possible to see from the console response whether a particular tape unit supports IDRC, because this information is overlaid by the L indicating that the unit belongs to a library.

  The output of DEVSERV is not changed in this way.

- `MOUNT devnum, VOL=(NL/SL/AL,serial)`

  The processing of MOUNT has been modified to accommodate automated tape libraries and the requirement to verify that the correct volume has been mounted.

- `UNLOAD devnum`

  The UNLOAD command allows you to unload a drive, if the rewind/unload process was not successful in the first place.

7.4 Basic operations

In the next few sections, we explain some of the basic tasks that might be needed during the operation of a TS7700.

7.4.1 Clock and time setting

The TS7700 time can be set from a network time protocol (NTP) server or by the IBM SSR. It is set to UTC, also called GMT. Refer to “Time coordination” on page 22 for more details about time coordination.

The Library Manager time should be set to your local time. The TS7700 logs its UTC time in the Library Manager logs every 10 minutes. This provides a means to line up the Library Manager and TS7700 logs.

7.4.2 Library online/offline

The vary online/offline command for a IBM TS7700 always uses the composite LIBRARY-ID either in a Single Cluster or in a Multi Cluster Grid environment.

7.4.3 Library in Pause Mode

In a Multi Cluster Grid environment, one distributed library can enter the pause mode, like it is possible for a standalone VTS. Reasons for the pause can include an enclosure door being opened for clearing a device after a load/unload failure or removing cartridges from the high capacity I/O station. The following message is displayed at the host when a library is in pause or manual mode:

```
CBR3757E Library library-name in {paused | manual mode} operational state
```

During Pause mode all recalls and physical mounts are held up and queued by the Library Manager for later processing when the library leaves the Pause mode.

Because both scratch mounts and private mounts with data in the cache are allowed to execute, but not physical mounts, no more data can be moved out of the cache after the currently mounted stacked volumes are completely filled. The cache is filling up with data that has not been copied to stacked volumes. This results in significant throttling and finally in the stopping of any mount activity in the library. For this reason it is important to minimize the amount of time spent with the library in Pause mode condition.

7.4.4 Preparing for service

When an element of the TS7700 needs to be serviced, it must be prepared prior to taking it away; otherwise, continued host access to data might not be possible. The service preparation task is an administrator responsibility, and will remove a whole site from a grid activity. More details on service preparation can be found in Chapter 2, “Architecture, components, and functional characteristics” on page 7.

Note: Before invoking service preparation at the TS7700, all virtual devices must be varied offline from the host. Make sure the customer is aware of this. All logical volumes must be dismounted, all devices associated with the cluster varied offline and jobs moved to other clusters in the grid before service preparation is invoked. After service is complete, when the TS7700 is ready for operation, you must vary the devices online at the host.

Preparing a TS7700 for service

When an operational TS7700 needs to be taken offline for service, the TS7700 grid must first be prepared for the loss of the resources involved, to provide continued access to customer data. The controls to prepare a TS7700 for service (Service Prep) are provided through the MI. The menu is described in Figure 7-50 on page 364.

Here is the message posted to all hosts when the TS7700 grid is in this state:

```
CBR3788E Service preparation occurring in library library-name.
```

Preparing the tape library for service

If the 3494 Tape Library or a TS3500 Tape Library in a TS7700 grid require to be serviced, the TS7740 associated with it must first be prepared for service. After the TS7740 has completed service preparation, the normal procedures for servicing the tape library can continue. Refer to 7.2.9, “Service & Troubleshooting” on page 364 on how to set the TS7700 in service preparation mode.

7.4.5 IBM TS3500 Tape Library inventory

You can perform a physical inventory by frame or an inventory for the total TS3500 Tape Library.

Perform the following steps:

1. Type the Ethernet IP address on the URL line of the browser and press Enter. The Welcome Page displays.
2. Select **Manage Library** → **by Frame**. The Physical Library Summary panel displays with options for performing an inventory on a single frame or on all frames.
3. Select **All Frames**, then select **Inventory**.

Important: During the inventory, all tasks in the work queue are delayed until the library inventory is performed. Depending on whether you requested a frame inventory or a library inventory, this can impact your production.

7.4.6 Inventory upload

The Request Inventory Upload window on the Library Manager console or the ETL Specialist panel (Figure 7-56 on page 387) allows you to manually update the logical volume inventory from the tape library. This procedure is necessary after insertion of a volume into the tape library that does not match any defined tape library Cartridge Assignment Policies (CAP). If you do not respond to the insert notification at the tape library operator panel, the inserted volume will not be assigned to a logical library.

Before you use the Request Inventory Update window, you must first specify the correct logical library for the inserted volume using the Web specialist of the TS3500 Tape Library.

Then select this function in the 3953 Library Manager console using **Commands** → **Request inventory upload**.

The Library Manager might first prompt you for the system administrator password. To upload the most recent tape library inventory to the Library Manager, select **Request Upload**. A confirmation window will open. Select **Yes** to confirm the request and initiate an inventory upload. Select **No** to close the confirmation window and return to the Request Inventory Upload window without initiating an inventory upload. After selecting **Yes** on the confirmation window, a notification window opens to inform you that the requested action completed successfully. Select **OK** to close the notification window. Note that the Request Inventory Upload window remains open until **Exit** is selected. The other two control push buttons on the Request Inventory Upload window are: Exit closes the Request Inventory Upload window. Help provides help about the Request Inventory Upload window.

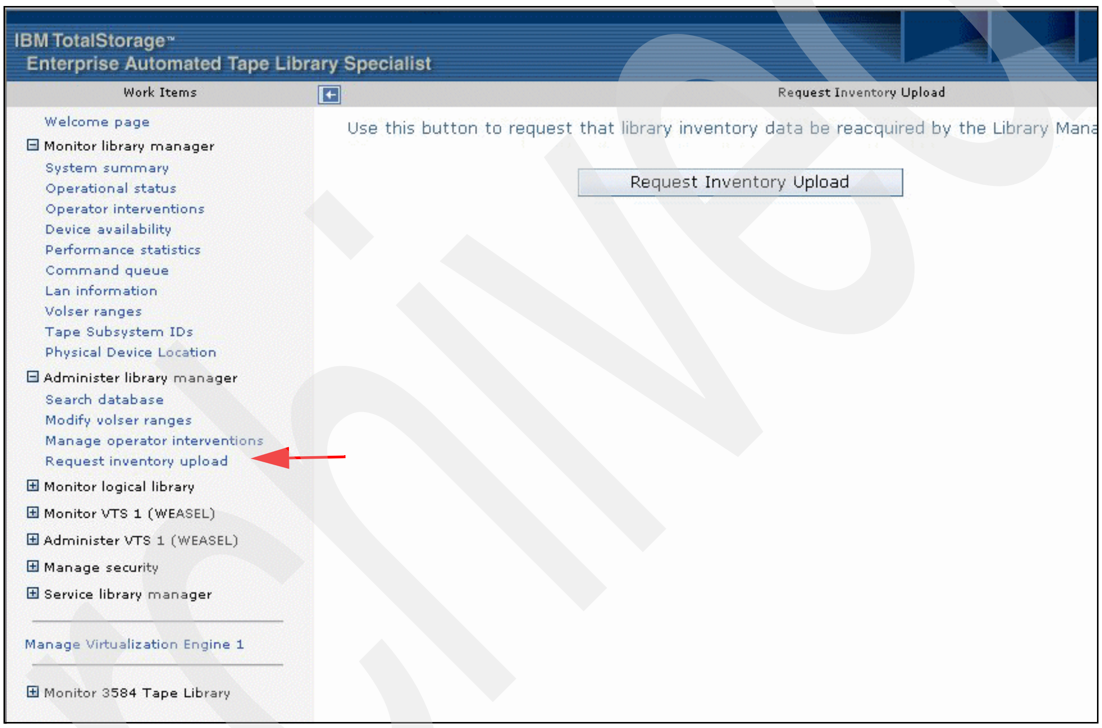

Figure 7-56 Request Inventory Upload through ETL Specialist

Refer to *IBM System Storage TS3500 Tape Library Operator Guide*, GA32-0560, for more information about inventory functions of the IBM TS3500.

7.5 Tape cartridge management

Most of the tape management operations are described in 7.1, “User interfaces” on page 304.

Here we give information about tape cartridges and labels, inserting and ejecting stacked volumes, and Exception Conditions.

7.5.1 Tape cartridges and labels

The data tape cartridge used in a 3592 contains the following:

- Single reel of magnetic tape
- Leader pin 1
- Clutch mechanism 2
- Cartridge write-protect mechanism 3
- Internal cartridge memory (CM)

Figure 7-57 depicts a J-type data cartridge, and Table 7-4 provides information about the various 3592 J-type data cartridges.

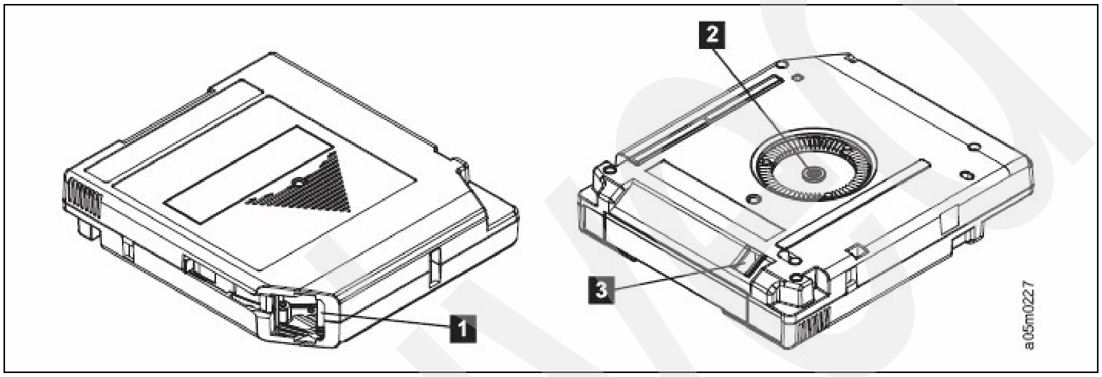

Figure 7-57 Tape cartridge

Table 7-4 Types of data tape cartridges. used with Figure 7-57

Identifier	Case Color	Tab, Door, and Label Color	J1A Native Capacity	E05 Native Capacity	Description
JA	Black	Dark Blue	300 GB	500 GB	Standard read/write
JB	Black	Dark Green	Not supported	700 GB	Extended length read/write
JJ	Black	Light Blue	60 GB	100 GB	Economy read/write

Labels

The cartridges use a media label to describe the cartridge type, as shown in Figure 7-58 (JA example). In tape libraries, the library vision system identifies the types of cartridges during an inventory operation. The vision system reads a volume serial number (VOLSER), which appears on the label on the edge of the cartridge. The VOLSER contains from one to six characters, which are left-aligned on the label. If fewer than six characters are used, spaces are added. The media type is indicated by the seventh and eighth characters.

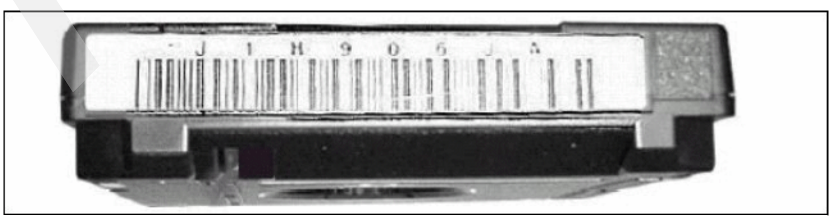

Figure 7-58 Cartridge label

7.5.2 Inserting stacked volumes

There are two methods for physically inserting a stacked volume into the IBM 3494 or into the IBM TS3500 Tape Libraries:

- Opening the library doors and directly inserting into tape library storage cells
- Using the tape library I/O station

Inserting directly into storage cells

You can insert cartridges in the 3494 and TS3500 tape libraries as described here:

- **IBM 3494**: When the cartridges are inserted into the Tape Library, if the barcode is readable and the VOLSER is unique, a record of the volume is added to the Library Manager database. In the IBM 3494, a VOLSER range is used to help determine a VOLSER's media type when the cartridge is inserted into the Tape Library. When you insert stacked cartridges, these are the steps you perform:

  a. Define the correct VOLSER ranges for the media types.
  b. Pause the Tape Library.
  c. Assure that Enable Inventory Update is enabled.
  d. Open the door where there is an empty cell.
  e. Insert the VTS stacked cartridge into an empty cell.
  f. Close the door and set the Library to Auto Mode.
  g. Select **Commands** → **System Management** → **Manage Insert Volume** → **Select Volume** → **Take action**.

  For more information, refer to *IBM TotalStorage Automated Tape Library (3494) Operator Guide,* GA32-0449.

- **TS3500**: The IBM TS3500 Cartridge Assignment Policy defines which volumes are assigned to which logical library partition; if the VOLSER is included in the System z range, then it will be assigned to the TS3500 logical library partition.

After the doors on the library are closed and the tape library has performed inventory, the upload of the inventory to the 3953 Tape System will be processed before the IBM TS3500 tape library reaches the READY state. The Library Manager updates its database accordingly.

> **Note:** The inventory is performed only on the frame where the door is opened and not on the frames to either side. If you insert cartridges into a frame adjacent to the frame that you opened, then you must perform a manual inventory of the adjacent frame, using the operator panel on the IBM TS3500 itself.

Inserting cartridges using the I/O station

The IBM 3494 and the IBM TS3500 Tape Libraries detect volumes in the I/O station, and then move the volumes to empty cells.

- **IBM 3494**: A cartridge is placed in the Convenience I/O Station for insertion into the IBM 3494. When the Convenience I/O door is closed, the Library Manager senses the presence of a cartridge and locks the door. The Library Manager then instructs the accessor to remove the cartridge from the Convenience I/O station and place it into an empty cell. At this time the bar code is read, and the VOLSER and media type are checked. If the barcode is readable and the VOLSER is unique, a record of the volume is added to the Library Manager database. The cartridge is placed in the Library Manager INSERT category. Then the operator must select **Commands** → **System Management** → **Manage Insert Volume** → **Select Volume** → **Take action**, from the Library Manager console.

Refer to *IBM TotalStorage Automated Tape Library (3494) Operator Guide,* GA32-0449 for details.

- **IBM TS3500**: The IBM TS3500 Cartridge Assignment Policy defines which volumes are assigned to which logical library; if the VOLSER is included in the System z range, it will be assigned to the IBM TS3500 logical library partition. If any VOLSER is not in a range defined by the CAP, then the operator identifies a System z logical library as the destination using the Insert Notification process.

  Under certain conditions, cartridges will not be assigned to a logical library partition in the TS3500 or in the IBM 3494 or 3953 Library Managers.

> **Note:** Unassigned cartridges can exist in the TS3500, as well as in the IBM 3494 or 3593 Library Managers. But “unassigned” has different meanings and needs different actions from the operator.

7.5.3 Ejecting stacked volumes

To remove a cartridge from an IBM 3494 or an IBM TS3500 Tape Library you can use various methods. If you want to eject cartridges that are owned by a TS7700 Virtualization Engine, you should only use the options provided through the Library Manager operator panels or the ETL Specialist.

From the ETL Specialist panel from the task pane on the left side, select **Move/Eject Stacked Volumes**. You will see the panel shown in Figure 7-59.

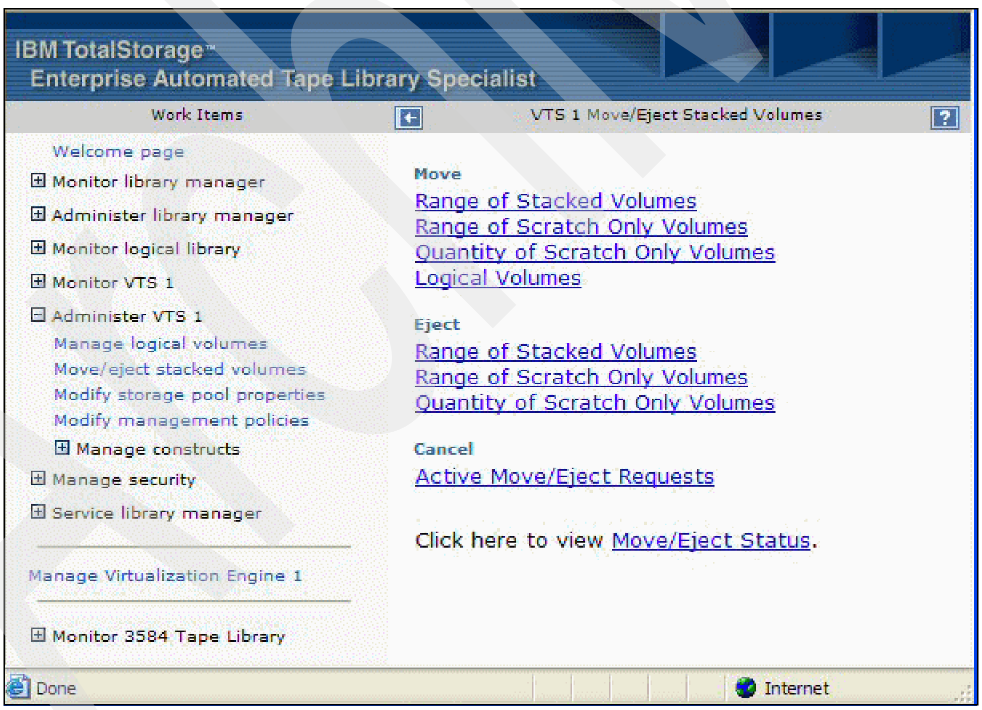

Figure 7-59 Move/Eject Stacked Volumes panel

You can eject a range of stacked volumes (both scratch and private), a range of scratch-only volumes, or a number of scratch-only volumes. To select a Range of Stacked Volumes, click the text. The panel shown in Figure 7-60 is displayed.

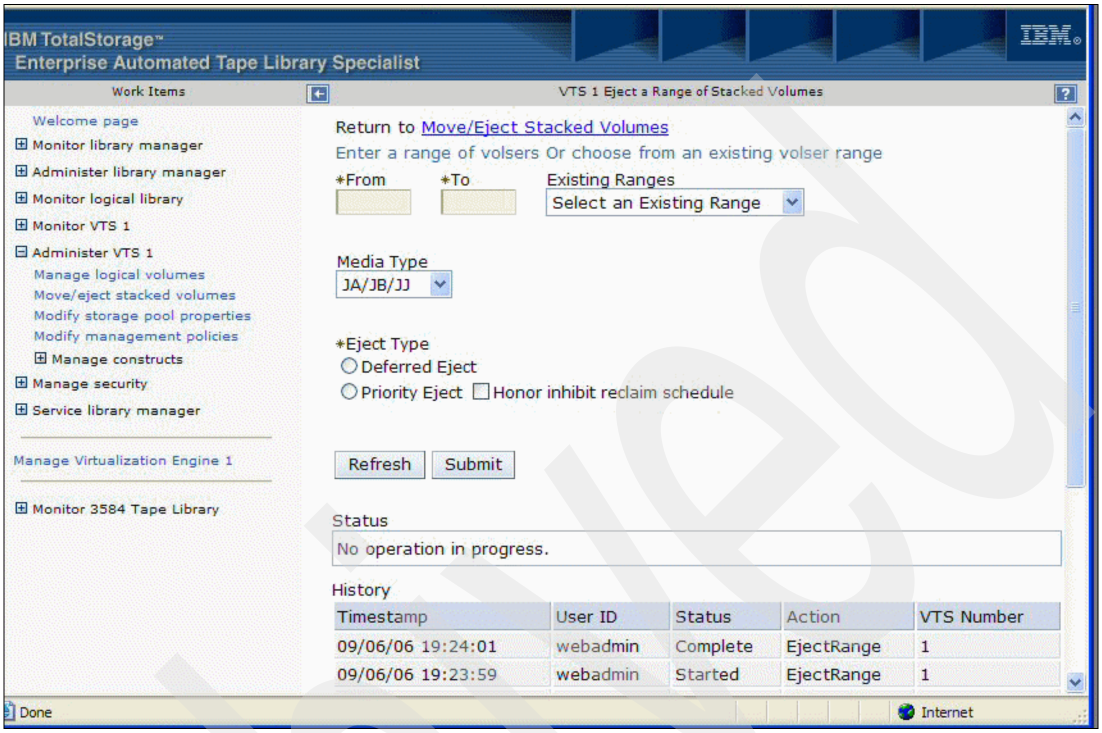

Figure 7-60 Move/Eject Range of Stacked Volumes panel

Fill in the range you want to eject and click **Submit**.

Be aware that the removing of TS7700 cartridges out of a library can take some time, depending on the amount of active data on a cartridge. TS7700 cartridges can only be ejected when they are empty. Therefore, all active data must be transferred to a different cartridge before the TS7700 cartridge can be ejected. There is no possibility to eject a TS7700 physical cartridge from the host interface; you must always use the appropriate 3953 Library Manager console or ETL Specialist panel.

7.5.4 Exception conditions

On a physical volume, one important exception condition is the Out of Physical Volumes condition.

Out of Physical Volumes

When a distributed library associated with a cluster runs out of scratch stacked physical volumes, operations of the TS7740 are impacted.

As part of normal processing, data is copied from cache to physical volumes in a primary pool managed by the virtualization engine. A copy can also be made to a physical volume in a secondary pool if the dual copy function is specified using Management Class. Empty physical volumes are needed in a pool, or if a pool is enabled for borrowing, in the common scratch pool, for operations to continue. If a pool runs out of empty physical volumes and there are no volumes that can be borrowed or borrowing is not enabled, operations that can use that pool on the distributed library must be suspended. If one or more pools run out of empty physical volumes, the distributed library enters the Out of Physical Scratch state. The Out of Physical Scratch state is reported to all hosts attached to the cluster associated with the distributed library and, if included in a Grid configuration, to the other clusters in the Grid. The following MVS console message is generated to inform an operator of this condition:

```
CBR3789E VTS library library-name is out of empty stacked volumes.
```

`Library-name` is the name of the distributed library in the state. The CBR3789E message will remain on the MVS console until empty physical volumes have been added to the library or the pool that is out has been enabled to borrow from the common scratch pool and there are empty physical volumes to borrow. Intervention required conditions are also generated for the out of empty stacked volume state and for the pool that is out of empty physical volumes. If the option to send intervention conditions to attached hosts is set on the Library Manager associated with the distributed library, the following console messages are also generated to provide specifics about the pool that is out of empty physical volumes:

```
CBR3750I Message from library library-name: OP0138 The Common Scratch Pool
(Pool 00) is out of y media volumes.
CBR3750I Message from library library-name: OP0139 Storage pool xx is out of
scratch volumes.
```

The 0P0138 message indicates the media type that is out in the common scratch pool. These messages will not remain on the MVS console. The intervention conditions can be viewed through the Library Manager console or the Web specialist.

If the TS7740 is in a grid configuration, if its associated distributed library enters the out-of-empty-stacked-volume state, there are other impacts to operations:

- All copy operations are immediately suspended into the cluster (regardless of which pool has become empty).
- If the cluster has a copy consistency point of RUN, the grid enters the Immediate Mode Copy Operations Deferred state. This generates an MVS console message:

  ```
  CBR3787E One or more immediate mode copy operations deferred in library
  library-name.
  ```

- If another cluster attempts to copy a logical volume that is not resident in the cache, the copy attempt fails.
- In choosing a TVC cluster, the grid prefers clusters that are not in the out-of-empty-stacked-volume state, but could still select a remote TVC whose cluster is in that state. If the data needed is not in the remote cluster's TVC, the recall of the data will fail. If data is being written to the remote cluster's TVC, the writes will be allowed, but because there might not be any empty physical volumes available to copy the data to, the cache can become full of data that cannot be copied and all host I/O using that cluster's TVC will become throttled to prevent a cache overrun.

Note: Because having a distributed library in the out-of-empty-stacked-volume state impacts operations in a TS7740, it is something that should be avoided if at all possible!

It is highly recommended that you monitor the number of empty stacked volumes in a library. If the library is close to running out of a physical volume media type, action should be taken to either expedite the reclaim of physical stacked volumes or add additional ones. You can use the Bulk Volume Information Retrieval function to obtain the physical media counts for each library. The information obtained includes the empty physical volume counts by media type for the common scratch pool and each defined pool.

Above Threshold Warning state

The TS7700 enters the Above Threshold Warning state when the amount of data to copy exceeds the threshold for the installed cache capacity for 5 consecutive sample periods (amount of data to copy is sampled every 30 seconds). The TS7700 leaves the Above Threshold Warning state when the amount of data to pre-migrate is below the threshold capacity for 30 consecutive sample periods. The consecutive sampling criteria is to prevent excessive messages being created.

```
CBR3750I Message from library library-name:OP0160 Above threshold for uncopied
data in cache, throttling possible
```

7.6 Managing logical volumes

In addition to the tasks described in Chapter 4, “Hardware implementation” on page 131 and in 7.2, “Virtualization Engine TS7740 Management Interface” on page 315, we cover a few more management tasks and considerations for logical volumes in the following sections.

7.6.1 Scratch volume recovery for logical volumes

The advantage for this method of managing data is that if you determine that a volume was mistakenly returned to scratch, you only have to return the volume to private status to recover from the mistake, as long as you have not re-used the volume or the “grace period” has not expired. The method to recover depends on the tape management system used. For example, for DFSMSrmm, the following command will return the volume to private status and increase its retention period, including communicating the change to the TS7700 and Library Manager (see *z/OS DFSMSrmm Guide and Reference*, SC26-7404 for details of the command):

```
RMM CHANGEVOLUME yyyyyy STATUS(USER) RETPD(days) OWNER(userid)
where yyyyyy is the VOLSER
```

7.6.2 Ejecting logical volumes

Logical volumes are not physical entities that can be individually removed from the library. They reside on stacked volumes with many other logical volumes. If you issue an EJECT for a logical volume, all data on that volume will be lost.

> **Note:** You cannot recover the data on a logical volume after the EJECT command is processed.

Due to the permanent nature of the EJECT, the TS7700 only allows you to EJECT a logical volume that is in either the INSERT or SCRATCH (defined with fast- ready attribute) category. If a logical volume is in any other status, the EJECT fails.

> **Note:** This fact has proven to be cumbersome for volumes that happen to be in the ERROR category (000E). An easy way to eject such volumes is to use ISMF panels to set these volumes to the PRIVATE status. The volume status is propagated to DFSMSrmm. You can use DFSMSrmm to subsequently assign the volume to the SCRATCH status and eject it.

Ejecting large numbers of logical volumes can have a performance impact on the host and the library.

Tapes that are in INSERT status can be ejected by the resetting of the return code using the CBRUXENT exit. This exit is usually provided by your tape management system vendor.

When the tape is in SCRATCH status, follow the procedure for EJECT processing based on whether your environment is system-managed tape or BTLS. You also need to follow the procedure specified by your tape management system vendor. For DFSMSrmm, issue the RMM CHANGEVOLUME volser EJECT command. If your tape management system vendor does not specify how to do this, you can use one of the following commands:

- z/OS command LIBRARY EJECT,volser
- IDCAMS command LIBRARY EJECT,volser (for BTLS)
- ISMF EJECT line operator for the tape volume

The eject process fails if the tape is in another status or category. For libraries managed under DFSMS system-managed tape, the system command LIBRARY EJECT,volser issued to a logical volume in PRIVATE status fails with this message:

```
CBR3726I Function incompatible error code 6 from library <library-name>
         for volume <volser>.
```

> **Note:** In a DFSMS system-managed tape environment, if you try to eject a logical volume and get this error, OAM notifies the tape management system. This is done through the OAM eject exit CBRUXEJC before the eject request is sent to the tape library. The Library Manager will eventually fail the eject, but the tape management system has already marked the volume as ejected. Prior to APAR OW54054 there was no notification back that the eject failed.
>
> Failed Eject Notification was added to OAM with APAR OW54054 and is currently in all supported releases of DFSMS. Any tape management system supporting this notification can use this function.

If your tape management system is DFSMSrmm, you can use the following commands to clean up the RMM CDS for failed logical volume ejects and to re synchronize the TCDB and RMM CDS:

```
RMM SEARCHVOLUME VOL(*) OWN(*) LIM(*) INTRANSIT(Y) LOCATION(vts) -
  CLIST('RMM CHANGEVOLUME ',' LOC(vts)')

EXEC EXEC.RMM
```

The first RMM command asks for a list of volumes that RMM thinks it has ejected and writes a record for each in a sequential data set called *prefix.EXEC.RMM.CLIST*. The CLIST then checks that the volume is really still resident in the VTS library and, if so, it corrects the RMM CDS.

> **Note:** Limiting the number of outstanding ejects to a couple of thousand total per system will limit exposure to performance problems.

There are considerations to be aware of when ejecting large numbers of logical volumes. APAR OW42068 introduced Peer-To-Peer toleration support. This in effect treats all libraries as though they had the same eject processing restrictions as a Peer-to-Peer. This support is currently in all supported releases of DFSMS. There is a logical volume eject limit of 1000 that the Peer-to-Peer VTS can handle at any one time. Because libraries can be shared among systems, this limit can be reached quickly if many ejects are issued from multiple hosts.

OAM helps by restricting the number of ejects sent to each library at a given time and manages all the outstanding requests. This management requires storage on the host, and a large number of ejects can force OAM to reserve large amounts of storage. Additionally, there is a restriction on the number of eject requests on the device services' queue. All of these conditions can have an impact on the host's performance.

So the recommended limit for the number of outstanding eject requests is no more than a couple thousand per system. Additional ejects can be initiated when others complete. Further information can be obtained within APAR OW42068. The following command can be used on the System z hosts to list the outstanding and the active requests.

F OAM,QUERY,WAITING,SUM,ALL

F OAM,QUERY,ACTIVE,SUM,ALL

7.6.3 Searching the Library Manager database for volser, category, devices

Using the 3953 or 3494 Library Managers, you can search the Library Manager database and view selected volumes in the TS3500 and TS7700 according to specified search criteria. You can search for both logical and physical volumes.

Keep in mind, that you are looking in the Library Manager database and that this database does not include information from a logical point of view. For example, to obtain information about logical-stacked volume relationship, copy status, or cache status, you must use the TS7700 Management Interface. Note that you can only obtain the information for one cluster at a time and not for multiple clusters together.

Figure 7-61 shows the ETL Specialist Welcome panel to query the Library Manager database. Select **Search Database** from the Administer Library Manager work items.

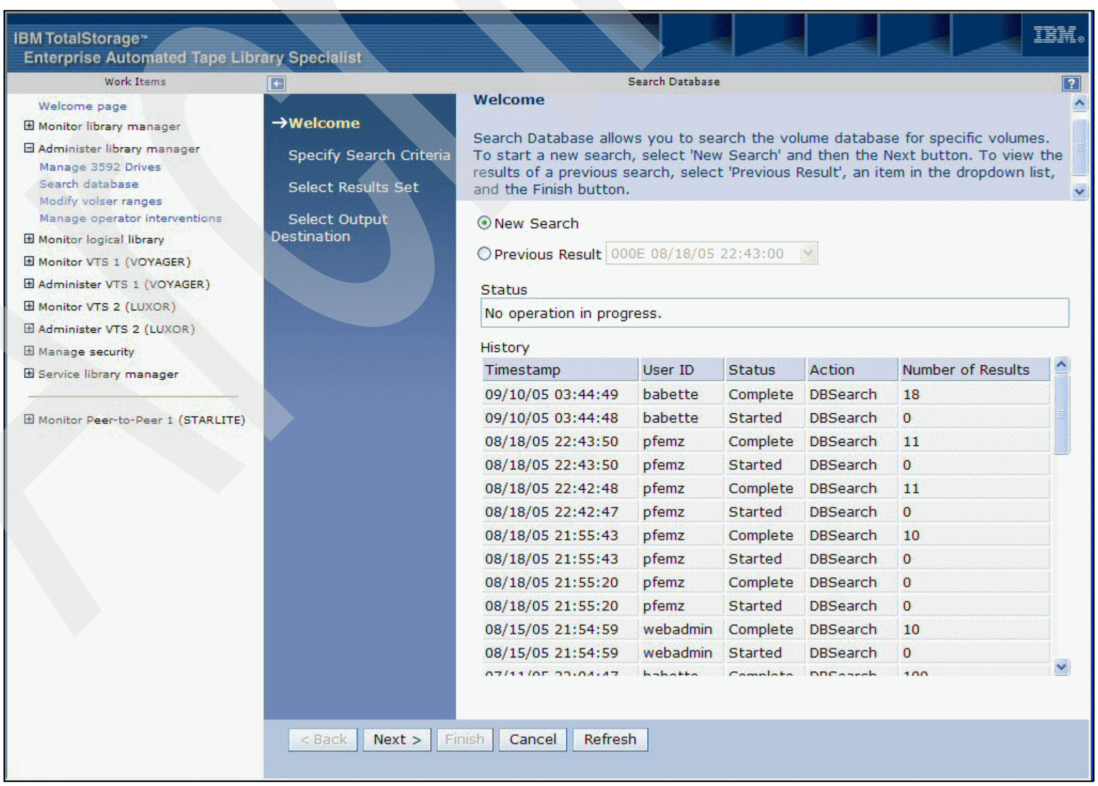

Figure 7-61 ETL Specialist Search Database for Volumes window

When clicking **Next**, the panel shown in Figure 7-62 on page 396 displays, which allows you to make your selections for the search criteria. You can specify the following search criteria:

- Volser
- Category
- Media type
- Device
- Partition
- Expire time
- Home rack
- Home row
- Home column
- Home pool
- Current rack
- Current row
- Current column
- Current pool
- APM construct names
  - Storage group
  - Management class
  - Storage class
  - Data class
- Volume flags
  - Misplaced
  - Unreadable
  - Mounted
  - Inaccessible
  - Manual mode

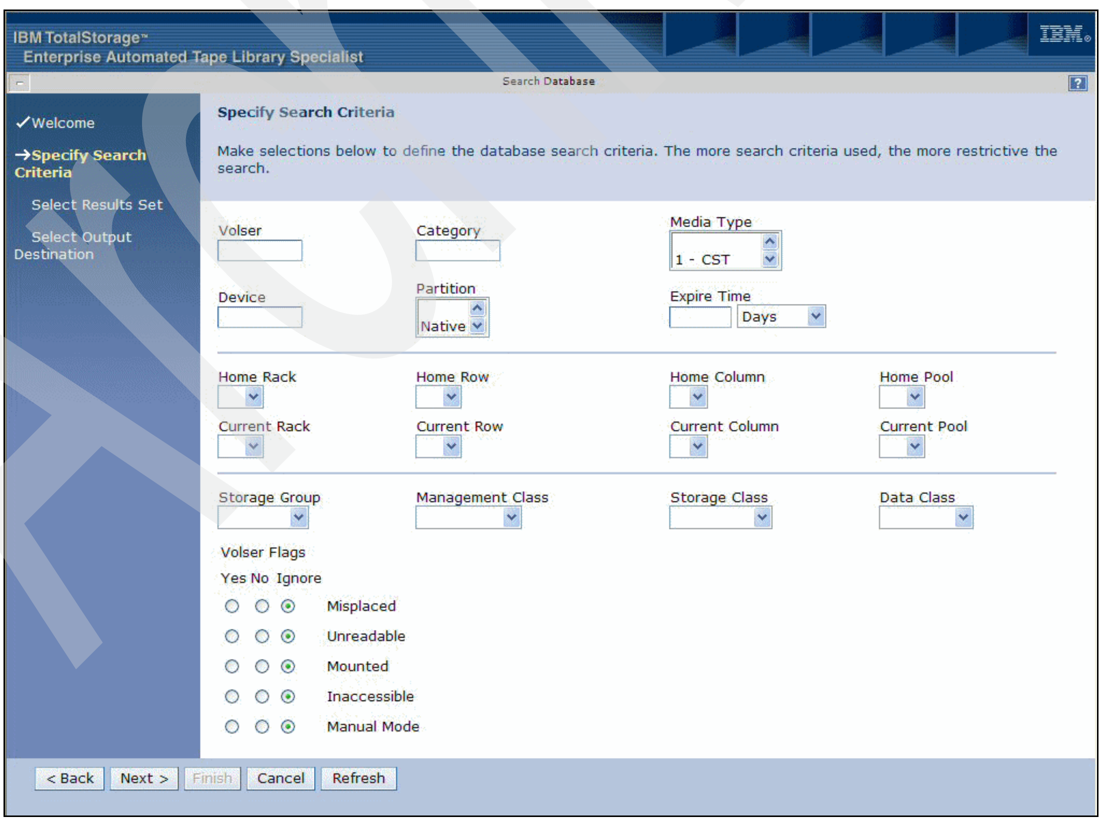

Figure 7-62 Search Database window: Select Search Criteria

After specifying your search criteria and clicking **Next**, the window shown in Figure 7-63 on page 397 is displayed. In this panel you can select which values you want to have displayed.

You can select the following values to be displayed for the volumes that meet the search criteria specified before:

- Volser
- Category
- Media type
- Device
- Partition
- Expire time
- Home rack
- Home row
- Home column
- Home pool
- APM construct names
  - Storage group
  - Management class
  - Storage class
  - Data class
- Volume flags
- Category order
- Mount date
- Number of mounts

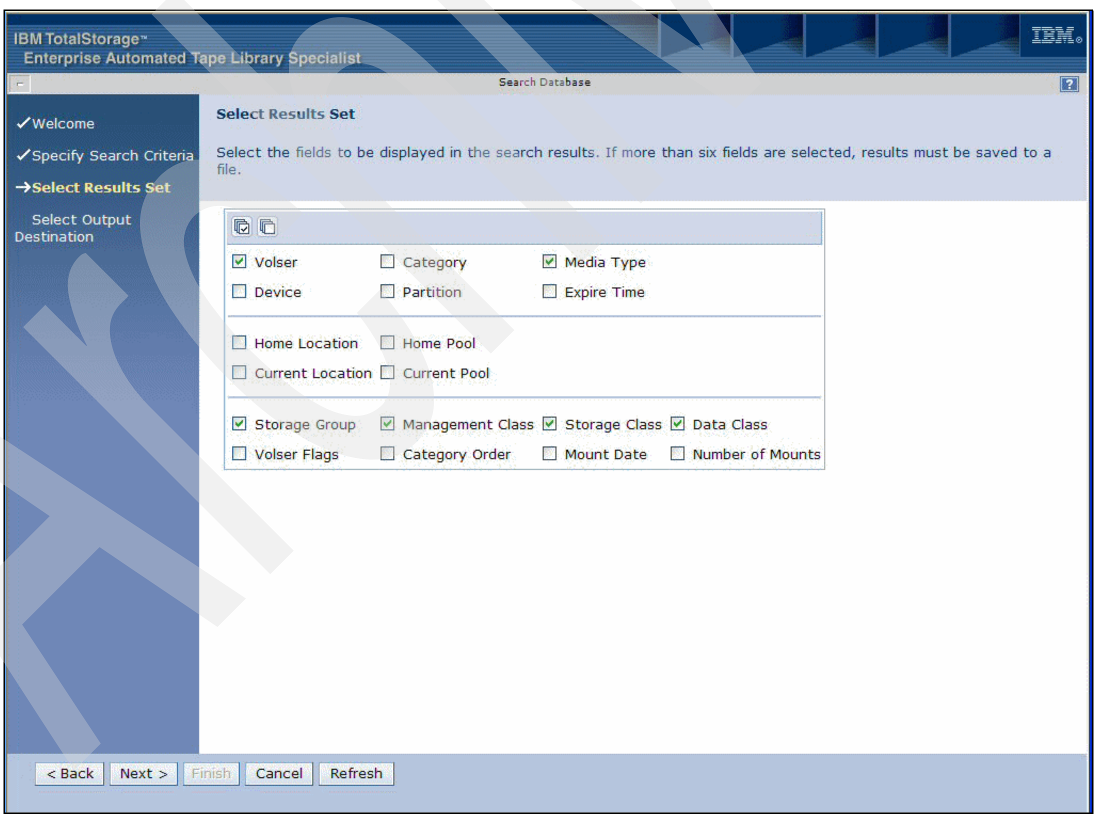

Figure 7-63 Search Database window: Select Results Set

Depending on the number of columns selected in the window shown in Figure 7-63, you can select whether you want to have your output displayed on the panel or written to a file (see Figure 7-64).

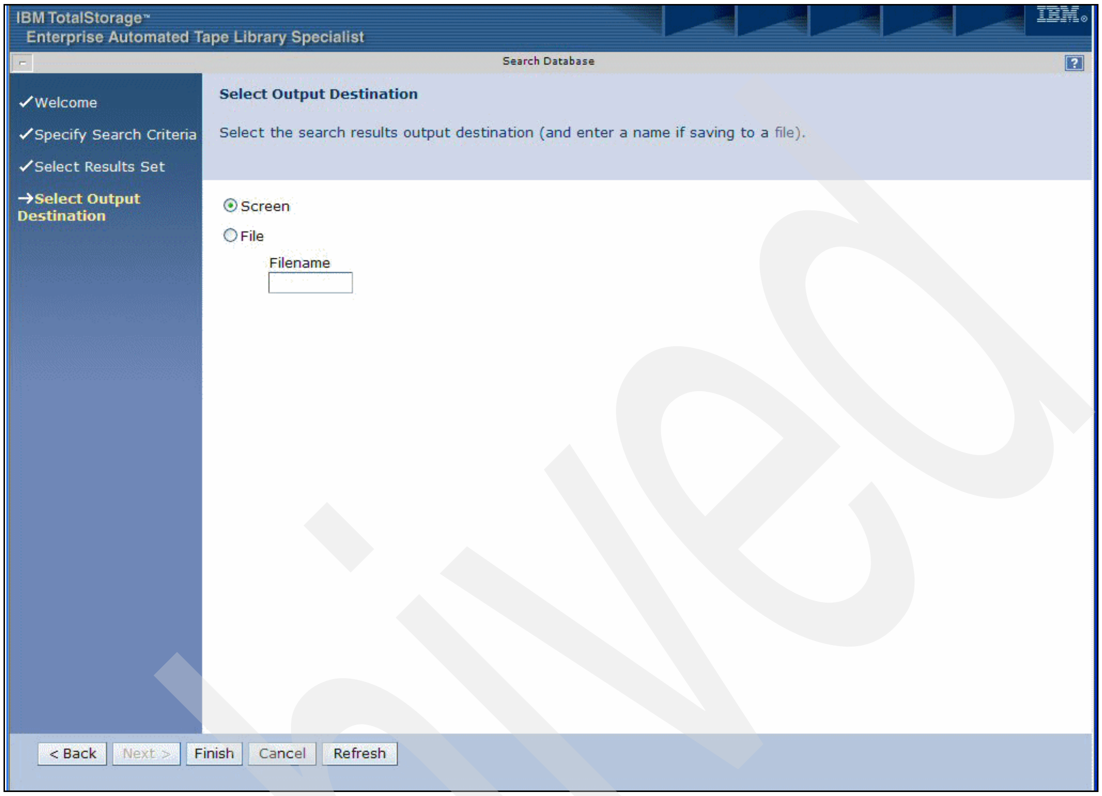

Figure 7-64 Search Database window: Select Output Destination

If you select between one and six columns for the Results Set, you can select between Screen and File on the panel shown in Figure 7-65. If you select seven or more columns, the output will be written to a file.

Figure 7-65 shows the sample output of a search selecting all VOLSERs.

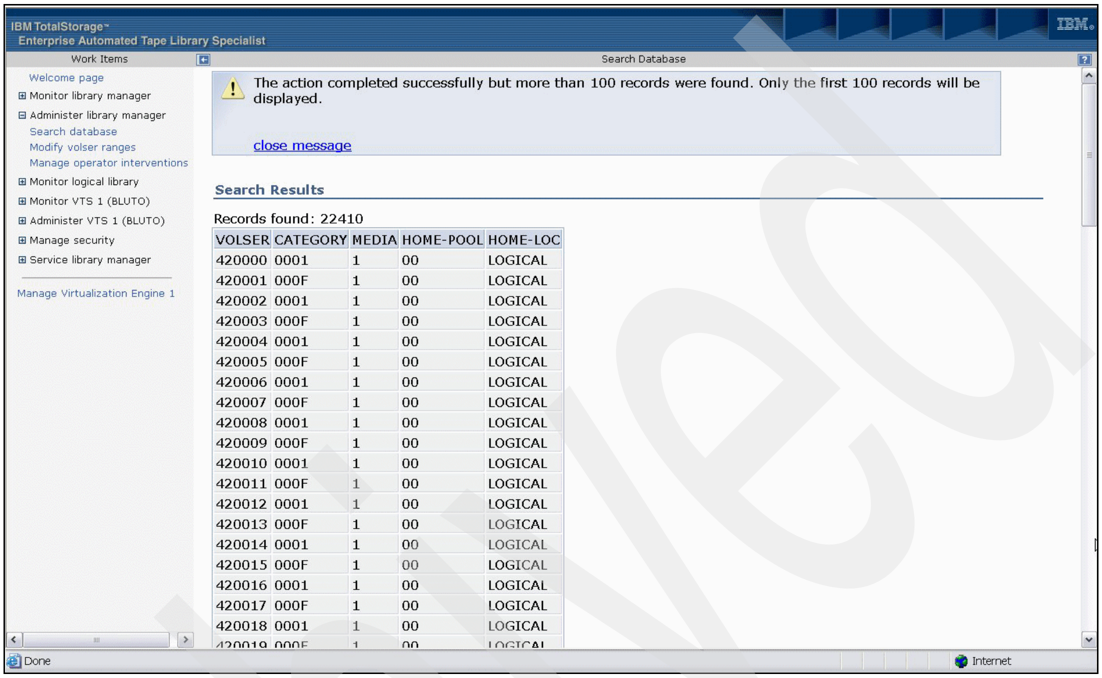

Figure 7-65 Search Database window: Search Results

As you can see from the categories listed in the example shown in Figure 7-65, these are all TS7700 volumes (example: Category 0001 is z/OS DFSMS scratch MEDIA1; Category 000F indicates a private volume).

7.6.4 Find logical volume

Is it not possible to search a logical volume to determine on which a stacked volume resides from the ETL Specialist. For the logical volumes that belong to the TS7700 Virtualization Engine, you can initiate the search from the TS7700 Virtualization Engine Management Interface.

For more information, refer to “Logical Volume Details” on page 321. See Figure 7-66 and Figure 7-67 on page 400 as reference. This is an extract from the display of the Logical Volume detail.

You can see at the bottom of the panel that the logical volumes on which we are looking for information reside on the two physical volumes, one on the current owner cluster and the other on the cached copy cluster.

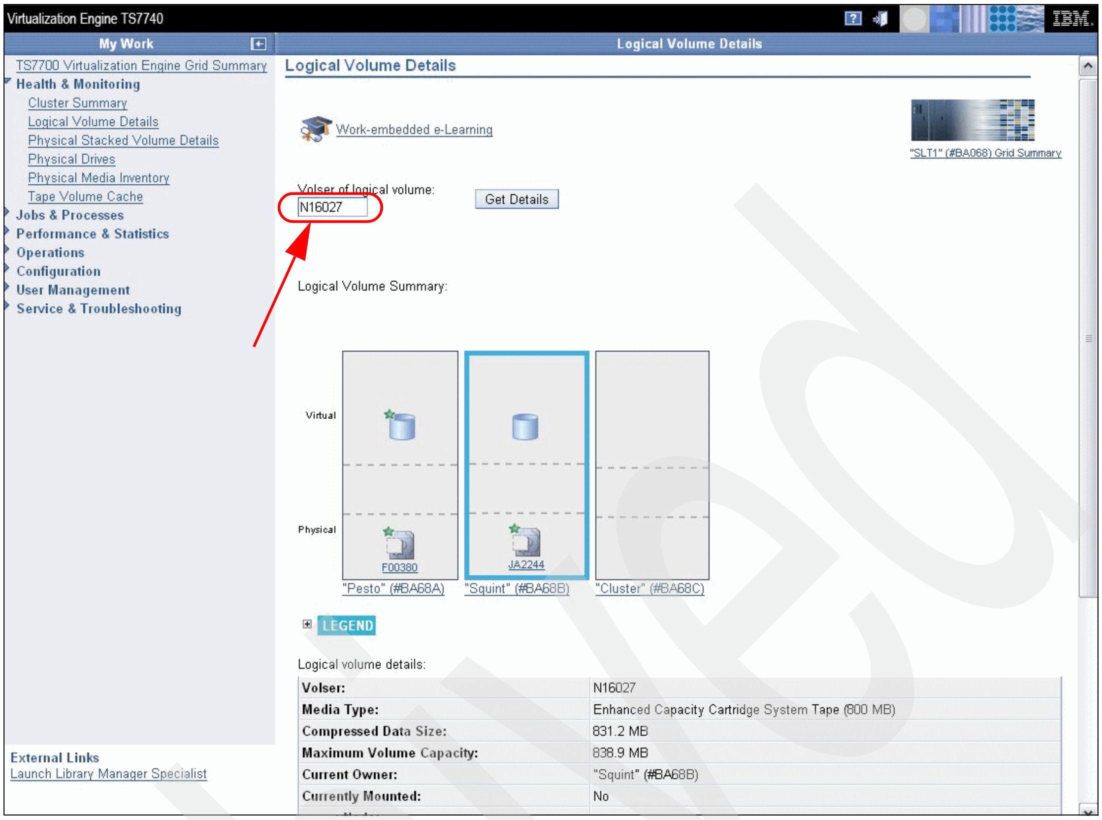

Figure 7-66 TS7700 Virtualization Engine Logical Volume Detail part1

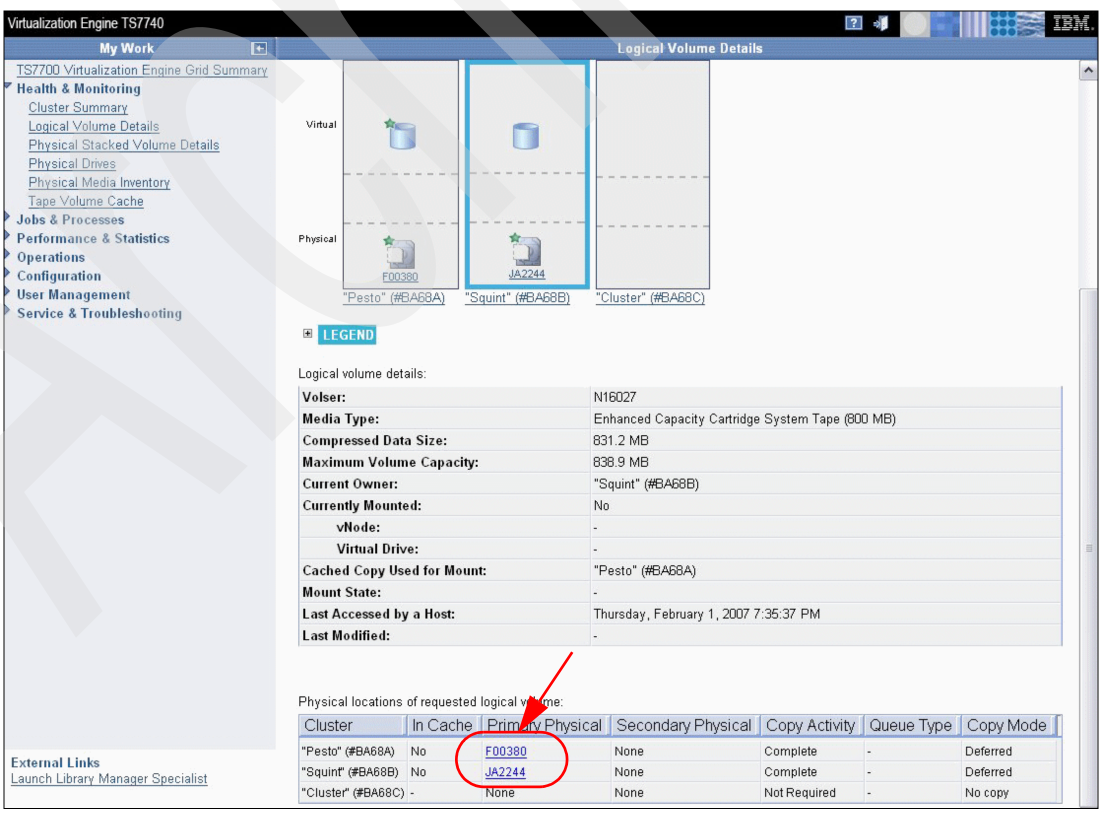

Figure 7-67 TS7700 Virtualization Engine Logical Volume Detail part 2

7.7 Messages and displays

This section describes the enhanced message support and some relevant messages related to the TS7700.

7.7.1 Console name message routing

Today, with console name message routing support, many of the library-specific messages are only issued to the specified library console (if defined) and not to the specified routing codes.

Although this is not specific to a TS7700, the following critical, action-related messages will now be issued using the specified library console and routing codes, providing maximum visibility.

```
CBR3759E Library x safety enclosure interlock open.
CBR3764E Library x all storage cells full.
CBR3765E No cleaner volumes available in library x.
CBR3753E All convenience output stations in library x are full.
CBR3754E High capacity output station in library x is full.
CBR3755E {Input|Output} door open in library x.
CBR3660A Enter {list of media inserts} scratch volumes into x.
```

7.7.2 Grid messages

This section lists some of the TS7700 grid-specific messages that you might see. For a complete and current list, see the appropriate volume of *z/OS MVS System Messages*.

Incompatibility error message

In case of an incompatible function error, you might see the message CBR3726I.

```
CBR3726I Function incompatible error code error-code from library library-name
for volume volser.
```

Explanation: An error has occurred during processing of volume *volser* in library *library-name*. The library returned a unit check with an error code *error-code*, which indicates that an incompatible function has been requested. A command has been issued that requests an operation that is understood by the subsystem microcode, but cannot be performed due to one of the following errors:

X'00' The function requested is not supported by the subsystem to which the order was issued.

X'01' Library attachment facility not installed and allowed.

X'02' Not currently used.

X'03' High capacity input/output facility is not configured.

X'04' Reserved

X'05' Volume requested to be mounted is not compatible with the device allocated.

X'06' The logical volume can only be ejected if it is in the insert category and has a mount count of zero, or is assigned to a category that has the Fast-Ready attribute set.

X'07' There is no pending import or export operation to cancel.

X'08' There are not enough (four are needed) physical drives available to initiate the import or export operation.

X'09' Reserved

X'0A' Reserved

X'0B' Reserved

X'0C' Reserved

X'0D' The Grid TS7700 subsystem is either in service preparation mode or has an unavailable component within the subsystem such as an unavailable distributed library. Audit, eject, or entry-related commands are not being accepted at this time.

X'0E' The Grid TS7700 subsystem already has one thousand eject requests queued and is not accepting any more eject requests at this time.

X'0F' An inappropriate library function was issued to the Grid TS7700 subsystem.

X'10' The VTC in the Peer-to-Peer VTS subsystem or the distributed library in a TS7700 grid configuration that the command was issued to is in read-only or write-protect mode and is not accepting requests that change the category or attributes of a volume. This mode of operation is provided to support disaster recovery operations in a configuration where the configuration is split between two physical sites.

X'12' The volume specified has a non-zero expire time associated with it. A volume in this state cannot be mounted, moved, or have its attributes modified until the expire time has elapsed.

X'30' The TS7700 cluster that the command was received or does have an available path to the cluster that currently owns the volume and ownership takeover is not enabled.not obtain ownership of the volume from the cluster that owns it and ownership takeover is not enabled.

X'32' There is more than one valid copy of the specified export list volume in the TS7700 grid configuration.

X'33' An export operation was issued to a TS7700 that is performing a global operation. Global operations include volume inserts, volume deletions through the management interface, damaged volume recovery and disaster recovery. Export operations are not being accepted at this time.

7.7.3 Display grid status

The following messages can be issued for the Grid TS7700.

CBR1100I OAM status

The message is issued in response to the following operator command:

DISPLAY SMS,OAM

Example 7-14 shows the complete message text.

Example 7-14 DISPLAY SMS,OAM command

```
CBR1100I OAM status: 618
 TAPE TOT   ONL   TOT   TOT   TOT   TOT    TOT    ONL    AVL    TOTAL
      LIB   LIB   AL    VL    VCL   ML     DRV    DRV    DRV   SCRTCH
        5     2     0     0     2     0    192    192    127    46079
 There are also 2 VTS distributed libraries defined.
```

```
CBRUXCUA processing ENABLED.
CBRUXEJC processing ENABLED.
CBRUXENT processing ENABLED.
CBRUXVNL processing ENABLED.
```

If there are TS7700 subsystems defined to the system, the following status line is displayed, reflecting the number of distributed libraries that are associated with the composite libraries.

```
There are also numvdl-lib VTS distributed libraries defined.
```

Additional status lines might appear containing one or more of the following messages:

```
Copy operations disabled.
VTS operations degraded.
Immediate Mode Copy operations deferred.
Service preparation occurring.
Library is out of empty stacked volumes.
```

CBR1110I OAM library status

This message is issued in response to

```
DISPLAY SMS,LIBRARY(library-name),DETAIL
```

Example 7-15 shows the complete message text.

Example 7-15 DISPLAY SMS, LIBRARY command

```
CBR1110I OAM library status: 334
 TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY SCRTCH  ON OP
 LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS   VOLS
 HYDRAO    VCL  3957-V06  128  128   64      0      0  44365  Y  Y
 ------------------------------------------------------------------
 MEDIA      SCRATCH       SCRATCH       SCRATCH
 TYPE         COUNT     THRESHOLD      CATEGORY
 MEDIA2       44365           100          0002
 ------------------------------------------------------------------
 DISTRIBUTED LIBRARIES:  HYDRAD
-------------------------------------------------------------------
 LIBRARY ID:  10001
 OPERATIONAL STATE:  AUTOMATED
 ERROR CATEGORY SCRATCH COUNT:            0
 CORRUPTED TOKEN VOLUME COUNT:            0
-------------------------------------------------------------------
```

Note: Library type VCL indicates the Composite Library as opposed to VDL for the distributed library.

CBR1180I OAM tape volume status

This message is presented in response to the following operator command:

```
DISPLAY SMS,VOLUME(volser)
```

Example 7-16 lists the complete message text for logical volume A01878.

Example 7-16 DISPLAY SMS,VOLUME command

```
CBR1180I OAM tape volume status: 933
VOLUME  MEDIA     STORAGE   LIBRARY   USE  W  C  SOFTWARE  LIBRARY
        TYPE      GROUP     NAME      ATR  P  P  ERR STAT  CATEGORY
A01878  MEDIA2    SGHYP02   HYDRAO    P    N  N  NOERROR   PRIVATE
-------------------------------------------------------------------
RECORDING TECH:     36 TRACK        COMPACTION:         YES
SPECIAL ATTRIBUTE:  NONE            ENTER/EJECT DATE:   2006-08-03
CREATION DATE:      2006-08-03      EXPIRATION DATE:    2006-09-23
LAST MOUNTED DATE:  2006-09-20      LAST WRITTEN DATE:  2006-09-20
SHELF LOCATION:
OWNER: BHAEUSS
LM SG: SGHYP02   LM SC: SCHYPG1   LM MC: MCHYNCOP  LM DC: DCHY1GB
-------------------------------------------------------------------
Logical volume.
Volume is cache resident.
```

7.8 Recovery scenarios

In this section we discuss some of the potential recovery scenarios you might have to perform. You are notified of most of the errors that require operator attention through an `Intervention Required` message on the IBM 3953 Library Manager or at the IBM 3494 Tape Library and the hosts.

7.8.1 Hardware conditions

Here we describe potential hardware failure scenarios.

IBM 3592 tape drive failure

When the TS7700 determines that one of its tape drives is not operating correctly and requires service (it is likely that the drive has excessive read or write errors), the drive is marked offline and an SSR must be called. The following Intervention Required message is displayed on the Library Manager Console.

```
CBR3750I MESSAGE FROM LIBRARY lib: Device xxx made unavailable by a VTS. (VTS z)
```

Operation of the TS7700 continues with a reduced number of drives until the repair action on the drive is complete. To recover, the SSR repairs the failed tape drive and makes it available for the TS7700 to use it again.

Power failure

User data is protected in the event of a power failure, as it is stored on the TVC. Any host jobs reading or writing to virtual tapes will fail as they would with a real IBM 3490E and will need to be restarted after the TS7700 is available again. When power is restored, the TS7700 recovers access to the TVC, using information available from the TS7700 database and logs. After power is restored, restart the IBM TS3500 Tape Library. The IBM 3953 Library Manager and the IBM 3494 Tape Library server restart the TS7700.

TS7700 database corruption

To recover from corruption of the TS7700 database, you have to perform a so-called disaster recovery. This is a TS7700 process further described in Chapter 9, “Disaster recovery and failover scenarios” on page 489.

The TS7700 maintains a database of information about the location and status of logical volumes on the stacked volumes it manages. When a stacked volume has been filled with logical volumes, a backup of the entire database is placed at the end of the filled stacked volume. The database contains a time and date stamp that identifies when the backup was performed.

When the database copy operation is complete, a message is sent to the attached hosts:

```
CBR3750I MESSAGE FROM LIBRARY lib: VTS Database Backup written to Physical Tape
xxxxxx.
```

The disaster recovery process causes the TS7700 to load the stacked volumes, locate the latest version of the database and restore it. Any logical volumes written after the last backup of the TS7700 database are lost. When the restoration is complete, a message is displayed on the Library Manager console informing you of the date and time when the TS7700 database was restored.

Library Manager failure

If the IBM 3953 or 3494 Library Managers fail in a TS7700 environment, the situation is the same as with any IBM 3953 or IBM3494. Current tape tasks complete, but no new tasks can start until the Library Manager switchover is complete (on HA1) or the problem has been solved. To recover, fix the problem that caused the Library Manager to fail. If the HA1 Frames are not installed, the IBM 3593 or IBM 3494 and the TS7700 are unavailable until the repair action is complete and the Library Manager is restored. If HA1 Frames are installed, the second Library Manager takes over. The library will be taken away from the hosts and will be unavailable until takeover is complete.

For detailed scenarios on how to handle failure scenarios with HA1 Frames, refer to *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632.

> **Note:** On the TS3500/3953/3494 Tape Library the HA1 and the second Library Manager are two optional, separate features. Ordering the HA1 feature for the IBM TS3500 does not give you a second Library Manager but does give you the second accessor.

Accessor failure and manual mode

If the HA1 Frames are not installed, failure of the accessor results in the library being unable to automatically mount physical volumes. When a host requests a logical volume for recall, the request is converted into a message to mount the stacked volume where it is located. Fast-Ready scratch mounts do not require a physical mount until the data is to be copied to a stacked tape. The message is displayed on the Library Manager Console for manual mode operation.

If HA1 Frames are installed, the second accessor takes over. Then you can call your IBM service representative to repair the failed accessor.

Gripper failure

Here we show the gripper failure scenario on the two different libraries that can be connected to the TS7700 Virtualization Engine.

- **IBM 3494 Tape Library**: If the IBM 3494 Tape Library has dual grippers installed, and a gripper fails, the library operations will continue with the other one. While the gripper is being repaired, the accessor is not available, so the library should be run in manual mode until the repair is complete.

  If the HA1 Frames are installed, the second accessor will be used until the gripper is repaired. For detailed information about running the IBM 3494 Tape Library system in manual mode, see *TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation,* SG24-4632.

- **IBM TS3500 Tape Library**: Because the IBM TS3500 Tape Library has dual grippers, if a gripper fails, the library operations will continue with the other one. While the gripper is being repaired, the accessor is not available, so the library should be run in manual mode until the repair is complete.

  If the HA1 Frames are installed, the second accessor will be used until the gripper is repaired. For detailed information about running the IBM TS3500 Tape Library system in manual mode see *IBM System Storage TS3500 Tape Library Operator Guide,* GA32-0560.

Out of stacked volumes

If the Tape Library runs out of stacked volumes, then copying to the 3592 tape drives will fail and an Intervention Required message is sent to the Library Manager. All further logical mount requests are delayed by the Library Manager until more stacked volumes are added to the IBM TS3500 Tape Library or to the IBM 3494 Tape Library if is connected to the TS7700 Virtualization Engine. To recover, insert more stacked volumes. Copy processing can then continue.

Damaged cartridge pin

The 3592 it has a metal pin that is grabbed by the feeding mechanism in the 3592 tape drive to load the tape on to the take-up spool inside the drive. If this pin gets dislodged or damaged, then follow the information in *IBM TotalStorage Enterprise Tape System 3592 Operators Guide,* GA32-0465.

> **Important:** Repairing 3592 tape should only be done for data recovery. After the data has been moved to a new volume, you should replace the repaired cartridge.

Broken tape

If a 3592 tape cartridge is physically damaged and unusable (the tape is crushed, or the media is physically broken, for example), the TS7700 cannot recover the contents. This is the same for any tape drive media cartridges today. You can generate a list of logical volumes that are on that stacked volume. Consult with your SSR, because further ways might be available to recover a broken tape from within IBM.

Logical mount failure

When a mount request is received for a logical volume, the TS7700 determines whether the mount request can be satisfied and if so, tells the host that it will process the request. Unless an error condition is encountered in the attempt to mount the logical volume, the mount operation completes and the host is notified that the mount was successful. With the TS7700, the way that a mount error condition is handled is different then with the prior generations of VTSs. With the prior generation of VTSs, the VTS always indicated to the host that the mount completed, even if a problem had occurred. When the first I/O command was issued, the VTS would then fail that I/O because of the error. This would result in a failure of the job, without the opportunity to try and correct the problem and retry the mount.

With the TS7700 subsystem, if an error condition is encountered during the execution of the mount, instead of indicating that the mount was successful, the TS7700 returns completion and reason codes to the host indicating that a problem was encountered. With DFSMS, the logical mount failure completion code results in the console messages shown in Example 7-17.

Example 7-17 Unsuccessful mount completion and reason codes

```
CBR4195I LACS RETRY POSSIBLE FOR JOB job-name
CBR4171I MOUNT FAILED. LVOL=logical-volser, LIB=library-name,
PVOL=physical-volser, RSN=reason-code
...
CBR4196D JOB job-name, DRIVE device-number, VOLSER volser, ERROR CODE error-code.
REPLY 'R' TO RETRY OR 'C' TO CANCEL
```

The following reason codes provide information about the condition that caused the mount to fail:

X'10' Internal Error Detected
X'11' Resend Special Case
X'20' Specific Volume In Use On Another Cluster
X'21' Scratch Volume Selected In Use On Another Cluster
X'22' Valid Volume Inaccessible
X'23' Local Cluster Path to Volume's Data No Longer Available
X'24' Remote Cluster Path To Volume's Data No Longer Available
X'25' Copy Required, but Cluster Copying Inhibited
X'30' Local Cluster Recall Failed, Stacked Volume Misplaced
X'31' Local Cluster Recall Failed, Stacked Volume Inaccessible
X'32' Local Cluster Recall Failed, Stacked Volume Unavailable
X'33' Local Cluster Recall Failed, Stacked Volume No Longer In Library
X'34' Local Cluster Recall Failed, Stacked Volume Load Failure
X'35' Local Cluster Recall Failed, Stacked Volume Access Error
X'38' Remote Cluster Recall Failed, Stacked Volume Misplaced
X'39' Remote Cluster Recall Failed, Stacked Volume Inaccessible
X'3A' Remote Cluster Recall Failed, Stacked Volume Unavailable
X'3B' Remote Cluster Recall Failed, Stacked Volume No Longer In Library
X'3C' Remote Cluster Recall Failed, Stacked Volume Load Failure
X'3D' Remote Cluster Recall Failed, Stacked Volume Access Error

When a mount is failed on a TS7700, you can attempt to resolve the underlying problem indicated by the reason code and then have the mount retried. For example, if the failure was because a recall was required and the stacked volume was unavailable because it was accidentally removed from the library, recovery involves returning the volume to the library and then replying to the CBR4196D message with RETRY.

Orphaned logical volume

This occurs when the TS7700 database has a reference to a logical volume but no reference to its physical location. This could result from hardware or internal software errors. When it does occur, any data on the logical volume is lost. When this error occurs, contact your System Service Representative.

Internal-external label mismatch

If a label mismatch occurs, the stacked volume is ejected to the convenience Input/Output station and the Intervention Required condition is posted at the 3953 Library Manager console (see Example 7-18).

Example 7-18 Label mismatch

```
CBR3750I MESSAGE FROM LIBRARY lib: A stacked volume has a label mismatch and has
been ejected to the Convenience Input/Output Station.
Internal: xxxxxx, External: yyyyyy
```

The host is notified that Intervention Required conditions exist. Investigate the reason for the mismatch. If possible, relabel the volume to use it again.

Failure during reclaim

If there is a failure during the reclamation process, the process is managed by the TS7700 microcode. No user action is needed; recovery is managed internally.

Excessive temporary errors on stacked volume

When a stacked volume is determined to have an excessive number of temporary data errors, to reduce the possibility of a permanent data error, the stacked volume is placed in read-only status.

7.8.2 TS7700 software failure

If a problem develops with the TS7700 software, the TS7700 issues an Intervention Required message to the Library Manager and attempts to recover. In the worst case, this can involve a reboot of the TS7740 itself. If the problem persists, you need to contact your IBM service representative. The Intervention Required message (Example 7-19) is sent to the Library Manager.

Example 7-19 VTS software failure

```
CBR3750I MESSAGE FROM LIBRARY lib: Virtual Tape System z has a CHECK-1 (xxxx)
failure
```

The TS7700 internal recovery procedures handle this situation and restart the TS7740. Refer to Chapter 9, “Disaster recovery and failover scenarios” on page 489.

8

Performance and monitoring

In this chapter, we describe the factors that determine and influence the performance of the IBM TS7700 Virtualization Engine. We also describe what actions to take, when necessary, to improve the TS7700 Virtualization Engine's performance.

This chapter includes:

- An overview of the shared tasks that are running in the TS7700 server
- A description of a TS7700 monitoring and performance evaluation methodology
- A walk-through of a TS7700 capacity planning case study
- A review of BVIR and VEHSTATS reporting

We discuss the TS7700 shared resources so that you can understand the impact that contention for these resources has on the performance of the TS7700.

The monitoring section can help you to understand the important values recorded in the TS7700, as well as discussing in-depth investigations on performance issues that might arise with the TS7700. It can also help you recognize the symptoms that indicate the VTS is at or near its maximum capacity. The information provided can help you to evaluate the options available to improve the throughput or performance of the TS7700 Virtualization Engine.

The capacity planning case study illustrates some guidelines and techniques for the management of virtual and stacked volumes associated with the TS7700.

8.1 TS7700 performance characteristics

The TS7700 Virtualization Engine can provide significant benefits to the tape processing environment. In general, performance depends on such factors as total system configuration, Tape Volume Cache capacity, the number of physical tape drives available to the TS7700, the number of channels, the read/write ratio as well as data characteristics such as blocksize and mount pattern. These performance factors and numbers for different configurations are discussed in 8.3, “Monitoring TS7700 performance” on page 420.

A comparison of bandwidth numbers with previous IBM Virtual Tape Server (VTS) generations compared to a fully configured TS7700, based on initial modeling and measurements, and assuming a 2.66:1 compression ratio, is shown in Figure 8-1 and described in more detail in the Performance White Papers.

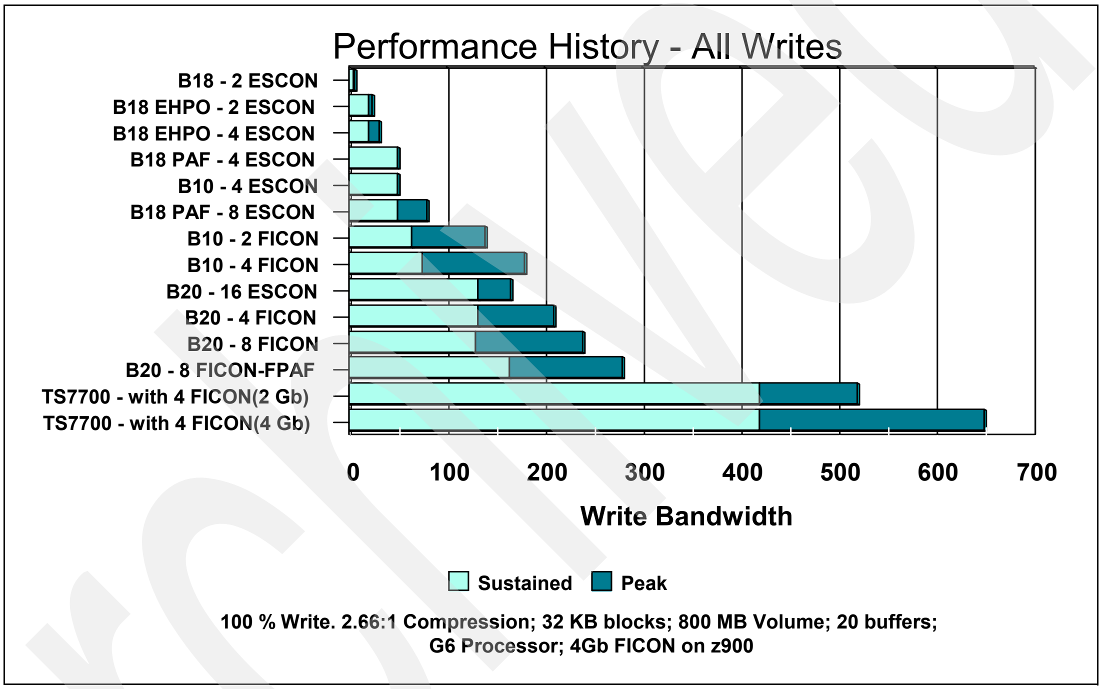

Figure 8-1 Performance comparison

In addition to faster performing hardware components such as faster FICON channel adapters, the more powerful TS7740 Controller, and faster disk cache, the new TS7700 architecture provides for improved performance and throughput characteristics of the TS7700 Virtualization Engine. From a performance aspect, important characteristics of the new architecture are:

- With the selection of DB2 as the central repository, the TS7700 provides a standard SQL interface to the data and all data is stored and managed in one place. DB2 also allows for more control over backend performance.
- The cluster design with vNode and hNode provides increased configuration flexibility over the monolithic design of the Virtual Tape Server.
- The use of TCP/IP instead of FICON for site-to-site communication eliminates the requirement to use vendor channel extenders.

The TS7700, compared to the VTS, also allows for faster job run times. Laboratory measurements have shown significant reduction in single job run times; see Figure 8-2.

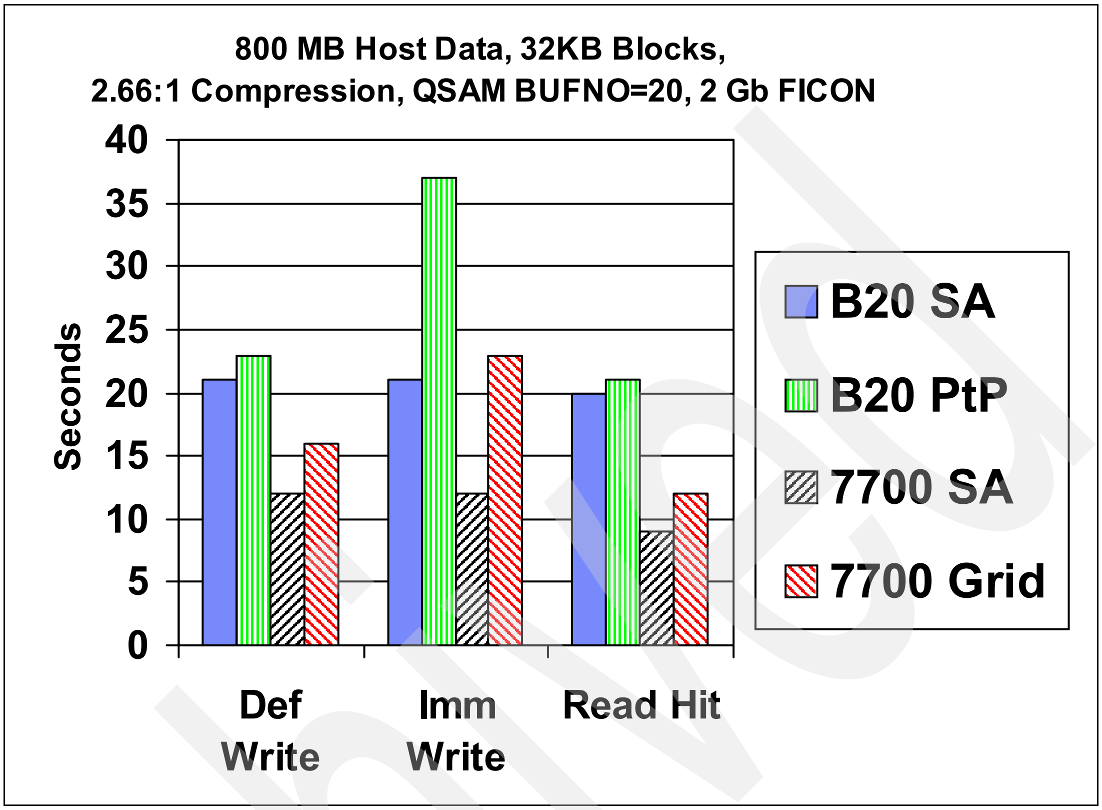

Figure 8-2 Single-job time comparison TS7700 versus B20

The same measurements as in Figure 8-2 were taken on TS7700 with a 4 Gb FICON connection running on z900. The first results are as follows:

- Single Cluster Grid with Deferred Write 7 seconds
- Single Cluster Grid with Immediate Write 7 seconds
- Single Cluster Grid Cache read hit 6 seconds
- Two-Cluster Grid with Deferred Write 12 seconds
- Two-Cluster Grid with Immediate Write 19 seconds
- Two-Cluster Grid Cache read hit 10 seconds

8.1.1 Single cluster performance

Laboratory measurements have shown the following results:

- TS7700 logical fast ready and read hit mount times have been typically in the range of 2 to 5 seconds.
- TS7700 recall performance has been on a par with B20 recall performance.
- TS7700 pre-migration rates have been at least twice the B20 pre-migration rates.
- TS7700 return-to-scratch rates have been 3 to 4.5 returns per second with no other TS7700 activity.

Refer to the TS7700 Performance White Paper for more details on Single Cluster performance characteristics.

8.1.2 Multi Cluster Grid performance

Laboratory measurements have shown the following results:

- Read/write data rates for a Two-Cluster Grid are somewhat reduced from those for a Single Cluster Grid and reduced further for a Three-Cluster Grid.
- TS7700 cluster-to-cluster copy rates have equalled or exceeded the B20 copy rates.
- TS7700 fast ready and read hit mount times have been typically in the 3 to 7 second range.
- TS7700 recall performance has been on a par with B20 recall performance.
- TS7700 return-to-scratch rates have been between 2 and 2.5 returns per second when returned from just one site, and 2.5 to 3 per second when returned concurrently from dual sites, with no other TS7700 activity and with less than 100,000 logical scratch volumes in the Two or Three-Cluster Grid.
- Impact of distance on write data rates is similar to the impact on VTS. Read performance at distance has been significantly improved with Release 1.3.

Refer to the TS7700 Performance White Paper for more details on performance characteristics.

8.1.3 Performance delay dependent on distance of a Two-Cluster Grid

There has been some testing where simulation of distance will show the impact of placing the two clusters many miles apart from each other. Figure 8-3 and Figure 8-4 show examples of read hit rates in terms of MB/s and how they are affected based on distance. Both 2 Gbps and 4 Gbps FICON adapters are shown.

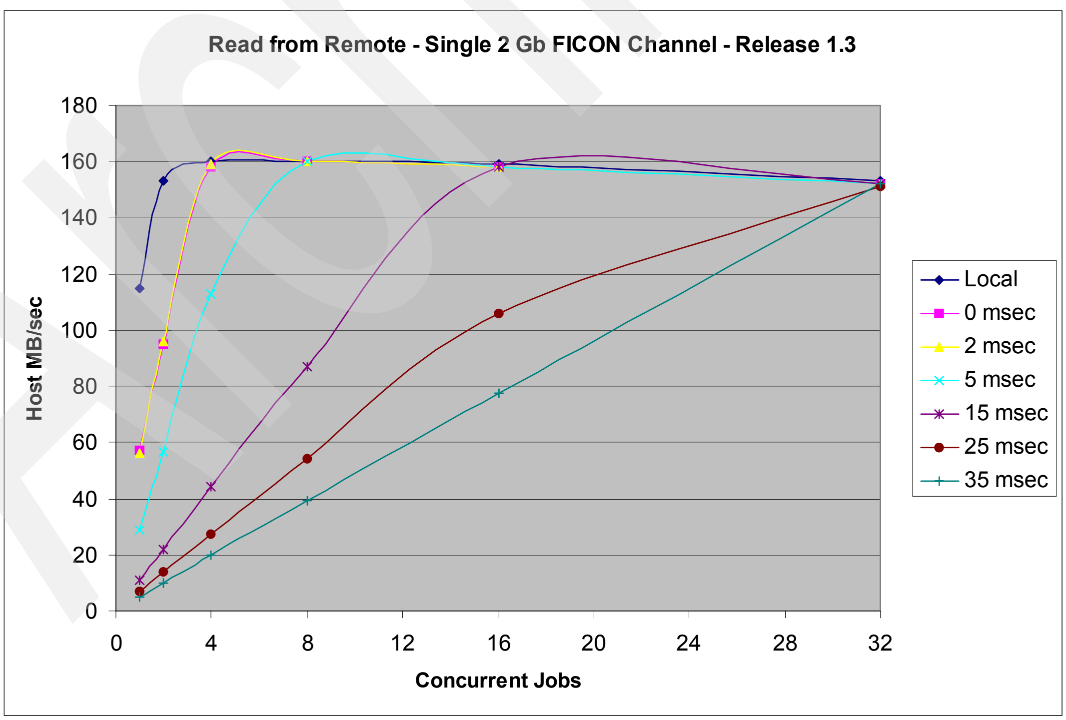

Figure 8-3 Read hit for remote TS7700 in Two-Cluster Grid on a single 2 Gb FICON channel

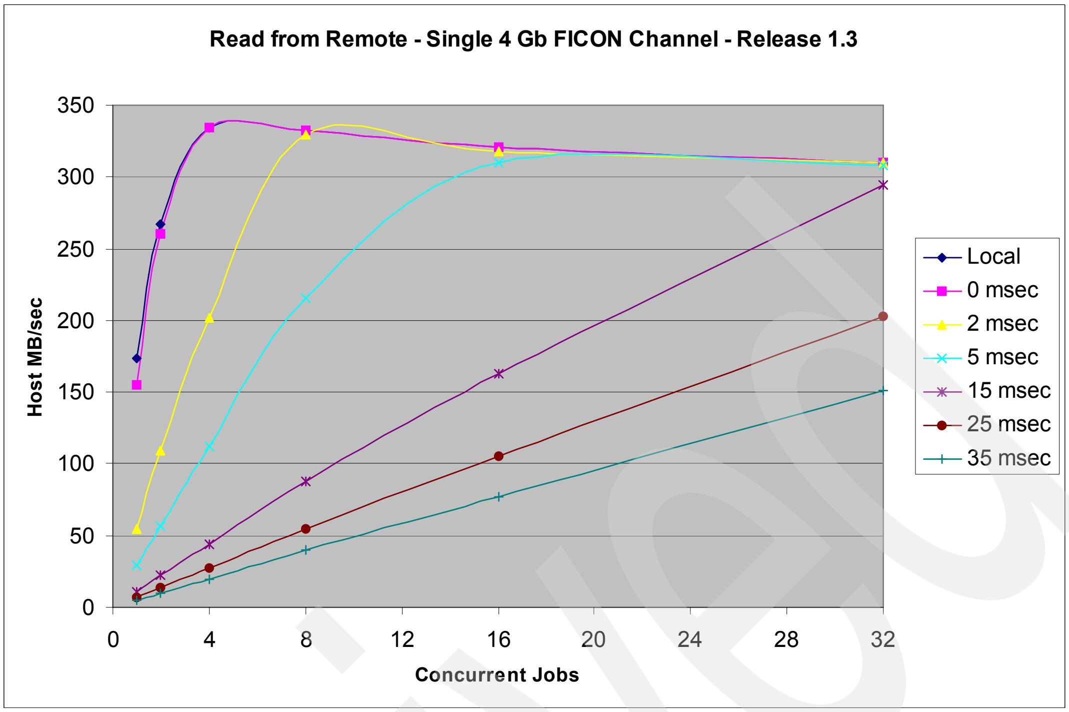

Figure 8-4 Read hit for remote TS7700 in Two-Cluster Grid on a single 4 Gb FICON channel

If the distance between the local and remote TS7700 is 500 miles, then choose the 5 ms curve. Then choose the compressed data rate for the number of concurrent jobs. In this example, done on a 2 Gbps FICON connection, this would give a maximum read data rate of approximately 150 MB/s for 32 concurrent jobs and about 28 MB/s for a single job. Note that communication distance might be greater than physical distance, so allow for this when choosing a delay value.

8.1.4 Performance for different cache sizes

With Release 1.4 you can physically get one, two, or four cache drawers,

The maximum data rate achieved with two cache drawers is about 2/3 of that with four cache drawers. The data rate with a single cache drawer will be less. The total number of FC5267 will affect the percentage of read hits. FC5267 is described in 3.1.1, "TS7700 configuration requirements" on page 78.

8.2 Shared resources in the TS7700 subsystem

In the process of writing scratch volumes or pre-migrating and recalling virtual volumes on physical stacked volumes, components are shared by tasks running on the TS7700 server. Some of these tasks represent users' work (such as scratch mounts) and other tasks are associated with the internal operations of the TS7700 (such as reclamation). All these tasks must share the same resources, especially the TS7700 Server processor, the Tape Volume Cache (TVC) and physical tape drives. Contention might occur for these resources when heavy demands are placed on the TS7700 subsystem. To manage the use of shared resources, the TS7700 uses various resource management algorithms, which can have a significant impact on the level of performance achieved for a specific workload.

In this section we discuss the effects on performance of the following shared resources:

- TS7740 Model V06 Server processor cycles
- Tape Volume Cache (TVC) management
- Physical drives
- Physical stacked volumes

8.2.1 Processor cycles

All tasks running in the TS7700 server require a share of processor cycles. These tasks include the emulation of virtual drives, copy tasks, and recall tasks. For example, if there were eight active virtual drives, two copy tasks, and one recall task active at one point, then there would be a total of eleven concurrent tasks. The processor cycles would be shared by these tasks through a time-slicing multiprocessing algorithm.

The TS7700 monitors the utilization of the processor and uses processor utilization or idle time in many of its internal task management algorithms. High processor utilization can adversely affect channel throughput and logical volume mount times.

The TS7700 does not report processor utilization as part of its performance statistics; therefore, you cannot monitor this metric.

8.2.2 Tape Volume Cache management

All virtual volumes are written and read by the host into and out of the Tape Volume Cache (TVC). The process of copying virtual volumes to stacked volumes and the recall of the virtual volumes from stacked volumes is transparent to the host.

The primary objectives of TVC management are to ensure that sufficient free space is available for new or old virtual volumes, and to maximize the number of read cache hits. The primary mechanisms which are used by the TS7700 to manage the cache are the internal allocation of physical drives, logical volume fragmenting, and throttling.

Note: The software management algorithms used by the TS7700 for TVC management, including throttling values, are internal processes that cannot be controlled by the user.

Also see 2.3.2, "Tape Volume Cache Management" on page 36.

Figure 8-5 shows the Tape Volume Cache (TVC) processes.

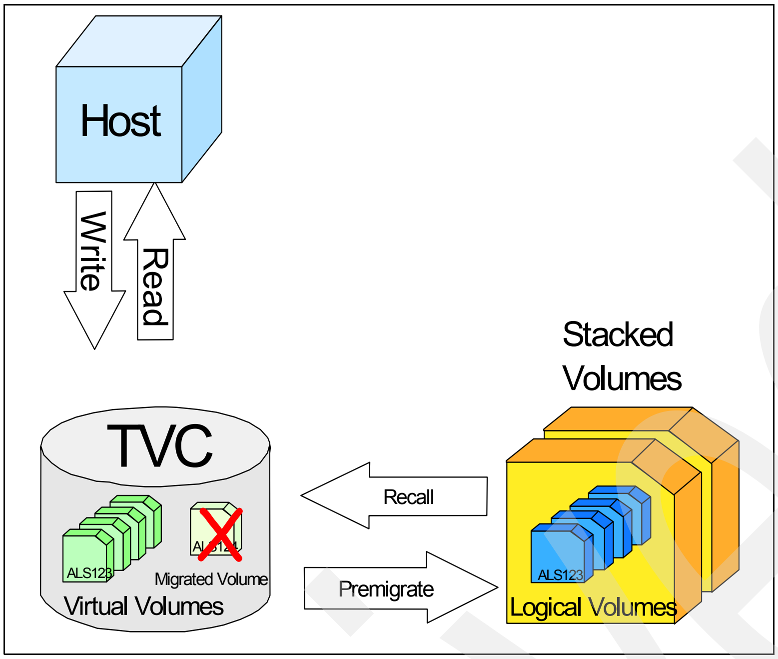

Figure 8-5 Tape Volume Cache processes

Physical drive allocation

Write operations for new volumes, appending to old volumes and read operations requiring the recall of a volume into TVC, all require free space in the TVC for storage of data. When free space becomes too low, or the space occupied by volumes which are closed and ready to be copied to stacked volumes (copy queue) becomes too high, the TS7700 will, if possible, increase the number of copy tasks allowed. In addition, if free space becomes too low, the TS7700 reduces the number of recall tasks allowed, therefore making more physical drives available for copying. The TS7700 always reserves at least one drive for the copy function and one drive for recalls.

Logical volume fragmenting

Copied volumes are eligible for fragmenting, which is the process of reducing the data portion of the virtual volume in cache down to 0 KB in size. The information recorded on the beginning of the volume (such as the volume header) is now being provided from the internal DB2 database. Fragmenting generally works on a first in, first out basis. The process of fragmenting reduces used space within the TVC and therefore “creates” the space required for new volumes to be written and old volumes to be recalled. The TS7700 will fragment logical volumes in the cache only when cache free space becomes low.

Before a logical volume can be fragmented, it must have been copied to a stacked volume. The fragmenting process is also called *Migration* of a logical volume. After a logical volume is migrated, it disappears from the TVC.

Host write throttling

In order to make sure to give other tasks such as pre-migrate and grid replication to reach an equilibrium with the host activity, the TS7700 might introduce a delay in response to host write operations. This is called throttling. High levels of throttling can cause a reduction in host write activity, which might result in elongated or erratic job run times.

The TS7700 calculates throttling values for each of the following at regular intervals:

- Amount of TVC freespace
- Amount of TVC non-pre-migrated data
- Number of TVC logical volumes
- Size of the Two-Cluster Grid copy queue

The actual amount of throttling that the TS7700 applies at any given time depends on the above factors and might vary depending on the LIC level due to changes in the throttling algorithm. There are three types of write throttling, as described below.

Freespace throttling

When the low freespace threshold is crossed, the amount of freespace throttling is inversely proportional to the amount of TVC freespace remaining. As long as the amount of freespace is above the threshold, freespace throttling will not occur.

Figure 8-6 shows TS7700 actions when cache freespace is low:

- Delay the response to EACH host write I/O operation until free space is above trigger.
- Increase the number of drives available for pre-migration processes.
- Reduce number of drives available for recall processes.

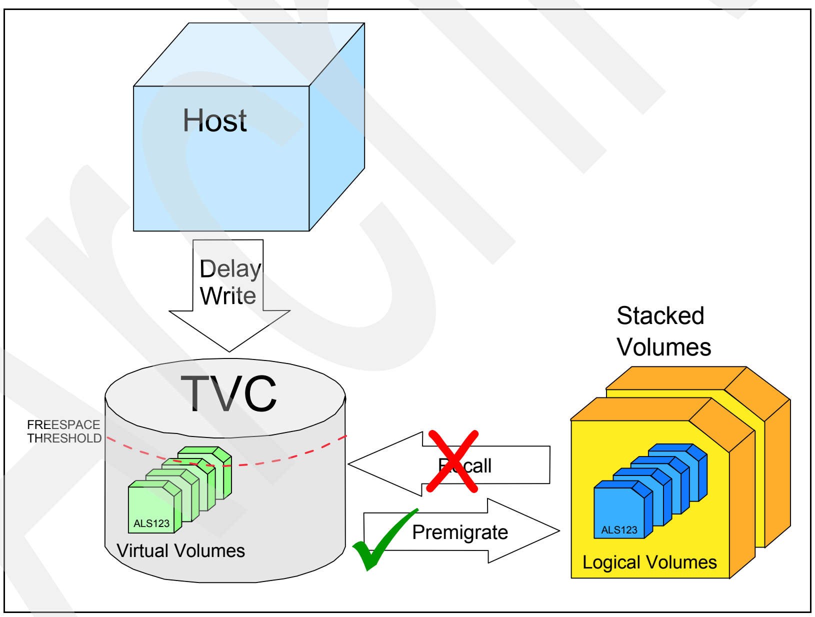

Figure 8-6 Cache freespace throttling

Copy queue throttling

This applies only to a TS7700 in a Multi Cluster Grid configuration. The TS7700 introduces host write throttling and slows the copy activity to the other cluster if either the number of copy tasks or the amount of data waiting to be copied exceeds certain thresholds.

Figure 8-7 shows actions when the copy queue value is high. A delay response to EACH host write I/O operation will take place until the copy queue is below trigger.

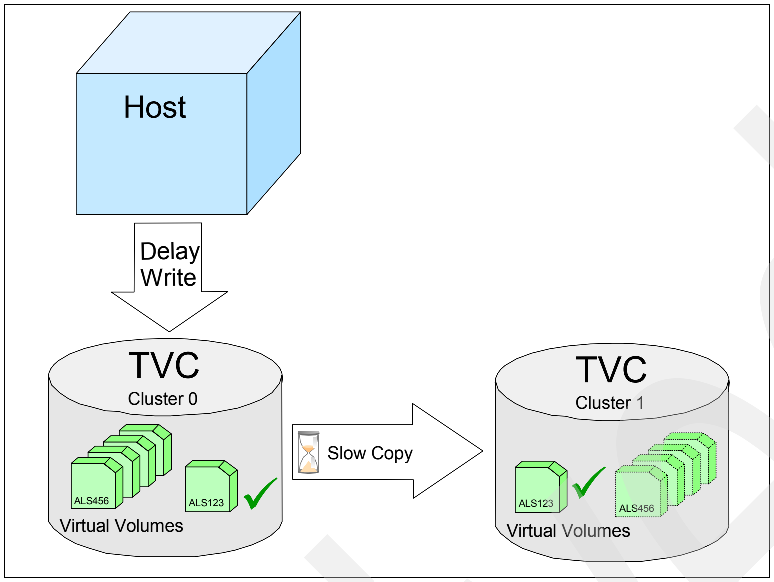

Figure 8-7 Copy queue high actions, copy throttling

Pre-migrate throttling

This type of write throttling will occur when the pre-migrate threshold is reached. This threshold varies between 40% to 50% of the total TVC space. Hitting this threshold means that there are a lot of volumes in the TVC that have not been destaged to physical tape, and the TS770 needs to slow down the host write rate until a balance is achieved between what comes in from the host and what gets written out to physical tape (Figure 8-8).

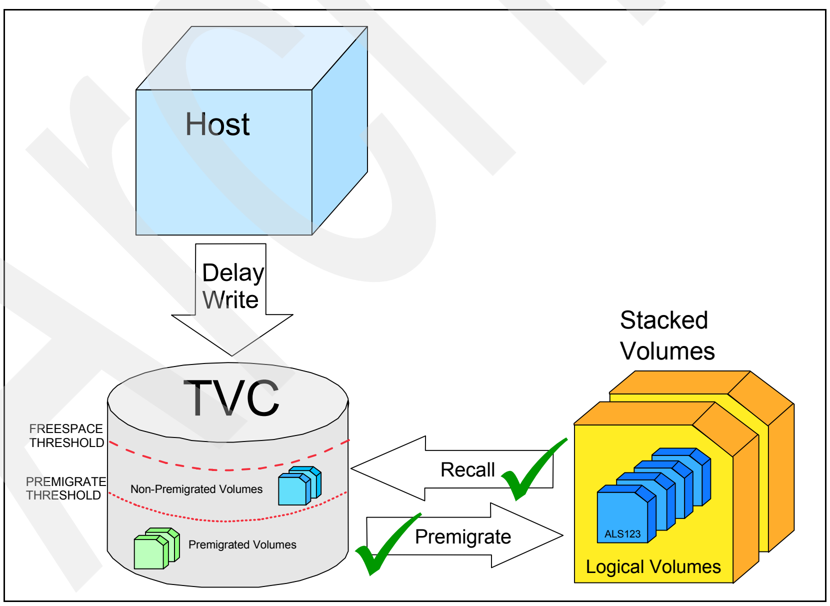

Figure 8-8 Pre-migration throttling

8.2.3 TS7700 physical drive aspects

The physical drives are managed by the TS7700 internal management software and cannot be accessed from any other attached host. These drives are used exclusively by the TS7700 for the mounts required for copying virtual volumes to stacked volumes, recalling virtual volumes into the cache and reclaiming stacked volume space.

The availability of TS7700 physical drives for certain functions can significantly affect TS7700 performance. The TS7700 manages the internal allocation of these drives as required for various functions, but it always reserves at least one physical drive for recall and one drive for pre-migration.

Tape Volume Cache management algorithms will also influence the allocation of physical drives. For example:

- **Cache freespace low**: The TS7700 will increase the number of drives available to the pre-migration function and reduce the number of drives available for recalls.
- **Pre-migration threshold crossed**: The TS7700 will reduce the number of drives available for recall down to a minimum of one drive in order to make drives available for the pre-migration function.

The number of drives available for recall or copy are also reduced during reclamation.

If the number of drives available for pre-migration is restricted, this can lead to limiting the number of virtual volumes in the cache that are eligible to be migrated. This can lead to freespace or copy queue throttling being applied.

If the number of drives for recall is restricted, this can lead to elongated virtual mount times for logical volumes being recalled.

Recall performance is highly dependent on both the placement of the recalled logical volumes on stacked volumes and the order in which the logical volumes are recalled. To minimize the effects of volume pooling on sustained write performance, volumes are pre-migrated using a different distribution algorithm.

This algorithm chains several volumes together on the same stacked volume for the same pool. This can change recall performance, sometimes making it better, possibly sometimes making it worse. Other than differences in performance due to differences in distribution over the stacked volumes, recall performance should be unchanged.

Reclaim policies must be set in the LM and for each volume pool. Reclaim occupies drives and can affect performance. The Inhibit Reclaim Schedule is also set from LM and can prevent reclaim from running at specified time frames during the week. If Secure Data Erase is used, less physical drives might be available even during times when you use inhibited reclamation. If used, IBM recommends to do that for a limited group of data. Inhibit Reclaim specifications only partially apply to Secure Data Erase. Figure 8-8 on page 417 shows that, although you do not allow reclamation during your peak hours to have all your drives available for recall and pre-migration, Secure Data Erase will not honor your settings and thus will run erasure operations as long as there are physical volumes to be erased. Find more about reclaim policies in 4.3.7, “Define TS7700 management policies” on page 157.

The use of Copy Export and Selective Dual Copy will also increase the use of physical drives. Both are used to create two copies of a logical volume in a TS7700.

JCL itself can affect use of physical drives. If you need to read/recover a large amount of non-cache resident data written as single file multivolume, that data requires physical activity.

```
//RESTORE  EXEC PGM=ADRDSSU
//SYSPRINT DD  SYSOUT=*
//TAPE     DD  DISP=SHR,DSN=A.B.C,DISP=OLD,
//         UNIT=(,2)
//SYSIN    DD  *
 RESTORE INDD(TAPE) .......................
```

Figure 8-9 JCL example with use of UNIT(,2)

By using UNIT=(,2), as in Figure 8-9, on a cataloged tape data set, the job will allocate two logical drives and request the TS7700 to allocate both. If the needed physical drives are available for cache recall, this JCL will speed up the time to get logical volumes to cache and turnaround time will decrease. In the job output there will be switching from drive to drive that shows that virtual volume to one drive will be pre-staged while the job is doing I/O to the other virtual volume. The true value of using UNIT=(,2), must be measured on a specific job.

8.2.4 Physical and logical volume performance aspects

There are two instances where the physical stacked volumes can impact TS7700 performance:

- When multiple concurrent recalls occur for logical volumes that reside on the same stacked volume, these recalls will be processed serially, and the second and subsequent recalls will experience elongated mount times. Even though the recalls are serialized, the mount times are likely shorter than if the logical volumes were separate physical volumes, because there is no physical mount of a tape. Also, the TS7700 reorders the recalls so that they are sequential on the physical media to minimize search times.
- When the minimum amount of scratch stacked volumes falls below a certain level, this value is two. At this point, the TS7700 automatically raises the priority of the reclamation process and assigns two physical drives and reclaims until it reaches 2 available scratch stacked volumes. If this happens, the total number of drives available for pre-migration or recall activity is reduced.

The TS7700 performance can also be impacted during return-to-scratch operations of logical volumes.

Return-to-scratch processing involves running a set of tape management tools that identify the logical volumes that no longer contain active data and then communicating with the TS7700 Virtualization Engine to change the status of those volumes from private to scratch. The amount of time the process takes depends on the type of tape management system being employed as well as how busy the TS7700 Virtualization Engine is when it is processing the volume status change requests and whether a grid configuration is being used. You should expect that when the TS7700 Virtualization Engine is not handling a workload peak, up to 5,000 logical volumes per hour can be returned to scratch for a single cluster and up to 2,500 logical volumes per hour can be returned to scratch in a Two-Cluster Grid configuration.

Note: For more information about planning for logical and physical volumes, see 3.5, “Planning for logical and physical volumes” on page 107.

8.3 Monitoring TS7700 performance

To help you monitor the performance of a virtual tape subsystem, various methods were developed over time. Statistical and monitoring information appeared in several forms. For example, hourly records sent to the hosts known as SMF94 records, periodic real-time statistics available through the Library Manager's Web specialist or console, and the Bulk Volume Information Retrieval function (BVIR).

For the next generation of enterprise tape virtualization, the TS7700 Virtualization Engine, statistics reporting design has been revised. Useful statistics from the VTS have been retained and new statistics relevant for the TS7700 have been added. Also, both point-in-time (PIT) and historical statistics are recorded and a new user Interface, the TS7700 Management Interface, has been introduced.

You can use the following interfaces, tools, and methods to monitor the TS7700 Virtualization Engine:

- TS3500 Tape Library Specialist
- Enterprise Tape Library Specialist (or the Library Manager console)
- TS7700 Management Interface
- Bulk Volume Information Retrieval function (BVIR)
- VEHSTATS

The Specialists and Management Interfaces are Web-based. With the BVIR function different types of monitoring and performance-related information can be requested through a host logical volume out of the TS7700 Virtualization Engine. Finally, the VEHSTATS tool can be used to format those BVIR responses that are in a binary format to create human readable statistical reports.

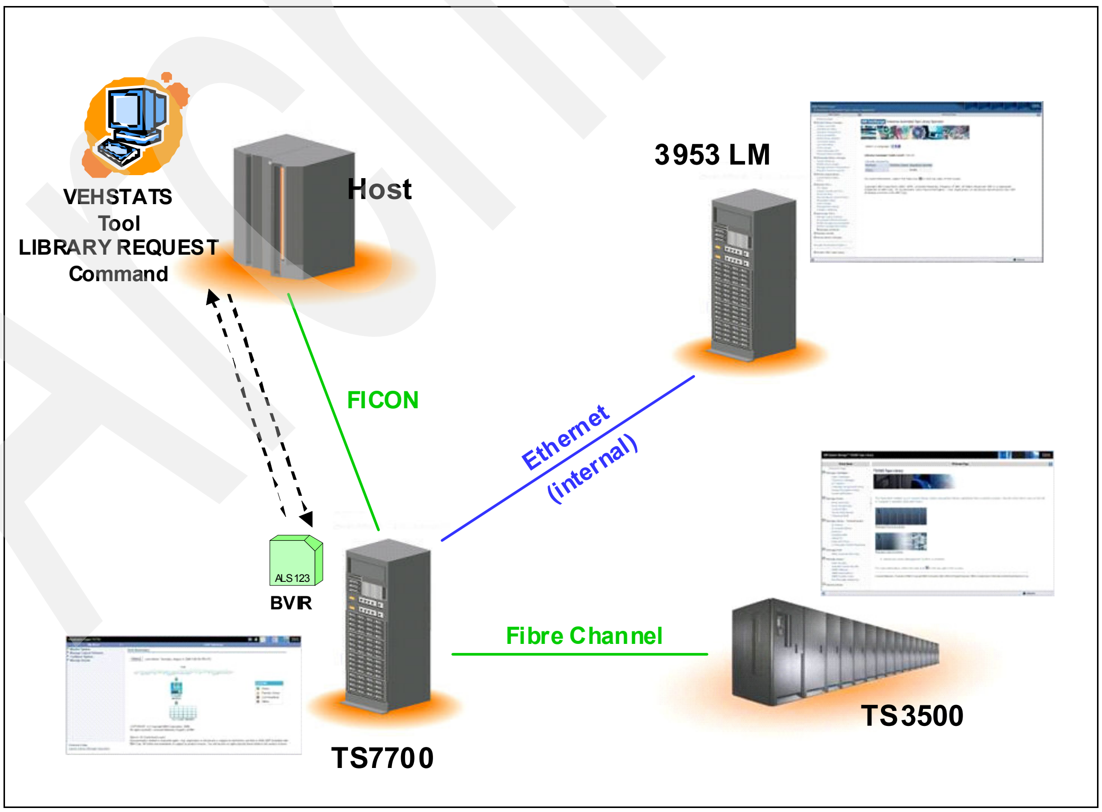

Figure 8-10 Interfaces, tools, and methods to monitor the TS7700 Virtualization Engine

All of the above interfaces, tools, and methods to monitor the TS7700 Virtualization Engine are explained in detail in the following sections.

8.3.1 Using the TS3500 Tape Library Specialist for monitoring

The IBM System Storage Tape Library Specialist (IBM TS3500 Specialist) provides management and monitoring of the TS3500 Tape Library-related items. Initially, the Web user interface to the IBM TS3500 Tape Library only supported a single user at any given time. Now each Ethernet-capable frame on the TS3500 Tape Library allows five simultaneous users of the Web user interface. This allows multiple users to simultaneously access the library Web user interface. Figure 8-11 displays the TS3500 Welcome Page.

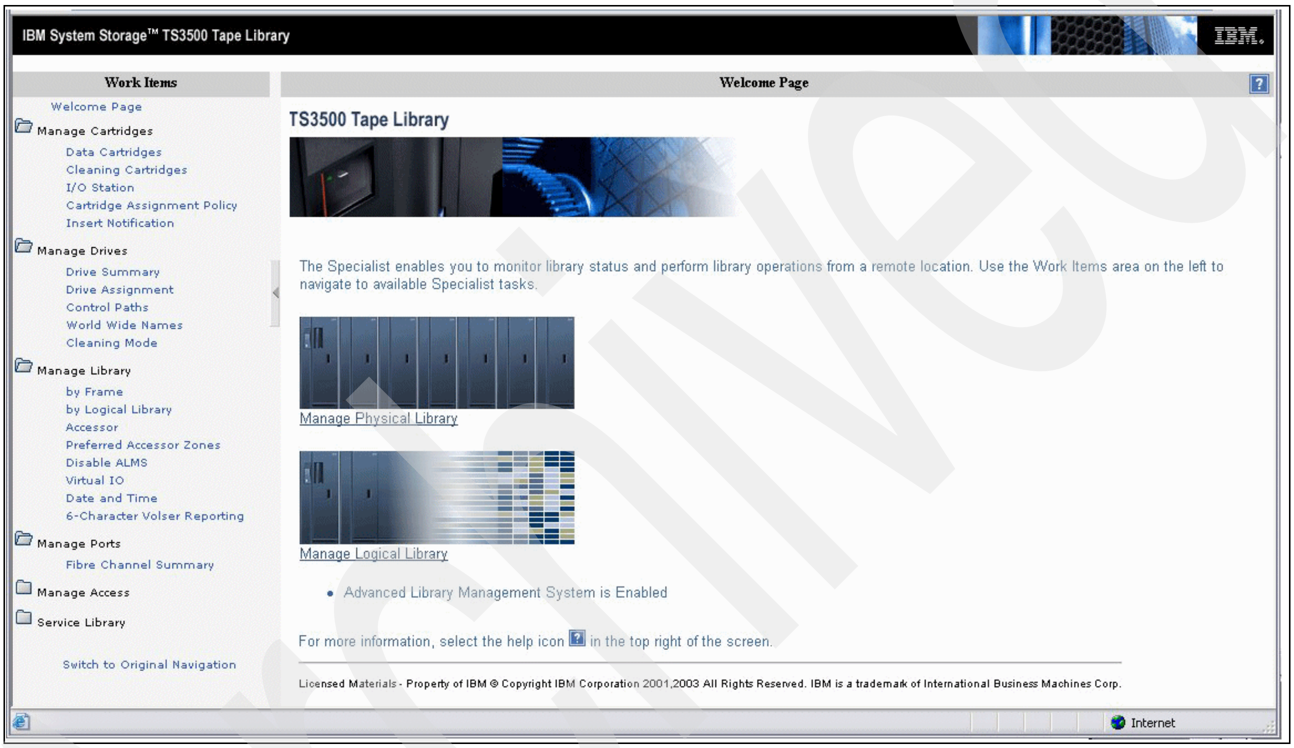

Figure 8-11 TS3500 Specialist Welcome Page

Note that, different from the ETL Specialist and the TS7700 Management Interface, the TS3500 Specialist will time out its session after a default time of ten minutes. You can change these default values from the TS3500 Specialist home page in the panel Manage Access - Web Security, shown in Figure 8-12.

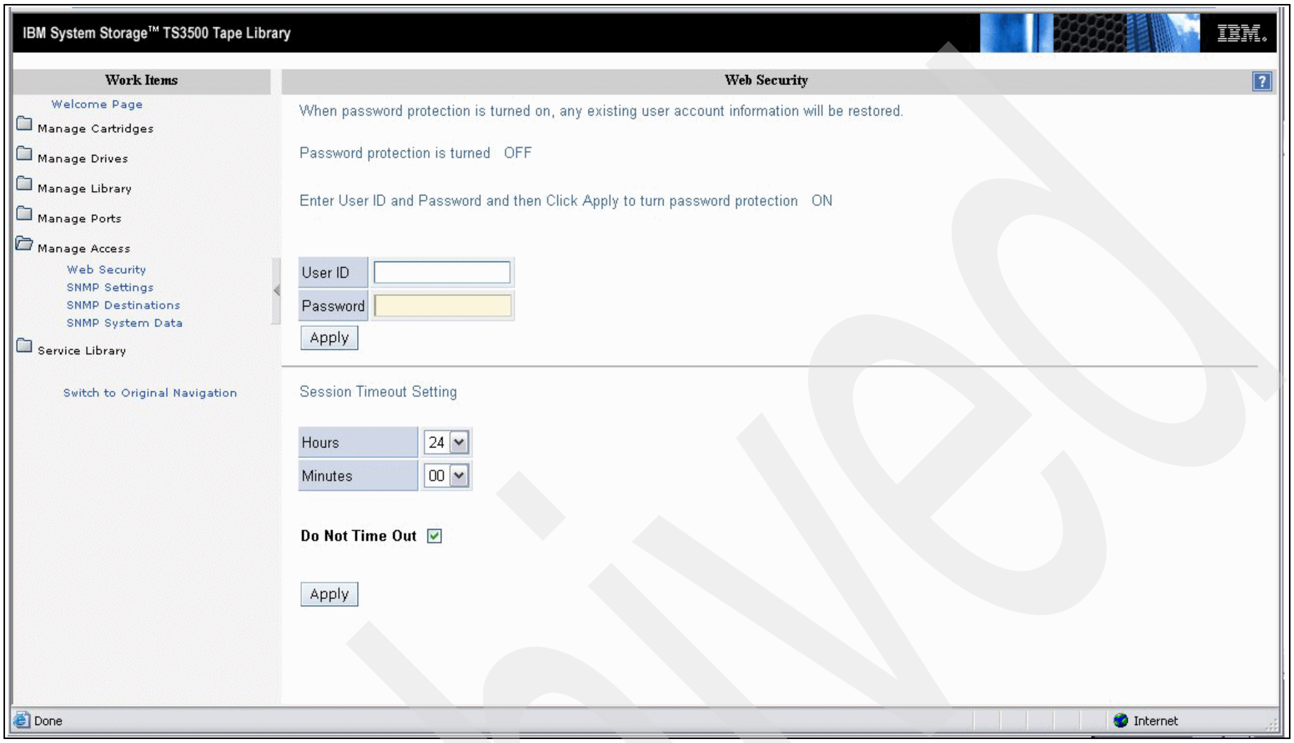

Figure 8-12 TS3500 Specialist Manage Access - Web Security

Some of the information in the TS3500 Specialist is only in a display format. There is no possibility to download it in a CSV format. Other information is only provided in a download format. It can be downloaded directly to your PC and then used as input for a snapshot analysis—or if historical data is stored, for a forecast.

For further information, how to request and use this data, refer to *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

- Accessor Usage (display only)
  - Activity of each Accessor and gripper
  - Travel meters of Accessors
- Drive Status and Activity (display only)
- Drive Statistics (display only)
  - Last Volser on this drive
  - Write and Read MB per drive
  - Write and Read errors corrected per drive
  - Write and Read errors uncorrected per drive
- Mount History for cartridges (download only)
  - Last Tape Alert
  - Number of Mounts of a specific cartridge
  - Number of Write and Read retries of a specific cartridge in the life cycle
  - Number of Write and Read permanent Errors of a specific cartridge in the life cycle

- Fibre Port statistics (download only)
  - Report Errors, Aborts, Reset and Recoveries between the TS7700 and the drives or, for native attached drive, between the controller and the drive

> **Restriction:** This statistic does not provide information from the host to the TS7700 or controller.

- Library statistics, on a hourly basis - download only
  - Total Mounts
  - Total Ejects
  - Total Inserts
  - Average and Maximum amount of time a drive was mounted on a drive (residency)
  - Average and Maximum amount of time was needed to perform a single mount
  - Average and Maximum amount of time was needed to perform an eject

  These statistics can be downloaded; however, they are not included in the Bulk Volume Information Retrieval (BVIR) records, processed by the TS7700 Virtualization Engine.

> **Restriction:** The information are only available on the TS3500 Tape Library.

For more information about the meaning of the statistics, refer to *IBM System Storage TS3500 Tape library data gathering - Introduction* in the Library Statistics White Paper.

8.3.2 Using the 3953 Library Manager for monitoring

In this section we describe a few windows that can be used to display information about performance and capacity of the TS7700 Virtualization Engine subsystem. We do not cover all windows available on the 3953 Library Manager here. See *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473 for more detailed information about Library Manager windows. If the TS7700 is connected to a 3494 Library see *IBM Totalstorage 3494 Tape Library Operator Guide*, GA32-0449.

You can access the windows through the 3953 Library Manager ETL Specialist (Figure 8-13) from the menu selection Monitor VTS x.

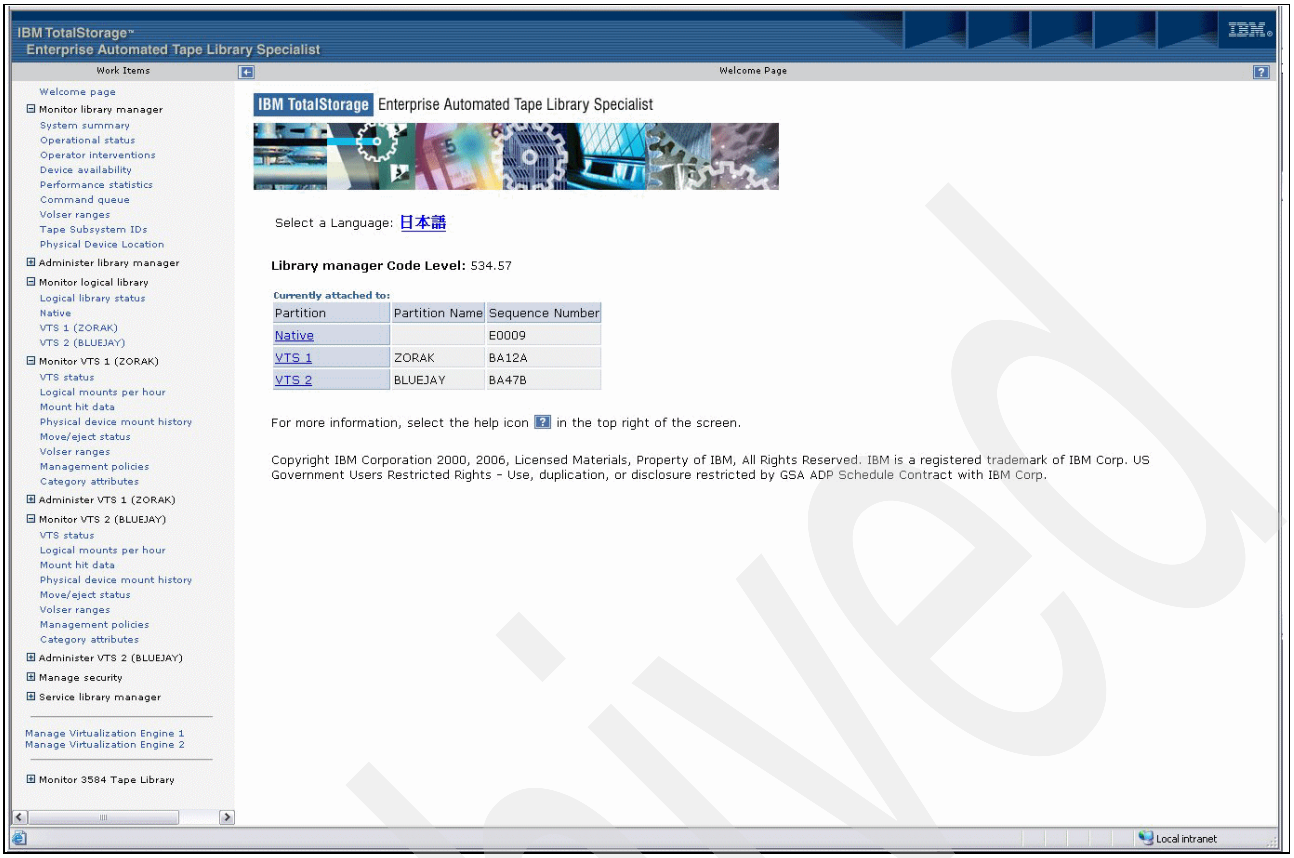

Figure 8-13 ETL Specialist for monitoring Library and TS7700/VTS

TS7700 Status window from the ETL Specialist

The TS7700 Status window shown in Figure 8-14 gives you an overview of the TS7700 and its capabilities. For the TS7700 in this example, you can see that:

- The library sequence number (Distributed LIBRARY-ID).
- It has the Advanced Policy Management feature installed (default for the TS77000).
- IBM 3592 tape drives are installed.
- Larger logical volumes are supported up to 4,000 MB support.
- TS7700 does not support Secure Data Erase for the initial product release.
- TS7700 does not support Import/Export for the initial product release.

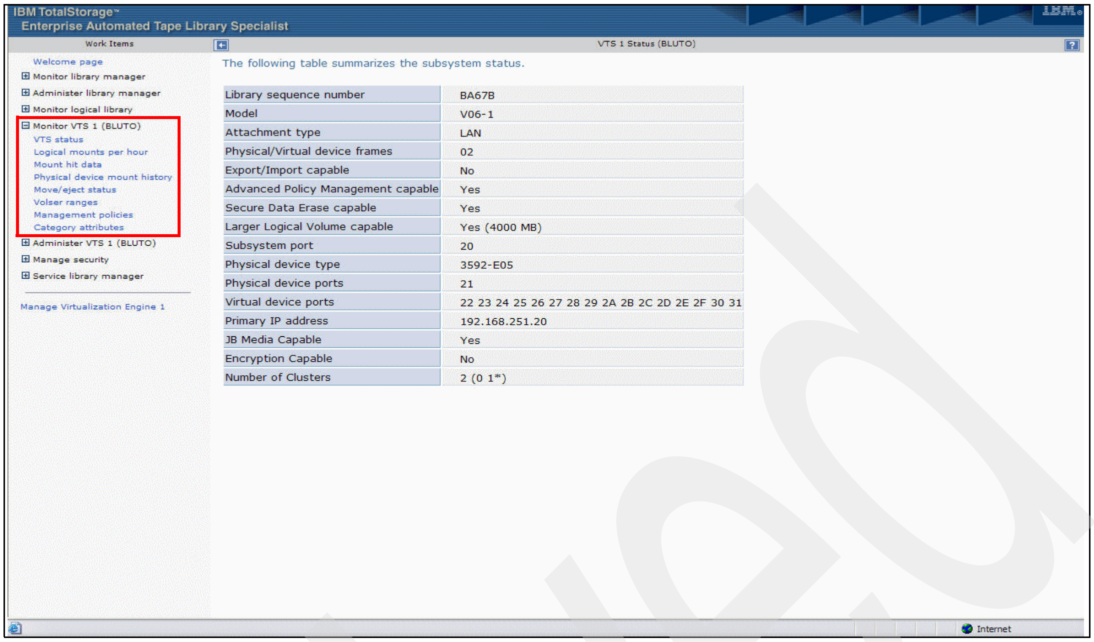

Figure 8-14 TS7700/VTS Status panel through the ETL Specialist

On the left side of the panel shown in Figure 8-14, you can see the work items to monitor the TS7000:

Logical mounts per hour
Displays the number of logical mounts per hour, which includes the sum of fast ready mounts, cache hit mounts, and physical mounts (recalls). Data is displayed for the previous 24 hours.

Mount hit data
Displays the distribution in percentage of three types of logical mounts: fast ready hits, cache hits, and physical mounts required. Data is displayed for the previous 24 hours.

Physical device mount history
Displays the maximum, average, and minimum numbers of physical tape drives used at one time to mount stacked volumes. Data is displayed for the previous 24 hours.

Move/eject status
Displays two tables which show the status of in progress eject and move stacked volume requests.

Volser ranges
Summarizes the current defined Volser ranges for the TS7700 stacked volumes.

Management policies
Displays the current TS7700 management policy settings for inhibit reclaim and the free storage threshold. These values can be changed at the Library Manager console, but not from the Specialist panel.

Category attributes
Displays the Library Manager categories that have the Fast Ready attribute set. The Fast Ready attribute can only be changed from the Library Manager console.

General library performance statistics

The Performance Statistics window displays the statistics involving physical movement of the accessors in the physical tape library. This includes a brief mount summary for the physical drives managed by this 3953 Library Manager. Actions involving TS7700 logical volumes are not part of these statistics. They show the performance through the current hour, over the previous 24 hours, and for the last 7 days. The 7-day and 24-hour statistics are updated hourly.

Figure 8-15 shows this ETL Specialist panel. You can access the display selecting **Monitor Library Manager → Performance statistics**. Note that the graph scales dynamically on the maximum number of mounts per hour.

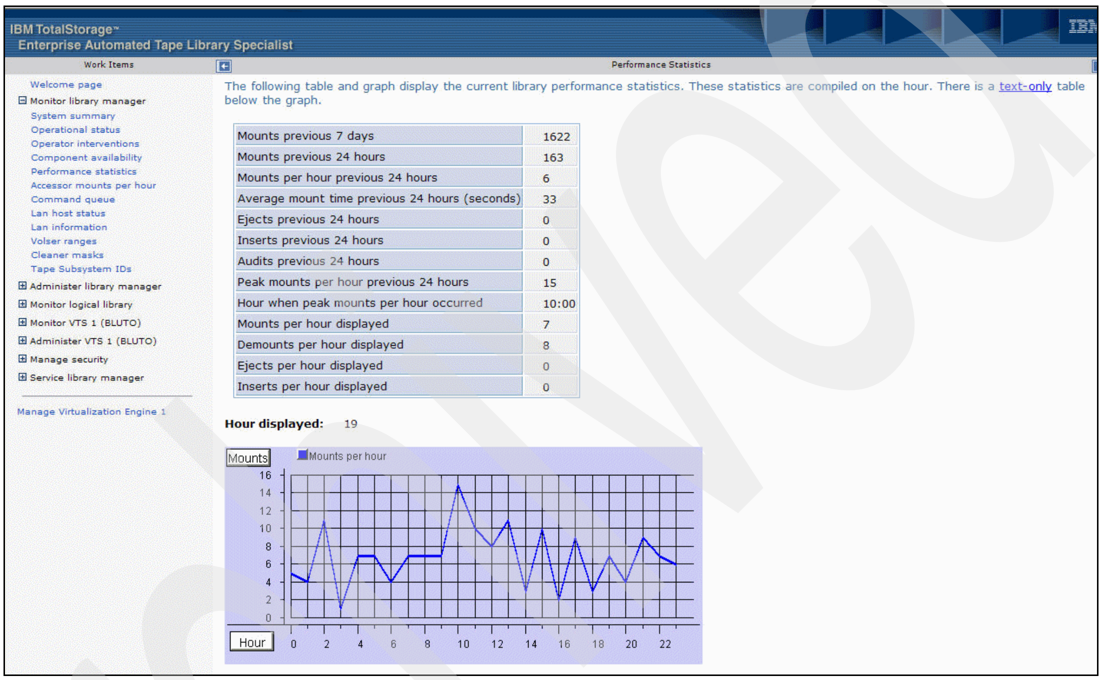

Figure 8-15 Performance Statistics through the ETL Specialist panel

TS7700 Mount hit data

The Mount Hit Data window (Figure 8-16) shows the distribution in percentage of three types of logical mounts:

- Fast-ready hits
- TVC Cache hits
- TVC Cache misses, mounts requiring a recall

Data is displayed for the previous 24 hours. The current hour's data is designated by a diamond-shaped marker.

The graph displays three lines, one for each type of mount, indicating its percentage of the total number of mounts for an hour. To display the data, from the ETL Specialist select **Monitor VTS x → Mount hit data**.

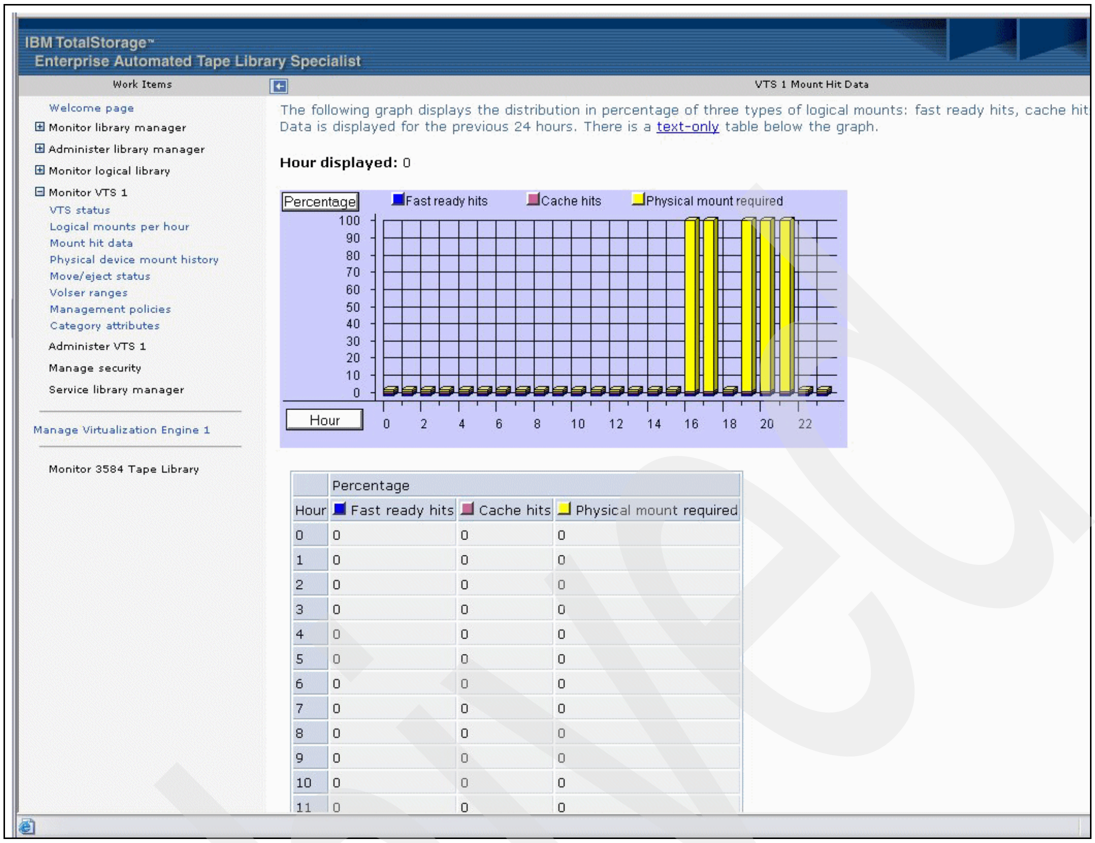

Figure 8-16 TS7700 Mount hit data through the ETL Specialist panel

The data shown in Figure 8-16 means:

- **Fast Ready hits**: The number of times the host requested a mount from a specific volume category, and the category was designated as “fast-ready”. Typically, these are scratch mounts. Because this type of mount does not require any recall of data from tape, it is the fastest. To benefit from Fast Ready hits, assign the Fast Ready attribute to your scratch volume categories. If this value is 0, you should check your Fast Ready categories settings described in 4.3.5, “Define Fast Ready categories” on page 154.
- **Cache hits**: The number of virtual mounts, where the volume to be mounted still resides in the TVC. This type of mount does not require any recall of data from tape.
- **Physical mount required**: The number of times a logical volume had to be recalled from a stacked volume into the TVC to satisfy a virtual mount request. While a stacked volume is mounted, the VTS can recall more than one logical volume from it. Therefore, the graph does not represent the number of times a stacked volume actually had to be mounted in response to a virtual mount request. This is the slowest type of virtual volume mount.

TS7700 Physical device mount history

The Physical Device Mount History window (Figure 8-17 on page 428) displays a graph showing the average, maximum, and minimum number of physical tape drives used at one time to mount stacked volumes. Data is displayed for the previous 24 hours. The current hour's data is designated by a diamond-shaped marker.

To access the data from the ETL Specialist, select **Monitor VTS x** → **Physical device mount history**.

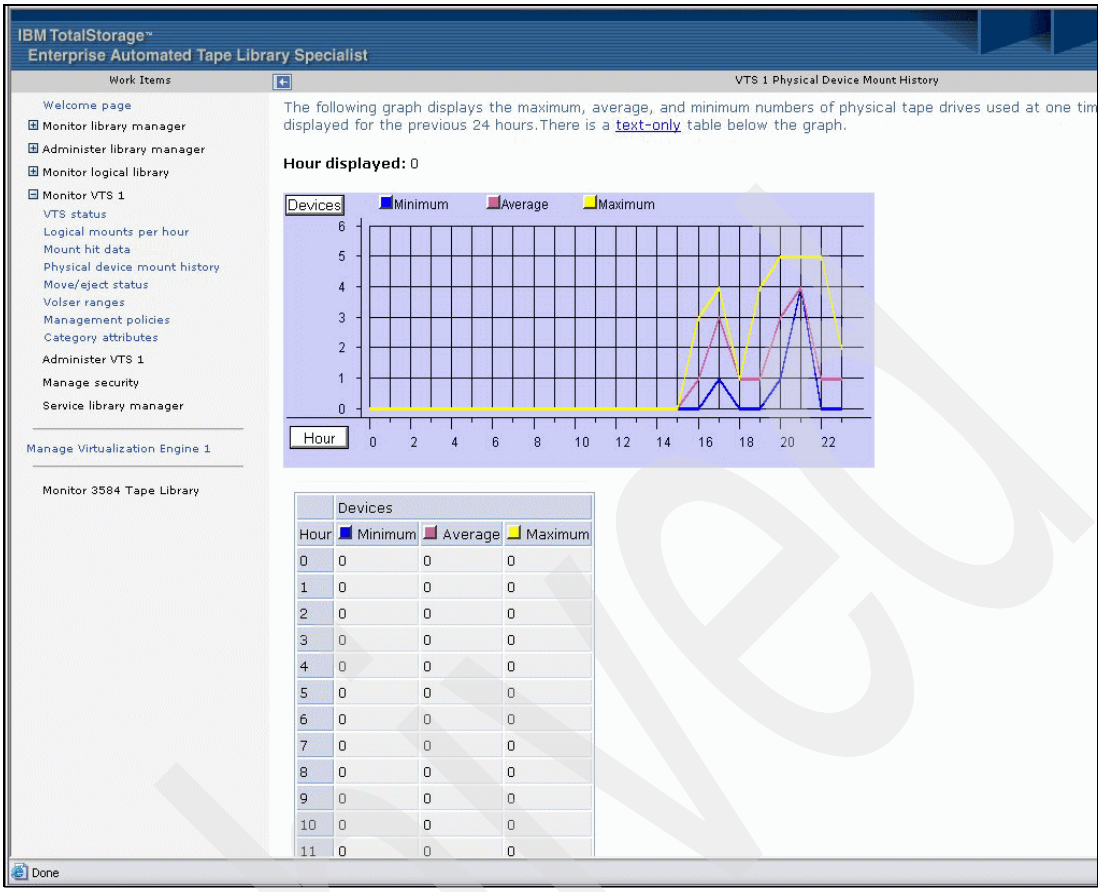

Figure 8-17 TS7700 Physical device mount history through the ETL Specialist panel

The data shown in Figure 8-17 is:

- **Maximum**: The maximum number of physical drives used concurrently to mount stacked volumes during the hour.
- **Average**: The average number of physical drives that were used concurrently to mount stacked volumes during the hour. The sampling rate for calculating the average is 10 seconds.
- **Minimum**: The minimum number of physical drives used concurrently to mount stacked volumes during the hour.

TS7700 Logical mounts per hour

The Logical Mounts Per Hour window (Figure 8-20 on page 432) shows the number of logical mounts per hour. The figure is the sum of Fast Ready mounts, cache hit mounts and cache misses (Recall). Data is displayed for the previous 24 hours. The current hour's data is designated by a diamond-shaped marker.

To display the percentage distribution of logical mounts, use the Mount Hit Data window (Figure 8-16).

To access the information, from the ETL Specialist, select **Monitor VTS x** → **Logical mounts per hour**.

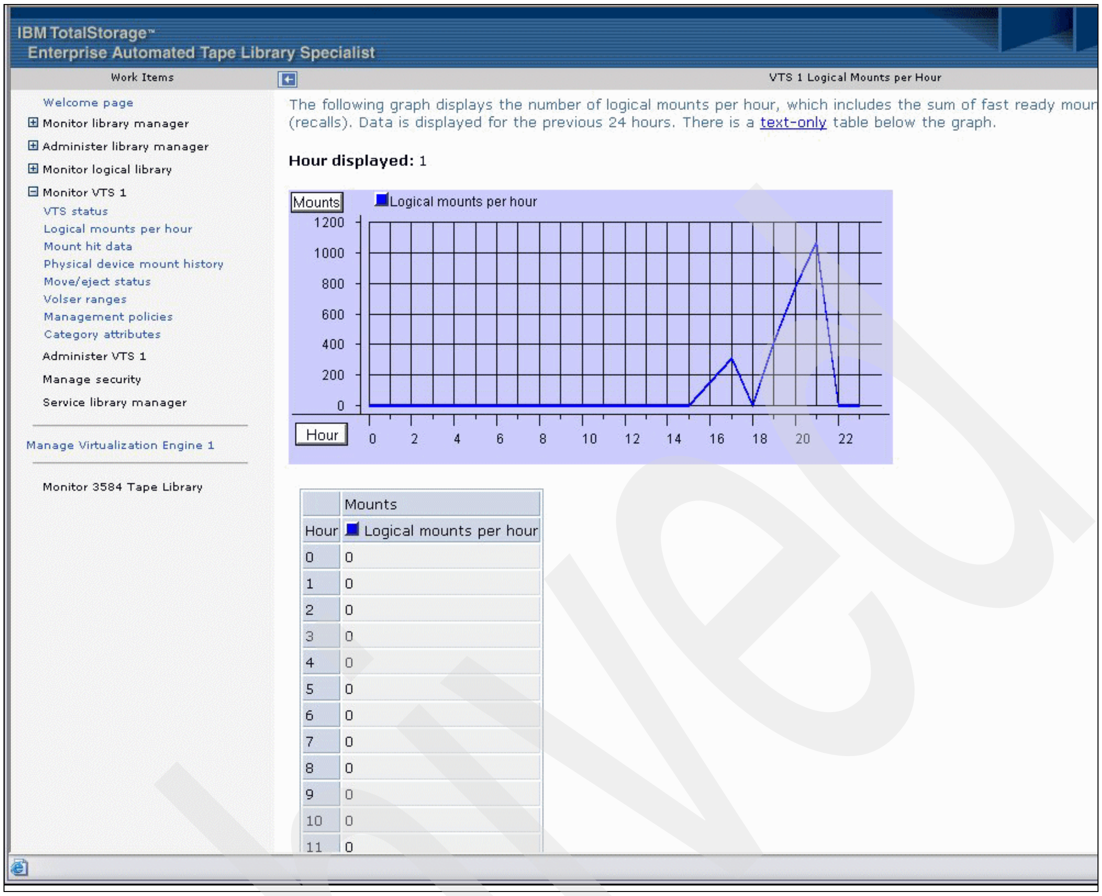

Figure 8-18 TS7700 Logical Mounts Per Hour panel using the ETL Specialist

TS7700 Active Data, Data Distribution, and Data Flow

Although the queries for Active Data, Data Distribution and Data Flow are still available for selection through the 3953 Library Manager console, for the attached TS7700 Virtualization Engine the underlying information has been enhanced and is available now using a different method.

To get the required information, use the Bulk Volume Information Retrieval (BVIR) functions described in 8.6, "Bulk Volume Information Retrieval (BVIR)" on page 442.

If you select one of these functions from the 3953 Library Manager console, you will get the informational message box shown in Figure 8-19.

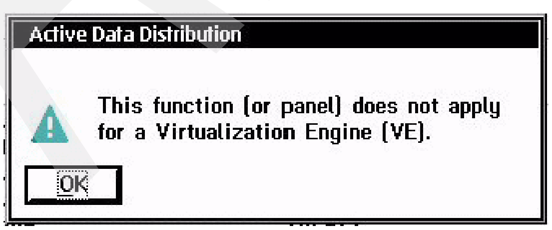

Figure 8-19 TS7700 query for Active Data Distribution through the 3953 Library Manager console

Active Data Distribution

The Active Data Distribution information displays the distribution of the data on the stacked volumes. There is a number for each 10% increment in active data for the number of used stacked tapes. You use that output to identify the number of tapes that will be freed up if you were to change the reclaim threshold.

With prior VTSs, the active data distribution displayed on the LM was a combination of all media types and pools. It was difficult to determine if changing the reclamation thresholds would yield a specific media type needed. With TS7740 this information is not in the LM itself anymore. You must use BVIR function.

In TS7700, the Bulk Volume Information Retrieval function, PHYSICAL MEDIA POOLs, provides the active data distribution by media type and pool so that you can determine exactly what would happen if you changed the reclaim threshold. This output provides counts of physical media by media type, pool and use for a specific distributed library in a human readable formatted response (Example 8-1).

Example 8-1 Bulk Volume Information Retrieval for PHYISCAL MEDIA POOLS

```
VTS BULK VOLUME DATA REQUEST
PHYSICAL MEDIA POOLS
05/20/2006 08:45:29 VERSION 02
S/N: 1B245  LIB ID: BAR09

R POOL MEDIA  EMPTY  FILLING   FULL  ERASE    ROR  UNAVAIL
C   00    JA  00083
C   00    JB  00051
C   00    JJ  00002
C   01    JA  00134    00003  00609  00000  00000    00000
C   01    JB  00001    00002  00001  00006  00000    00000
R POOL MEDIA    0+   10+   20+   30+   40+   50+   60+   70+   80+   90+
D   01    JA 00000 00002 00034 00056 00092 00078 00084 00065 00195 00183
D   01    JB 00000 00000 00000 00000 00000 00000 00000 00002 00003 00001
```

At the upper area of the BVIR response for the PHYSICAL MEDIA POOLS query you can identify the number of empty and active volumes per pool, that is, the number of stacked volumes that have active data on them for each pool. This includes stacked volumes that are in the process of being filled with new data. Pool 0 is the Common Scratch Pool (CSP) and therefore it can only contain EMPTY stacked volumes.

The lower section of the output shows the active data distribution per pool. The output should not show any volumes below your Reclaim Threshold Percentage value. If it does, perhaps you need to decrease, as possible, the Inhibit Reclaim Schedule time frame to allow the TS7740 to “reclaim” more volumes.

8.4 Using the TS7700 Management Interface for monitoring

The TS7700 Management Interface belongs to the family of tools used for reporting and monitoring IBM storage products. These tools do not provide reports, but can be used for online queries about the status of the TS7700 Virtualization Engine, its components, and the distributed libraries. They also provide information about the copies that have not been completed yet and on the amount of data to be copied.

The TS7700 Management Interface is based on a Web server that is installed in each TS7700 Virtualization Engine. You can access this interface with any standard Web browser.

This TS7700 Management Interlace is also required for implementation and operation purposes. You will insert logical volumes, invoke Service Mode, control cluster failover, and more.

In a TS7700 Virtualization configuration, you have three possible Web interfaces available: the ETL Specialist, the TS3500 Tape Library Specialist, and the TS7700 Management Interface. You can access the TS7700 Management pages from the ETL Specialist and vice versa, as there is a link between the products that enables you to switch between them seamlessly and hence easily find the required information. From the ETL Specialist you can navigate further to the attached physical TS3500 Tape Library Specialist.

In this section we introduce the TS7700 Management Interface to be used for performance and statistics. The My Work items of the TS7700 are covered in 7.2, "Virtualization Engine TS7740 Management Interface" on page 315.

8.4.1 Performance and statistics

This section presents information that is related to viewing performance information and statistics for the IBM TS7700 Virtualization Engine for Single and Multi Cluster Grid configurations. The graphical views are display snapshots of the processing activities from the last 15 minutes. You can access the following selections by navigating to the Performance & Statistics section in the TS7700 Management Interface. The examples are taken from different configurations—some are from a Two-Cluster Grid and some from a Three-Cluster Grid.

Logical mounts

Use this page to view logical mount statistics for the TS7700 Virtualization Engine. The logical mount statistics for each cluster are displayed in two bar graphs and tables: one for the number of mounts and one for average mount time. The example provided in Figure 8-20 is taken from a TS7700 cluster that is part of a Multi Cluster Grid configuration (Three-Cluster Grid).

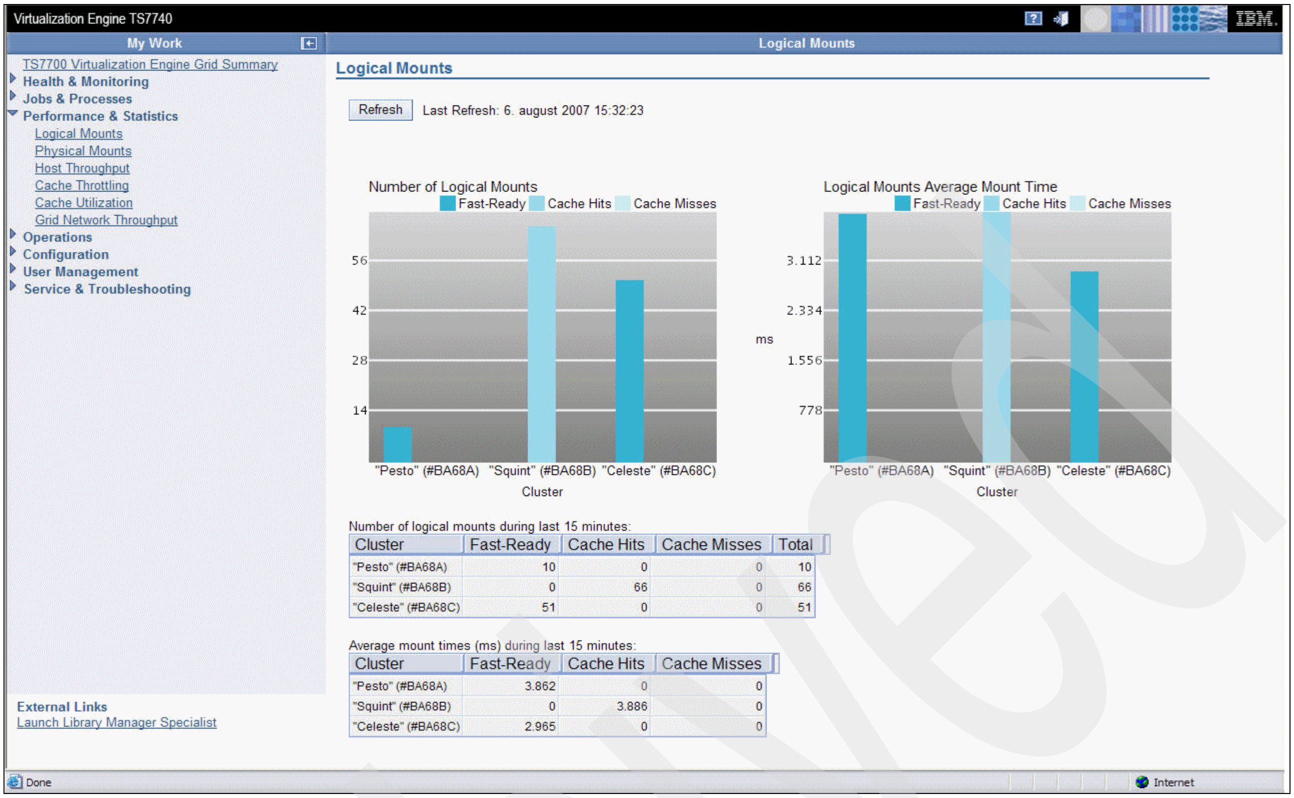

Figure 8-20 TS7700 Management Interface Logical Mounts

Under “Number of logical mounts during last 15 minutes” you can find:

Cluster The cluster name
Fast-Ready Amount of logical mounts completed using the Fast-Ready method
Cache Hits Amount of logical mounts completed from cache
Cache Misses Amount of mount requests that are unable to be fulfilled from cache
Total Total amount of logical mounts

Under “Average mount times (ms) during last 15 minutes” you see:

Cluster The cluster name.
Fast-Ready Average mount time for logical mounts completed using the Fast-Ready method.
Cache Hits Average mount time for logical mounts completed from cache.
Cache Misses Average mount time for mount requests that are unable to be fulfilled from cache.

Physical mounts

Use this page to view physical mount statistics for the TS7700 Virtualization Engine. The physical mount statistics for each cluster are displayed in two bar graphs and tables: one for the number of mounts by category and one for average mount time per cluster. The example provided in Figure 8-21 is taken from a TS7700 cluster that is part of a Multi Cluster Grid configuration (Two-Cluster Grid).

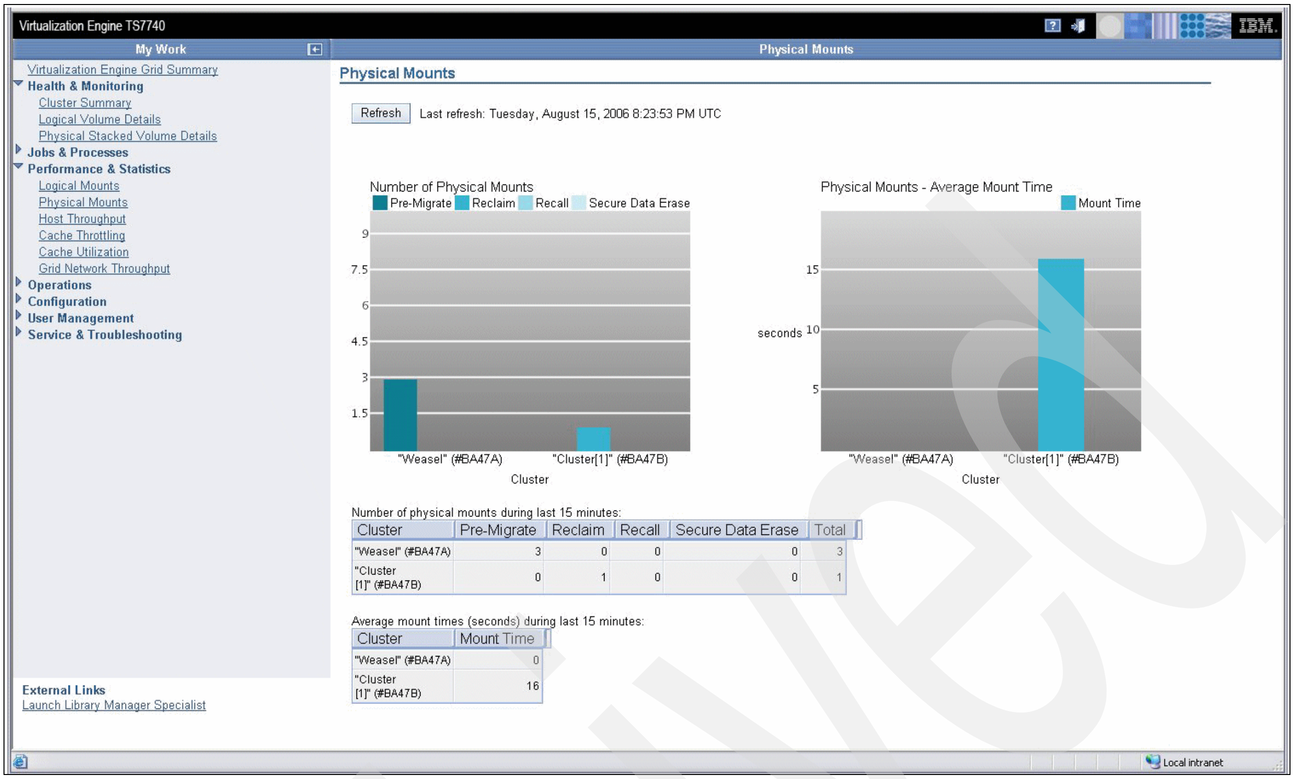

Figure 8-21   TS7700 Management Interface Physical Mounts

The table cells show:

Cluster The cluster name
Pre-Migrate Number of pre-migrate mounts
Reclaim Number of reclaim mounts
Recall Number of recall mounts
Secure Data Erase Number of Secure Data Erase mounts
Total Total physical mounts
Mount Time Average mount time for physical mounts

Host throughput

You can use this page to view host throughput statistics for the IBM TS7700 Virtualization Engine. Note that the information is provided in 15 second intervals, not in 15 minute intervals like the other performance data.

This page allows you to view statistics for each cluster, vNode, host adapter, and host adapter port in the grid. At the top of the page is a collapsible tree that lets you view statistics for a specific level of the grid and cluster. Clicking the grid hyperlink will display information for each cluster. Clicking the cluster hyperlink will display information for each vNode. Clicking the vNode hyperlink will display information for each host adapter. Clicking a host adapter link will display information for each of its ports.

The example provided in Figure 8-22 is taken from a TS7700 cluster that is part of a Multi Cluster Grid configuration (Three-Cluster Grid).

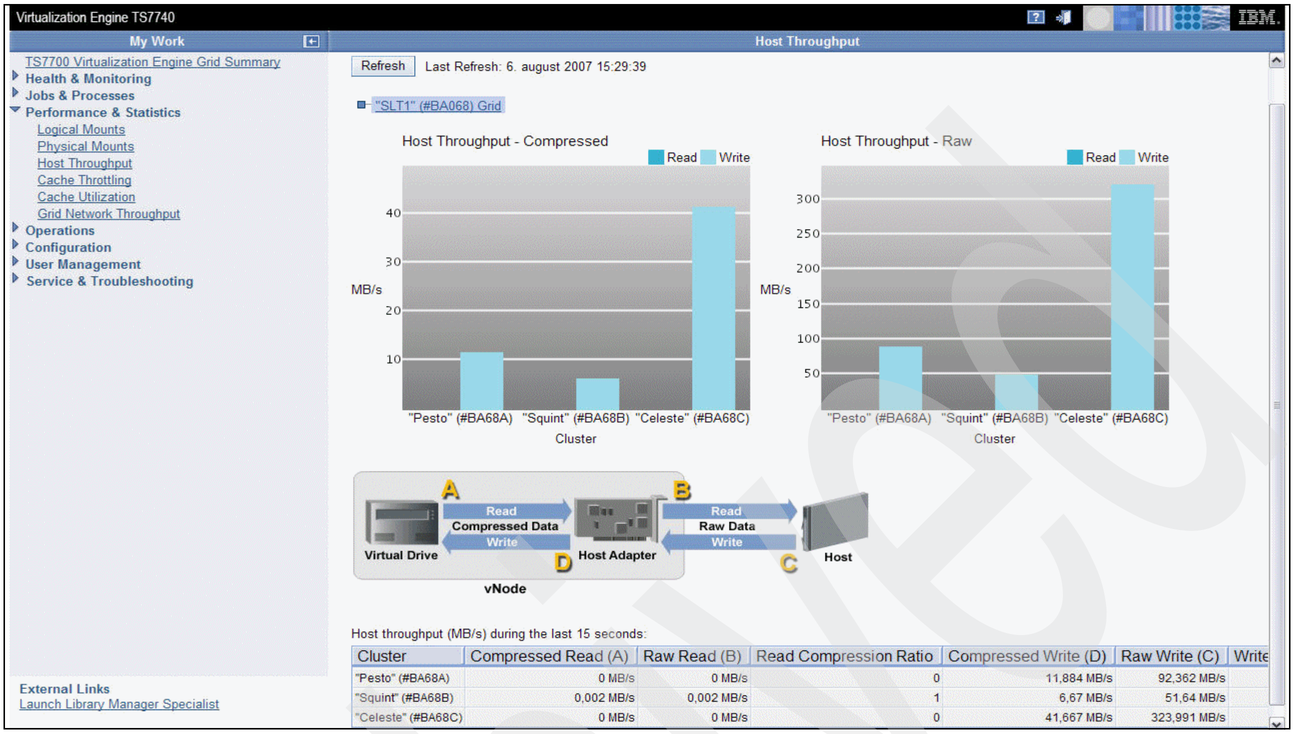

Figure 8-22 TS7700 Management Interface Host Throughput

The host throughput data is displayed in two bar graphs and one table. The bar graphs are for raw data coming from the host to the host bus adapter (HBA) and for compressed data going from the HBA to the virtual drive on the vNode.

The letter next to the table heading corresponds with the letter step in the diagram shown above the table. Data is available for a cluster, vNode, host adapter or host adapter port. The table cells include:

- Cluster/vNode/Host Adapter/Host Adapter Port: The cluster or cluster component data is being displayed for.
- Compressed Read (A): Amount of data read between the virtual drive and HBA.
- Raw Read (B): Amount of data read between the HBA and host.
- Read Compression Ratio: Ratio of compressed data read to raw data read.
- Compressed Write (D): Amount of data written from the HBA to the virtual drive.
- Raw Write (C): Amount of data written from the host to the HBA.
- Write Compression Ratio: Ratio of compressed data written to raw data written.

Cache throttling

You can use this page to view cache throttling statistics for the IBM TS7700 Virtualization Engine. The example provided in Figure 8-23 is taken from a TS7700 cluster that is part of a Multi Cluster Grid configuration (Three-Cluster Grid).

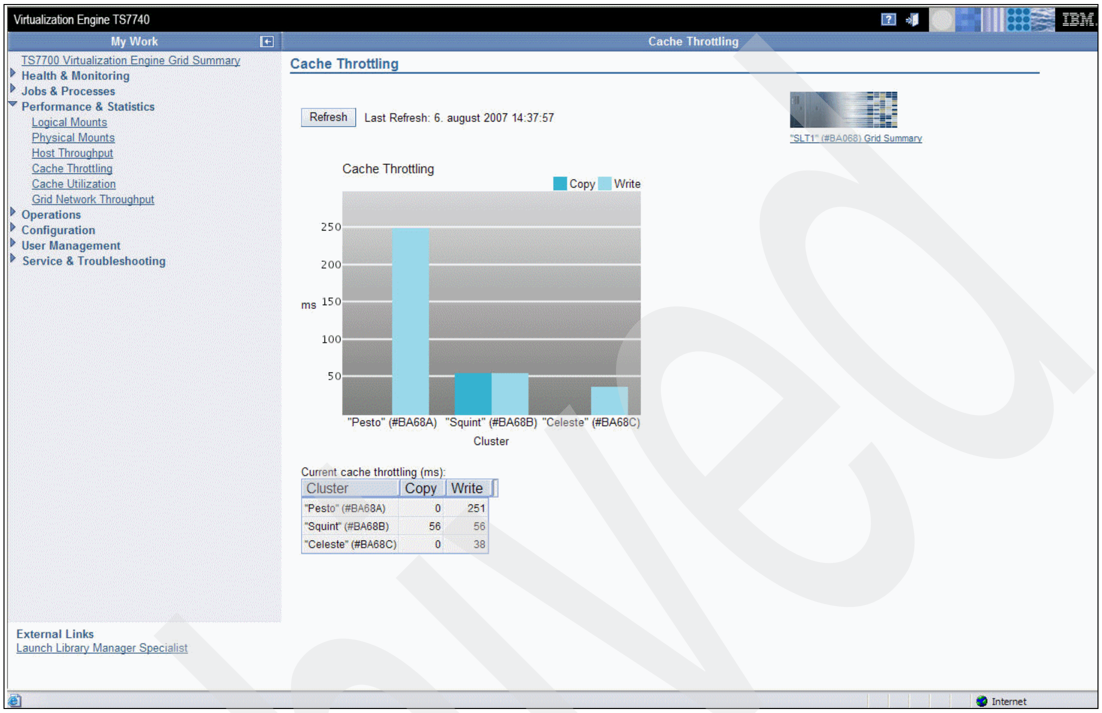

Figure 8-23 TS7700 Management Interface Cache Throttling

Cache throttling is a time interval between internal TS7700 job execution that allows for higher access performance for the host. The cache throttling statistics for each cluster in regards to copy and write are displayed both in a bar graph form and in a table. The table shows:

Cluster The cluster name

Copy The amount of time inserted between internal copy operations

Write The amount of time inserted between internal write operations

Cache utilization

You can use this page to view cache utilization statistics for the IBM TS7700 Virtualization Engine. The example provided in TS7700 Management Interface Cache Utilization is taken from a TS7700 cluster that is part of a Multi Cluster Grid configuration (Two-Cluster Grid).

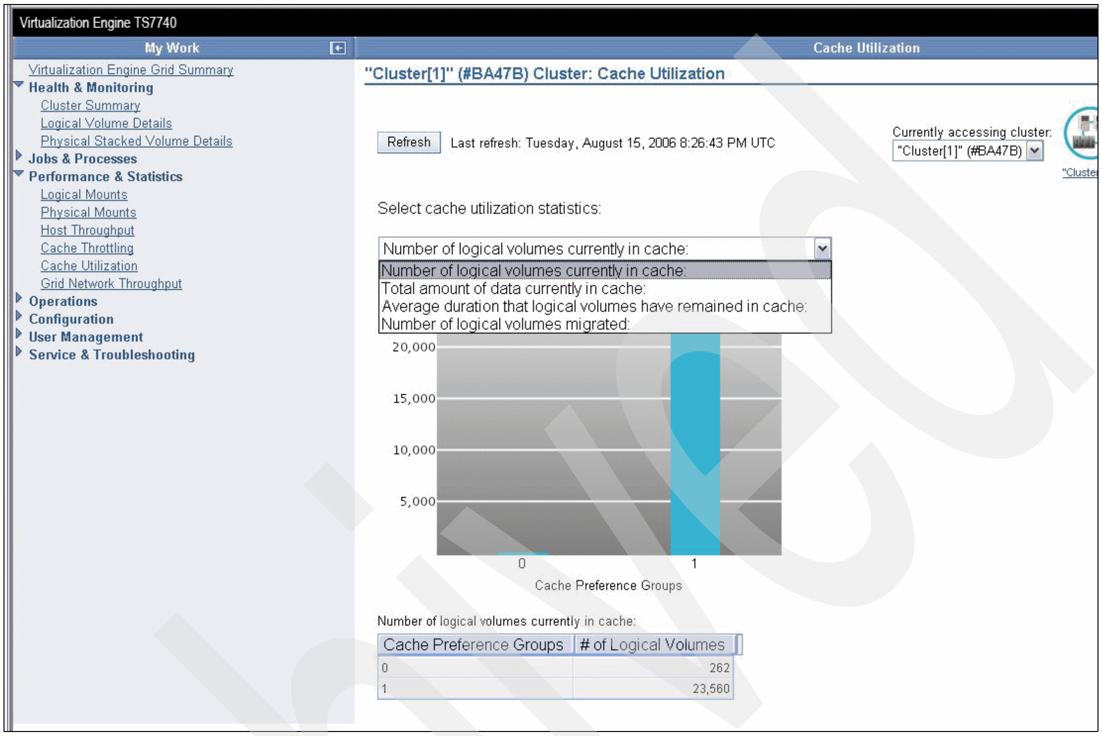

Figure 8-24 TS7700 Management Interface Cache Utilization

The cache utilization statistics for each cluster with regards to copy and write are displayed both in a bar graph form and in a table. Each statistic has its own bar graph and table selected from the View drop-down.

Average Duration in Cache and Number of Logical Volumes Migrated have a table column for each of the time periods mentioned in parenthesis.

- Cache Preference Group: Possible values:

  0: Volumes in this group have preference to be removed from cache over other volumes.
  1: Volumes in this group have preference to be retained in cache over other volumes.

- # of Logical Volumes: The number of logical volumes present in the cache preference group
- Amount of data: Total amount of data present in volumes assigned to the cache preference group
- Average Duration in Cache (Last 4 hours, 48 hours, 35 days): Rolling average of the cache age of volumes migrated out of this cache preference group for the specified amount of time
- Number of Logical Volumes Migrated (4 hours, 48 hours, 35 days): Rolling average of the number of volumes migrated to this cache preference group

Grid network throughput

Use this page to view network path utilization (Grid Network Throughput) statistics for the IBM TS7700 Virtualization Engine Cluster.

This page presents information about cross-cluster data transfer rates. This selection will be present only in a Multi Cluster Grid configuration. If the TS7700 grid only has one cluster, there is no cross-cluster data transfer using the Ethernet adapters.

The example provided in Figure 8-25 is taken from a TS7700 cluster that is part of a Multi Cluster Grid configuration (Three-Cluster Grid).

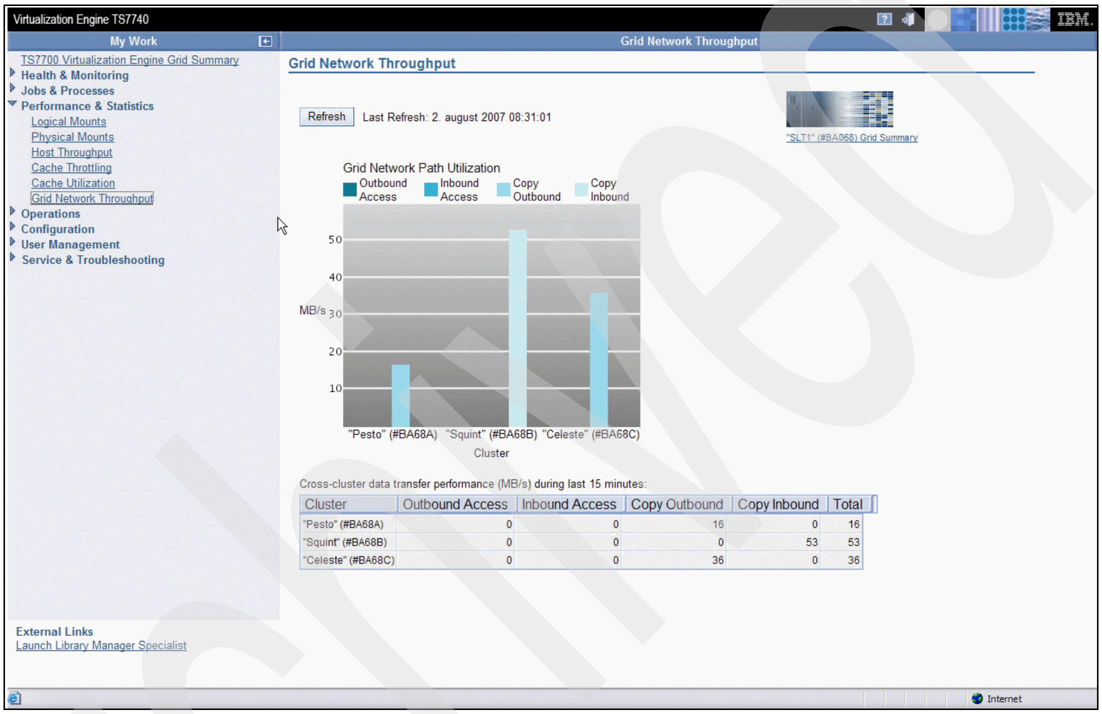

Figure 8-25 TS7700 Management Interface Grid Network Throughput in a Three-Cluster Grid

The table displays data for cross-cluster data transfer performance (MB/s) during last 15 minutes. The table cells show:

Cluster	The cluster name
Access Read	Data transfer rate for read operations
Access Write	Data transfer rate for write operations
Copy Read	Data transfer rate for copy read operations
Copy Write	Data transfer rate for copy write operations

8.5 TS7700 statistical data

The IBM TotalStorage 3494 Virtual Tape Server (VTS) revolutionized the way System z customers utilized their tape resources. To help them monitor the performance of the VTS various statistics were developed. These statistics appeared in two forms: Hourly records sent to the hosts known as SMF94 records, and periodic real-time statistics available using the Library's Web specialist.

For the next generation of VTS, the TS7700 Virtualization Engine, the statistics design has been revised. The TS7700 does not send SMF Type 94 records to be stored at the host, but rather writes statistical records to a TS7700 database. Various statistics from the VTS have been retained and new statistics relevant for the TS7700 have been added.

Both point-in-time (PIT) and historical statistics are recorded. The point-in-time records present data from the most recent interval providing speedometer like statistics. The historical statistics provide statistics where historical trends can be observed.

These statistics records are available to a host through the Bulk Volume Information Retrieval (BVIR) facility. Refer to 8.6, “Bulk Volume Information Retrieval (BVIR)” on page 442 for more information about how to retrieve the statistics records and to 8.7, “Using VEHSTATS for monitoring and reporting” on page 474 about how to format and analyze these records.

Each cluster in a grid has its own set of vNode and hNode point-in-time (PIT) and historical statistics.

8.5.1 Type of statistical records

The TS7700 Virtualization Engine provides two types of statistics:

- Point-in-Time Statistics: Report operations over the last 15-second interval. The main intent is to help answer the question: What is going on now?
- Historical Statistics: Report operations in 15 minute intervals. Intervals are accumulated on a daily basis (96 intervals/day) and 90 days of historical statistics is retained in the TS7700 subsystem. The main intent is to answer the question: How are you using the resources?

Point-in-time statistics

The data provided by this type of statistic is a snapshot of the TS7700 activity over the last 15 second interval. Each new 15 second interval data overlays the prior interval’s data. But not all data is updated every 15 seconds (primarily hNode data). Those statistics contain Single and Multi Cluster Grid information.

You can obtain the point-in-time statistics using the appropriate BVIR request; the response returns the last PIT snapshot from the TS7700. The data returned is not in a human readable format, it is primarily binary data {use DCB=(RECFM=U,BLKSIZE=24000)}. The records are of different length depending on the record type.

Appendix I, “TS7700 Statistics record format” on page 631 and the White Paper *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide Version 1.0* provide you with details.

Basically, the point-in-time statistics provide these record types:

- vNode - Virtual Device
- vNode - Host Adapter
- hNode - HSM
- hNode - Grid

A variable number of records in data is returned, depending on the number of vNodes and hNodes and the number of clusters (if in a Multi Cluster Grid configuration).

Each record has a common header containing:

- Record length
- Record type
- Node & distributed library ID
- Timestamp
- Machine type, model and serial number
- Code level

Point-in-time statistics - vNode - virtual device

These point-in-time statistics provide the state of the reporting vNode:

- Overall state (Online, offline, and so forth)
- Maximum configured throughput
- Number of installed virtual devices

In addition, the state and usage of each virtual drive is presented:

- Volume mounted or last mounted
- Distributed library access point
- Mount state (unloaded, mount in progress, failed, mounted, and so forth.)
- Device state (ready, write mode, BOT, and so forth)
- Buffer wait condition count (view of whether the subsystem is pacing the channel or visa versa)
- Data transferred

Point-in-time statistics - vNode - host adapter

These kinds of point-in-time statistics provide the status for each of the four host adapter positions:

- Adapter type installed
- State (Online, offline, and so forth)
- Location (drawer/slot)

In addition, the state and usage of each adapter port is provided:

- Port interface ID
- Errors (loss of light, timeouts, bit errors)
- Data transferred before and after compression

Point-in-time statistics - hNode - HSM

These point-in-time statistics provide the status for the management tasks running in the hNode:

- Pre-migrate and recall queue counts
- Throttling value (write and copy)
- Library sequence number
- State and usage of the physical drives
  - Device type
  - Physical volume mounted & pool ID
  - Device state (offline/online)
  - Device role (idle/premig/recall/reclaim/etc)
  - Logical volume being processed
  - Data transferred

Note: HSM does not relate to the Hierarchical Storage Manager (DFSMShsm), the z/OS storage management software.

Point-in-time statistics - hNode - grid

These point-in-time statistics provide the grid status from this hNode perspective:

- Status and usage for the distributed library
  - Run and deferred copy queue counts for this distributed library
  - Active copies for this distributed library
  - Link usage %
  - Errors (network CRC)
- Data transferred to/from each cluster

Historical statistics

The data provided by this type of statistic is captured over a 15 minute interval in the TS7700. Each new 15 minute interval data does not overlay the prior interval's data. But not all data is updated every 15 minutes. Those statistics contain Single and Multi Cluster Grid information. Up to 96 intervals per day can be acquired and therefore 90 rolling days of historical statistics are kept in the TS7700 subsystem.

You can obtain the complete set or a subset of these historical statistics through the appropriate BVIR request (for more details, see "Historical statistics" on page 464). The request will contain the day or the days of needed historical statistic information. For the current day, records up to the last 15 minute interval are returned. The data returned is not in a human readable format, it is primarily binary data (use DCB=(RECFM=U,BLKSIZE=24000). The records are of different lengths depending on the record type. Refer to Appendix I, "TS7700 Statistics record format" on page 631 for more information about the format of the statistics records, and about a formatting tool called VEHSTATS. For more information about using VEHSTATS, see 8.7, "Using VEHSTATS for monitoring and reporting" on page 474.

Basically, the historical statistics provide six record types:

- vNode - Virtual Device
- vNode - Host Adapter
- hNode - HSM
- hNode - Library
- hNode - Grid

A number of records are returned, depending on the number of vNodes and hNodes and the number of clusters (if in a Multi Cluster Grid configuration).

Each record has a common header containing:

- Record length
- Record type
- Node & distributed library ID
- Timestamp
- Machine type, model and serial number
- Code level

Historical statistics - vNode - virtual device

These types of historical statistics will provide the following information about the usage for a vNode's virtual devices:

- Number of installed virtual devices
- Virtual device type
- Blocksizes being written
- Configured throughput
- Min, max, and average virtual devices mounted

Historical statistics - vNode - host adapter

You can use this type of historical statistics to receive the status for each of the four host adapter positions. The provided statistics contain:

- Adapter type installed
- State (Online, offline, and so forth)
- Location (drawer/slot)

Both state and usage of each adapter port are provided:

- Port interface ID
- Interface data transfer rate setting (capable and actual)
- Data transferred before and after compression
- Selective and system reset counts

Historical statistics - hNode - HSM

This data portion within the historical statistics will give you HSM information. You will get the VOLSER of the physical volume with the latest database copy on it.

You will also see state and status of the Tape Volume Cache:

- Usable size in GBytes
- Throttling values
- State of each cache partition
  – Partition size
  – Number of fast-ready, cache hit, cache miss mounts
  – Average fast-ready, cache hit, cache miss mount times
  – Number of volumes in cache by preference group
  – Space occupied by the volumes in cache by preference group
  – Volume aging by preference group

Historical statistics - hNode - library

This part of the historical statistics will give you library information such as:

- Information about the attached physical library
- Information about the physical tape devices in the library
- Information about the common scratch pool media in the library
- Information about each physical volume pool in the library

Historical statistics - hNode - grid

This last set of historical statistics will give you information in a Multi Cluster Grid environment for:

- Status and usage for the distributed library:
  – Number of logical volumes and data to copy for this distributed library
  – Average age of the deferred and immediate copy queues for this distributed library
  – Number of distributed libraries in the grid configuration.
- Data transferred to/from each cluster

8.5.2 Collecting and analyzing the TS7700 statistical records

As stated in the introduction of this section, you need to consider two aspects for working with the TS7700-provided statistics and reporting mechanisms.

- **Getting the statistics and reports**: The main interface to access statistic and reports from the TS7700 Virtualization Engine is the Bulk Volume Information Retrieval (BVIR) functions. Depending on the request, you will receive readable output or, for the TS7700 point-in-time and historical statistics, binary data. For the response of binary data further documentation and tools are needed. For details on how to use BVIR functions, refer to 8.6, “Bulk Volume Information Retrieval (BVIR)” on page 442.

- **Formatting and displaying the information**: Some of the response data of the BVIR functions is already in a readable format. For the remaining binary format provided by the point-in-time and historical statistics you need a formatting tool. IBM provides a tool called VEHSTATS. Further information about where to download this tool and how to use it is in 8.7, “Using VEHSTATS for monitoring and reporting” on page 474.

8.6 Bulk Volume Information Retrieval (BVIR)

Internal to the TS7700 Virtualization Engine, a large amount of information is captured and maintained about the state and operational aspects of the resources in the TS7700 Virtualization Engine. With Bulk Volume Information Retrieval (BVIR) you are now able to obtain information about all of the logical volumes a TS7700 manages:

- Volume Status Information
- Cache Contents Information
- Physical Volume to Logical Volume Mapping Information
- Point-in-Time Statistics
- Historical Statistics
- Physical Media Pools

> **Note:** The information about BVIR is current and correct at the time of this writing. However, to make sure that you are always using the most current information about BVIR functions and record formats, refer to
>
> - *IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide* at:
>
>   http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100828
>
> - *Virtualization Engine TS7700 Series Statistical Data Format* White Paper at:
>
>   http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100829

8.6.1 Overview of the BVIR function

With the potential to support hundreds of thousands of logical volumes in a TS7700 Virtualization Engine, providing a set of information for all of those volumes through normal channel control type commands is not very practical. Luckily, the functions of a TS7700 subsystem that allow it to virtualize a tape volume, also provide a simple and effective method to transfer the information to a requesting application. The TS7700 converts the format and storage conventions of a tape volume into a standard file managed by a file system within the subsystem.

The Bulk Volume Information Retrieval (BVIR) facility uses an IBM standard labeled tape volume to both initiate a request for information and return the results. By using a standard tape volume, no special interfaces or access methods are needed for an application to use this facility. In practice, no specific applications are required, as standard IBM utilities, such as IEBGENER, provide the function needed to request and obtain the information.

The BVIR function was first introduced on the prior generation of Virtual Tape Servers. Its use has been expanded with the introduction of the TS7700 Virtualization Engine.

There are two steps to obtain information using this facility:

1. A single data set with the information request is written to a logical volume. The logical volume can be any logical volume in the subsystem the information is to be obtained from. Either a scratch or specific volume request can be used. The data set contains a minimum of two records and a maximum of three records that specify the type of data being requested. The records are in human readable form, that is, lines of character data. The data set can be cataloged or uncataloged (although cataloging the data set can make it easier for subsequent access to the data). On closing the volume, the TS7700 Virtualization Engine server recognizes it as a request volume and "primes" the subsystem for the next step.

   Note: Some of the information obtained through this function is specific to the cluster the logical volume is written on, for example cache contents or logical-physical volume map. In a TS7700 grid configuration with multiple clusters, use a Management Class for the volume to obtain statistics for a specific cluster.

   Historical statistics for a Multi Cluster Grid can be obtained from any of the clusters.

2. The request volume is again mounted, this time as a specific mount. Seeing that the volume was primed for a data request, the TS7700 appends the requested information to the data set. The process of obtaining the information and creating the records to append can take up to several minutes, depending on the request and, from a host's viewpoint, is part of the mount processing time. When the TS7700 has completed appending to the data set, the host is notified that the mount has completed. The requested data can then be accessed like any other tape data set.

   In a JES2 environment, the JCL to perform the two steps can be combined into a single job. However, in a JES3 environment, they must be run in separate jobs. This is because the volume will not be demounted and remounted between job steps in a JES3 environment.

After the response data set has been written to the request logical volume, that logical volume functions identically to any other logical volume in the subsystem. Subsequent mount requests and read accesses to the logical volume should have no effect on its contents. Subsequent mount requests and write accesses to the logical volume will overwrite its contents. It can be returned to scratch status and reused by any application.

Figure 8-26 shows the process flow of BVIR.

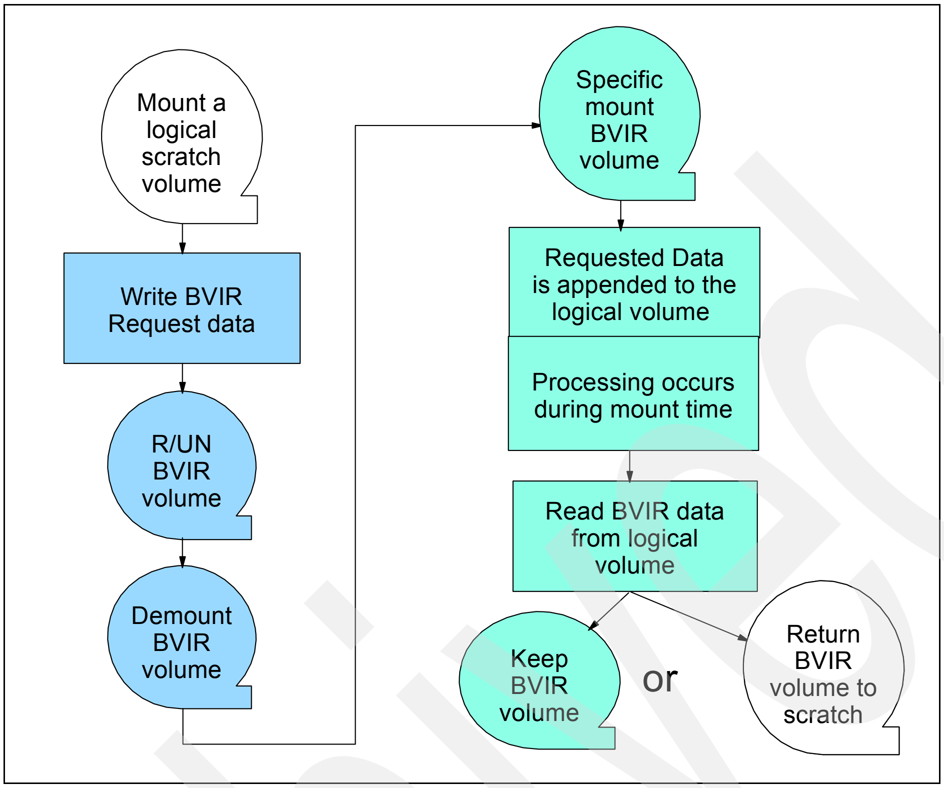

Figure 8-26 BVIR process flow

The building of the response information does require a small amount of the resources of the TS7700. It is recommended that the BVIR function not be used to “poll” for a specific set of information and that only one request be issued at a time. Some requests, for example the volume map, might take several minutes to complete and to prevent “locking” out another request during that time, the TS7700 is designed to handle two concurrent requests. If more than two concurrent requests are issued, they will be processed as prior requests are completed.

Whereas the request data is always in a human readable format, depending on the request, the data returned from the TS7700 can be in human readable or binary form. Refer to the response sections for the specifics of the returned data.

The general format for the request/response data set is shown in Example 8-2.

Example 8-2 BVIR output format

```
VTS BULK VOLUME DATA REQUEST
VOLUME MAP
11/20/2003 12:27:00 VERSION 02
S/N: 0F16F LIB ID: AB123

PHYSICAL LOGICAL P/B  ORDER   PART       SIZE
P00024   GK0000    P 000001 1 OF 1    23.45 M
P00024   GK0020    P 000002 1 OF 1    76.50 M
P00024   GK0010    P 000003 1 OF 1   134.24 M
```

Record 0 is identical for all requests and not part of the output but just for supporting Records 1 through 5. Records 6 and above contain the requested output, which differs depending on the request:

- Records 1 and 2 contain the data request commands.
- Record 3 contains the date and time when the report was created and the version of BVIR.
- Record 4 contains the serial number and the Library ID of the TS7700.
- Record 5 contains all blanks.
- Record 6 contains the title of the output columns, in our example for a Volume Map.
- Records 7 and above contain the requested data. The fields of these and the other reports are described in the next sections.

8.6.2 Prerequisites

Any logical volume defined to a TS770 Virtualization Engine can be used as the request/response volume. Logical volumes in a TS7700 are formatted as IBM Standard Labeled volumes. Although a user can reformat a logical volume with an ANSI Standard Label or as an unlabeled tape volume, those formats are not supported for use as a request/response volume.

There are no restrictions regarding the prior use of a volume used as a request/response volume and no restrictions regarding its subsequent use for any other application. It is recommended that normal scratch allocation methods are used for each request (i.e. DISP=(NEW,CATLG)). In this way, any of the available scratch logical volumes in the TS7700 can be used. Likewise, it is recommended that when the response volume's data is no longer needed, the logical volume is returned to scratch status through the normal methods (typically by deletion of the data set on the volume and a return to scratch policy based on data set deletion).

8.6.3 Request Data Format

Several types of data can be requested. The type of data requested is indicated in the request data set. The request data set must be the only data set on the volume and must be written with a record format of F and a logical record size of 80 bytes. Request information is in EBCDIC character form, beginning in the first character position of the record and padded with blank characters on the right to fill out the record.

> **Important:** The request fields must be as shown. Not beginning in the first character position of the record or extra blanks between words will result in the request failing.

Although the request data format uses fixed records, not all response records are fixed. For the point-in-time and historical statistics responses, the data records are of variable length and the record format used to read them is the Undefined (U) format. Refer to the sample JCL section.

In a multi-site TS7700 grid configuration, the request volume must be created on the cluster that the data is being requested for. The Management Class assigned to the volume needs to specify the particular cluster that is to have the copy of the request volume.

The following sections describe the format for the request data set records.

Record 1

Record 1 must contain the following command, exactly as shown:

VTS BULK VOLUME DATA REQUEST

The format for the request data set records is shown inTable 8-1.

Table 8-1 BVIR Request Record 1

Record 1: Request Identifier		
Bytes	**Name**	**Contents**
1 - 28	Request Identifier	VTS BULK VOLUME DATA REQUEST
29 - 80	Blanks	Blank padding
41 - 80	Blank	Blank padding

Record 2

With Record 2, you can specify which data you want to obtain. The following options are available:

- VOLUME STATUS zzzzz
- PHYSICAL VOLUME STATUS VOLUME zzzzzz
- PHYSICAL VOLUME STATUS POOL xx
- CACHE CONTENTS
- VOLUME MAP
- POINT IN TIME STATISTICS
- HISTORICAL STATISTICS FOR xxx
- HISTORICAL STATISTICS FOR xxx-yyy
- PHYSICAL MEDIA POOLS
- COPY AUDIT COPYMODE INCLUDE/EXCLUDE libids

The format for the request data set records is shown in Table 8-2.

Table 8-2 BVIR Request Record 2

Record 2: Request		
Bytes	**Name**	**Contents**
1 - 40	Request	VOLUME STATUS zzzzzz or PHYSICAL VOLUME STATUS VOLUME zzzzzz or PHYSICAL VOLUME STATUS POOL xx or CACHE CONTENTS or VOLUME MAP or POINT IN TIME STATISTICS or HISTORICAL STATISTICS FOR xxx-yyy or PHYSICAL MEDIA POOLS COPY AUDIT COPYMODE INCLUDE/EXCLUDE libids left justified, padded with blanks on the right
41 - 80	Blank	Blank padding

For the VOLUME STATUS request, “zzzzzz” specifies the volume serial number mask to be used. By using the mask, one to thousands of volume records can be retrieved for the request. The mask must be 6 characters in length, with the “_” character representing a positional wildcard mask.

For example, assuming that volumes in the range of ABC000 through ABC999 have been defined to the cluster, a request of VOLUME STATUS ABC1_0 would return database records that exist for ABC100, ABC110, ABC120, ABC130, ABC140, ABC150, ABC160, ABC170, ABC180, ABC190, ABC1A0, ABC1B0,...ABC1Z0.

- For the HISTORICAL STATISTICS request, “xxx” specifies the Julian day being requested. Optionally, “-yyy” can also be specified and indicates that historical statistics from xxx through yyy are being requested. Valid days are 001 through 366 (to account for the leap year).

  For leap years, February 29th is Julian day 060 and December 31st is Julian day 366. For other years, Julian day 060 is March 1st and December 31st is Julian day 365. If historical statistics do not exist for the days requested, that will be indicated in the response record (this would occur if a request is issued for a day prior to the day the system was installed, days the system was powered off, or after the current day before a rolling year has been accumulated).

  If a request spans the end of the year, for example a request that specified: HISTORICAL STATISTICS FOR 364-002, responses are provided for days 364, 365, 366, 001 and 002, regardless of whether the year was a leap year.

- For point-in-time and HISTORICAL STATISTICS requests, any additional characters provided in the request record past the request itself are retained in the response data, but otherwise ignored.

8.6.4 Response data format

When the request data set has been written to the volume and subsequently closed and demounted, when mounted again, the TS7700 Virtualization Engine validates the contents of the request volume and appends the requested data records to the data set.

Human readable appended records are 80 bytes in length. Binary data appended records can be variable in length of up to 32760 bytes. The data set is now a response data set. The appropriate block counts in the end of file (EOF) records will be updated to reflect the total number of records written to the volume.

After appending the records and updating the EOF records, the host that requested the mount is signaled that the mount is complete and can read the contents of the volume. If the contents of the request volume is not valid, either one or more error description records will be appended to the data set or the data set will be unmodified prior to signaling the host that the mount completed, depending on the problem encountered.

All human readable response records begin in the first character position of the record and are padded with blank characters on the right to fill out the record. All binary records are variable in length and are not padded.

> **Note:** In the response records, the dates and times presented are all based on the internal clock of the TS7700 handling the request. The internal clock of a TS7700 is not synchronized to the host, but is synchronized with any other TS7700.

The response data set contains both request records that we describe in 8.6.3, “Request Data Format” on page 445, and the response data set contains three explanatory records (Records 3 - 5) and, starting with Record 6, the actual response to the data request.

Record 3

This record contains the date and time the response data set was created as well as a format version number for the results, as shown in Example 8-3 and Table 8-3.

Example 8-3 BVIR Response Record 3

```
11/20/2003 12:27:00 VERSION 02
```

Table 8-3 BVIR Response Record 3

Bytes	Name	Description
1-10	Date	The date the response was generated in the format of MM/DD/YYYY
11	Field Delimiter	Blank character
12-19	Time	The time the response was generated in the format HH:MM:SS
20	Field Delimiter	Blank character
21-27		"VERSION"
28	Field Delimiter	Blank character
29-30	Version Number	The Version Number of this response right justified and zero filled. Original VTS supports Version 01, TS7700 supports Version 02
31-80	Blanks	Blank character padding

Record 4

This record contains the 5 character sequence number of the TS7700 cluster that generated the response shown in Example 8-4 and Table 8-4.

Example 8-4 BVIR Response Record 4

```
S/N: 0F16F  LIB ID: AB123
```

Table 8-4 BVIR Response Record 4

Bytes	Name	Description
1-4	Date	"S/N:"
5	Field Delimiter	Blank character
6-10	Time	Sequence number of the TS7700 that generated the response
11-12	Field Delimiter	Blank character
13-19		"LIB ID:"
20	Field Delimiter	Blank character
21-25	Version Number	Sequence number of the library attached to the TS7700
26-80	Blanks	Blank character padding

The TS7700 sequence number field (S/N) is the 5 character identifier assigned to the TS7700 cluster. The S/N is defined by the IBM CE during installation of the TS7700 Virtualization Engine. The Library ID field (LIBID) is defined by the IBM CE during installation of the library. It is also called Library ID on the DFSMS Tape Library Define panel, or as shown on the Tape Library Display panel.

Record 5

This record contains all character blanks.

Record 6-N

These records contain the specific response records based on the request. If the request could not be understood or was invalid, that will be indicated.

Refer to the following sections for the response record definitions.

8.6.5 Interpreting the BVIR Response Data

Here we explain how to interpret each BVIR Response Data set for the specific request information, such as:

- Volume Status Information
- Physical volume status information
- Cache Contents Information
- Physical Volume to Logical Volume Mapping Information
- Point in Time Statistics
- Historical Statistics
- Physical Media Pools

When records are listed in this chapter there will be a record showing ‘123456789012’. This record does not exist but is shown to improve readability.

Volume Status Information

A database is maintained on each individual TS7700 Cluster that contains information about the state and data validity of each logical volume on that TS7700 Cluster. The database also contains other information related to the management of the logical volumes on the cluster and copy and resynchronization processes when the TS7700 Virtualization Engines are in a grid configuration. Several of the database fields returned can be useful in handling operational exceptions at one or more of the clusters.

Data Inconsistent

This field indicates whether the cluster has a valid version of the data or not. If it indicates that the data on the logical volume is not valid, this means that the same volume on another TS7700 Virtualization Engine in the Grid has been modified and it has not yet been copied. For customers who use the deferred copy consistency point (which is typically when there is some significant distance between the TS7700 Virtualization Engines in the Grid configuration), there will be some number of volumes that are not consistent between the TS7700s at any point in time. If a situation occurs that renders the site inoperable where the source data resides by issuing the Volume Status request to an operable TS7700, this field can be used to identify the volumes that were not copied prior to the situation so that appropriate recovery steps can be performed for them.

MES Volume

This field indicates that the logical volume was created in the TS7700 Cluster or even created within a VTS, prior to being merged into a grid configuration. Volumes that existed in a TS7700 Cluster prior to being included in a grid configuration are not automatically copied to the other TS7700 Clusters in the configuration until they have been accessed and closed. This field could be used to determine which volumes in each TS7700 Cluster have not been copied, to build a set of jobs to access them, and force the copy. The PRESTAGE program from the TAPETOOL FTP site could support you to do that job in an efficient way. The VEHSYNC job can be used to identify volumes needing copies.

Copy Required for Cluster n

This field indicates that a copy to another TS7700 Cluster in a grid configuration is required. In cases where deferred mode copy is used, this field can be used to determine whether a critical set of volumes has completed their copy operations to specific clusters.

Volume Ownership and Volume Ownership Taken

At any point in time a logical volume is owned by a specific cluster. Ownership is transferred as part of mount processing. Ownership can transfer in one of two ways, either through communication with the current owning cluster or through a recovery process called ownership takeover. Normally, the cluster receiving a mount command requests that the current owning cluster transfer ownership (assuming that the cluster receiving the mount request for the volume does not already have ownership of the volume) and ownership is transferred.

However, if the cluster receiving the mount request cannot communicate with the owning cluster, that method does not work. In this case the requesting clusters cannot determine whether the owning cluster has failed or just the communication paths to it have failed. Operator intervention is required to indicate that the owning cluster has failed and that ownership takeover by the other clusters is allowed.

There are two types of ownership takeover, Write and Read-only. With Write ownership takeover (WOT), the cluster taking over ownership of the volume has complete freedom to modify the contents of the volume or modify any of the properties associated with the volume. With Read-only takeover (ROT), the cluster taking over ownership of the volume is restricted to reading the volume's data only. Scratch mounts are failed in a cluster in Read-only Takeover mode.

Current and Pending Category

One of the key properties associated with a volume is the *category* it is assigned. The primary usage for category is to group together scratch volumes. A volume's category assignment changes as the volume is used. The current category field indicates the category the volume is assigned to in the Library Manager associated with the cluster. The pending category field indicates that a new category assignment is in progress for the volume. These fields can be used to determine whether the category assignments are in sync between the clusters and the host databases.

Data Deleted

As part of normal processing in a TS7700 Cluster, you can specify that after a certain period of time after being returned to scratch, the contents of a volume can be deleted. This field indicates whether or not the data associated with the volume has been deleted on the cluster.

Volume Status

The volume status information returned represents the status of the volume on the cluster the request volume is written. In a TS7700 grid configuration, separate requests must be issued to each cluster to obtain the volume status information for the individual clusters. A response record is written for each logical volume, selected based on the volume serial number mask specified in the request, that exists in the cluster.

A response record consists of the database fields defined in the following table. Fields are presented in the order defined in the table and are comma (,) separated. The overall length of each record is 320 bytes with blank padding after the last field as needed. The first few fields of the record returned for VOLSER ABC123 would be as shown in Example 8-5.

Example 8-5 BVIR Volume Status Information

```
12345678901234567890123456789012345678901234567890123456789012345
ABC123,0,2006-04-22-11.56.45.871263,0,0,32,0,N,2548,N,8719,N...
```

Record 6-N

If the request is for a specific volume and it does not exist in the cluster's database, or the request is for a range of volumes and none of the volumes exist in the cluster's database, the record shown in Example 8-6 is returned.

Example 8-6 BVIR request for a non-existing specific volume

```
NO VOLUME RECORD(S) EXIST FOR THE REQUESTED VOLUME(S)
```

For the requested volumes that do exist in the cluster's database, each of these records provides information for one logical volume as shown in Example 8-6, according to Table 8-5.

Table 8-5 Decoding Volume Status Information

Field Name	Description
Volser	6 character volume serial number (VOLSER)
domain_lock_cluster	The cluster identifier of the cluster which currently has ownership of the volume. 0 is the identifier for the first cluster in a Grid.
domain_lock_timestamp	Time stamp of the last volume ownership change for the volume. The format of the time stamp is: Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2006-05-23-19.34.23.876129
mounted_cluster	If the volume is mounted, the cluster identifier of the cluster the mount request for the volume was issued to. 0 is the identifier for the first cluster in a Grid. If the volume is not mounted, -1 is indicated.
mounted_vnode	If the volume is mounted, the vNode identifier in the cluster that the mount request for the volume was issued to. 0 is the identifier for the first VNode in a cluster. If the volume is not mounted, -1 is indicated.
mounted_device	If the volume is mounted, the virtual device number in the vNode of the cluster that the mount request for the volume was issued on. 0 is the identifier for the first virtual device in a VNode. If the volume is not mounted, -1 is indicated.
tvc_cluster	The cluster identifier associated with the tape volume cache being used for the mounted volume. With remote mounting support, this cluster might not be the same cluster as the mounted_cluster. 0 is the identifier for the first cluster in a Grid.
data_inconsistent	This field indicates whether the cluster has an inconsistent copy of the volume. Y indicates it does not have a valid copy, N indicates that it has a valid copy
data_level	The current data level value for the volume on the cluster. Every time a volume's data is modified, its data level is incremented.

Field Name	Description
properties_inconsistent	This field indicates whether the cluster has inconsistent properties for the volume. Y indicates it has one or more properties that are down level, N indicates that it does not. The properties consist of the category and storage constructs assigned to the volume.
properties_level	The current properties level value for the volume on the cluster. Every time a volume's properties are modified, its data level is incremented.
Cluster0_copy_required	This field indicates whether Cluster 0 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
Cluster1_copy_required	This field indicates whether Cluster 1 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
Cluster2_copy_required	This field indicates whether Cluster 2 needs to make a copy of the volume from this cluster. Y indicates it does, N indicates that either a copy is not required or one has already been made.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.
Cluster0_copy_mode	This field indicates whether Cluster 0 is to have a copy of the volume and the copy consistency point defined for the volume. The values are: I Rewind unload (RUN) copy consistency point. D Deferred copy consistency point. N No copy E The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
Cluster1_copy_mode	This field indicates whether Cluster 1 is to have a copy of the volume and the copy consistency point defined for the volume. The values are: I Rewind unload (RUN) copy consistency point. D Deferred copy consistency point. N No copy E The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
Cluster2_copy_mode	This field indicates whether Cluster 2 is to have a copy of the volume and the copy consistency point defined for the volume. The values are: I Rewind unload (RUN) copy consistency point. D Deferred copy consistency point. N No copy E The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.

Field Name	Description
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.
mes_flag	This field indicates whether the volume was part of a MES merge operation. Y indicates the volume existed prior to merging the cluster into a Grid configuration. N indicates that the volume has been successfully copied to the clusters specified in the copy_mode fields.
recall_error	This field indicated whether or not an error occurred while recalling this volume from a physical tape associated with the tvc_cluster. Y indicates that an error occurred the last time a recall was attempted with the volume. N indicates that the last recall of the volume was successful.
read_error	This field indicates whether the data associated with the volume is corrupted. Y indicates that the volume's data is corrupted. N indicates that the data is not corrupted. A volume's data is determined to be corrupted if an error is detected when reading the data.
disaster_rec	This field indicates the disaster recovery state for the volume. Y indicates that the volume has been through the disaster recovery process. N indicates that no disaster recovery processing has been done or is required for the volume.
data_deleted	This field indicated whether the data associated with the volume has been deleted by the delete expired volume data function. Y indicates that the data has been deleted. N indicates that the data has not been deleted.
volume_damaged	The volume/token is determined to be damaged or unrecoverable without human intervention. Y indicates that the volume/token has been damaged. N indicates that the volume/token has not been damaged.
Mount_operation_time	Timestamp of when the volume was last mounted or unloaded. The format of the time stamp is: Year-Month-Day-Hour.Minute.Second.Microsecond for example: 2006-05-23-19.34.23.876129
Media_type	The media type defined when the volume was inserted. The values are: 0 Cartridge System Tape (400 MB) 1 Enhanced Capacity Cartridge System Tape (800 MB)
steal_active	This field indicates whether the volume's ownership has been taken over. 0 indicates that the volume's ownership has transferred normally. A non-zero value indicates that the volume's ownership was transferred by one of the ownership takeover methods. ► bits 0:3 Cluster ID that stole this volume ► bits 4:7 Cluster ID that was the victim of the steal. ► bit 8 - WOT/ROT 0 ROT or Read Ownership Takeover 1 WOT or Write Ownership Takeover ► bit 9 - Set if SOT or Service Ownership Takeover
pending_category	The category the volume is to be set to within the Library Manager associated with the cluster.

Field Name	Description
Current_category	The category the volume is currently assigned to in the Library Manager associated with the cluster. When this value and the pending_category are equal, then the category is consistent
category_time	Time stamp of when the current_category or pending_category fields were last changed. The format of the time stamp is: Year-Month-Day-Hour.Minute.Second.Microsecond For example: 2006-05-23-19.34.23.876129 Prior to the first change in category, this field is set to 1970-01-01-00.00.00.000000.
takeover_timestamp	Time stamp of when the volume's ownership was last changed due to write or read ownership takeover. The format of the time stamp is: Year-Month-Day-Hour.Minute.Second.Microsecond For example: 2006-05-23-19.34.23.876129. If the volume has not had its ownership taken, this field is set to 1970-01-01-00.00.00.000000.
Blanks	As needed to pad record to 320 bytes

Physical volume status information

A database is maintained on each individual TS7700 Cluster that contains information related to the management of the physical volumes on the cluster.

The physical volume status information returned represents the status of the volumes on the cluster the request volume is written. In a TS7700 Grid configuration, separate requests must be issued to each cluster to obtain the physical volume status information for the individual clusters. A response record is written for each physical volume, selected based on the volume serial number mask or pool number specified in the request, that exists in the cluster.

A response record consists of the database fields defined in Table 8-6. Fields are presented in the order defined in the table and are comma (,) separated. The overall length of each record is 400 bytes with blank padding after the last field as needed. For example, the first few fields of the record returned for VOLSER A03599 would be:

```
A03599,2,FULL,READ-WRITE,2007-05-05-06.40.08.030061,2007-05-04-13.45.15.918473,..
```

> **Note:** The generation of the response might take several minutes to complete depending on the number of volumes requested and how busy the TS7700 cluster is at the time of the request.

Record 6-N

If the request is for a specific physical volume and it does not exist in the cluster's database or the request is for a range of physical volumes and none of the volumes exist in the cluster's database, or the request is for a pool and there are no physical volumes currently resident in the pool, the following record is returned.

```
NO VOLUME RECORD(S) EXIST FOR THE REQUESTED VOLUME(S)
```

The structure of the record is shown in Table 8-6.

Table 8-6 Decoding Volume not found record

Bytes	Name	Description
1-53		NO VOLUME RECORD(S) EXIST FOR REQUESTED VOLUME(S)
54-80	Blanks	Blank character padding

For the requested volumes that do exist in the cluster's database, each of these records provides information for each physical volume, as shown in Table 8-7.

Table 8-7 Decoding Volume record Information

Field Name	Description
VOLSER	6 character volume serial number
CURRENT_POOL	This field indicates the pool that the volume is currently assigned to. Pool 0 is the common scratch pool. Pools 1-32 are the specific data pools.
VOLUME_STATUS	This field indicates the volume's current capacity state. The following are the values that can be indicated: ► EMPTY: The volume contains no data and is available for use as a physical scratch volume. ► FILLING: The volume contains valid data, but is not yet full. It is available for additional data. ► FULL: The volume contains valid data. At some point it was marked as full and additional data cannot be added to it. A volume can be marked full in some cases short of the volume capacity limit. ► UNKNOWN: The volume's capacity state is unknown.
VOLUME_ACCESS	This field indicates the volume's current accessibility state. The following are the values that can be indicated: ► READ-WRITE: The volume can be read from or written to. ► READ-ONLY: The volume contains valid data, but had encountered read or write errors that exceeded an allowable threshold. It is only accessible for read. ► UNAVAILABLE : The volume is unavailable to the TS7700. ► DAMAGED: The volume has been damaged physically or logically such that it cannot be mounted or read from. ► COPY-EXPORTED: The volume has been exported using the Copy Export function and is not currently resident in the library.
STARTED_EXPIRING_TIMESTAMP	Timestamp of when the data first started expiring on the physical volume. The format of the timestamp is: Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2006-05-23-19.34.23.876129. This field is set to 1970-01-01-00.00.00.000000 when the volume becomes empty (returned to scratch) or is inserted into the library.
BECAME_EMPTY_TIMESTAMP	Timestamp of when all of the data on the volume last became invalid and the volume is empty. The format of the timestamp is: Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2006-05-23-19.34.23.876129. When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.

Field Name	Description
BECAME_FULL_TIMESTAMP	Timestamp of when the volume is last marked as full. The format of the timestamp is: Year-Month-Day-Hour.Minute.Second.Microsecond, for example: 2006-05-23-19.34.23.876129. When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.
RECORDING_FORMAT	This field indicates the format that was used in recording the data on the volume. The following are the values that can be reported: ► 4: The volume is written in the 3592 J1A format. ► 5: The volume is written in the TS1120 E05 format. ► 6: The volume is written in the TS1120 E05E format (encrypted). Values not defined are reserved.
PCT_UTILIZED	This field indicates the percent of the volume that contains active data. The percentage is determined using the active_size and the total_bytes_written fields. The number is reported in 1/10th of a percent. This field is updated hourly (if needed). It is reset to 0 when the volume becomes empty (scratch).
TOTAL_BYTES_WRITTEN	This field indicates the number of bytes that were written to the volume when it was filled. It is reset to 0 when the volume becomes empty (scratch).
ACTIVE_SIZE	This field indicates the number of bytes of the active data on the volume. It is reset to 0 when the volume becomes empty (scratch).
ERROR_TYPE_FLAG	This field indicates whether the TS7700 has detected an error with the use of the volume. The error values are used by IBM service.
SCRATCH_COUNT	This field indicates the number of times the volume has been returned to scratch status since it was inserted into the library.
MOUNT_COUNT	This field indicates the number of times the volume has been mounted since it was inserted into the library.
MEDIA_TYPE	This field indicates the media type for the volume. The following are the values that can be reported: ► 5: IBM JA ► 6: IBM JJ ► 7: IBM JB Values not defined are reserved.

Field Name	Description
BATE_ROR_REASON	If the volume has a volume access value of Read-Only, this field provides additional information about the reason for the access value. A volume might be marked as Read-Only for reasons other than a media error. The following are the values that can be reported: ► 0: BTE_ROR_BATE_NONE. ► 1: BTE_ROR_BATE_RECLAIM. ► 2: BTE_ROR_BATE_MOVE_PEND ► 3: BTE_ROR_BATE_ADSMEXIT ► 4: BTE_ROR_BATE_EJECT_STACKED ► 5: BTE_ROR_BATE_PROTECTED_EMPTY_VOL ► 6: BTE_ROR_BATE_DEFERRED ► 7: BTE_ROR_BATE_IN_PROGRESS ► 8: BTE_ROR_BATE_FF08 ► 9: BTE_ROR_BATE_EJECT_CLEANUP ► 10: BTE_ROR_BATE_ERASE ► 11: BTE_ROR_BATE_PFE_SET ► FF: BTE_ROR_BATE_UNKNOWN Values not defined are reserved.
DVM_ROR_REASON	If the volume has a volume access value of Read-Only, this field provides additional information about the reason for the access value. A volume might be marked as Read-Only for reasons other than a media error. The following are the values that can be reported: ► 0: BTE_ROR_ADSM_NONE. ► 1: BTE_ROR_ADSM_DBRESTORE. ► 2: BTE_ROR_ADSM_DBBACKFAIL. ► 3: BTE_ROR_ADSM_DEFAULT ► 4: BTE_ROR_ADSM_8447_STRIP ► 5: BTE_ROR_ADSM_8447_NO_STRIP ► FF: BTE_ROR_ADSM_UNKNOWN Values not defined are reserved.
ERASE_FLAG	This field indicates whether or not the volume needs to be secure data erased.
MOUNT_FAILURE_COUNT	This field indicates the number of time a mount failure occurred with this volume since it was inserted into the library.
LAST_DEVICE_MOUNTED	This field contains the device number the volume was last mounted on.
LAST_WRITE_TIMESTAMP	Timestamp of when the volume was last written to. The format of the timestamp is: Year-Month-Day-Hour.Minute.Second.Microsecond for example: 2006-05-23-19.34.23.876129 When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.
LAST_READ_TIMESTAMP	Timestamp of when the volume was last read from. The format of the timestamp is: Year-Month-Day-Hour.Minute.Second.Microsecond for example: 2006-05-23-19.34.23.876129 When the volume is inserted into the library, this field is set to 1970-01-01-00.00.00.000000.

Field Name	Description
KEKLABEL1	If the volume has been encrypted, this field contains the first key label. A key label can be up to 64 characters in length. If the volume is not encrypted, this field is null. It is set to null when the volume becomes empty (scratch).
KEKLABEL2	If the volume has been encrypted, this field contains the second key label. A key label can be up to 64 characters in length. If the volume is not encrypted, this field is null. It is set to null when the volume becomes empty (scratch).
INSERT_TIMESTAMP	Timestamp of when the volume was inserted into the library. The format of the timestamp is: Year-Month-Day-Hour.Minute.Second.Microsecond for example: 2006-05-23-19.34.23.876129
ADSM_FORMAT_FLAG	This field indicates whether the physical volume is written using the pre-TS7700(B10/b20 VTS) or TS7700 format. The following are the values that can be indicated: ► 0: The volume is written in the TS7700 format. ► 1: The volume is written in thepre-TS7700 format.
Blanks	As needed to pad a record to 400 bytes

Cache content information

Volumes accessed by a host are maintained in the tape volume cache managed by each cluster. The cache might be partitioned in later code releases. The TS7700 Virtualization Engine controls the movement of logical volumes out of a cache partition as space is needed for newly created or recalled volumes for that partition.

The primary goal of the cache management algorithms in the TS7700 is to maximize the utilization of its cache for volumes that have some likelihood to be accessed again. The cache management function of the TS7700 arranges the volumes in a cache partition in the anticipated order they are to be removed when space is needed. In order to remove a volume from cache it must first have been pre-migrated (which means copied to a physical tape).

For this reason, it is possible that volumes with a higher order number are removed from cache first. As part of the Outboard Policy Management functions of the TS7700, the Storage Class construct provides for customer control of the partition for a volume's data and cache preferencing policies for the management of the volume in cache. Two preferencing policies are supported:

Preference Group 0 (PG0)

When space is needed in the cache, pre-migrated volumes assigned to preference group 0 are removed from cache before volumes assigned to preference group 1. Within preference group 0, the volumes are ordered for removal from cache by largest volumes first.

> **Note:** Volumes assigned to preference group 0 might also be removed from the cache, independent of the need for cache space, as a background task within the TS7700 Virtualization Engine. For more details, see 2.3.2, “Tape Volume Cache Management” on page 36.

Preference Group 1 (PG1)

When space is needed in the cache and there are no pre-migrated preference group 0 volumes to remove, pre-migrated volumes assigned to preference group 1 are removed.

Within preference group 1, the volumes are ordered for removal from cache based on time since last access (LRU).

The order of removal of a volume from cache might also be influenced by other storage constructs settings for a volume, so do not rely on the order presented in the response data to be exact.

The contents of the cache associated with the specific cluster the request volume is written to are returned in the response records. In a TS7700 Multi Cluster Grid configuration, separate requests must be issued to each cluster to obtain the cache contents of all the clusters.

The response records are written in 80 byte fixed format.

Record 6

This record provides a heading for the data records to follow as shown in Example 8-7 and Table 8-8.

Example 8-7 BVIR heading for the data records for Cache Content information

```
ORDER VOLSER DATE/TIME IN CACHE    PG  PART
```

Table 8-8 Decoding heading for the data records for Cache Content information

Bytes	Name	Description
1-3		Blank characters
4-8		ORDER Heading for the volume order column
9	Field Delimiter	Blank character
10-15		VOLSER Heading for the volume serial number column
16	Field Delimiter	Blank character
17-34		DATE/TIME IN CACHE Heading for the data/time stamp columns
35-38	Field Delimiter	Blank character
39-40		PG Heading for the preference group column
41-42	Field Delimiter	
43-46		PART Heading for the partition column
47-80	Blanks	Blank character padding

Record 7-N

Each of these records provides information for one logical volume. The cache content information is provided by partition, then in the order in which the volumes are anticipated to be removed from the cache partition. Example 8-8 and Table 8-9 show an example.

Example 8-8 BVIR data records for cache content information

```
12345678901234567890123456789012345678901234567890123456
       1 VOL020 11/30/2005 11:57:00    0     0
       2 VOL019 11/29/2005 03:00:00    0     0
       3 VOL023 11/20/2005 09:57:00    1     0
       4 VOL016 11/20/2005 10:01:00    1     0
       1 ABC309 11/20/2005 17:31:00    0     1
       2 ABC333 11/20/2005 11:44:00    1     1
       3 ABC789 11/20/2005 10:51:00    1     1
       4 ABC234 11/20/2005 04:27:00    1     1
       5 ABC045 11/19/2005 21:45:00    1     1
```

Table 8-9 Decoding for the data records for Cache Content information

Bytes	Name	Description
1-8	Order in Cache Partition	The order in which volumes are to be removed from the cache partition, right justified and blank filled.
9	Field Delimiter	Blank character
10-15	Logical Volser	6 character volume serial number
16	Field Delimiter	Blank characters
17-26	Date in Cache	The date the volume was created or recalled into the cache in the format of MM/DD/YYYY.
27	Field Delimiter	Blank character
28-35	Time in Cache	The time the volume was created or recalled into the cache in the format of HH:MM:SS.
36-39	Field Delimiter	Blank characters
40	Preference Group	Contains the preference group the volume is assigned.
41-45	Field Delimiter	
46	Partition	Contains the partition the volume is resident in
47-80	Blanks	Blank character padding

Note: The contents of the cache typically are all private volumes; however, it is possible that some might have been returned to scratch status soon after being written. The TS7700 does not filter the cache contents based on the private or scratch status of a volume.

Physical volume to logical volume mapping information

The TS7700 Virtualization Engine maintains the mapping between logical and physical volumes in a database on each cluster. It is possible that there are inconsistencies in the mapping information provided with this function. This results when a logical volume is being moved from one physical volume to another. For a period of time, the volume is shown on more than one physical volume. This can result in a small number of logical volumes reported as being on physical volumes which they were located on in the past, but presently are not.

Even with some inconsistencies, the mapping data is useful to customers that want to design jobs that recall data efficiently off of physical volumes. If the logical volumes reported on a physical volume are recalled together, the efficiency of the recalls will be increased. If a logical volume with an inconsistent mapping relationship is recalled, it will recall correctly, but an additional mount of a different physical volume might be required.

The physical volume to logical volume mapping associated with the physical volumes managed by the specific cluster the request volume is written to are returned in the response records. In a TS7700 Multi Cluster Grid configuration, separate requests must be issued to each cluster to obtain the mapping for all physical volumes.

The response records are written in 80 byte fixed format.

Record 6

This record provides a heading for the data records to follow (Example 8-9 and Table 8-10).

Example 8-9 BVIR heading for the data records for Volume Mapping information

```
123456789012345678901234567890123456789012345678
PHYSICAL  LOGICAL P/B  ORDER   PART        SIZE
```

Table 8-10 Decoding heading for the data records for Cache Content information

Bytes	Name	Description
1-8		PHYSICAL Heading for the physical volume serial number column
9-10	Field Delimiter	Blank characters
11-17		LOGICAL Heading for the logical volume serial number column
18	Field Delimiter	Blank characters
19-21		P/B Heading for the primary or backup pool indicator column
22-23	Field Delimiter	Blank characters
24-28		ORDER Heading for the logical volume order column
29-31	Field Delimiter	Blank characters
32-35		PART Heading for the logical volume spanning indicator column
36-44	Blanks	Blank characters
45-48		SIZE Heading for the logical volume compresses size
49-80	Blanks	Blank character padding

The primary and backup pool indicator column indicates whether the volume being reported resides on a primary pool volume or the secondary pool volume as defined using the Selective Dual Copy function that is part of the TS7700 Virtualization Engine's advanced policy management function.

The size field reports the number of MBs, rounded to two places after the decimal point, a logical volume occupies on the physical volume. This includes the effect of the compression performed on the data by the TS7700, but does not include any effect of the compression performed by the physical drive. Any volume with a size of less than 5 KB reports a size of

0.00. When a volume indicates that it spans, the size shown represents the amount of the logical volume data that is stored on each physical volume.

Note: Only data created on prior generation B10 or B20 VTSs and migrated under the control of a TS7700 Virtualization Engine will indicate that they span from one physical volume to another. Volumes created or recopied by the TS7700 *do not span* physical volumes.

Record 7-N

Each of these records provides information for a logical to physical volume association. The records are ordered alphanumerically (0-9, A-Z) by physical VOLSER, then by logical sequence of the active logical volumes on the physical volume. Example 8-10 provides an example and Table 8-11 shows the decoding.

Example 8-10 BVIR data records for Volume Mapping information

```
1234567890123456789012345678901234567890123456789O1
P00024    GK0000   P  000001 1 OF 1        23.45 M
P00024    GK0020   P  000002 1 OF 1        76.50 M
P00024    GK0010   P  000003 1 OF 1       145.70 M
P00024    GK0030   P  000004 1 OF 1       670.32 M
P00024    GK0040   P  000005 1 OF 1      1934.12 M
P00024    GK0060   P  000006 1 OF 1         0.00 M
P00024    GK0050   P  000007 1 OF 2       540.12 M
P00467    GK0050   P  000001 2 OF 2       540.12 M
```

Table 8-11 Decoding for the data records for Volume Mapping information

Bytes	Name	Description
1-6	Physical Volser	Physical VOLSER the logical VOLSER is located on, left justified and padded with blanks.
7-10	Field Delimiter	Blank character
11-16	Logical Volser	Logical VOLSER, left justified and padded with blanks.
17-19	Field Delimiter	Blank characters
20	Pool Indicator	P indicates the logical volume is the primary copy, B indicates that the volume is the backup copy.
21-22	Field Delimiter	Blank character
23-28	Order	The relative order of the logical volume on the physical volume.
29	Field Delimiter	Blank characters
30-35	Spanning	Indicates whether the logical volume spans to another physical volume or not and if so, which part.
36-41	Blanks	Blank characters
42-45	Size	Integer part of the size, right justified and padded with leading blanks
46	Decimal Point	"."
47-48	Hundredths	Decimal part of the size
49	Field Delimiter	Blank character

Bytes	Name	Description
50	Unit Indicator	M indicates the size is in megabytes, which is 1024x1024 bytes
51-80	Blanks	Blank character padding

Point-in-time statistics

A TS7700 Virtualization Engine is continually logging information regarding the activities within it. The logged information is referred to as statistical information and is recorded in two forms, Point-in-time and Historical. Point-in-time statistics indicate the state and operation aspects of the TS7700 over a short interval of time. The time interval is currently approximately 15 seconds. A request for Point-in-time statistics will respond with the data accumulated in the interval completed just prior to the request being processed. Because of this, the state information reported might lag the actual state of the TS7700 by an interval.

Other than an information header, Point-in-time statistics are provided in a mixture of character and binary format fields. The format of the statistical records is defined in the *Virtualization Engine TS7700 Series Statistical Data Forma*t White Paper.

The Point-in-time statistics for all clusters are returned in the response records. In a TS7700 grid configuration, this means that the request volume can be written to any cluster to obtain the information for the entire configuration.

> **Note:** The request records are written in F format. To read the response records, use the Undefined (U) format with a maximum blocksize of 24000. The response records are variable in length.

Record 6

This record provides a human readable header for the response record. (Example 8-11 and Table 8-12).

Example 8-11 BVIR heading for the data records for Point-in-time statistics

```
POINT IN TIME STATISTICS
```

Table 8-12 Decoding heading for the data records for Point-in-time statistics

Bytes	Name	Description
1-24		POINT IN TIME STATISTICS
25-80	Blanks	Blank character padding

Record 7-N

Each of these records provides Point-in-time information for the nodes of the TS7700 Cluster. Character data is encoded using EBCDIC and non-character data is in binary format. Records vary in length. The first four bytes of the record identify the length, version, and data type of the record.

Because the Point-in-time information is provided for all nodes in the configuration, including nodes in all clusters in a grid configuration, the number of records returned depends on the configuration. There are two records for each vNode and two records for each hNode. The application processing the response volume records should not assume that the records for a vNode or hNode are presented in any specific order.

Table 8-13 lists the decoding of the response command.

Table 8-13 Decoding the data records for Point-in-time statistics

Bytes	Name	Description
1-2	Length	Total length of the response record in binary
3	Version	Binary 1
4	Data Type	Identifies the type of response record data
5-length	Data	Response data

Data type

The statistics for a node in the subsystem are subdivided into different data types. Refer to the *Virtualization Engine TS7700 Series Statistical Data Format* White Paper.

Binary response data

This record contains the actual response data. Refer to the *Virtualization Engine TS7700 Series Statistical Data Format* White Paper for the format of the response data for each data type.

Historical statistics

A TS7700 Virtualization Engine is continually logging information regarding the activities within it. The logged information is referred to as statistical information and is recorded in two forms, Point-in-time and Historical. Historical statistics indicate the operational aspects of the TS7700 accumulated over a 15 minute interval. The data from each 15 minute interval is maintained and logged in the TS7700. A request for Historical statistics results in a response file that contains all of the data logged up to that point for the requested Julian day.

Other than an information header, Historical statistics are provided in character and binary format fields. The format of the statistical records are defined in the *Virtualization Engine TS7700 Series Statistical Data Format* White Paper.

The Historical statistics for all clusters are returned in the response records. In a TS7700 grid configuration, this means that the request volume can be written to any cluster to obtain the information for the entire configuration.

Note: The TS7700 Virtualization Engine retains 90 days worth of historical statistics. If you want to keep statistics for a longer period of time, it is recommended that you retain the logical volumes used to obtain the statistics.

The request records are written in F format. To read the response records, use the Undefined (U) format with a maximum blocksize of 24,000. The response records are variable in length.

Record 6-N

The Historical statistical response records for each day requested are preceded by the following human-readable header. If Historical statistics are requested for a day for which the TS7700 does not have historical data, the “No Historical Statistics for xxx” header is the only response record provided for the day (Example 8-12 and Table 8-14).

Example 8-12 BVIR heading for the data records for Historical statistics

```
HISTORICAL STATISTICS FOR xxx or NO HISTORICAL STATISTICS FOR xxx
```

Table 8-14 Decoding heading for the data records for Point-in-time statistics

Bytes	Name	Description
1-46		HISTORICAL STATISTICS for xxx or NO HISTORICAL STATISTICS FOR xxx, left justified, padded with blanks on the right.
47-80	Blanks	Blank character padding

Data response records

Each of these records provides historical information for the nodes of the TS7700 Cluster for the day indicated in the header record. Character data is encoded using EBCDIC and non-character data is in binary format. Records vary in length. The first four bytes of the record identify the length, version, and data type of the record. Because the historical information is provided for all nodes in the configuration, including nodes in all clusters in a grid configuration, the number of records returned depends on the configuration. There are two records for each vNode and from 4 to 7 records for each hNode (depending on the number of physical libraries attached to the cluster). The application processing the response volume records should not assume that the records for a V or H node are presented in any specific order. Below we provide the table for decoding the response command (Table 8-15).

Table 8-15 Decoding the data response for Historical statistics

Bytes	Name	Description
1-2	Length	Total length of the response record in binary
3	Version	Binary 1
4	Data Type	Identifies the type of response record data
5-length	Data	Response data

Data type

The statistics for a node in the subsystem are subdivided into different data types. Refer to the *Virtualization Engine TS7700 Series Statistical Data Format* White Paper for the description of the data types.

Binary response data

Refer to the *Virtualization Engine TS7700 Series Statistical Data Format* White Paper for the format of the response data for each data type.

Physical media pools

The TS7700 Virtualization Engine supports separating the physical volumes it manages into pools. The supported pools include a pool that contains scratch (empty) volumes that are common and up to 32 pools that might contain scratch (empty) and data (filling/full) volumes.

Pools can borrow and return volumes from the common scratch pool. Each pool can contain several types of media.

For pool 00 (common scratch pool), because it only contains empty volume, only the empty count is returned. Volumes that have been borrowed from the common pool are not included.

For pools 1-32, a count of the physical volumes that are empty, are empty and waiting for erasure, are in the process of being filled, or have been marked as full is returned. The count for empty includes physical volumes that have been specifically assigned to the pool as well as volumes that were borrowed from the common scratch pool but have not yet been returned. The count of volumes that are marked as Read Only or Unavailable (including destroyed volumes) are returned. Also, the full data volumes contain a mixture of valid and invalid data. Response records are provided for the distribution of active data on the data volumes marked as full for a pool.

Information is returned for the common pool and all other pools that are defined and have physical volumes associated with them.

The physical media pool information managed by the specific cluster the request volume is written to are returned in the response records. In a TS7700 grid configuration, separate requests must be issued to each cluster to obtain the physical media pool information for all clusters.

The response records are written in 80 byte fixed format. Counts are provided for each media type associated with the pool (up to a maximum of 8) and the response lines are formatted as follows:

Record 6

This record provides a heading for the data records to follow (Example 8-13 and Table 8-16).

Example 8-13 BVIR heading for the data records for Physical Media Pooling

```
R POOL MEDIA  EMPTY  FILLING   FULL  ERASE    ROR  UNAVAIL
```

Table 8-16 Decoding heading for the data records for Physical Media Pooling

Bytes	Name	Description
1		R Heading for the record type column
2	Field Delimiter	Blank character
3-6		POOL Heading for the pool number column
7	Field Delimiter	Blank character
8-12		MEDIA Heading for the media type column
13-14	Field Delimiter	Blank character
15-19		EMPTY Heading for the empty count column
20-21	Field Delimiter	Blank character
22-28		FILLING heading for the filling count column
29-31	Field Delimiter	Blank character
32-35		FULL Heading for the full count column
36-37	Field Delimiter	Blank character

Bytes	Name	Description
38-42		ERASE Heading for the waiting to erase count column
43-46	Field Delimiter	Blank characters
47-49		ROR Heading for the read only recovery count column
50-51	Field Delimiter	Blank characters
52-58		UNAVAIL Heading for the unavailable count column
59-80	Blanks	Blank character padding

Record 7-N

Each of these records provides the counts for a specific media type for each pool. Up to eight media types might be reported for a pool, as shown in Example 8-14.

Example 8-14 BVIR response for the data records for Physical Media Pooling

```
12345678901234567890123456789012345678901234567890123456789012345678

C   00    JA  00083
C   00    JB  00051
C   00    JJ  00002
C   01    JA  00134     00003  00609  00000  00000     00000
C   01    JB  00001     00002  00001  00006  00000     00000
```

Use Table 8-17 to decode the data records shown in Example 8-14.

Table 8-17 Decoding for the data records for Physical Media Pooling

Bytes	Name	Description
1	Record Type	Record type identifier. C for media count.
2-4	Field Delimiter	Blank characters
5-6	Pool Number	The pool number. Pool number is right justified and padded with a leading zero.
7-10	Field Delimiter	Blank characters
11-12	Media Type	The media type defined for the pool. Media types with a single character are padded on the left with a blank.
13-14	Field Delimiter	Blank characters
15-19	Empty Count	The count of the physical volumes that are empty for the media type. The count is right justified and padded with leading zeroes.
20-23	Field Delimiter	Blank characters
24-28	Filling Count	The count of the physical volumes that are in the filling state for the media type. The count is right justified and padded with leading zeroes. This field is all blanks for pool 00.
29-30	Field Delimiter	Blank characters
31-35	Full Count	The count of the physical volumes that have been marked full for the media type. The count is right justified and padded with leading zeroes. This field is all blanks for pool 00.

Bytes	Name	Description
36-37	Field Delimiter	Blank characters
38-42	Erase Count	The count of the physical volumes that have been reclaimed, but need to be erased before they become empty. The count is right justified and padded with leading zeroes. This field is all blanks for pool 00.
43-44	Field Delimiter	Blank characters
45-49	ROR Count	The count of the physical volumes that are in the read only recovery state.
50-53	Field Delimiter	Blank characters
54-58	Unavailable Count	The count of the physical volumes that are in the unavailable or destroyed state.
59-80	Blanks	Blank character padding

Record N+1

For pools 1-32, this record provides a heading for the data records to follow; see Example 8-15.

Example 8-15 BVIR response for the data records N+1 for Physical Media Pooling

```
R POOL MEDIA    0+   10+   20+   30+   40+   50+   60+   70+   80+   90+
```

Table 8-18 describes the layout of the record.

Table 8-18 Decoding for the data record N+1 for Physical Media Pooling

Bytes	Name	Description
1		R Heading for the record type column
2	Field Delimiter	Blank character
3-6		POOL Heading for the pool number column
7	Field Delimiter	Blank characters
8-12		MEDIA Heading for the media type column
13-16	Field Delimiter	Blank characters
17-18		0+ Heading for the 0-10% count column
19-21	Field Delimiter	Blank characters
22-24		10+ Heading for the 10-20% count column
25-27	Field Delimiter	Blank characters
28-30		20+ Heading for the 20-30% count column
31-33	Field Delimiter	Blank characters
34-36		30+ Heading for the 30-40% count column
37-39	Field Delimiter	Blank characters
40-42		40+ Heading for the 40-50% count column

Bytes	Name	Description
43-45	Field Delimiter	Blank characters
46-48		50+ Heading for the 50-60% count column
49-51	Field Delimiter	Blank characters
52-54		60+ Heading for the 60-70% count column
55-57	Field Delimiter	Blank characters
58-60		70+ Heading for the 70-80% count column
61-63	Field Delimiter	Blank characters
64-66		80+ Heading for the 80-90% count column
67-69	Field Delimiter	Blank characters
70-72		90+ Heading for the 90+% count column
73-80	Blanks	Blank character padding

Record N+2-M

Each of these records provides the active data distribution counts for a specific media type. Up to eight media types might be reported for a pool. Example 8-16 lists the first three pools.

Example 8-16 BVIR response for the data records N+2-M for Physical Media Pooling

```
12345678901234567890123456789012345678901234567890123456789012345678901234567890

D   01    JA 00000 00002 00034 00056 00092 00078 00084 00065 00195 00183
D   01    JB 00000 00000 00000 00000 00000 00000 00000 00002 00003 00001
```

Table 8-19 describes the layout of the records.

Table 8-19 Decoding for the data records N+2-M for Physical Media Pooling

Bytes	Name	Description
1	Record Type	Record type identifier.D for active media data distribution count.
2-4	Field Delimiter	Blank characters
5-6	Pool Number	The pool number, right justified and padded with a leading zero.
7-10	Field Delimiter	Blank characters
11-12	Media Type	The media type defined for the pool. Media types with a single character are padded on the left with a blank.
13	Field Delimiter	Blank character
14-18	0-9 Count	The count of the physical volumes that are marked full and have an active data percentage greater than 0 and less than 10. The count is right justified and padded with leading zeroes.
19	Field Delimiter	Blank character
20-24	10s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 10 and less than 20. The count is right justified and padded with leading zeroes.

Bytes	Name	Description
25	Field Delimiter	Blank character
26-30	20s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 20 and less than 30. The count is right justified and padded with leading zeroes.
31	Field Delimiter	Blank character
32-26	30s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 30 and less than 40. The count is right justified and padded with leading zeroes.
37	Field Delimiter	Blank character
38-42	40s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 40 and less than 50. The count is right justified and padded with leading zeroes.
43	Field Delimiter	Blank character
44-48	50s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 50 and less than 60. The count is right justified and padded with leading zeroes.
49	Field Delimiter	Blank character
50-54	60s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 60 and less than 70. The count is right justified and padded with leading zeroes.
55	Field Delimiter	Blank character
56-60	70s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 70 and less than 80. The count is right justified and padded with leading zeroes.
61	Field Delimiter	Blank character
62-66	80s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 80 and less than 90. The count is right justified and padded with leading zeroes.
67	Field Delimiter	Blank character
68-72	90s Count	The count of the physical volumes that are marked full and have an active data percentage greater than 90. The count is right justified and padded with leading zeroes.
73-80	Blanks	Blank character padding

Copy audit request

A database is maintained on each individual TS7740 Cluster that contains status information about the logical volumes defined to the grid. Two key pieces of information are whether the cluster contains a valid copy of a logical volume and whether the copy policy for the volume indicates that it should have a valid copy.

This request performs an audit of the databases on a set of specified TS7700 distributed libraries to determine if there are any volumes that do not have a valid copy on at least one of them. If you specify the COPYMODE option, whether the volume is supposed to have a copy on the distributed library is taken into account in determining whether that distributed library has a valid copy. If COPYMODE is specified and the copy policy for a volume on a specific cluster is R or D, then that cluster is considered during the audit. If COPYMODE is specified and the copy policy for a volume on a specific cluster is N, then the volume's validity state is ignored because that cluster does not need to have a valid copy. The request then returns a list of any volumes that do not have a valid copy, subject to the copy mode if the COPYMODE option is specified, on the TS7700s clusters that are specified as part of the request.

The specified clusters might not have a copy for several reasons:

- The copy policy associated with the volume did not specify that any of the clusters specified in the request were to have a copy, and the COPYMODE option was not specified. This might be because of a mistake in defining the copy policy or it might be really intended. For example, volumes used in a disaster recovery test need to reside only on the disaster recovery TS7700 and not on the production TS7700s. If the request specified only the production TS7700s, all of the volumes that are used in the test are returned in the list.
- The copies have not yet been made from a source TS7700 to one or more of the specified clusters. This could be because the source TS7700 or the links to it are unavailable or because a copy policy of deferred was specified and a copy had not been completed when the audit was performed.

The Copy Audit request is intended to be used for the following circumstances:

- A TS7700 is to be removed from a Grid configuration. Prior to its removal you want to ensure that the TS7700s that are to remain in the Grid configuration have a copy of all important volumes that were created on the TS7700 that is to be removed.
- A condition has occurred (could be a site disaster or as part of a test procedure) where one of the TS7700s in a Grid configuration is no longer available and you want to determine which, if any, volumes on the remaining TS7700s do not have a valid copy.

In the Copy Audit request, you need to specify which TS7700 clusters are to be audited. You specify the clusters using their associated distributed library ID (this is the 5 character library sequence number that was defined when the TS7700 cluster was installed). If you specify more than one distributed library ID, separate them when a comma. Use the following rules to determine which TS7700 clusters to include in the audit:

- Specify the INCLUDE parameter to include all specified distributed library IDs in the audit. All clusters that are associated with these IDs must be available, or the audit will fail.
- Specify the EXCLUDE parameter to exclude all specified distributed library IDs from the audit. All other clusters in the Grid configuration must be available, or the audit will fail.
- Specify distributed library IDs to check for the validity of those IDs in the Grid configuration. If one or more of the specified distributed library IDs are invalid, the Copy Audit fails, and the response indicates the IDs that are considered invalid.
- Specify distributed library IDs, or the Copy Audit will fail.

The following examples show valid requests (these examples assume a Three-Cluster Grid configuration with distributed library IDs of BA45A, BA45B, and BA45C):

- COPY AUDIT INCLUDE BA45A

  Audits the copy status of all volumes only on the cluster that is associated with distributed library ID BA45A.
- COPY AUDIT COPYMODE INCLUDE BA45A

  Audits the copy status of volumes that also have a valid copy policy only on the cluster that is associated with distributed library ID BA45A.

- COPY AUDIT INCLUDE BA45B,BA45C

  Audits the copy status of volumes on the clusters that are associated with distributed library IDs BA45B and BA45C.
- COPY AUDIT EXCLUDE BA45C

  Audits the copy status of volumes on the clusters in the Grid configuration that are associated with distributed library IDs BA45A and BA45B.

When the audit completes, a response record is written for each logical volume that did not have a valid copy on any of the specified clusters. Volumes that have never been used, that have had their associated data deleted, or that have been returned to scratch are not included in the response records. The record includes the volume serial number and the copy policy definition for the volume. The volser and the copy policy definitions are comma separated, as shown in the following example:

```
L00001,R,D,D,N,N,N,N,N
```

Note: The output for copy audit includes copy consistency points for up to eight TS7700 clusters to provide for future expansion of the number of clusters that are supported in a TS7700 Grid to the architected maximum.

Copy Audit might take more than an hour to complete, depending on the number of logical volumes that are defined, how many clusters are configured in the Grid configuration, and how busy the TS7700s are at the time of the request.

Record 6

This record lists the distributed libraries IDs by their 5 character library sequence number separated by a comma. The distributed library IDs that are listed are either the IDs used for a successful Copy Audit operation or the IDs with which the TS7700 had a problem. For example:

```
BA45A,BA45B
```

Record 7-N

When there are logical volumes that do not have a copy on any of the specified TS7700 clusters, each of these records provide information for each volume. Table 8-20 lists the field descriptions.

Table 8-20 Decoding of Record 7-N for Copy Audit

Field name	Description
VOLSER	6 character volume serial number
cluster0_copy_mode	This field indicates whether cluster 0 is to have a copy of the volume and the copy consistency point defined for the volume. The values are: ► R: Rewind unload (RUN) copy consistency point. ► D: Deferred copy consistency point. ► N: No copy ► E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.

Field name	Description
cluster1_copy_mode	This field indicates whether cluster 1 is to have a copy of the volume and the copy consistency point defined for the volume. The values are: ► R: Rewind unload (RUN) copy consistency point. ► D: Deferred copy consistency point. ► N: No copy ► E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
cluster2_copy_mode	This field indicates whether cluster 2 is to have a copy of the volume and the copy consistency point defined for the volume. The values are: ► R: Rewind unload (RUN) copy consistency point. ► D: Deferred copy consistency point. ► N: No copy ► E: The volume was previously assigned a copy consistency point of rewind unload or deferred, but was changed to no copy.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.
reserved	This field is reserved for future use and is set to N.

If the audit finds no invalid logical volumes on the specified TS7700 clusters, the following record is returned.

```
NO INVALID VOLUMES ON SPECIFIED CLUSTERS
```

If the TS7700 encounters errors or operational conditions while processing the copy audit request, the audit cannot be performed, and a record is returned, depending on the condition. The message starts at Byte 1 of the record and is followed with padding blanks to Byte 80. The following records can be returned based on the error condition:

- `CLUSTERS NOT IN VALID STATE`

  One or more of the clusters that are associated with the specified distributed library IDs are not available because they are in service or offline. Record 6 lists the distributed library IDs specified that are associated with the unavailable clusters.

- `NO DISTRIBUTED IDS SPECIFIED`

  No distributed library IDs are included in the request. At least one is required. Record 6 will be all blanks.

- `DISTRIBUTED ID NOT VALID`

  One or more of the specified distributed library IDs are not defined to the Grid configuration. Record 6 lists the distributed library IDs specified that are not defined.

- `ALL CLUSTERS HAVE BEEN EXCLUDED`

  The EXCLUDE parameter was specified and the distributed library IDs specified includes all defined in the Grid configuration. Record 6 lists the distributed library IDs specified.

- `INCLUDE OR EXCLUDE MUST BE SPECIFIED`

  Neither the INCLUDE or the EXCLUDE parameter was specified.

- `GLOBAL RESOURCE CURRENTLY BUSY - TRY AGAIN LATER`

  For copy audit to be performed, it requires that no other major database update functions are being performed. If the TS7700 Grid is in the process of inserting logical volumes, performing a copy export operation, or is already executing a copy export operation.

- `INTERNAL ERROR CAUSED NO RESPONSE`

  An internal error in the Grid resulted in no response being generated.

Unknown or Invalid Request

If the request file does not contain the correct number of records or if the first record is incorrect, the request file on the volume is unchanged, and no error is indicated.

If the request file contains the correct number of records and if the first record is correct but the second is not, the response file indicates in Record 6 that the request is unknown as follows:

```
UNKNOWN REQUEST TYPE
```

The response starts at Byte 1 of the record and is followed with padding blanks to Byte 80.

If the request file contains the correct number of records, if the first record is correct, and if the second is recognized but includes a variable that is not within the range supported for the request, the response file indicates in record 6 that the request is invalid as follows:

```
INVALID VARIABLE SPECIFIED
```

The response starts at Byte 1 of the record and is followed with padding blanks to Byte 80.

8.7 Using VEHSTATS for monitoring and reporting

In this section, we show how to work further with the binary reports for Point-in-time statistics as well the Historical statistics. You have collected these with the BVIR functions described earlier in 8.6, “Bulk Volume Information Retrieval (BVIR)” on page 442.

To convert the binary response record from historical BVIR data, you can use the IBM-provided tool VEHSTATS. IBM also provides the white paper with the record layout of the binary BVIR response data which you can use to decode the binary file or which you might use to program your own tool for creating statistical reports. The VEPSTATS job is used to process the point-in-time statistics.

8.7.1 Tools download

You have access to the IBM Tape Tools library, which contains various tools that help you analyze your tape environment. This set of tools also includes the VEHSTATS tool, which uses the BVIR reports for a comprehensive performance analysis.

Figure 8-27 shows part of the tools catalog as accessible by customers. To obtain VEHSTATS, download the ibmtools.exe file. See the ibmtools.txt file for information about how to download the tools.

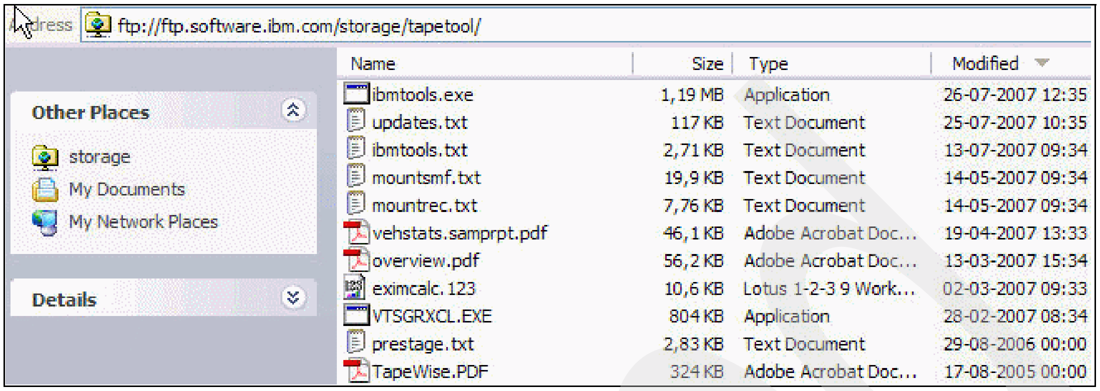

Figure 8-27 Tape Tools catalog

You can locate all the TS7700 monitoring and evaluation tools that we discuss in this chapter at the following site:

`ftp://ftp.software.ibm.com/storage/tapetool/`

The overview.pdf file provides a complete list of all tape tools that are available. Figure 8-28 shows one page of this document as an example.

IBM Tape Analysis Tools Overview April 18, 2008

ftp://ftp.software.ibm.com/storage/tapetool <=== as IBMTOOLS.EXE

Tool	Major Use	Benefit	Inputs	Outputs	Product pre-reqs
BADBLKSZ	Identify small VTS blksizes	Improve VTS performance, make jobs run faster	Logrec MDR & CA1, TLMS, RMM, ZARA,CTLT	Volser, Jobname, Dsname for VTS volumes with small blksizes	none
BIRTHDST	Find old cartridges	Improve data reliability	CA1, TLMS, RMM, ZARA, CTLT	List of active data on old carts and cartridge birth date distribution	none
BVIRHIST	Get stats from TS7740	U, VB, or to SMF	TS7740	Statistics file	TS7740
BVIRRPT	Identify VTS virtual volumes by owner	Determine which applications or users have virtual volumes	BVIR data & CA1, TLMS, RMM, ZARA, CTLT	Logical volumes by jobname or dsname, logical to physical rpts	VTS 7.4 for size
BVPITRPT	Point in Time stats as WTO	Immediately available	TS7740	Point in Time stats displayed as WTO	TS7740
CHGDATE	Change dates in SMF 14,15,21,30,40 records	Allows analyst to simulate combining workloads from two different periods.	SMF 14,15,21,30,40	Modified SMF records	none
CHKDUPS	Identify duplicate volsers	Identify duplicate volsers across multiple tape management catalogs	CA1, TLMS, RMM, ZARA, CTLT	List of duplicate volsers	none
CRTDIST	Show volumes and GBs created over time	Determine customer use patterns	CA1, TLMS, RMM, ZARA, CTLT	Distribution of create dates	none
DCOLLRPx	Report number of migrate or backup data sets	Identify size of TMM data if considering VTS	DFHSM MCDS file	Number of files and GBs migrated and dsns not recalled since damonyr	none
DIFFEXP	Identify multi-file volumes with different expiration dates	Prevent single file from not allowing volume to return to scratch	CA1, TLMS, RMM, ZARA, CTLT	List of files not matching file 1 expiration date	none
EXPDIST	Quantify number of volumes expiring n days from now	Determine rate of return to scratch	CA1, TLMS, RMM, ZARA, CTLT	Volume count distribution by media	none
EXPORT	Quantify VTS export/import processing times	Determine if E/I is viable	CA1, TLMS, RMM, ZARA, CTLT	Projected off-site volumes, MBs, amount of time for E/I process	none

Figure 8-28 Overview of the tools on the Tapetool FTP site

8.7.2 VEHSTATS tool overview

The IBM TS7700 Virtualization Engine's activity is recorded in the subsystem. There are two different types of statistics:

- Point-in-time statistics - Snapshot™ about activity in the last 15 seconds
- Historical statistic - Up to 90 days in 15 minute increments

Both can be obtained using the BVIR functions (for more information see 8.6, "Bulk Volume Information Retrieval (BVIR)" on page 442).

Because both types of statistical data are delivered in binary format from the BVIR functions, you need to translate the content into a readable format. You either can do this manually using the information provided in the *Virtualization Engine TS7700 Series Statistical Data Format* White Paper, or you can use an existing automated tool.

IBM provides a tool called VEHSTATS. Like the other IBM tape tools, the program is provided as-is, without official support, for the single purpose of showing how the data might be reported. There is no guarantee of its accuracy, and there is no additional documentation available for this tool. We provide some guidance for interpretation of the reports in 8.7.3, "VEHSTATS reports" on page 478.

The ibmtools.exe self-extracting zipped file expands into three separate files:

IBMJCL.BIN Execution JCL for current tape analysis tools

IBMCNTL.BIN Parameters needed for job execution, but do *not* need to be modified by the user

IBMLOAD.BIN Load library for executable load modules

IBMPAT.BIN Data pattern library, only needed if you are running the QSAMDRVR utility

If your VEHSTATS job stops running because the expiration date has passed, you need to download ibmtools.exe again to get the latest IBMTOOLS.JCL(EXPIRE) member. The file UPDATES.TXT contains all fixes and enhancements made to the tools. Our suggestion is that you look at this file regularly to see whether any programs that you run have been modified.

In addition to the VEHSTATS tool, sample BVIR jobs are included in the IBMTOOLS libraries. These jobs help you obtain the input data from the TS7700 Virtualization Engine. With these jobs, you can control where the historical statistics are accumulated for long term retention. The TS7700 still maintains historical statistics for the previous 90 days but you can have the pulled statistics recorded directly to the SMF log file or continue to use the disk flat file method. The flat files can be recorded as either RECFM=U or RECFB=VB.

There are three specific jobs in IBMTOOLS.JCL designed to fit your particular needs:

BVIRHSTS To write statistics to the SMF log file

BVIRHSTU To write statistics to a RECFM=U disk file

BVIRHSTV To write statistics to a RECFM=VB disk file

The VEHSTATS reporting program accepts any or all of the different formats. You define which input is to be used through a DD statement in the VEHSTATS job. The three input DD statements are optional, but at least one of the statements shown in Example 8-17 must be specified.

Example 8-17 VEHSTATS Input DD statements

```
//* ACTIVATE ONE OR MORE OF THE FOLLOWING DD STATEMENTS FOR YOUR DATA
//*STATSU  DD DISP=SHR,
//*           DSN=&USERHLQ..#&VTSID..BVIRHIST.D070205.D070205
//*STATSVB DD DISP=SHR,
//*           DSN=&USERHLQ..#&VTSID..BVIRHIST.D070206.D070206
//*STATSMF DD DISP=SHR,      RECORDS WILL BE SELECTED BASED ON SMFNUM
//*           DSN=&USERHLQ..#&VTSID..SMF194
```

The SMF input file can contain all SMF record types kept by the user. The SMFNUM parameter defines which record number is processed when you specify the STATSMF statement.

You can use VEHSTATS to monitor TS7700 drive and Tape Volume Cache (TVC) activity and to do trend analysis reports, based on BVIR binary response data. The tool summarizes TS7700 activity on a specified time basis up to 90 days in time sample intervals of 15 minutes or 1 hour, depending upon the data item. VEHSTATS reports the following information:

- Virtual drive activity
  - Number of drives allocated
  - Number of mounts (fast ready, read hits, read misses)

- Physical drive activity
  - Number of drives allocated
  - Mounts for recall, pre-migration, reclaim
- Tape volume cache activity
  - MBs read and written to Tape Volume Cache
  - MBs read and written to physical tape
  - Virtual volume size
  - Logical volumes managed
- Recall and write overrun statistics
- Percentage of throttling and average throttle value
- Number of scratch stacked volumes (available cartridge storage)
- Number of stacked private volumes (cartridges containing active data)
- Active cartridge data distribution
  - 20 values with number of volumes with 0-5%, 5-10%, and so forth of active data
  - Reclaim threshold
- Min/Max/Avg mount times for fast-ready, read hits and read misses
- Compression ratio achieved in the Host Adapter cards
- Compression ratio achieved between TVC and Drive
- Block sizes of data written to TS7700
  - Average channel block size
  - Number of blocks written
  - Distribution of block sizes
    - 0 to 2 K, 2 K to 4 K, 4 K to 8 K, 8 K to 16 K,
    - 16 K to 32 K, 32 K to 64 K, greater than 64 K
- Reports by physical volume pool
- Reports by cache preference level

8.7.3 VEHSTATS reports

VEHSTATS can be used to monitor TS7700 Drive and Tape Volume Cache (TVC) activity and for doing trend analysis to see where the bottlenecks are and whether an upgrade, such as, for example, adding additional physical tape drives might improve the overall performance of the TS7700 Virtualization Engine. VEHSTATS is not a projection tool, but it provides the basis for an overall health check of the TS7700.

The following reports are available with initial release of the TS7700:

- Virtual drive activity
- vNode Adapters 0/1 activity
- vNode Adapters 2/3 activity
- Cache Partition 1 Activity
- Cache Partition 2 Activity
- Import/Export Activity (reserved for future use)
- Physical Tape Types 1/2 Activity
- Physical Tape Types 3/4 Activity
- Common Scratch Pool
- General Use Pools 01/02

- General Use Pools 03/04
- General Use Pools 05/06
- General Use Pools 07/08
- General Use Pools 09/10
- General Use Pools 11/12
- General Use Pools 13/14
- General Use Pools 15/16
- General Use Pools 17/18
- General Use Pools 19/20
- General Use Pools 21/22
- General Use Pools 23/24
- General Use Pools 25/26
- General Use Pools 27/28
- General Use Pools 29/30
- General Use Pools 31/32
- Distributed Library Statistics

Virtual Device Activity

Example 8-18 lists the report for Virtual Drive activity. This report gives an overview per 15 minute interval of the taken time frame about:

- The min/avg/max mounted virtual drives
- The amount of channel blocks written based on blocksize

Note that the report provides the view for all distributed libraries.

Example 8-18 VEHSTATS report for Virtual Drive Activity

```
1(C) IBM  REPORT=H20VIRT (07215)       VNODE VIRTUAL DEVICE HISTORICAL RECORDS       RUN ON 15AUG2007 @ 11:37:40
PAGE   1
 GRID#=00000  DIST_LIB_ID= 0  VNODE_ID= 0  NODE_SERIAL=78-13840  VE_CODE_LEVEL=008.003.000.0098
UTCMINUS=07
  19JUL07  -VIRTUAL_DRIVES-
   RECORD      --MOUNTED--    MAX    ---------CHANNEL_BLOCKS_WRITTEN_FOR_THESE_BLOCKSIZES-----------------------------
     TIME  INST MIN AVG MAX  THRPUT     <=2048      <=4096      <=8192     <=16384     <=32768     <=65536      >65536
  4:15:00   256 114 124 127    MAX         630           0           0           0     2485298           0           0
  4:30:00   256 117 125 127    MAX         631           0           0           0     2026062           0           0
  4:45:00   256 113 124 127    MAX         530           0           0           0     2620099           0           0
  5:00:00   256 117 125 127    MAX         474           0           0           0     3118714           0           0
  5:15:00   256 118 125 127    MAX         606           0           0           0     3199644           0           0
  5:30:00   256 117 125 127    MAX         385           0           0           0     2069711           0           0
  5:45:00   256 112 124 127    MAX         635           0           0           0     2625621           0           0
  6:00:00   256 121 125 127    MAX         513           0           0           0     1764645           0           0
  6:15:00   256 120 126 127    MAX         410           0           0           0     2977421           0           0
  6:30:00   256 115 125 127    MAX         625           0           0           0     3100414           0           0
  6:45:00   256 115 124 127    MAX         382           0           0           0     2004155           0           0
  7:00:00   256 119 125 127    MAX         635           0           0           0     1448850           0           0
1(C) IBM  REPORT=H20VIRT (07215)       VNODE VIRTUAL DEVICE HISTORICAL RECORDS       RUN ON 15AUG2007 @ 11:37:40
PAGE   2
 GRID#=00000  DIST_LIB_ID= 1  VNODE_ID= 0  NODE_SERIAL=78-16000  VE_CODE_LEVEL=008.003.000.0098
UTCMINUS=07
  19JUL07  -VIRTUAL_DRIVES-
   RECORD      --MOUNTED--    MAX    ---------CHANNEL_BLOCKS_WRITTEN_FOR_THESE_BLOCKSIZES-----------------------------
     TIME  INST MIN AVG MAX  THRPUT     <=2048      <=4096      <=8192     <=16384     <=32768     <=65536      >65536
  4:15:01   256 115 124 127    MAX         860           0           0           0     3244251           0           0
  4:30:01   256 116 125 127    MAX         736           0           0           0     3199442           0           0
  4:45:01   256 116 125 127    MAX         651           0           0           0     4845938           0           0
  5:00:01   256 114 125 127    MAX         876           0           0           0     4240055           0           0
  5:15:01   256 116 125 127    MAX         700           0           0           0     3480676           0           0
  5:30:01   256 121 126 127    MAX         512           0           0           0     2972331           0           0
  5:45:01   256 116 125 127    MAX         600           0           0           0     2219093           0           0
  6:00:01   256 115 125 127    MAX         940           0           0           0     4374096           0           0
```

```
  6:15:01   256 114 124 127    MAX          760           0           0           0     3321785           0           0
  6:30:01   256 116 123 127    MAX          752           0           0           0     3372917           0           0
  6:45:01   256 110 125 127    MAX          689           0           0           0     4411612           0           0
  7:00:01   256 115 125 127    MAX          635           0           0           0     1979149           0           0
1(C) IBM   REPORT=H20VIRT  07215)        VNODE VIRTUAL DEVICE HISTORICAL RECORDS        RUN ON 15AUG2007 @ 11:37:40
PAGE   3
 GRID#=00000   DIST_LIB_ID= 2  VNODE_ID= 0  NODE_SERIAL=78-12345  VE_CODE_LEVEL=008.003.000.0098
UTCMINUS=07
  19JUL07  -VIRTUAL_DRIVES-
   RECORD       --MOUNTED--    MAX    ---------CHANNEL_BLOCKS_WRITTEN_FOR_THESE_BLOCKSIZES------------------------------
     TIME   INST MIN AVG MAX  THRPUT      <=2048      <=4096      <=8192     <=16384     <=32768      <=65536      >65536
  4:15:05   256 112 125 128    MAX          646           0           0           0     3555376           0           0
  4:30:05   256 111 126 128    MAX          683           0           0           0     3138833           0           0
  4:45:05   256 114 126 128    MAX          577           0           0           0     2524311           0           0
  5:00:05   256 117 125 128    MAX          642           0           0           0     2876831           0           0
  5:15:05   256 120 126 128    MAX          558           0           0           0     3205550           0           0
  5:30:05   256 120 126 128    MAX          392           0           0           0     2989166           0           0
  5:45:05   256 113 125 128    MAX          429           0           0           0     2127037           0           0
  6:00:05   256 119 126 128    MAX          574           0           0           0     1534035           0           0
  6:15:05   256 119 126 128    MAX          535           0           0           0     3148438           0           0
  6:30:05   256 119 127 128    MAX          483           0           0           0     3005242           0           0
  6:45:05   256 107 124 128    MAX          427           0           0           0     1741808           0           0
  7:00:05   256 117 126 128    MAX          540           0           0           0     1808489           0           0
```

Vnode Adapter Activity

Example 8-19 lists detail of the Vnode adapter activity. Although there is a lot information (one report per distributed library per FICON adapter), we think that Vnode Adaptor Historical Activity Combined provides enough information to give an overall view of the FICON channel performance. As always, there is one report per distributed library. In this case, we see a hourly report basis (shown in Example 8-19) with the following information:

- Total throughput per distributed library every hour
- Read and write channel activity
- Read and write device activity with compression rate achieved

Example 8-19 Vnode Adapter Historical Activity combined Report

```
1(C) IBM   REPORT=H21ADPSU(08028)         VNODE ADAPTOR HISTORICAL ACTVITY COMBINED        RUN ON 11MAR2008 @ 15:50:27 PAGE   1
 GRID#=?????   DIST_LIB_ID= 0  VNODE_ID= 0  NODE_SERIAL=13-10570  VE_CODE_LEVEL=008.002.000.0030 UTC NOT CHG
  03MAR08
   RECORD Total ----------CHANNEL-----------   ----------DEVICE----------
     TIME  MB/s   RD_GB  MB/s    WR_GB  MB/s     RD_GB  COMP  WR_GB  COMP
 01:06:11    72   178.4    49     81.7    22      39.2  4.54   29.3  2.78
 02:06:11   126   287.6    79    168.7    46      68.3  4.20   59.1  2.85
 03:06:16    67   128.3    35    114.6    31      40.1  3.19   34.2  3.34
 04:06:16    61   199.5    55     23.5     6      68.3  2.91    8.4  2.80
 05:06:16    36    82.0    22     49.9    13      31.2  2.62   23.5  2.11
 06:06:16    93   151.4    42    183.6    51      28.3  5.33   37.9  4.84
 07:06:16    73   147.6    41    116.4    32      34.5  4.27   30.3  3.83
 08:06:16   110   159.1    44    237.6    66      40.9  3.89   64.2  3.69
 09:06:17    19    59.4    16     10.5     2      14.4  4.12    2.3  4.42
 10:06:17    32    72.3    20     42.9    11      24.9  2.90   12.2  3.49
 11:06:17     9    32.6     9      0.0     0      10.7  3.03    0.0
 12:06:17     8    16.7     4     13.7     3       4.4  3.73    2.4  5.59
 13:06:17    10    32.7     9      5.1     1       8.5  3.80    0.7  7.06
 14:06:17    14    52.7    14      0.0     0      18.1  2.91    0.0
 15:06:17     3    11.5     3      0.7     0       5.4  2.13    0.2  2.65
 16:06:17   253   565.9   157    348.2    96     137.5  4.11   69.8  4.98
 17:06:17   163   312.0    86    276.6    76      91.3  3.41   77.8  3.55
 18:06:17   146   318.1    88    210.6    58      73.8  4.30   46.8  4.49
 19:06:17   255   495.8   137    424.0   117     116.4  4.25   83.6  5.07
 20:06:17   164   322.9    89    269.9    74      55.5  5.81   50.4  5.35
 21:06:17    79    75.8    21    210.9    58      21.5  3.52   48.3  4.36
 22:06:17    64   125.3    34    108.6    30      34.9  3.58   31.5  3.44
 23:06:17   104   289.8    80     86.6    24      61.0  4.74   27.0  3.21
 23:51:17   145   309.4   114     82.4    30      73.3  4.21   21.1  3.89
1(C) IBM   REPORT=H21ADPSU 08028)         VNODE ADAPTOR HISTORICAL ACTVITY COMBINED        RUN ON 11MAR2008 @ 15:50:27          PAGE   2
 GRID#=?????   DIST_LIB_ID= 1  VNODE_ID= 0  NODE_SERIAL=13-10520  VE_CODE_LEVEL=008.002.000.0030                           UTC NOT
CHG
  03MAR08
```

```
  RECORD Total ----------CHANNEL-----------    ----------DEVICE----------
    TIME  MB/s   RD_GB  MB/s     WR_GB  MB/s      RD_GB  COMP  WR_GB  COMP
01:05:43    60   130.1    36      86.7    24       41.0  3.16   25.8  3.35
02:05:43   137   263.3    73     231.9    64       63.0  4.17   52.5  4.41
03:05:43    46    90.5    25      76.0    21       22.6  3.99   23.7  3.19
04:05:43    20    64.8    18      10.2     2       24.2  2.67    4.3  2.33
05:05:43     8    29.5     8       0.0     0        9.3  3.14    0.0
06:05:43    83   168.0    46     133.2    37       34.0  4.93   26.8  4.96
07:05:48    37    72.1    20      61.9    17       19.9  3.61   12.5  4.94
08:05:48    71   161.8    44      97.2    27       46.8  3.45   29.9  3.24
09:05:48    15    44.9    12       9.6     2       13.1  3.42    3.0  3.17
10:05:48    27    67.8    18      29.4     8       18.4  3.66   11.0  2.66
11:05:48    27    85.8    23      13.7     3       20.5  4.17    2.3  5.78
12:05:48    30   108.6    30       0.0     0       30.6  3.53    0.0  7.00
13:05:48    22    82.1    22       0.0     0       25.1  3.26    0.0
14:05:48    25    92.9    25       0.0     0       30.6  3.03    0.0
15:05:48     0     0.7     0       0.0     0        0.3  2.23    0.0
16:05:48   253   625.4   173     286.4    79      165.3  3.78   71.3  4.01
17:05:48   115   201.3    55     214.4    59       46.2  4.35   51.7  4.14
18:05:48   145   322.2    89     200.6    55       62.1  5.18   49.0  4.09
19:05:48   175   384.8   106     245.9    68       94.4  4.07   62.7  3.91
20:05:48   259   587.4   163     345.7    96      105.0  5.59   73.2  4.72
21:05:48    73   193.5    53      70.9    19       37.4  5.16   17.0  4.16
22:05:48   107   191.2    53     194.0    53       52.1  3.66   43.1  4.49
23:05:48   139   181.6    50     318.8    88       40.3  4.49   54.8  5.81
23:50:48    70   133.3    49      56.5    20       34.3  3.87   17.9  3.14
```

Cache Partition Activity

The next report example lists the details about the TS7700 Cache Partition Activity (Example 8-20). You can identify the following information for each 15 minute interval:

- The percentage of read and write throttling
- The numbers for Fast Ready mounts, cache hits and cache misses
- The amount and number of logical volumes separated by preference group (0 or 1)

Example 8-20 HNODE HSM Historical Cache Partition Report

```
RID#=00000   DIST_LIB_ID= 0  VNODE_ID= 0  NODE_SERIAL=78-13840  VE_CODE_LEVEL=008.003.000.0098   TVC_SIZE= 6000  HNODE=ACTIVE UTCMINUS=07 PARTITION
SIZE= 6000
------------PREFERENCE_GROUP_0------------  ------------PREFERENCE_GROUP_1------------
  19JUL07 %WR %CP       FAST_RDY CACHE_HIT CACHE_MIS   VIRT    GB --ROLLING_AVG-                       VIRT    GB --ROLLING_AVG-
   RECORD OVR OVR PART  NUM  AVG  NUM  AVG  NUM  AVG   VOLS   RES -TIME_IN_CACHE -VIRT_VOLS_MIG-   VOLS   RES -TIME_IN_CACHE -VIRT_VOLS_MIG-
     TIME THR THR HIT% MNTS SECS MNTS SECS MNTS SECS  CACHE CACHE  4HR 48HR 35DA  4HR 48HR  35DA  CACHE CACHE  4HR 48HR 35DA  4HR 48HR  35DA
 UPDATE INTERVAL==>                                                -ON_THE_HOUR-- --ON_THE_HOUR--               -ON_THE_HOUR-- --ON_THE_HOUR--
  4:15:00  68   0  100  127 9.47    0  .00    0  .00   1605   470  17M   1M   7M    0    K   88K  17233  5060 118H  97H 114H 1451   1K   74K
  4:30:00  66   0  100  126 5.07    0  .00    0  .00   2022   593  17M   1M   7M    0    K   88K  16783  4928 118H  97H 114H 1451   1K   74K
  4:45:00  53   0  100  118 8.92    0  .00    0  .00   1879   550  17M   1M   7M    0    K   88K  16881  4956 118H  97H 114H 1451   1K   74K
  5:00:00  33   0  100   64 12.5    0  .00    0  .00   1863   546  49M   4M   7M  498    K   88K  16979  4985 115H  97H 114H 2547   2K   74K
  5:15:00  43   0  100   90 5.60    0  .00    0  .00   1799   527  49M   4M   7M  498    K   88K  16909  4965 115H  97H 114H 2547   2K   74K
  5:30:00  56   0  100  127 6.23    0  .00    0  .00   1600   469  49M   4M   7M  498    K   88K  17037  5002 115H  97H 114H 2547   2K   74K
  5:45:00  66   0  100  127 6.31    0  .00    0  .00   1562   458  49M   4M   7M  498    K   88K  17164  5039 115H  97H 114H 2547   2K   74K
  6:00:00  56   0  100  102 6.20    0  .00    0  .00   1355   397  70M   5M   7M 1987   1K   88K  17289  5076 111H  98H 114H 2657   2K   74K
  6:15:00  60   0  100   40 5.08    0  .00    0  .00   1148   336  70M   5M   7M 1987   1K   88K  17292  5077 111H  98H 114H 2657   2K   74K
  6:30:00  46   0  100  136 6.99    0  .00    0  .00   1096   321  70M   5M   7M 1987   1K   88K  17485  5134 111H  98H 114H 2657   2K   74K
  6:45:00  76   0  100  108 10.4    0  .00    0  .00   1108   324  70M   5M   7M 1987   1K   88K  17550  5153 111H  98H 114H 2657   2K   74K
  7:00:00  80   0  100  127 4.84    0  .00    0  .00   1046   306  90M   7M   7M 3309   3K   88K  17677  5190 108H  99H 114H 2657   2K   74K
1(C) IBM   REPORT=H30TVC1  07215)          HNODE HSM HISTORICAL CACHE PARTITION            RUN ON 15AUG2007 @ 11:37:40          PAGE   2
 GRID#=00000   DIST_LIB_ID= 1  VNODE_ID= 0  NODE_SERIAL=78-16000  VE_CODE_LEVEL=008.003.000.0098   TVC_SIZE= 6000  HNODE=ACTIVE UTCMINUS=07
 PARTITION SIZE= 6000                                        ------------PREFERENCE_GROUP_0------------  ------------PREFERENCE_GROUP_1------------
  19JUL07 %WR %CP       FAST_RDY CACHE_HIT CACHE_MIS   VIRT    GB --ROLLING_AVG-                       VIRT    GB --ROLLING_AVG-
   RECORD OVR OVR PART  NUM  AVG  NUM  AVG  NUM  AVG   VOLS   RES -TIME_IN_CACHE -VIRT_VOLS_MIG-   VOLS   RES -TIME_IN_CACHE -VIRT_VOLS_MIG-
     TIME THR THR HIT% MNTS SECS MNTS SECS MNTS SECS  CACHE CACHE  4HR 48HR 35DA  4HR 48HR  35DA  CACHE CACHE  4HR 48HR 35DA  4HR 48HR  35DA
 UPDATE INTERVAL==>                                                -ON_THE_HOUR-- --ON_THE_HOUR--               -ON_THE_HOUR-- --ON_THE_HOUR--
  4:15:01  73   0  100  222 6.25    0  .00    0  .00   1490   436  19M   2M   5M  116    K  131K    990   290  29D  37D  15D    0    K   32K
  4:30:01  70   0  100  128 5.38    0  .00    0  .00   1722   505  19M   2M   5M  116    K  131K   1118   327  29D  37D  15D    0    K   32K
  4:45:01  40   0  100  137 5.90    0  .00    0  .00   1797   527  19M   2M   5M  116    K  131K   1245   365  29D  37D  15D    0    K   32K
  5:00:01  43   0  100  131 7.27    0  .00    0  .00   1691   495  51M   5M   5M  530    K  131K   1499   439  19D  37D  15D    0    K   32K
  5:15:01  46   0  100  136 7.85    0  .00    0  .00   1630   478  51M   5M   5M  530    K  131K   1519   445  19D  37D  15D    0    K   32K
  5:30:01  56   0  100  140 4.28    0  .00    0  .00   1415   414  51M   5M   5M  530    K  131K   1709   501  19D  37D  15D    0    K   32K
  5:45:01  66   0  100  125 6.23    0  .00    0  .00   1368   401  51M   5M   5M  530    K  131K   1775   520  19D  37D  15D    0    K   32K
  6:00:01  30   0  100  174 6.43    0  .00    0  .00   1234   361  73M   6M   5M 1827   1K  131K   2033   596 9.9D  36D  15D    0    K   32K
  6:15:01  60   0  100  167 6.96    0  .00    0  .00   1154   338  73M   6M   5M 1827   1K  131K   2160   633 9.9D  36D  15D    0    K   32K
  6:30:01  30   0  100  159 7.16    0  .00    0  .00   1184   347  73M   6M   5M 1827   1K  131K   2287   670 9.9D  36D  15D    0    K   32K
  6:45:01  63  63  100  128 6.25    0  .00    0  .00   1204   352  73M   6M   5M 1827   1K  131K   2414   708 9.9D  36D  15D    0    K   32K
  7:00:01  96  75  100  126 5.47    0  .00    0  .00   1006   294  96M   9M   5M 2761   2K  131K   2541   745 158M  35D  15D    0    K   32K
1(C) IBM   REPORT=H30TVC1  07215)          HNODE HSM HISTORICAL CACHE PARTITION            RUN ON 15AUG2007 @ 11:37:40          PAGE   3
 GRID#=00000   DIST_LIB_ID= 2  VNODE_ID= 0  NODE_SERIAL=78-12345  VE_CODE_LEVEL=008.003.000.0098   TVC_SIZE= 6000  HNODE=ACTIVE UTCMINUS=07
 PARTITION SIZE= 6000                                        ------------PREFERENCE_GROUP_0------------  ------------PREFERENCE_GROUP_1------------
  19JUL07 %WR %CP       FAST_RDY CACHE_HIT CACHE_MIS   VIRT    GB --ROLLING_AVG-                       VIRT    GB --ROLLING_AVG-
   RECORD OVR OVR PART  NUM  AVG  NUM  AVG  NUM  AVG   VOLS   RES -TIME_IN_CACHE -VIRT_VOLS_MIG-   VOLS   RES -TIME_IN_CACHE -VIRT_VOLS_MIG-
     TIME THR THR HIT% MNTS SECS MNTS SECS MNTS SECS  CACHE CACHE  4HR 48HR 35DA  4HR 48HR  35DA  CACHE CACHE  4HR 48HR 35DA  4HR 48HR  35DA
```

```
 UPDATE INTERVAL==>                                                       -ON_THE_HOUR-- --ON_THE_HOUR--               -ON_THE_HOUR-- --ON_THE_HOUR--
  4:15:05  66   0  100  131 8.51   0  .00   0  .00  1457   427  18M   2M   9M  108    K   98K  17347  5093 101H  79H  14D 1512   1K   49K
  4:30:05  66   0  100  129 10.6   0  .00   0  .00  1842   540  18M   2M   9M  108    K   98K  16893  4960 101H  79H  14D 1512   1K   49K
  4:45:05  53   0  100  122 7.29   0  .00   0  .00  1848   542  18M   2M   9M  108    K   98K  16924  4969 101H  79H  14D 1512   1K   49K
  5:00:05  51   0  100  129 6.75   0  .00   0  .00  1824   535  49M   4M   9M  588    K   98K  16990  4988 101H  80H  14D 2596   2K   49K
  5:15:05  46   0  100  103 6.71   0  .00   0  .00  1800   528  49M   4M   9M  588    K   98K  16930  4970 101H  80H  14D 2596   2K   49K
  5:30:05  36   0  100   71 6.71   0  .00   0  .00  1734   509  49M   4M   9M  588    K   98K  17001  4991 101H  80H  14D 2596   2K   49K
  5:45:05  80   0  100   98 10.2   0  .00   0  .00  1785   524  49M   4M   9M  588    K   98K  16955  4978 101H  80H  14D 2596   2K   49K
  6:00:05  50   0  100  114 6.79   0  .00   0  .00  1725   506  74M   6M   9M 1640   1K   98K  17082  5015 101H  81H  14D 2964   2K   49K
  6:15:05  33   0  100   90 7.06   0  .00   0  .00  1596   468  74M   6M   9M 1640   1K   98K  17131  5030 101H  81H  14D 2964   2K   49K
  6:30:05  46   0  100   80 5.59   0  .00   0  .00  1503   441  74M   6M   9M 1640   1K   98K  17242  5062 101H  81H  14D 2964   2K   49K
  6:45:05  73  23  100  119 12.8   0  .00   0  .00  1615   474  74M   6M   9M 1640   1K   98K  17254  5066 101H  81H  14D 2964   2K   49K
  7:00:05  80  53  100  114 4.52   0  .00   0  .00  1504   441  95M   8M   9M 2821   2K   98K  17301  5079 101H  82H  14D 3138   3K   49K
```

Physical Drive Activity

Another important report to look at is the report for Physical Drive Activity, shown in Example 8-21. From this report, you can identify the following for each 15 minute interval when the report was taken:

- How many physical tape drives were installed
- How many physical tape drives were available
- How many (min/avg/max) were mounted
- How long the mount (min/avg/max in seconds) took
- The number of physical mounts sorted by purpose:
  - RCL = Recalls of logical volume back to cache
  - MIG = pre-migration of logical volume from cache to physical tape
  - RCM = Reclamation
  - SDE = Secure Data Erase (not used within first release of the TS7700)

Example 8-21 VEHSTATS for Physical Tape Types Activity

```
1(C) IBM   REPORT=H32TDU21(07215)        HNODE LIBRARY HISTORICAL DRIVE ACTIVITY        RUN ON 15AUG2007 @ 11:37:40
PAGE   1
 GRID#=00000  DIST_LIB_ID= 0  VNODE_ID= 0  NODE_SERIAL=78-13840  VE_CODE_LEVEL=008.003.000.0098  3494  -L10  BA08A
UTCMINUS=07
  19JUL07 ------------PHYSICAL_DRIVES_3592-E05---------------      ----------PHYSICAL_DRIVES_NONE    ----------------
   RECORD          --MOUNTED--  -MOUNT_SECS-  --MOUNTS_FOR---               --MOUNTED--  -MOUNT_SECS-  --MOUNTS_FOR---
     TIME INST AVL MIN AVG MAX  MIN AVG  MAX  RCL MIG RCM SDE      INST AVL MIN AVG MAX  MIN AVG  MAX  RCL MIG RCM SDE
  4:15:00    8   8   0   1   1   61  61   61    0   1   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  4:30:00    8   8   0   3   7   30  36   63    0   7   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  4:45:00    8   8   7   7   7    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  5:00:00    8   8   7   8   8   31  31   31    0   1   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  5:15:00    8   8   5   6   8   31  31   32    0   2   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  5:30:00    8   8   7   8   8   61  61   61    0   1   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  5:45:00    8   8   7   7   7    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  6:00:00    8   8   7   7   7    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  6:15:00    8   8   7   7   7    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  6:30:00    8   8   5   7   8   31  41   62    0   3   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  6:45:00    8   8   7   8   8    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  7:00:00    8   8   7   7   8   33  33   33    0   1   0   0         0   0   0   0   0    0   0    0    0   0   0   0
1(C) IBM   REPORT=H32TDU21 07215)        HNODE LIBRARY HISTORICAL DRIVE ACTIVITY        RUN ON 15AUG2007 @ 11:37:40
PAGE   2
 GRID#=00000  DIST_LIB_ID= 1  VNODE_ID= 0  NODE_SERIAL=78-16000  VE_CODE_LEVEL=008.003.000.0098  3584  -L22  BA08B
UTCMINUS=07
  19JUL07 ------------PHYSICAL_DRIVES_3592-E05---------------      ----------PHYSICAL_DRIVES_NONE    ----------------
   RECORD          --MOUNTED--  -MOUNT_SECS-  --MOUNTS_FOR---               --MOUNTED--  -MOUNT_SECS-  --MOUNTS_FOR---
     TIME INST AVL MIN AVG MAX  MIN AVG  MAX  RCL MIG RCM SDE      INST AVL MIN AVG MAX  MIN AVG  MAX  RCL MIG RCM SDE
  4:15:01    6   3   1   2   2   32  32   32    0   1   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  4:30:01    6   3   2   2   3   32  32   32    0   1   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  4:45:01    6   3   3   4   5   31  31   31    0   2   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  5:00:01    6   3   5   5   5    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  5:15:01    6   3   5   5   5    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  5:30:01    6   3   5   5   5    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  5:45:01    6   3   5   5   5    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  6:00:01    6   3   5   5   5    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
  6:15:01    6   3   5   5   5    0   0    0    0   0   0   0         0   0   0   0   0    0   0    0    0   0   0   0
```

```
  6:30:01    6   3   5   5   6   31  31   31    0   1   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  6:45:01    6   3   5   6   6    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  7:00:01    6   3   5   5   6   31  31   31    0   1   0   0          0   0   0   0   0    0   0    0    0   0   0   0
1(C) IBM   REPORT=H32TDU21 07215)        HNODE LIBRARY HISTORICAL DRIVE ACTIVITY        RUN ON 15AUG2007 @ 11:37:40
PAGE   3
 GRID#=00000  DIST_LIB_ID= 2  VNODE_ID= 0  NODE_SERIAL=78-12345  VE_CODE_LEVEL=008.003.000.0098   3584  -L22  BA08C
UTCMINUS=07
  19JUL07 ------------PHYSICAL_DRIVES_3592-E05--------------UNKNOWN ----------PHYSICAL_DRIVES_NONE    ----------------
   RECORD         --MOUNTED--  -MOUNT_SECS-  --MOUNTS_FOR---                --MOUNTED--  -MOUNT_SECS-  --MOUNTS_FOR---
     TIME INST AVL MIN AVG MAX  MIN AVG  MAX  RCL MIG RCM SDE       INST AVL MIN AVG MAX  MIN AVG  MAX  RCL MIG RCM SDE
  4:15:05    8   8   0   1   1   31  31   31    0   1   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  4:30:05    8   8   0   2   7   30  30   31    0   7   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  4:45:05    8   8   7   7   7    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  5:00:05    8   8   7   7   7    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  5:15:05    8   8   7   7   7    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  5:30:05    8   8   7   7   7    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  5:45:05    8   8   7   7   7    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  6:00:05    8   8   7   7   7    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  6:15:05    8   8   7   7   7    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  6:30:05    8   8   7   7   7    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  6:45:05    8   8   7   7   7    0   0    0    0   0   0   0          0   0   0   0   0    0   0    0    0   0   0   0
  7:00:05    8   8   7   8   8   31  31   31    0   1   0   0          0   0   0   0   0    0   0    0    0   0   0   0
```

Common Scratch Pool

The report for Common Scratch Pool (Example 8-22) shows you the amount of scratch tapes per physical media type.

Example 8-22   VEHSTATS report for Common Scratch Pool

```
1(C) IBM   REPORT=H32CSP  (07215)        HNODE LIBRARY HIST SCRTCH POOL ACTIVITY        RUN ON 15AUG2007 @ 11:37:40
PAGE   1
 GRID#=00000   DIST_LIB_ID= 0  VNODE_ID= 0  NODE_SERIAL=78-13840  VE_CODE_LEVEL=008.003.000.0098
UTCMINUS=07
  19JUL07  ----------SCRATCH_STACKED_VOLUMES_AVAILABLE_BY_TYPE-----------
   RECORD  3590J   3590K   3592JA  3592JJ    NONE    NONE  3592JB    NONE
     TIME  MEDIA0  MEDIA1  MEDIA2  MEDIA3  MEDIA4  MEDIA5  MEDIA6  MEDIA7
  4:15:00       0       0      42       0       0       0       0       0
  4:30:00       0       0      42       0       0       0       0       0
  4:45:00       0       0      42       0       0       0       0       0
  5:00:00       0       0      41       0       0       0       0       0
  5:15:00       0       0      41       0       0       0       0       0
  5:30:00       0       0      41       0       0       0       0       0
  5:45:00       0       0      41       0       0       0       0       0
  6:00:00       0       0      41       0       0       0       0       0
  6:15:00       0       0      41       0       0       0       0       0
  6:30:00       0       0      40       0       0       0       0       0
  6:45:00       0       0      40       0       0       0       0       0
  7:00:00       0       0      40       0       0       0       0       0
1(C) IBM   REPORT=H32CSP   07215)        HNODE LIBRARY HIST SCRTCH POOL ACTIVITY        RUN ON 15AUG2007 @ 11:37:40
PAGE   2
 GRID#=00000   DIST_LIB_ID= 1  VNODE_ID= 0  NODE_SERIAL=78-16000  VE_CODE_LEVEL=008.003.000.0098
UTCMINUS=07
  19JUL07  ----------SCRATCH_STACKED_VOLUMES_AVAILABLE_BY_TYPE-----------
   RECORD  3590J   3590K   3592JA  3592JJ    NONE    NONE  3592JB    NONE
     TIME  MEDIA0  MEDIA1  MEDIA2  MEDIA3  MEDIA4  MEDIA5  MEDIA6  MEDIA7
  4:15:01       0       0      42       0       0       0       0       0
  4:30:01       0       0      42       0       0       0       0       0
  4:45:01       0       0      41       0       0       0       0       0
  5:00:01       0       0      41       0       0       0       0       0
  5:15:01       0       0      41       0       0       0       0       0
  5:30:01       0       0      41       0       0       0       0       0
  5:45:01       0       0      41       0       0       0       0       0
  6:00:01       0       0      41       0       0       0       0       0
  6:15:01       0       0      41       0       0       0       0       0
  6:30:01       0       0      41       0       0       0       0       0
```

```
  6:45:01       0       0      40       0       0       0       0       0
  7:00:01       0       0      40       0       0       0       0       0
1(C) IBM   REPORT=H32CSP   07215)        HNODE LIBRARY HIST SCRTCH POOL ACTIVITY        RUN ON 15AUG2007 @ 11:37:40
PAGE   3
 GRID#=00000   DIST_LIB_ID= 2  VNODE_ID= 0  NODE_SERIAL=78-12345  VE_CODE_LEVEL=008.003.000.0098
UTCMINUS=07
  19JUL07  ----------SCRATCH_STACKED_VOLUMES_AVAILABLE_BY_TYPE-----------
   RECORD  3590J   3590K   3592JA  3592JJ    NONE    NONE  3592JB    NONE
     TIME  MEDIA0  MEDIA1  MEDIA2  MEDIA3  MEDIA4  MEDIA5  MEDIA6  MEDIA7
  4:15:05       0       0      87       0       0       0       0       0
  4:30:05       0       0      87       0       0       0       0       0
  4:45:05       0       0      87       0       0       0       0       0
  5:00:05       0       0      87       0       0       0       0       0
  5:15:05       0       0      87       0       0       0       0       0
  5:30:05       0       0      87       0       0       0       0       0
  5:45:05       0       0      87       0       0       0       0       0
  6:00:05       0       0      87       0       0       0       0       0
  6:15:05       0       0      87       0       0       0       0       0
  6:30:05       0       0      87       0       0       0       0       0
  6:45:05       0       0      87       0       0       0       0       0
  7:00:05       0       0      86       0       0       0       0       0
```

General Pool Use

The next report (Example 8-23) is the report for General Pool Use. One report is always used to show two pools. In this example, the report shows Pool 01 and Pool 02. You can see in detail for the recorded time frame, per pool:

- The amount of active logical volume
- The amount of active data in GB
- The amount of data written in MB
- The amount of data read in MB
- The current reclamation threshold and target pool

Example 8-23 VEHSTATS report for General Pool Use

```
(1(C) IBM   REPORT=H32GUP01(07215)          HNODE LIBRARY HIST GUP/POOLING ACTIVITY         RUN ON 15AUG2007 @ 11:37:40          PAGE
01
 GRID#=00000   DIST_LIB_ID= 0  VNODE_ID= 0  NODE_SERIAL=78-13840  VE_CODE_LEVEL=008.003.000.0098   3494  -L10  BA08A
UTCMINUS=07
  19JUL07  POOL 01 3592-E05 3592JA                                READ  UN-    POOL 02 3592-E05                                      READ  UN-
   RECORD  ACTIVE  ACTIVE      MB      MB   VOL_COUNT  RECLAIM- ONLY AVAI   ACTIVE  ACTIVE      MB      MB   VOL_COUNT  RECLAIM- ONLY AVAI
     TIME  LVOLS       GB  WRITTN    READ   SCR  PRIV  PCT POOL VOLS VOLS   LVOLS       GB  WRITTN    READ   SCR  PRIV  PCT POOL VOLS VOLS
 UPD INT=>  -ON_THE_HOUR-                   ON_THE_HR                       -ON_THE_HOUR-                   ON_THE_HR
  4:15:00   65079   18052    5412       0     2    56   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  4:30:00   65079   18052   37888       0     2    56   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  4:45:00   65079   18052   83895       0     2    56   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  5:00:00   65630   18206   94721       0     2    57   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  5:15:00   65630   18206   98630       0     2    57   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  5:30:00   65630   18206  124490       0     2    57   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  5:45:00   65630   18206  119979       0     2    57   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  6:00:00   67069   18610  108854       0     2    57   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  6:15:00   67069   18610  108854       0     2    57   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  6:30:00   67069   18610   97126       0     2    57   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  6:45:00   67069   18610  104343       0     2    57   25   01   00   00       0       0       0       0     0     0   25   02   00   00
  7:00:00   68411   18985   78182       0     2    58   25   01   00   00       0       0       0       0     0     0   25   02   00   00
1(C) IBM   REPORT=H32GUP01 07215)          HNODE LIBRARY HIST GUP/POOLING ACTIVITY         RUN ON 15AUG2007 @ 11:37:40          PAGE  02
 GRID#=00000   DIST_LIB_ID= 1  VNODE_ID= 0  NODE_SERIAL=78-16000  VE_CODE_LEVEL=008.003.000.0098   3584  -L22  BA08B
UTCMINUS=07
  19JUL07  POOL 01 3592-E05 3592JA                                READ  UN-    POOL 02 3592-E05                                      READ  UN-
   RECORD  ACTIVE  ACTIVE      MB      MB   VOL_COUNT  RECLAIM- ONLY AVAI   ACTIVE  ACTIVE      MB      MB   VOL_COUNT  RECLAIM- ONLY AVAI
     TIME  LVOLS       GB  WRITTN    READ   SCR  PRIV  PCT POOL VOLS VOLS   LVOLS       GB  WRITTN    READ   SCR  PRIV  PCT POOL VOLS VOLS
 UPD INT=>  -ON_THE_HOUR-                   ON_THE_HR                       -ON_THE_HOUR-                   ON_THE_HR
  4:15:01   65188   18083    5412       0     2    66   10   01   00   00       0       0       0       0     0     0   10   02   00   00
  4:30:01   65188   18083   28566       0     2    66   10   01   00   00       0       0       0       0     0     0   10   02   00   00
  4:45:01   65188   18083   52622       0     2    66   10   01   00   00       0       0       0       0     0     0   10   02   00   00
  5:00:01   65584   18194   83595       0     2    67   10   01   00   00       0       0       0       0     0     0   10   02   00   00
  5:15:01   65584   18194   84497       0     2    67   10   01   00   00       0       0       0       0     0     0   10   02   00   00
  5:30:01   65584   18194  114266       0     2    67   10   01   00   00       0       0       0       0     0     0   10   02   00   00
  5:45:01   65584   18194  120581       0     2    67   10   01   00   00       0       0       0       0     0     0   10   02   00   00
  6:00:01   66903   18561   66755       0     2    67   10   01   00   00       0       0       0       0     0     0   10   02   00   00
  6:15:01   66903   18561   66455       0     2    67   10   01   00   00       0       0       0       0     0     0   10   02   00   00
```

```
  6:30:01   66903   18561   63147       0     2    67   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  6:45:01   66903   18561   67958       0     2    67   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  7:00:01   67820   18820  118175       0     2    68   10   01   00   00        0       0       0       0     0     0   10   02   00   00
1(C) IBM   REPORT=H32GUP01 07215)         HNODE LIBRARY HIST GUP/POOLING ACTIVITY         RUN ON 15AUG2007 @ 11:37:40           PAGE  03
 GRID#=00000   DIST_LIB_ID= 2  VNODE_ID= 0  NODE_SERIAL=78-12345  VE_CODE_LEVEL=008.003.000.0098   3584  -L22  BA08C
UTCMINUS=07
  19JUL07  POOL 01 UNKNOWN  3592JA                               READ  UN-    POOL 02 UNKNOWN                                  READ  UN-
  RECORD  ACTIVE  ACTIVE      MB      MB    VOL_COUNT  RECLAIM- ONLY AVAI    ACTIVE  ACTIVE      MB      MB    VOL_COUNT  RECLAIM- ONLY AVAI
    TIME   LVOLS      GB  WRITTN    READ    SCR  PRIV  PCT POOL VOLS VOLS     LVOLS      GB  WRITTN    READ    SCR  PRIV  PCT POOL VOLS VOLS
 UPD INT=>  -ON_THE_HOUR-                   ON_THE_HR                        -ON_THE_HOUR-                     ON_THE_HR
  4:15:05   65189   18076    5412       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  4:30:05   65189   18076   13832       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  4:45:05   65189   18076   91112       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  5:00:05   65676   18212   73972       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  5:15:05   65676   18212   76378       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  5:30:05   65676   18212   77280       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  5:45:05   65676   18212  109455       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  6:00:05   66745   18511   60140       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  6:15:05   66745   18511   96825       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  6:30:05   66745   18511  111861       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  6:45:05   66745   18511   76678       0     2    69   10   01   00   00        0       0       0       0     0     0   10   02   00   00
  7:00:05   67936   18845   78783       0     2    70   10   01   00   00        0       0       0       0     0     0   10   02   00   00
```

Peer to Peer Activity

The final example (Example 8-24) shows a report for the collection of the Peer to Peer Activity. This report is very useful in installations where copy mode is deferred to know the amount of data (and to calculate the time delay) that is pending to copy at the part of the day with the highest workload (this should be during batch window).

You can identify there, for each 15 minute interval of the time frame the report was taken:

- Number of logical volumes to be copied (valid only for a Multi Cluster Grid configuration)
- Amount of data to be copied (MB)
- Average age of Copy Jobs on the deferred and immediate copy queue
- Amount of data (in MB) to and from the Tape Volume Cache driven by copy activity
- MB to copy from cluster to cluster, seen from each cluster in a Multi Cluster Grid. If you go through the entire report, you see an example with three active clusters.

Example 8-24 VEHSTATS report for Peer to Peer Activity

```
1(C) IBM   REPORT=H33GRID (07215)        HNODE HISTORICAL PEER-TO-PEER ACTIVITY        RUN ON 15AUG2007 @ 11:37:40
PAGE   1
 GRID#=00000   DIST_LIB_ID= 0  VNODE_ID= 0  NODE_SERIAL=78-13840  VE_CODE_LEVEL=008.003.000.0098
UTCMINUS=07
                                                                                 MB_FR          MB_FR
  19JUL07   LVOLS      MB  AV_DEF AV_RUN   MB_TO  CALC  V_MNTS  MB_XFR  MB_XFR   1-->0  CALC    2-->0  CALC
               TO      TO  QUEAGE QUEAGE  TVC_BY   MB/ DONE_BY   FR_DL   TO_DL  TVC_BY   MB/   TVC_BY   MB/
             COPY    COPY  ---SECONDS---    COPY   SEC OTHR_DL  RMT_WR  RMT_RD    COPY   SEC     COPY   SEC
  4:15:00      34    9750       0    115       0             0       0       0   35489  39.4    58168  64.6
  4:30:00       6    1720       0    193       0             0       0       0   12444  13.8    62238  69.1
  4:45:00      30    8603       0    311       0             0       0       0   34255  38.0    33934  37.7
  5:00:00     153   43876       0    137       0             0       0       0   13125  14.5    37852  42.0
  5:15:00     125   35846       0     92       0             0       0       0    8962   9.9    12030  13.3
  5:30:00      15    4301       0     16       0             0       0       0   10868  12.0    34302  38.1
  5:45:00       2     573       0     10       0             0       0       0   11154  12.3    36608  40.6
  6:00:00     102   29250       0    173       0             0       0       0   14608  16.2    39468  43.8
  6:15:00     106   30397       0    213       0             0       0       0   19834  22.0    12996  14.4
  6:30:00      87   24949       0     89       0             0       0       0   29597  32.8    25557  28.3
  6:45:00       6    1720       0     11       0             0       0       0    7215   8.0    29229  32.4
  7:00:00       0       0       0      0       0             0       0       0   32564  36.1    48906  54.3
1(C) IBM   REPORT=H33GRID  07215)        HNODE HISTORICAL PEER-TO-PEER ACTIVITY        RUN ON 15AUG2007 @ 11:37:40
PAGE   2
 GRID#=00000   DIST_LIB_ID= 1  VNODE_ID= 0  NODE_SERIAL=78-16000  VE_CODE_LEVEL=008.003.000.0098
UTCMINUS=07
                                                                                 MB_FR          MB_FR
  19JUL07   LVOLS      MB  AV_DEF AV_RUN   MB_TO  CALC  V_MNTS  MB_XFR  MB_XFR   0-->1  CALC    2-->1  CALC
               TO      TO  QUEAGE QUEAGE  TVC_BY   MB/ DONE_BY   FR_DL   TO_DL  TVC_BY   MB/   TVC_BY   MB/
             COPY    COPY  ---SECONDS---    COPY   SEC OTHR_DL  RMT_WR  RMT_RD    COPY   SEC     COPY   SEC
```

```
  4:15:01       1      286       0       4       0             0       0       0    60987  67.7    25980  28.8
  4:30:01       1      286       0       5       0             0       0       0    36608  40.6     9724  10.8
  4:45:01       4     1147       0       5       0             0       0       0    45925  51.0    41363  45.9
  5:00:01      77    22081       0     337       0             0       0       0    41379  45.9    45230  50.2
  5:15:01     148    42442       0     138       0             0       0       0    25666  28.5    44594  49.5
  5:30:01      64    18353       0      37       0             0       0       0    29780  33.0    28066  31.1
  5:45:01       3      860       0      18       0             0       0       0    31676  35.1    29844  33.1
  6:00:01      26     7456       0     333       0             0       0       0    50647  56.2    46424  51.5
  6:15:01      91    26096       0     241       0             0       0       0    43393  48.2    48387  53.7
  6:30:01     152    43589       0      94       0             0       0       0    43846  48.7    50998  56.6
  6:45:01       8     2294       0      22       0             0       0       0    19548  21.7    37178  41.3
  7:00:01       4     1147       0       4       0             0       0       0    34792  38.6    20880  23.2
1(C) IBM   REPORT=H33GRID  07215)        HNODE HISTORICAL PEER-TO-PEER ACTIVITY         RUN ON 15AUG2007 @ 11:37:40
PAGE   3
 GRID#=00000   DIST_LIB_ID= 2  VNODE_ID= 0  NODE_SERIAL=78-12345  VE_CODE_LEVEL=008.003.000.0098
UTCMINUS=07
                                                                                         MB_FR            MB_FR
  19JUL07   LVOLS       MB  AV_DEF AV_RUN   MB_TO   CALC   V_MNTS  MB_XFR  MB_XFR    0-->2  CALC    1-->2  CALC
               TO       TO  QUEAGE QUEAGE  TVC_BY    MB/  DONE_BY   FR_DL   TO_DL   TVC_BY   MB/   TVC_BY   MB/
             COPY     COPY  ---SECONDS---    COPY    SEC OTHR_DL  RMT_WR  RMT_RD     COPY   SEC     COPY   SEC
  4:15:05      69    19787       0     191       0             0       0       0    35178  39.0    35412  39.3
  4:30:05      44    12617       0     447       0             0       0       0    36622  40.6    59493  66.1
  4:45:05      34     9750       0     237       0             0       0       0    30492  33.8    35368  39.2
  5:00:05     143    41008       0      69       0             0       0       0    33470  37.1    43146  47.9
  5:15:05      29     8316       0      20       0             0       0       0    43867  48.7    49081  54.5
  5:30:05       1      286       0       1       0             0       0       0    30812  34.2    45424  50.4
  5:45:05       0        0       0       0       0             0       0       0    31048  34.4    49305  54.7
  6:00:05     116    33265       0     132       0             0       0       0    32499  36.1    51124  56.8
  6:15:05     172    49324       0     116       0             0       0       0    28853  32.0    35713  39.6
  6:30:05      60    17206       0      84       0             0       0       0    25635  28.4    23833  26.4
  6:45:05       0        0       0       0       0             0       0       0    48758  54.1    44521  49.4
  7:00:05       5     1433       0     531       0             0       0       0    33282  36.9    35562  39.5
```

8.8 Other TS7700 control and monitor options

In addition to the previously introduced methods and possibilities for monitoring the TS7700 Virtualization Engine, we add here some more points for further subsystem monitoring.

8.8.1 MVS console commands

There are several SMS display commands to display the OAM status, the Composite Library, and volume details. New and very interesting are also the changes to the SMS operator command LIBRARY. This command has been extended with new parameters. The command could be subject for a thorough consideration. For more information, refer to Chapter 7, "Operation" on page 303 and to *z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427.

8.8.2 MVS console messages

With the new Dynamic Link Load Balancing function, you have a tool to know whether your links are working correctly or whether they are unbalanced. This function evaluates the link capabilities every 5 minutes. When the workload of one link drops below 75% of the workload of the other link, the following warning message is issued:

```
CBR3796E Grid links degraded in library library_name
```

When the situation is solved, this message is shown:

```
CBR3797E Grid links in library_name are no longer degraded
```

When this situation arises, you can use the Host Console Request command Status Gridlink (see 7.3.3, "Host Console Request" on page 375) to obtain additional information that can show you which link is having problems.

8.8.3 IOSTATS

IOSTATS is part of the IBMTOOLS.EXE file, which is available from:

ftp://ftp.software.ibm.com/storage/tapetool/

You can use IOSTATS to measure job execution times, especially before and after hardware upgrades.

IOSTATS can be run for a subset of job names for a certain period of time before the hardware installation. SMF type 30 records are required as input. The reports provided list the number of disk and tape I/O operations that were done for each job step, as well as the elapsed job execution time.

With the TS7700 running in a Multi Cluster Grid configuration, IOSTATS can be used for different purposes:

- To evaluate the effect of the Multi Cluster Grid environment, compare job execution times before implementation of the Multi Cluster Grid to those after having migrated, especially if you are operating in RUN/RUN cluster copy data consistency point (immediate copy) mode of operation.
- To evaluate the effect of hardware upgrades, compare job execution times before upgrading components of the TS7700 Virtualization Engine; for example, if you are increasing the size of the Tape Volume Cache (TVC) or the number of TS1120/3592 tape drives.
- To evaluate the effect of changing the copy mode of operation on elapsed job execution time.

8.8.4 TAPEWISE

Like IOSTATS, TAPEWISE is from the TAPETOOL FTP site. TAPEWISE can, based on input parameters, generate a number of reports that will help answer questions such as:

- Tape activity analysis
- Mounts and MBs processed by hour
- Input and output mounts by hour
- Mounts by SYSID during an hour
- Concurrent open drives used
- Long VTS mounts (recalls)

8.8.5 MOUNTMON and MOUNTRPT

Like IOSTATS, MOUNTMON is from the TAPETOOL FTP site.

MOUNTMON is running as a started task or batch job and monitors all tape activity on the system. The program must be APF-authorized and if it runs continuously it write statistics for each tape volume allocation to SMF or to a flat file.

Based on data gathered from MOUNTMON, the program MOUNTRPT can report on:

- How many tape mounts are we doing?
- How many scratch?
- How many specific?
- How many by system?
- How many by device type?
- How long does it take to mount a tape?
- How long are tapes allocated?
- How many drives are being used at any time?
- Most accurate reporting of concurrent drive usage
- Which jobs are allocating too many drives?

8.8.6 ST@S3494

If you attach a TS7700 to an IBM TotalStorage 3494 Tape Library, you can get reports that summarize the 3494 statistics for each hour of the analysis interval by use of ST@S3494, based on SMF94, which is the interval record used on the 3494 Library.

9

Disaster recovery and failover scenarios

In this chapter, we cover some TS7700-specific Disaster Recovery (DR) considerations with or without a Geographical Dispersed Parallel Sysplex (GDPS). We also explain some failure scenarios.

We discuss the following topics:

- Planning for disaster recovery
- TS7700 Grid failover principles
- Failover scenarios
- Disaster recovery using Copy Export
- Geographically Dispersed Parallel Sysplex (GDPS)
- Disaster recovery testing considerations

9.1 Planning for disaster recovery

Although it is hoped that a disaster does not happen, it is important to plan for such an event. This section provides information that can be of use in developing a disaster recovery plan as it relates to the TS7700.

There are many aspects of disaster recovery planning that must be considered. How critical is the data in the TS7700? Can the loss of some of the data be tolerated? How much time can be tolerated to resume operations after a disaster? What are the procedures for recovery and who will execute them? How will you test your procedures?

With the TS7700, you have two types of configuration that can be installed, standalone and grid. With a standalone system, a single TS7700 cluster is installed. If the site that system is installed at is destroyed, the data associated with the TS7700 might also have been destroyed. In the event that a TS7700 is not usable, because of interruption of utility or communication services to the site or significant physical damage to the site or the TS7700 itself, access to the data managed by the TS7700 is restored through automated processes designed into the product.

The recovery process assumes that the only elements available for recovery are the stacked volumes themselves. It further assumes that only a subset of the volumes is undamaged after the event. If the physical cartridges have been destroyed or irreparably damaged, recovery is not possible, as with other cartridge types. It is important that you integrate the TS7700 recovery procedure into your current disaster recovery procedures.

> **Note:** The disaster recovery process is a joint exercise that requires your involvement as well as that of your SSR.

For many customers, the potential data loss or the recovery time required with a standalone TS7700 is not acceptable. For those customers, the TS7700 Grid provides a near zero data loss and expedited recovery time solution. With a TS7700 Multi Cluster Grid configuration, two or three TS7700 clusters are installed, typically at two or three different sites, and interconnected so that data is replicated between them. The way the two or three sites are used then differs depending on customer requirements.

In a Two-Cluster Grid the typical use will be that one of the sites is the local production center and the other is a backup or disaster recovery center, separated by a distance dictated by your company's requirements for disaster recovery.

In a Three-Cluster Grid the typical use will be that two sites are connected to a host and the workload is spread even between them. The third site is strictly for disaster recovery and there will probably not be connections from the production host to the third site. Another use of a Three-Cluster Grid could be that there are three production sites, all interconnected and holding the backups of one another.

The only connection between the production sites and the disaster recovery site is the grid interconnection. There is normally no host connectivity between the production hosts and the disaster recovery site's TS7700. When data needed for recovery is created at the production sites, it is replicated to the disaster recovery site as defined through outboard policy management definitions and SMS settings.

As part of planning a TS7700 Grid configuration to address this solution, you need to consider the following items:

- Plan for the necessary WAN infrastructure and bandwidth to meet the copy requirements that you need. You generally will need more bandwidth if you are primarily using a copy consistency point of RUN because any delays in copy time caused by bandwidth limitations will result in an elongation of job run times. If you have limited bandwidth available between sites, use the deferred copy consistency point or only copy the data that is critical to the recovery of your key operations. The amount of data sent through the WAN can justify the establishment of a separate, redundant, and dedicated network only for the Multi Cluster Grid.
- If you use a consistency point of deferred copy, and the bandwidth is the limiting factor, it is possible that some of the data has not been replicated between the sites, and the jobs that created that data will need to be rerun.
- Plan for host connectivity at your disaster recovery site with sufficient resources to perform your critical workloads. If the local TS7700 Cluster becomes unavailable, there is no local host access to the data in the disaster recovery site's TS7700 Cluster through the local cluster.
- Design and code the Data Facility System Managed Storage (DFSMS) Automatic Class Selection routines to control what data gets copied and by which copy consistency point. You might need to consider management policies for testing of your procedures at the disaster recovery site that are different from the production policies.
- Prepare procedures that your operators would execute in the event the local site becomes unusable. The procedures would include such tasks as bringing up the disaster recovery host, varying the virtual drives online, as well as placing the disaster recovery TS7700 Cluster in one of the ownership takeover modes.
- Do a periodic capacity planning of your tape setup to evaluate whether the disaster setup is still capable of handling the production in case of a disaster.
- If encryption is used in production, make sure that the disaster site supports encryption as well. The Key Encrypting Keys (KEK) used for production must be available at the disaster recovery site to enable the data key to be decrypted. On the tape setup, the TS1120, the TS7700 and the Library Manager (LM) itself must support encryption. Validate that theTS7700 can communicate with the Encryption Key Manager (EKM) and that the keystore itself is available.
- Consider how you will test your disaster recovery procedures. Many different scenarios can be set up:
  - Will it be based on all data from an existing TS7700?
  - Based on use of the Copy Export function and an empty TS7700?
  - Based on stopping production of one TS7700 and running production to the other during a period of time when one cluster is down for service?

9.2 TS7700 Grid failover principles

In order to better understand and plan for the actions to be performed with the TS7700 Grid configuration in the event of failures, this section describes the key concepts for the grid operation and many of the failure scenarios that the grid has been designed to handle. A TS7700 Grid configuration provides the following data access and availability characteristics:

- In order to access the data on a particular cluster, a host mount request must be issued on a virtual device address defined for that cluster. The virtual device addresses for each cluster are independent. This is different from the prior generation's PTP VTS where the mount request was issued on a virtual device address defined for a virtual tape controller and the virtual tape controller then decided which VTS to use for data access.

- All logical volumes are accessible through any of the virtual device addresses on the TS7700 clusters in the grid configuration. The preference will be to access a copy of the volume in the tape volume cache associated with the TS7700 cluster the mount request is received on. If a recall is required to place the logical volume in the tape volume cache on that TS7700 cluster, it will be done as part of the mount operation. If a copy of the logical volume is not available at that TS7700 cluster (either because it does not have a copy or the copy it does have is inaccessible due to an error), and a copy is available at another TS7700 cluster in the grid, the volume is accessed through the tape volume cache at the TS7700 cluster that has the available copy. If a recall is required to place the logical volume in the tape volume cache on the other TS7700 cluster, it will be done as part of the mount operation.
- Whether a copy is available at another TS7700 cluster in a Multi Cluster Grid depends on the copy consistency point that had been assigned to the logical volume when it was written. The copy consistency point is set through the Management Class storage construct. It specifies if and when a copy of the data is made between the TS7700 clusters in the Grid configuration. There are three copy consistency policies that can be assigned:
  - Rewind Unload (RUN) Copy Consistency Point: If a data consistency point of RUN is specified, the data created on one TS7700 cluster is copied to the other TS7700 cluster as part of successful rewind unload command processing, meaning that for completed jobs, a copy of the volume will exist on both TS7700 clusters. Access to data written by completed jobs (successful rewind/unload) prior to the failure is maintained through the other TS7700 cluster. Access to data of incomplete jobs that were in process at the time of the failure is not provided.
  - Deferred Copy Consistency Point: If a data consistency point of Deferred is specified, the data created on one TS7700 cluster is copied to the other TS7700 cluster after successful rewind unload command processing. Access to the data through the other TS7700 cluster is dependent on when the copy completes. Because there will be some delay in performing the copy, access might or might not be available when a failure occurs.
  - No Copy Copy Consistency Point: If a data consistency point of No Copy is specified, the data created on one TS7700 cluster is not copied to the other TS7700 cluster. If the TS7700 cluster data was written to fails, the data for that logical volume is inaccessible until that TS7700 cluster's operation is restored.
  - With the introduction of the Three-Cluster Grid, a Logical Volume Copy Consistency Override has been enabled. By settings on each library it overrides existing RUN consistency points. Figure 9-1 shows an example of the setting from the Management Interface (MI) where a third copy will be forced to Deferred Copy, but only if three copies are requested from the MC definitions on all three clusters.

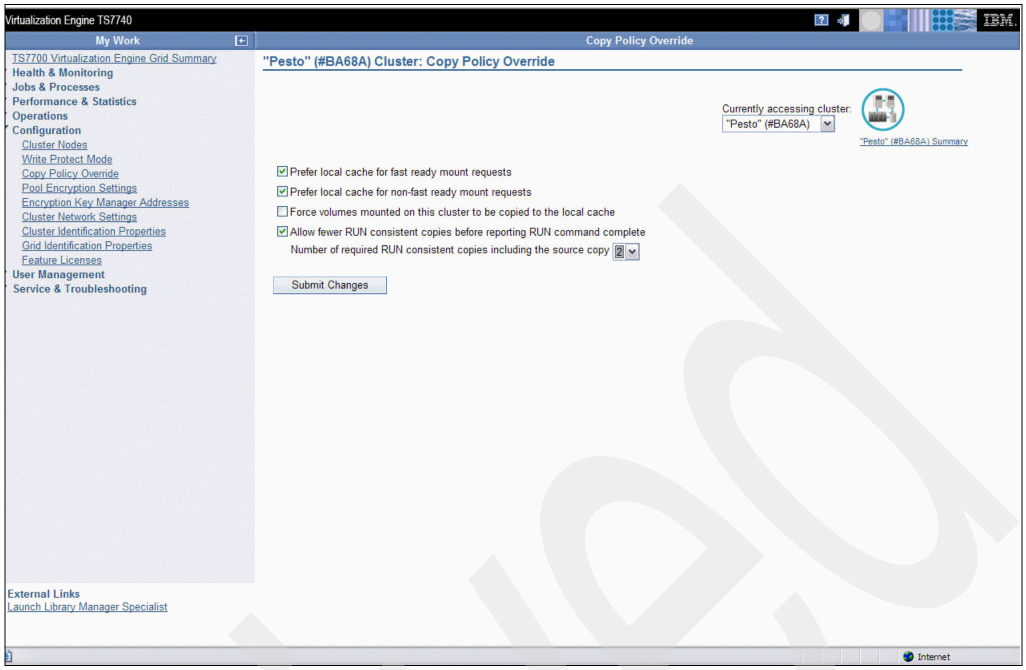

Figure 9-1 Logical Volume Copy Consistency Override

- If a logical volume is written on one of the TS7700 clusters in the grid configuration and copied to the other TS7700 cluster, the copy can be accessed through the other TS7700 cluster. This is subject to the so-called volume ownership.

  At any point in time a logical volume is "owned" by a cluster. The owning cluster has control over access to the volume and for changes to the attributes associated with the volume (such as category or storage constructs). The cluster that has ownership of a logical volume can change dynamically based on which cluster in the grid configuration is requesting a mount of the volume.

  When a mount request is received on a virtual device address, the TS7700 cluster for that virtual device must have ownership of the volume to be mounted or must obtain the ownership from the cluster that currently owns it. If the TS7700 clusters in a grid configuration and the communication paths between them are operational, the change of ownership and the processing of logical volume related commands are transparent in regards to the operation of the TS7700 cluster.

  However, if a TS7700 cluster that owns a volume is unable to respond to requests from other clusters, the operation against that volume will fail, unless some additional direction is given. In other words, clusters will not automatically assume or take over ownership of a logical volume, without being directed. This is done to prevent the failure of the communication paths between the TS7700 clusters resulting in both clusters thinking they have ownership of the volume. If more than one cluster has ownership of a volume, that could result in the volume's data or attributes being changed differently on each cluster, resulting in a data integrity issue with the volume.

  If a TS7700 cluster has failed or is known to be unavailable (for example, it is being serviced) its ownership of logical volumes is transferred to the other TS7700 cluster through one of the following modes. These modes are set through the management interface:

  – Read Ownership Takeover: When Read Ownership Takeover (ROT) is enabled for a failed cluster, ownership of a volume is allowed to be taken from a TS7700 cluster that has failed. Only read access to the volume is allowed through the other TS7700 cluster in the grid. After ownership for a volume has been taken in this mode, any operation attempting to modify data on that volume or change its attributes is failed. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster has been restored.

– Write Ownership Takeover: When Write Ownership Takeover (WOT) is enabled for a failed cluster, ownership of a volume is allowed to be taken from a cluster that has been marked as failed. Full access is allowed through the other TS7700 cluster in the grid. The mode for the failed cluster remains in place until a different mode is selected or the failed cluster has been restored.

– Service Preparation/Service Mode: When a TS7700 cluster is placed in service preparation mode or is in service mode, ownership of its volumes is allowed to be taken by the other TS7700 cluster. Full access is allowed. The mode for the cluster in service remains in place until it has been taken out of service mode.

In addition to the manual setting of one of the ownership takeover modes, an optional automatic method named Autonomic Ownership Takeover Manager (AOTM) is available when each of the TS7700 clusters is attached to a TSSC and there is a communication path provided between the TSSCs.

> **Note:** The links between the TSSCs should not be the same links as the ones used for cluster links. If they are, AOTM will have problems to act automatically if cluster links are failing.

If enabled by the CE, if a TS7700 cluster cannot obtain ownership from the other TS7700 cluster because it does not get a response to an ownership request, a check is made through the TSSCs to determine whether the owning TS7700 cluster is inoperable or just the communication paths to it are not functioning. If the TSSCs have determined that the owning TS7700 cluster is inoperable, then they will enable either read or write ownership takeover, depending on what was set by the CE.

- A failure of a TS7700 cluster will cause the jobs using its virtual device addresses to abend. To run the jobs again, you must enable the host connectivity to the virtual device addresses in the other TS7700 cluster (if not enabled already) and select an appropriate ownership takeover mode. As long as the other TS7700 cluster has a valid copy of a logical volume, the jobs can be retried.

  If a logical volume is being accessed in a remote cache through a 1 Gb Ethernet link and that link fails, the job accessing that volume will also fail. If the failed job is attempted again, the TS7700 Virtualization Engine will use the remaining 1 Gb Ethernet link. If both 1 Gb links fail, access to any data in a remote cache is not possible.

9.3 Failover scenarios

As part of a total systems design, you must develop business continuity procedures to instruct I/T personnel in the actions that they need to take in the event of a failure. We recommend that you test those procedures either during initial installation of the system or at some interval.

The scenarios described in the following sections are from some of the scenarios described in the White Paper *IBM Virtualization Engine TS7700 Series Grid Failover Scenarios,* which was written in an effort to assist IBM specialists and customers in developing such testing plans. The White Paper is available through the following links:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100831
http://www-1.ibm.com/support/docview.wss?uid=tss1wp100831

The paper documents a series of TS7700 Grid failover test scenarios for z/OS which were run in an IBM laboratory environment. Single failures of all major components and communication links and some multiple failures are simulated.

9.3.1 Test configuration

The hardware configuration used for the laboratory test scenarios is shown in Figure 9-2.

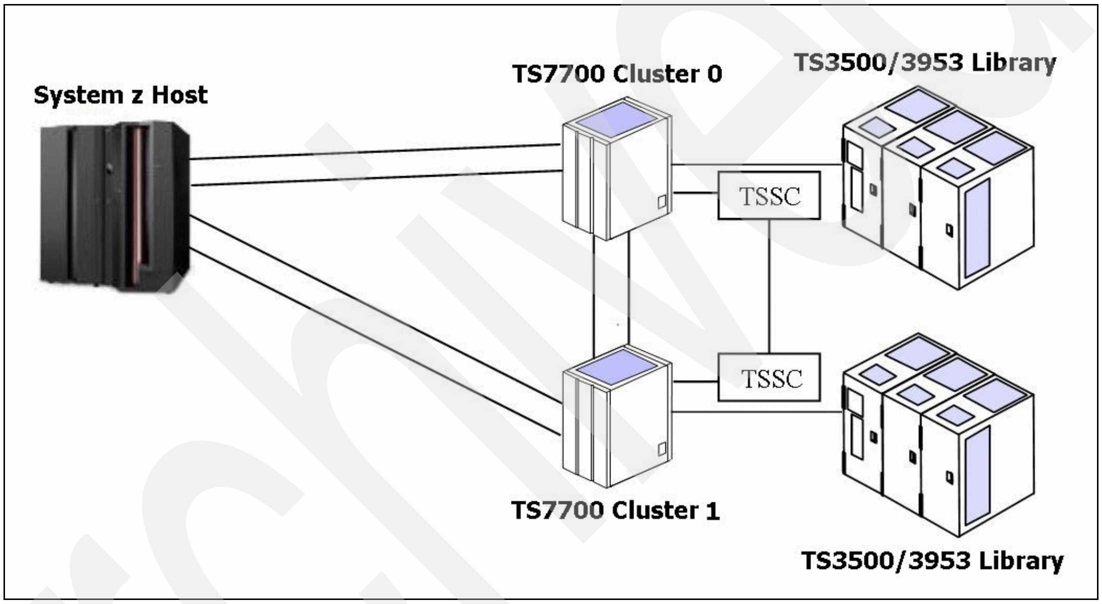

Figure 9-2 Grid test configuration

For the automatic takeover scenario, a TSSC attached to each of the TS7700 clusters is required as well as an Ethernet connection between the TSSCs. Although all the components tested were local, the results of the tests should be similar, if not the same, for remote configurations. All FICON connections were direct, but again, the results should be valid for configurations utilizing FICON directors. Any supported level of z/OS software, and current levels of TS7700, IBM 3953, and TS3500 microcode should all provide similar results. The test environment was MVS/JES2. Failover capabilities are the same for all supported host platforms, although host messages will differ and host recovery capabilities might not be supported in all environments.

For the tests, all host jobs are routed to the virtual device addresses associated with TS7700 Cluster-0. The host connections to the virtual device addresses in TS7700 Cluster-1 are used in testing recovery for a failure of TS7700 Cluster-0.

It is highly recommended that IBM SSRs be involved in the planning and execution of any failover tests. In some scenarios, CE intervention might be required to initiate failures, and restore “failed” components to normal operation.

Test job mix

The test jobs running during each of the failover scenarios consisted of 10 jobs that mounted single specific logical volumes for input (read), and five jobs that mounted single scratch logical volumes for output (write). The mix of work used in the tests was purely arbitrary, and any mix would be suitable, but in order for recovery to be successful, some logical drives must be available for a swap, and for that reason it is recommended that less than the maximum number of virtual drives be active during testing. Also, a large number of messages are generated during some scenarios, and fewer jobs will reduce the number of host console messages

9.3.2 Failover Scenario # 1

In this scenario, shown in Figure 9-3, we assume a failure of one link between TS7700 clusters.

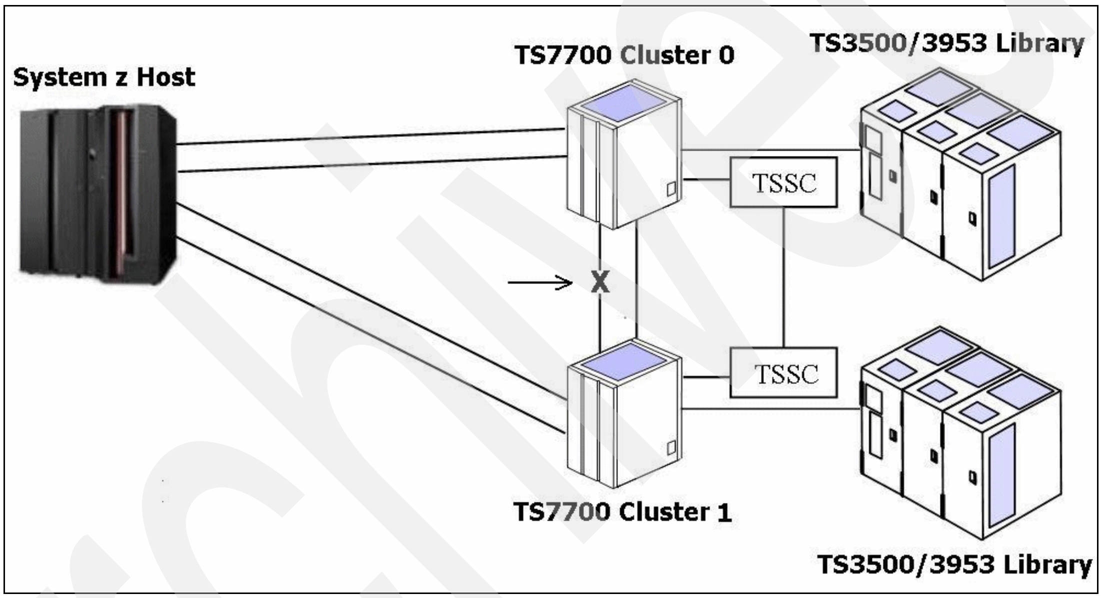

Figure 9-3 Failure of one link between TS7700 clusters

Effects of the failure

You will see the following effects of the failure:

- All grid components continue to operate.
- All channel activity on the failing host link is stopped.
- Host channel errors are reported or error information becomes available from the intermediate equipment.
- If there are alternate paths from the host to the TS7700, host I/O operations might continue.
- All data remains available.

Recovery from failure

To recover from the failures, you must:

- Normal error recovery procedures and repair will apply for the host channel and the intermediate equipment
- Contact your service representative for repair of the failed connection.

9.3.3 Failover Scenario # 2

In this scenario, shown in Figure 9-4, we assume a failure of both links between TS7700 clusters.

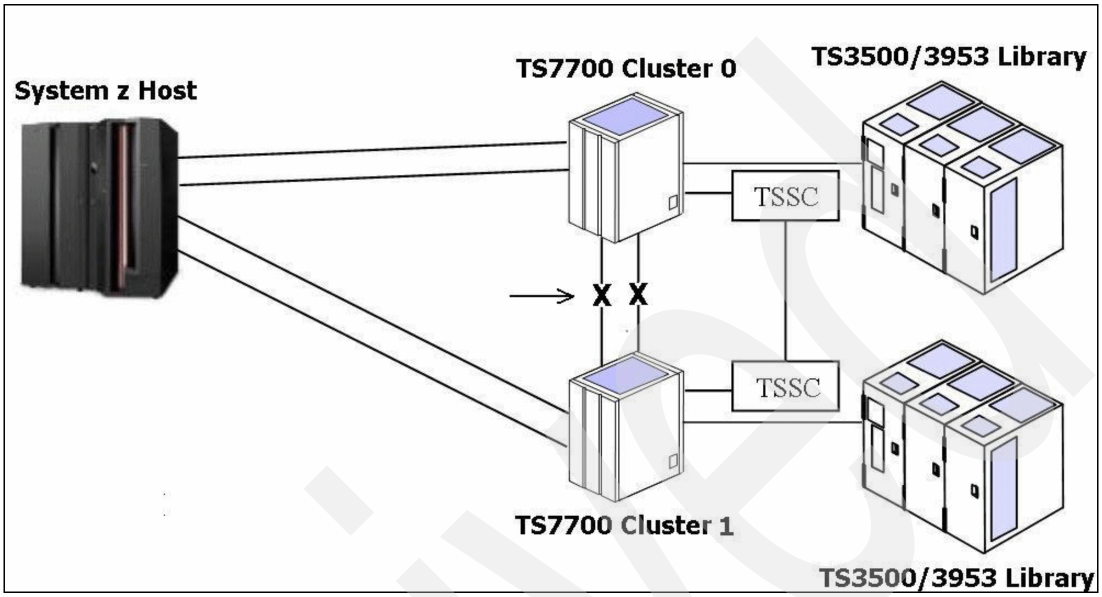

Figure 9-4 Failure of both links between TS7700 clusters

Effects of the failure

You will see the following effects of the failure:

- Jobs on virtual device addresses on TS7700 Cluster 0 will continue to run, assuming that the logical volumes are in that cluster.
- All copy operations are stopped.
- The grid enters the Copy Operation Disabled state.
- If the RUN copy consistency point is being used, the grid also enters the Immediate Mode Copy Completion's Deferred state.
- Call home support is invoked.

Recovery from failure

To recover from the failures, you must contact your service representative for repair of the failed connections.

9.3.4 Failover Scenario # 3

In this scenario, shown in Figure 9-5, we assume a failure of a link between TS7700 clusters with remote mounts.

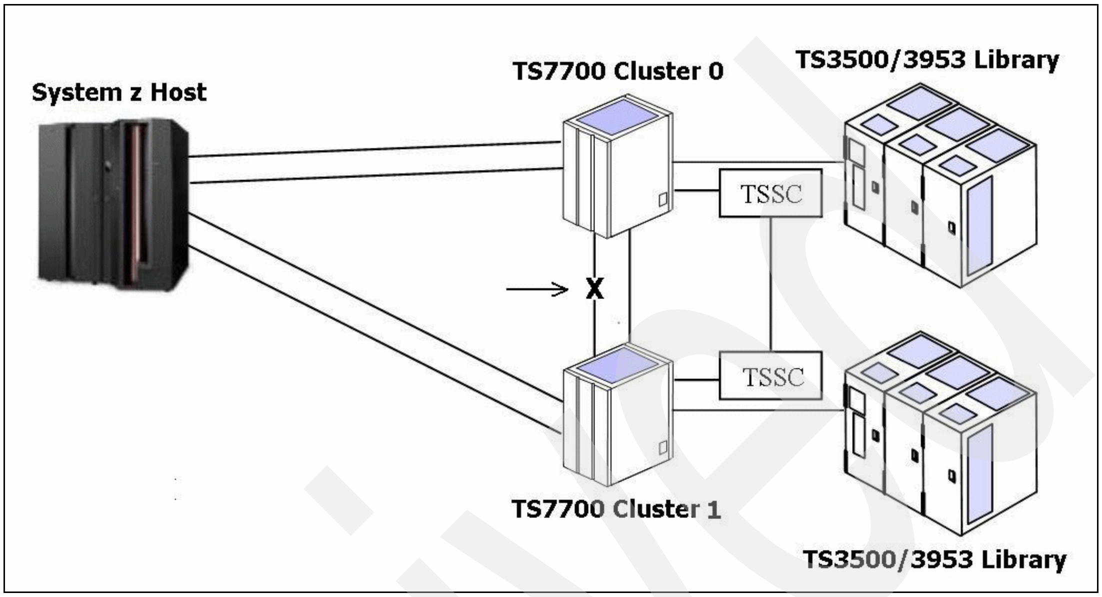

Figure 9-5 Failure of a link between TS7700 clusters with remote mounts

Effects of the failure

You will see the following effects of the failure:

- Any job using the remote link between TS7700 Cluster 0 and TS7700 Cluster 1 that was disconnected will fail.
- If the job is resubmitted, it will succeed using the other link.
- Call home support is invoked.

Recovery from failure

To recover from the failures, you must contact your service representative for repair of the failed connections.

9.3.5 Failover Scenario # 4

In this scenario, shown in Figure 9-6, we assume a failure of both links between TS7700 clusters with remote mounts.

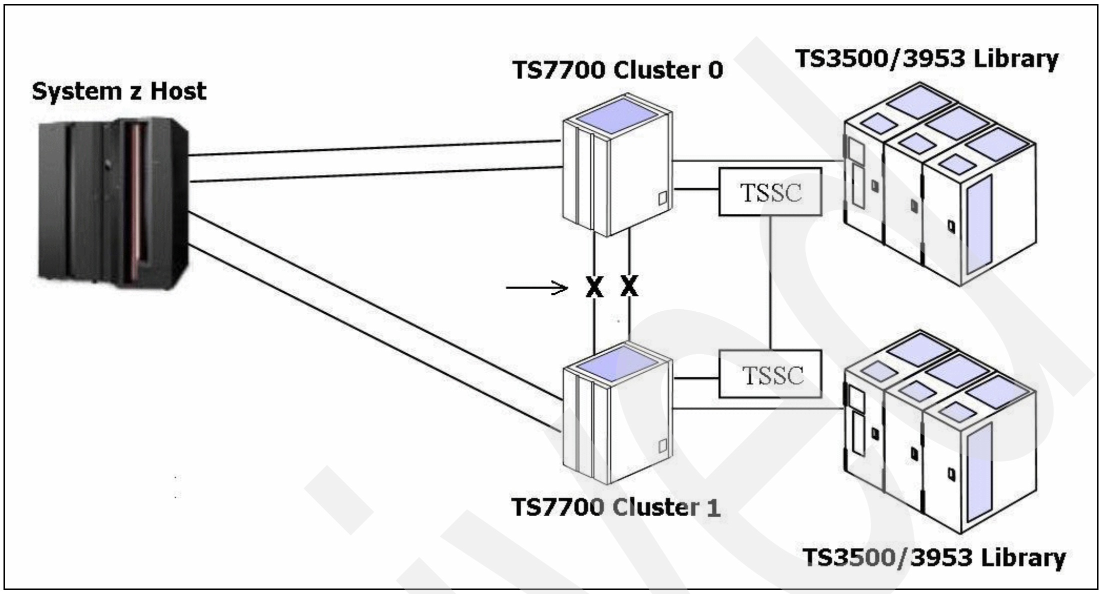

Figure 9-6 Failure of both links between TS7700 clusters with remote mounts

Effects of the failure

You will see the following effects of the failure:

- Jobs on virtual device addresses on TS7700-0 that are using TS7700-1 as the TVC cluster will fail.
- Subsequent jobs that attempt to access the data through TS7700-0 that only exists on TS7700-1 will fail.
- All copy operations are stopped.
- The grid enters the Copy Operation Disabled state.
- If the RUN copy consistency point is being used, the grid also enters the Immediate Mode Copy Completion's Deferred state.
- Call home support is invoked.

> **Note:** Even though the data resides on TS7700-1, if it was mounted on TS7700-0 when the failure occurred, it is not accessible through the virtual device addresses on TS7700-1 because ownership transfer cannot occur.

Recovery from failure

To recover from the failures, you must contact your service representative for repair of the failed connections.

9.3.6 Failover Scenario # 5

In this scenario, shown in Figure 9-7, we assume a failure of the local TS7700 Cluster 0.

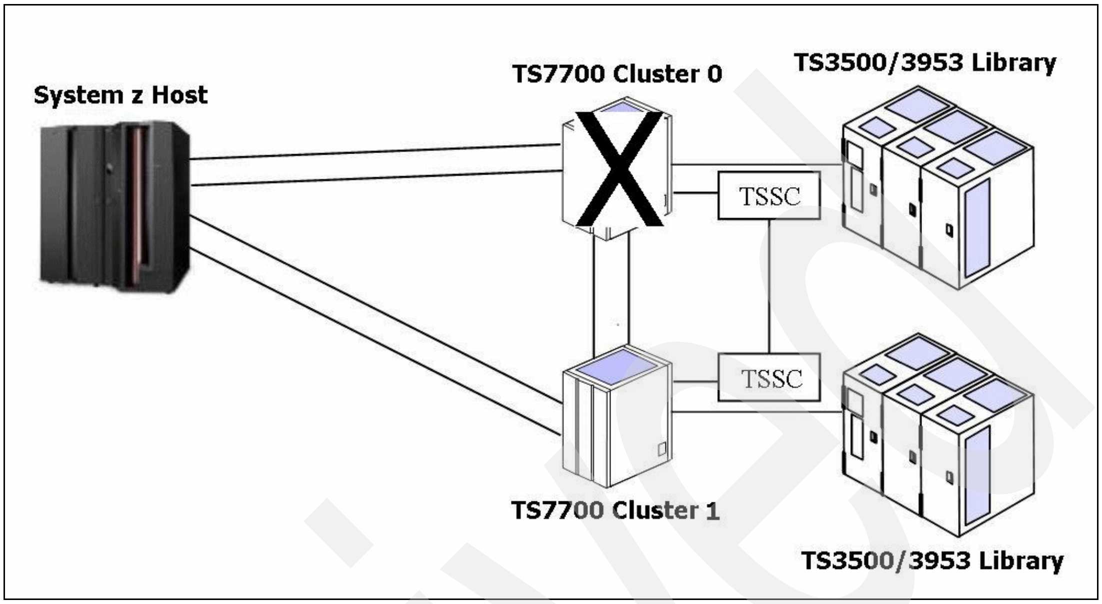

Figure 9-7 Failure of the local TS7700 Cluster 0

Effects of the failure

You will see the following effects of the failure:

- Virtual tape device addresses for TS7700 Cluster 0 are unavailable.
- All channel activity on the failing host links are stopped.
- Host channel errors are reported or error information becomes available from the intermediate equipment.
- Jobs that were using the virtual device addresses of TS7700 Cluster 0 will fail.
- The grid enters the Copy Operation Disabled and the VTS Operations Degraded states.
- If the RUN copy consistency point is being used, the grid also enters the Immediate Mode Copy Completion's Deferred state.
- All copied data can be made accessible through TS7700 Cluster 1 through one of the takeover modes. If a takeover mode for TS7700 Cluster 0 is not enabled, data will not be accessible through TS7700 Cluster 1 even if it has a valid copy of the data.

Recovery from failure

To recover from the failures, you must:

- Enable write or read ownership takeover through the management interface.
- Rerun the failed jobs using the virtual device addresses associated with TS7700 Cluster 1.
- Normal error recovery procedures and repair will apply for the host channels and the intermediate equipment.
- Contact your service representative for repair of the failed TS7700 Cluster.

9.3.7 Failover Scenario # 6

In this scenario, shown in Figure 9-8, we consider a failure of the remote TS7700 Cluster 1.

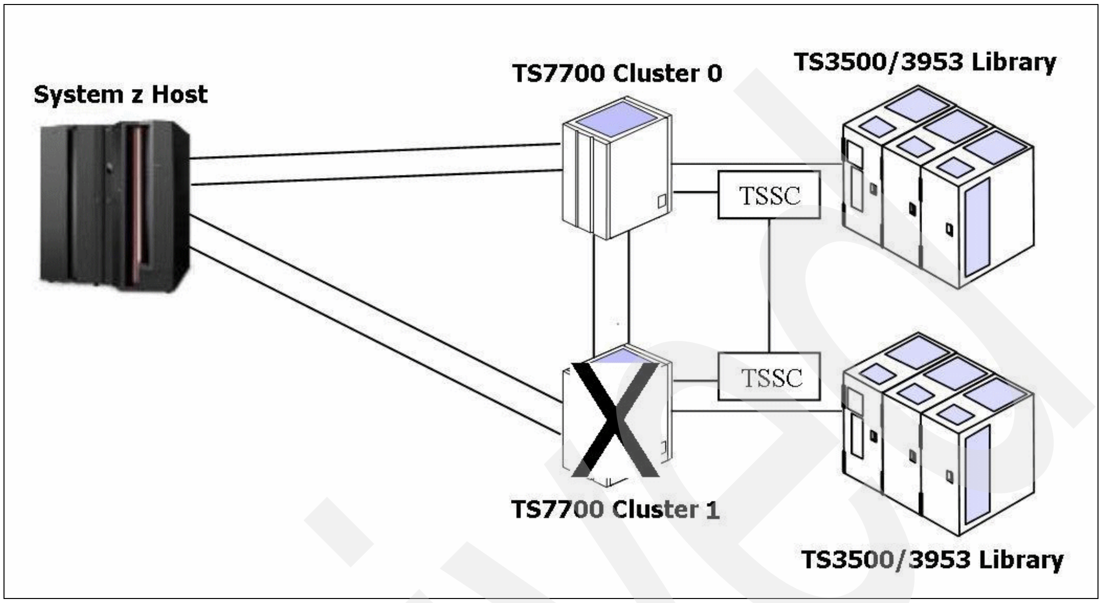

Figure 9-8 Failure of the remote TS7700 Cluster 1

Effects of the failure

You will see the following effects of the failure:

- All jobs continue to run.
- All copy operations are stopped.
- The grid enters the Copy Operation Disabled and VTS Operations Degraded states.
- If the RUN copy consistency point is being used, the grid also enters the Immediate Mode Copy Completion's Deferred state.
- Call home support is invoked.

Note: While TS7700 Cluster 1 is unavailable, data is not being copied. Even when TS7700 Cluster 1 is again available, it will take some time to copy the data from TS7700 Cluster 0.

Recovery from failure

To recover from the failures, you must contact your service representative for repair of the failed TS7700 Cluster.

9.3.8 Failover Scenario # 7

In this scenario, shown in Figure 9-9, we consider a failure of both links between TS7700 clusters with Automatic Takeover.

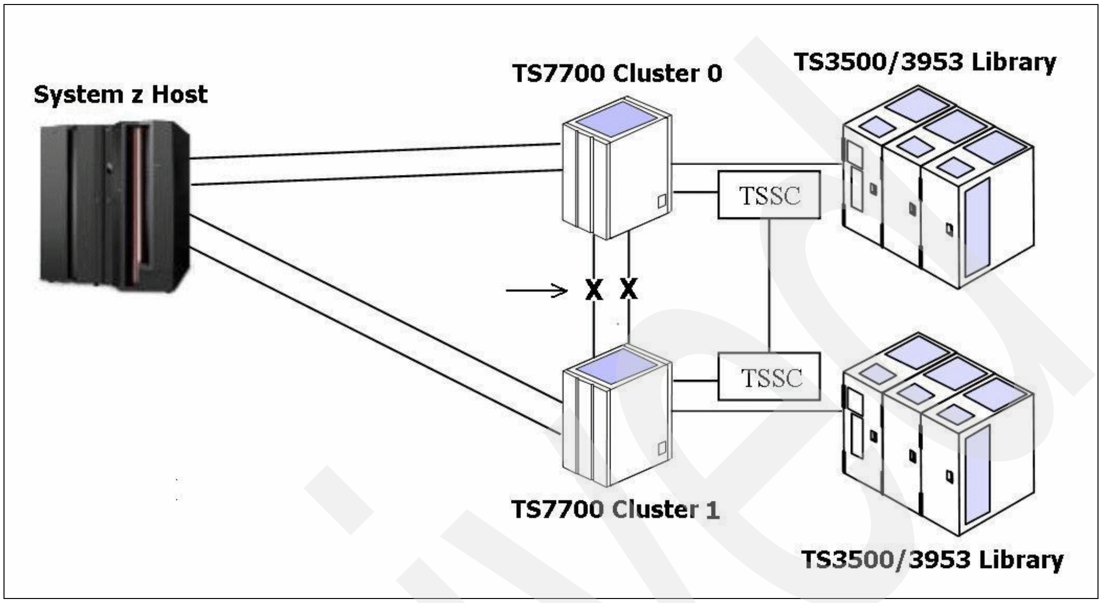

Figure 9-9 Failure of both links between TS7700 clusters with Automatic Takeover

Effects of the failure

You will see the following effects of the failure:

- Jobs subsequent to the failure using virtual device addresses on TS7700 Cluster 0 that need to access volumes that are owned by TS7700 Cluster 1 will fail (even if the data is local to TS7700 Cluster 0). Jobs using virtual device addresses on TS7700 Cluster 1 that need to access volumes that are owned by TS7700 Cluster 0 will also fail.
- All copy operations are stopped.
- The Grid enters the Copy Operation Disabled state.
- If the RUN copy consistency point is being used, the grid also enters the Immediate Mode Copy Completion's Deferred state.
- Call home support is invoked.
- The TSSCs will determine that TS7700 Cluster 0 is still operable and that takeover is not allowed.

Recovery from failure

To recover from the failures, you must contact your service representative for repair of the failed connections.

> **Note:** For more failover scenarios and details, refer to the White Paper *IBM Virtualization Engine TS7700 Series - Grid Failover Scenarios*.

9.3.9 Failover considerations with Three-Cluster Grid

The main focus is to look into what happens if access to the third cluster in a Three-Cluster Grid fails. The example is based on one customer having two TS7700 clusters (Cluster 0 and Cluster 1) active in production. Both are FICON-connected to the host for read and write of logical volumes. The third TS7700 cluster (Cluster 2) is placed far away at a remote location (Cross state) attached to Cluster 0 and Cluster 1 by the network connections and without any FICON connections to the production hosts. The third cluster will have backup host connectivity, devices varied offline, and no active host. Refer to Figure 9-10 for more information.

> **Note:** This is not a tested scenario. It is just an example.

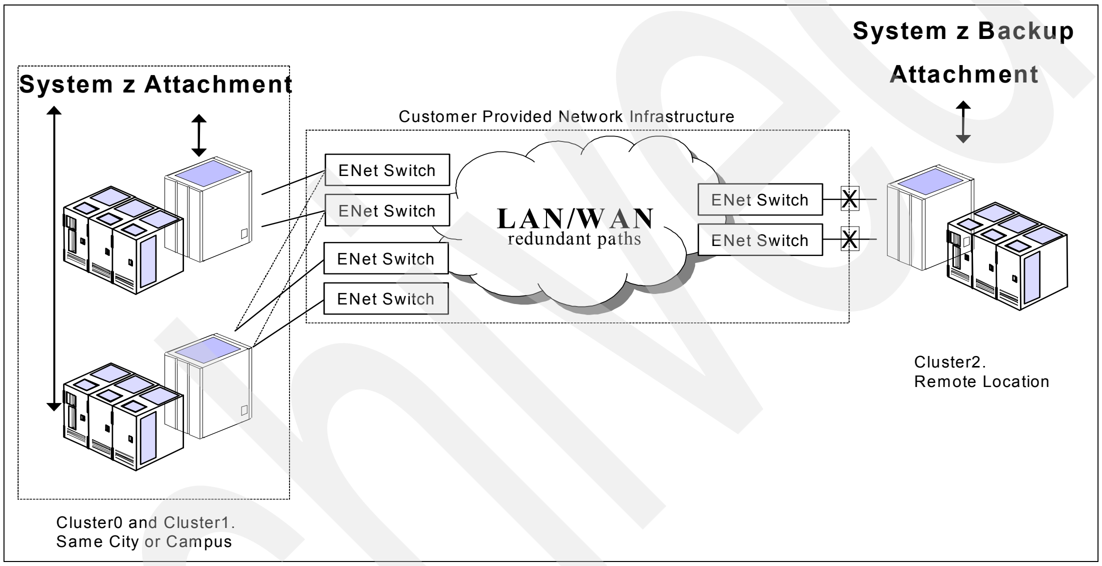

Figure 9-10 Three-Cluster Grid with failure on two links to Cluster 2

Different failures related to Cluster 0 and Cluster 1 is already described in the previous scenarios. In Figure 9-10, we consider what to do when both links to Cluster 2 has failed and the only shared component from Cluster 0 and Cluster 1 to Cluster 2 is the network.

Effects of the failure

You will see the following effects of the failure:

- All jobs will continue to run unaffected.
- All copy operations to Cluster 2 are stopped.
- Cluster 0 and Cluster 1 enter the Copy Operation Disabled state.
- If the RUN copy consistency point is being used from Cluster 0 and Cluster 1 to Cluster 2, the grid also enters the Immediate Mode Copy Completion's Deferred state. But if needed, the copy consistency point could then be changed to RUN between the two remaining clusters.
- Call home support is invoked.
- No AOTM should be used because all host tape activity is using Cluster 0 and Cluster 1.
- The TSSCs will determine that TS7700 Cluster 0 and Cluster 1 are still operable and that takeover is not allowed.

Recovery from failure

To recover from the failures, you must contact your service representative for repair of the failed connections.

9.4 Disaster recovery using Copy Export

Copy export provides a function to allow a copy of selected logical volumes written to the TS7700 to be removed and taken offsite for disaster recovery purposes. All the original logical or physical volumes are intact and usable for the production system. In 5.3.2, "Implementing Copy Export" on page 212, we explain how to set up and use the Copy Export function. In the following sections, we give you an example of how to use Copy Export Recovery.

We offer step-by-step instructions on how to implement and execute Copy Export Recovery. Note that these also apply if you are running a DR test. If it is a test it is specified in each step. For more details and error messages related to the Copy Export function, refer to the White Paper *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide,* which is available at:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101092

9.4.1 Execute Copy Export Recovery

Before executing the recovery, be sure that the described tasks from "Planning and considerations for testing Copy Export Recovery" on page 516 and "Cleanup after testing Copy Export Recovery" on page 517 have been tested with success. That brings us to the last tasks of the recovery. Some tasks are only verifications depending on the current settings in the LM:

1. Verify or define the VOLSER ranges for the physical volumes that are to be used for recovery. The recovery TS7700 must be aware of which VOLSER ranges to use. This is done through the LM associated with the recovery TS7700.
2. If you are going to have the TS7700 apply the storage management constructs from the original production settings as volumes are accessed by the host, then verify or define the same storage management constructs to the DR TS7700. But, if this is a test, then two copies would probably be overdoing it. Changes are performed through the LM associated with the recovery TS7700. If you write data and it is a test, then you would probably also accept only one copy of your test data as well.
3. If you leave the "Disable premigration on construct change" option selected, then the TS7700 ignores the storage management construct setting as logical volumes are accessed and closed. This only applies to existing volumes that are accessed but unchanged. Figure 9-12 on page 505 shows how to set it. If you create new volumes or modify an existing volume, then the storage management construct setting will be followed.
4. If the copy exported physical volumes are encrypted, you must set up the recovery TS7700 for encryption support and have it connected to an Enterprise Key Manager (EKM) that has access to the keys used to encrypt the physical volumes. When writing data to the recovery TS7700, you must also define the pools to be encrypted and set up their key labels.
5. Verify or define the reclamation settings. Make sure that no reclaim will be active on the physical volumes that have been imported if it is a test. Perform the verification or definitions for reclaim through the LM. Be aware that there are two panels that you will need to use, one to define the pool properties, the Stacked Volume Pool Properties

window; and one to define the reclaim inhibit schedule and reclaim percentages for each pool, the VTS Management Policies window. Make sure that all of the reclaim settings are set to 0 in a test. In a real disaster you probably set the values as defined in production.

6. Insert the physical volumes. They are all listed in the Manage Unassigned Volumes window on the LM. Move all volumes to the insert category and the TS7700 will add them to the database. Make sure that all the needed volumes are in the database and known to the LM.

> **Note:** If volumes are missing here and you continue, the entire process has to be rerun. You cannot run Copy Export Recovery without an erased TS7700.

7. With the TS7700 online and all virtual tape drives varied offline to any attached hosts, log in to the management interface and then select Copy Export Recovery from the Service & Troubleshooting menu, as shown in Figure 9-11. The menu will only be available if you have been given access to the function by the system administrator. If the TS7700 is part of a Multi Cluster Grid, the menu is not available.

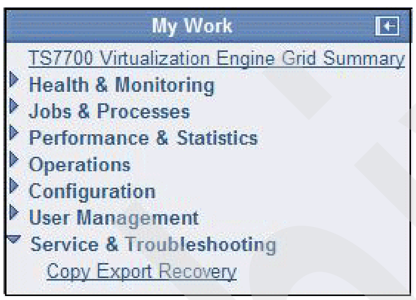

Figure 9-11 Menu with Copy Export Recovery

a. From Figure 9-11, select **Copy Export Recovery** and the menu in Figure 9-12 displays.

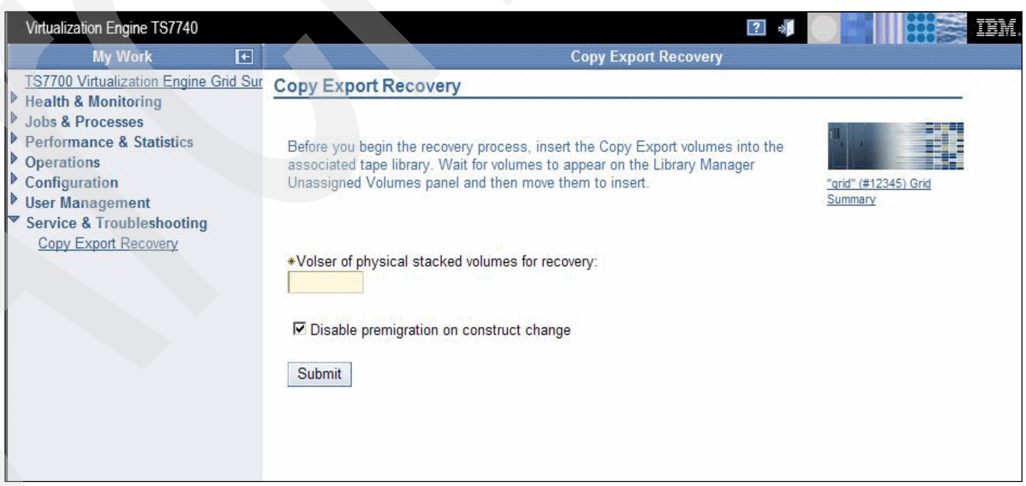

Figure 9-12 Menu with Copy Export Recovery on empty TS7700

b. Make sure that the VOLSER you type in Figure 9-12 is correct and has the newest database backup. Which VOLSER has the newest should be registered when doing the Copy Export on the production system.

If you leave the "Disable premigration on construct change" option enabled, then the TS7700 ignores the storage management construct setting as logical volumes are accessed and closed. This only applies to existing volumes that are accessed to read the data.

c. Click **Submit**. Then, confirm the submission as shown in Figure 9-13.

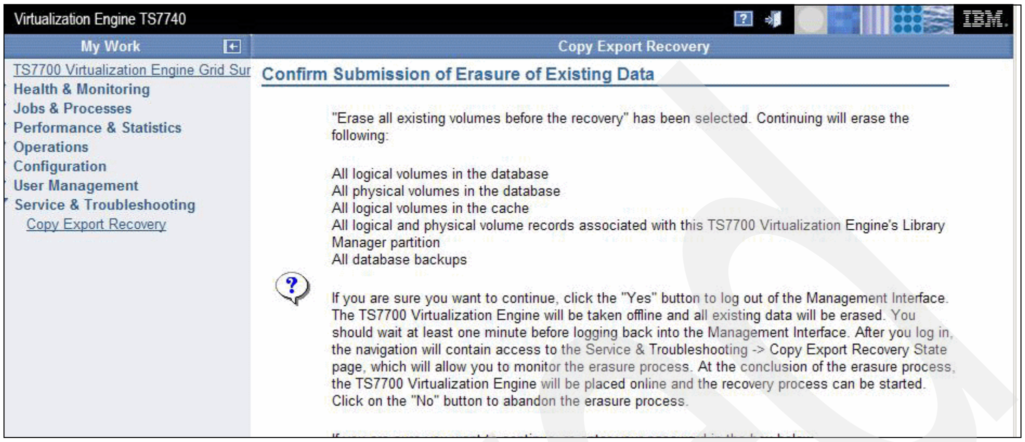

Figure 9-13 Menu to confirm Copy Export Recovery

8. The TS7700 will begin the recovery process. As part of this step, you will be logged off from the management interface.
9. After waiting about one minute, log in to the management interface. Because the TS7700 is performing the recovery process, the only selection that will be available on the Service & Troubleshooting menu is the Copy Export Recovery State page. Select that page to follow the progress of the recovery process. The page provides information about the process including the total number of steps required, the current step, when the operation was initiated, run duration, and overall status, as shown in Figure 9-14.

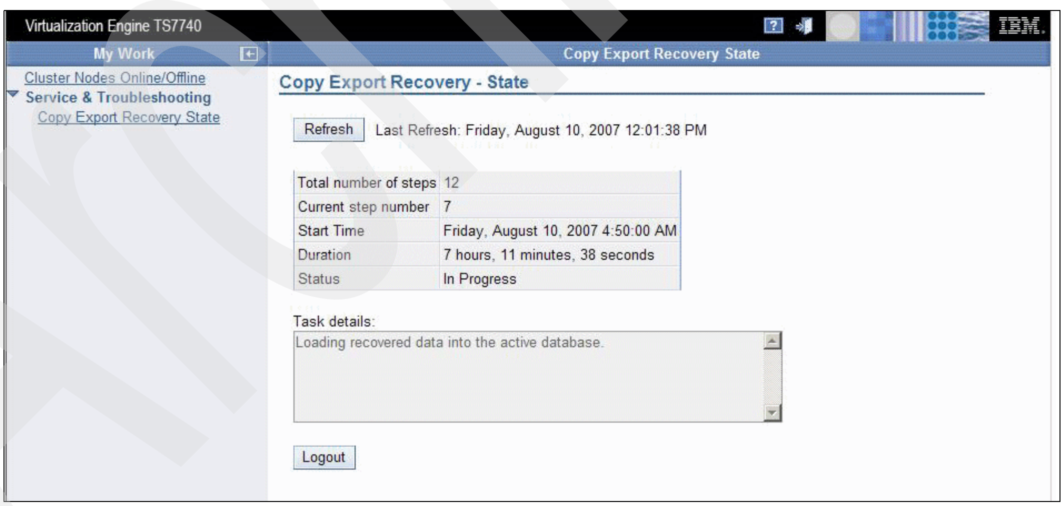

Figure 9-14 Menu with Copy Export Recovery State

The task detail window lists the following tasks as the copy export recovery steps are performed (the *X* and *Y* notation is replaced with appropriate values):

– Taking the TS7700 offline.
– The requested recovery tape *XXXXXX* is being mounted on device *YYY*.
– The database backup is being retrieved from the specified recovery tape *XXXXXX*.
– The requested recovery tape is being demounted following retrieval of the database backup.

- The database backup retrieved from tape is being restored on the TS7700.
- The restored database is being updated for this hardware.
- The restored database volumes are being filtered to contain the set of logical volumes which were copy exported.
- Token ownership is being set to this cluster from previous cluster.
- The restored database is being reconciled with the contents of cache, *XX* of *YY* complete.
- The restored database is being reconciled with the contents of cache, *XX* of *YY* complete.
- Logical volumes are being restored on the Library Manager, *XX* of *YY* complete.
- Copy Export Recovery is complete.
- Copy Export Recovery from physical volume *XXXXXX*.
- Requesting the TS7700 go online.
- Preparing the logical volume information to send to the Library Manager.
- Loading recovered data into the active database.
- In progress.
- Copy Export Recovery is complete.

If an error occurs, there are a number of possible error texts with detailed error descriptions that can help you to solve the problem. For more details and error messages related to the Copy Export Recovery function, refer to the White Paper *IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide* which is available at:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101092

If everything is completed, you can vary the virtual devices online, continue the disaster recovery or DR testing, and the tapes are ready to read.

Note: For more general considerations about DR testing, refer to 9.6, "Disaster recovery testing considerations" on page 510.

9.4.2 Restore the Host and Library environment

Before you can use the recovered logical volumes, you must restore the host environment as well. Only a minimum needed to continue the recovery process of your applications is listed:

1. Define the I/O definition including the library ID of the TS7700 used for DR.
2. Restore the ICF Catalogs including the Tape Configuration Database (TCDB).
3. Restore the Tape Management System (TMS).
4. Define and start a host and modify SMS library definitions with the library ID of the recovery TS7700.

To enable a base host setup on a DR site is complex and not detailed in this book. IBM provides services to support you to create procedures to be ready for a DR.

9.5 Geographically Dispersed Parallel Sysplex (GDPS)

The IBM System z multi-site application availability solution, the Geographically Dispersed Parallel Sysplex (GDPS), integrates Parallel Sysplex technology and remote copy technology to enhance application availability and improve disaster recovery.

GDPS topology is a Parallel Sysplex cluster spread across two sites, with all critical data mirrored between the sites. GDPS provides the capability to manage the remote copy configuration and storage subsystems, automates Parallel Sysplex operational tasks, and automates failure recovery from a single point of control, thereby improving application availability.

GDPS is a multi-site management facility incorporating a combination of system code and automation that utilizes the capabilities of Parallel Sysplex technology, storage subsystem mirroring, and databases to manage processors, storage, and network resources. It is designed to minimize and potentially eliminate the impact of a disaster or planned site outage. The GDPS provides the ability to perform a controlled site switch for both planned and unplanned site outages, with no data loss, maintaining full data integrity across multiple volumes and storage subsystems, and the ability to perform a normal DBMS restart (not DBMS recovery) at the opposite site.

9.5.1 GDPS functions

GDPS provides the following functions:

- Remote Copy Management Facility (RCMF), which automates management of the remote copy infrastructure
- Planned reconfiguration support, which automates operational tasks from one single point of control
- Unplanned reconfiguration support, which recovers from a z/OS, processor, storage subsystem, or site failure

Remote copy management facility

RCMF was designed to simplify the storage administrator's remote copy management functions by managing the remote copy configuration rather than individual remote copy pairs. This includes the initialization and monitoring of the PPRC or XRC volume pairs based upon policy and performing routine operations on installed storage subsystems.

Planned reconfigurations

GDPS planned reconfiguration support automates procedures performed by an operations center. These include standard actions to:

- Quiesce a system’s workload and remove the system from the Parallel Sysplex cluster (stop the system prior to a change window).
- IPL a system (start the system after a change window).
- Quiesce a system’s workload, remove the system from the Parallel Sysplex cluster, and re-IPL the system (recycle a system to pick up SW maintenance). The standard actions can be initiated against a single system or group of systems. Additionally, user-defined actions are supported (for example, a planned site switch in which the workload is switched from processors in site A to processors in site B).

Unplanned reconfigurations

GDPS was originally designed to minimize and potentially eliminate the amount of data loss and the duration of the recovery window in the event of a site failure; however, it will also minimize the impact and potentially mask an z/OS system or processor failure based upon GDPS policy. GDPS uses PPRC or XRC to help minimize or eliminate data loss. Parallel Sysplex cluster functions along with automation are used to detect z/OS system, processor, or site failures and to initiate recovery processing to help minimize the duration of the recovery window.

If a z/OS system fails, the failed system will automatically be removed from the Parallel Sysplex cluster, re-IPLed in the same location, and the workload restarted. If a processor fails, the failed system(s) will be removed from the Parallel Sysplex cluster, re-IPLed on another processor, and the workload restarted.

With PPRC, there will be limited or no data loss, based upon policy, because all critical data is being synchronously mirrored from site A to site B in the event of a site failure. There will be limited data loss if the production systems continue to make updates to the primary copy of data after remote copy processing is suspended (any updates after a freeze will not be reflected in the secondary copy of data) and there is a subsequent disaster that destroys some or all of the primary copy of data. There will be no data loss if the production systems do not make any updates to the primary PPRC volumes after PPRC processing is suspended.

Depending on the type of application or recovery options selected by the enterprise, multiple freeze options are supported by GDPS (the freeze is always performed to allow the restart of the software subsystems.

9.5.2 GDPS considerations on a TS7700 Grid configuration

One of the key principals of GDPS is to have all I/O be local to the system running production. Another is to provide a simplified method to switch between the primary and secondary site if needed. The TS7700 Virtualization Engine in a grid configuration provides a set of capabilities that can be tailored to allow it to operate efficiently in a GDPS environment. Those capabilities and how they might be used in a GDPS environment are described below.

Direct production data I/O to a specific TS7700

Unlike the prior generation's PTP VTS, the hosts are directly attached to the TS7700 local to it. So that is the first consideration in directing I/O to a specific TS7700. Host channels from each site's GDPS hosts are also typically installed to connect to the TS7700 at the site remote to a host to cover recovery Scenario 3 on page 510, but during normal operation, the remote virtual devices are set offline in each GDPS host.

The default behavior of the TS7700 Virtualization Engine in selecting which TVC will be used for the I/O is to follow the Management Class definitions as well as considerations to provide the best overall job performance. It will, however, use a logical volume in a remote TS7700's TVC if required to perform a mount operation unless override settings on a cluster are used. To direct the TS7700 to use its local TVC, the following are recommended.

- For the management class used for production data, ensure that the local cluster has a copy consistency point. If it is important to know that the data has been replicated at job close time, specify a copy consistency point of Rewind/Unload (RUN). If some amount of data loss after a job closes can be tolerated, then a copy consistency point of Deferred can be used. You might have production data with different data loss tolerance and if that is the case, you might want to define more than one management class with different copy consistency points. In defining the copy consistency points for a management class, it is important that you define the same copy mode for each site because in a site switch, the local cluster changes.
- Set Prefer Local Cache for Fast Ready Mounts in MI panel Copy Policy Override. This override will select the TVC local to the TS7700 the mount was received on as long as it is available and a copy consistency point other than No Copy is specified for that Cluster in the Management Class specified with the mount. The Cluster does not have to have a valid copy of the data for it to be selected for the I/O TVC.

- Set Prefer Local Cache for Non-Fast Ready Mounts in MI panel Copy Policy Override. This override will select the TVC local to the TS7700 the mount was received on as long as it is available and the cluster has a valid copy of the data, even if the data is only resident on a physical tape. Having an available, valid copy of the data overrides all other selection criteria. If the local cluster does not have a valid copy of the data, then without the next override, it is possible that the remote TVC would be selected.
- Set Force Volume Copy to Local. This override has two effects, depending on the type of mount requested. For a non-fast ready mount, if a valid copy does not exist on the cluster, a copy is performed to the local TVC as part of the mount processing. For a fast ready mount, it has the effect of 'ORing' the specified Management Class with a copy consistency point of Rewind/Unload for the Cluster and this will force the local TVC to be used. The override does not change the definition of the Management Class, it serves only to influence the selection of the I/O TVC or force a local copy.
- Make sure that these override settings are duplicated on both TS7700s.

Note: The MI panel Copy Policy Override is shown in Figure 9-1 on page 493.

Switch site production from one TS7700 to the other

Like the prior generation's PTP VTS, the way data is accessed by either TS7700 is based on the logical volume serial number. No changes are required in tape catalogs, JCL, or tape management systems. In case of failure in a TS7700 Grid environment together with GDPS there are three scenarios:

1. GDPS switches the primary host to the Remote Location; the TS7700 Grid is still fully functional.
   - No manual intervention required.
   - Logical volume ownership transfer is done automatically during each mount using grid.
2. A disaster happens at the primary site; the GDPS host and TS7700 Cluster are down or inactive.
   - Automatic ownership takeover of volumes, which then will be accessed from the remote host, is not possible.
   - Manual intervention is required. Using the TS7700 Management Interface, the administrator has to invoke a manual ownership takeover. To do so, use the TS7700 Management Interface and select Service and Troubleshooting, Ownership Takeover Mode.
3. Only the TS7700 Cluster in GDPS primary site is down. In this case, two manual interventions are required:
   - Vary online remote TS7700 Cluster devices from the primary GDPS host.
   - Because the down cluster cannot automatically take ownership of volumes that will then be accessed from the remote host, manual intervention is required. Using the TS7700 Management Interface, the administrator has to invoke a manual ownership takeover. To do so the Service and Troubleshooting menu in the TS7700 Management Interface has to be selected. There you will find the Ownership Takeover Mode selection.

9.6 Disaster recovery testing considerations

The TS7700 Grid configuration provides a solution for disaster recovery needs when data loss and the time for recovery are to be minimized. Although a real disaster is not something that can be anticipated, it is very important to have tested procedures in place in case one occurs. As you design a test involving the TS7700 Grid configuration, there are several capabilities designed into the TS7700 that you should consider.

9.6.1 The test environment represents a point in time

The test environment is typically a point in time, meaning that at the beginning of the test, the catalog, TCDB and tape management system control databases are all a snapshot of the production systems. Over the duration of the test, the production systems continue to run and make changes to the catalogs and TMS. Those changes are not reflected in the point-in-time snapshot. The main impact is that it is possible that a volume will be used in a test that has been returned to scratch status by the production system. The test system's catalogs and TMS will not reflect that change. If the links between the TS7700s remain connected, the TS7700 at the test location will be informed that a volume has been returned to scratch, it will not however, prevent the test host from accessing the data on that volume. What is important is that a volume returned to scratch needed during the test is not re-used on the production system during the test. See below for further suggestions on how to manage return-to-scratch handling during a test.

9.6.2 To break the interconnects between the TS7700s - or not

As with the prior generation of PTP VTS, there are two approaches to conducting the test: One where the site-to-site links are broken and one where they are left connected. A test can be conducted with either approach, but each presents trade-offs.

The main trade-off for breaking the links is:

- On the positive side:
  - You are sure that only the data that has been copied to the TS7700 connected to the test system is being accessed.
  - Logical volumes that are returned to scratch by the production system are not “seen” by the TS7700 under test.
  - Test data that is created during the test is not copied to the other TS7700.
- On the negative side:
  - If a disaster occurs while the test is in progress, data that was created by the production site after the links were broken is lost.
  - The TS7700 at the test site must be allowed to take over volume ownership (either read only or read write).
  - The TS7700 under test could select a volume for scratch that has already been used by the production system while the links were broken.

For most customers, the concern about losing data in the event of a disaster during a test is the major issue with using the link break method. The TS7700 has several design features that make valid testing possible without having to break the site-to-site links.

9.6.3 Writing data during the test

This test will typically include running a batch job cycle that will create new data volumes. This can be handled in two ways: 1) have a separate VTS or TS7700 available as the output target for the test jobs, or 2) have a separate logical volume range that is defined for use only by the test system.

The second approach is the most practical in terms of cost. It would involve defining the VOLSER range to be used, defining a separate set of categories for scratch volumes in the DFSMS DEVSUP parmlib and inserting the volume range into the test TS7700 prior to the start of the test. It is very important that the test volumes inserted using the management interface are associated with the test system so that the TS7700 at the test site will have ownership of the inserted volumes.

If the links are to be kept connected during the time when the volumes are inserted, it is also important to make sure that the tape management system at the production site does *not* accept use of the inserted volume range and that the tape management system at the test site does the following:

- Changes on production systems:
  – Use RMM parameter REJECT ANYUSE(tst*), meaning do not use VOLSERs named tst* here.
- Changes on the DR test systems:
  – Use the RMM parameter VLPOOL PREFIX(tst*) TYPE(S) to allow use of these volumes for default scratch mount processing.
  – Change DEVSUPxx to point to another scratch category. That would be the category of the tst* volumes.

After these settings are done it is time to insert the new tst* logical volumes. Any new allocations that are performed by the DR test system will only use the logical volumes defined for the test. At the end of the test the volumes can be returned to scratch status and left in the library, or deleted if desired.

Note: Be aware that one logical unit must have been or is online on the test system before entering logical volumes. For more information, see "Entering logical volumes after hardware changes" on page 210.

9.6.4 Protect production volumes from being written to / returned to scratch on the test system

While performing a test, you would not want the test system to inadvertently overwrite a production volume. If you will not be writing any new data in the test, putting the TS7700 at the test site into Write Protect mode (through the management interface) will prevent a host command from writing data to a logical volume or changing any attributes associated with it, for example, private to scratch. If you do plan on writing data to the TS7700 during the test and you do not plan on breaking the links between sites, then you should use the tape management system control to allow only read-access to the volumes in the production VOLSER ranges.

For example, with DFSMSrmm you would insert these extra statements into the EDGRMMxx parmlib member:

- For production volumes in a range of A00000-A09999, add:
  `REJECT OUTPUT(A0*)`
- For production volumes in a range of ABC000-ABC999, add:
  `REJECT OUTPUT(ABC*)`

With REJECT OUTPUT in effect, products and applications that append data to an existing tape with DISP=MOD must be handled manually to function properly. If the product is

DFSMShsm, tapes that are filling (seen as not full) from the test system Control Dataset, should be modified to full by issuing commands. If DFSMShsm then later needs to write data to tape, it will require a scratch volume related to the test system's logical volume range.

To protect production volumes from being returned to scratch accidently, we recommend:

- This procedure either runs automatically or can be started manually. You need to make sure that the procedure does not run during the entire test. Make sure that the RMM HSKP procedure is not run during the entire test to avoid returning production volumes to scratch. You would have to manually change the volume status back to MASTER to ensure specific mounts from the production system are honored.
- Make sure that the RMM "short on scratch" procedure is not started by providing enough virtual scratch volumes ahead of starting DR testing. The results could be the same as running a HSKP.

If you perform the test with the site-to-site links broken, you can us the Read Ownership Takeover mode to prevent the test system from modifying the production site's volumes. See 9.6.8, "Ownership takeover" on page 516 for further information about ownership takeover.

In addition to the protection options noted, you can also use the following RACF commands to protect the production volumes:

```
RDEFINE TAPEVOL x* UACC(READ) OWNER(SYS1)
SETR GENERIC(TAPEVOL) REFRESH
```

where *x* is the first character of the VOLSER of the volumes to protect.

9.6.5 Control of copies

One of the issues with not breaking the links is that data being created as part of the test might be copied to the production site, wasting space and inter-site bandwidth. This can be avoided by defining the copy mode for the management classes differently at the test TS7700 than at the production TS7700. Using a copy mode of No Copy for the production library site will prevent the test TS7700 from making a copy of the test data. It will not interfere with the copying of production data.

9.6.6 Preventing conflicts between production return to scratch processing and test use of older volumes

In a test environment where the links between sites are not used, having the production system return logical volumes to scratch status that are to be used during the test for input is not an issue, because the TS7700 at the test site will not be aware of the change in status.

In a test environment where the links are maintained, care must be taken to ensure that logical volumes that are to be in the test are not returned to scratch status and used by production applications to write new data. There are several ways that this can be accomplished:

1. Suspend all return to scratch processing at the production site. Unless the test is fairly short (hours not days), this is not likely to be acceptable because of the risk of running out of scratch volumes, especially for native tape workloads. If all tape processing uses logical volumes, the risk of running out of scratch volumes can be eliminated by making sure that the number of scratch volumes available to the production system is enough to cover the duration of the test.

2. If the suspension of all return-to-scratch processing is not feasible, we recommend that you:
   - Add sufficient new scratch volumes for the production system to cover the scratch volume needs during the DR test.
   - Remove the Fast-Ready attribute for the production scratch category at the DR TS7700. Do this for the duration of the DR test.
   - Use the expire-hold attribute for the scratch category on the production TS7700. This will ensure the production host does not mount and set to private a volume that might be needed by the DR host.

   Even if the number of volumes in the list is larger than the number of volumes that are needed per day times the number of days of the test, you still need to take steps to make it unlikely that a volume needed for test will be reused by the production.

   Let us have a look at the following example to clarify the reasons for the actions previously described:

   a. When a snapshot of the tape catalog is taken from the production host for the DR test, volume ABC123 is in volume category X'000F', or *private*.
   b. During the DR test, the production host returns volume ABC123 to category X'0002', or *scratch*. At this point the DR tape catalog has volume ABC123 in category X'000F' and the TS7700 has it in X'0002'.
   c. Then the DR host wants to mount volume ABC123 for read and issues a specific mount request. DFSMS OAM will allow this request because ABC123 is in a *private* category. However, if the TS7700 has ABC123 in category X'0002' with the fast ready attribute set, the TS7700 will mount a newly initialized tape volume and not recall the existing volume. If the fast ready attribute is turned off for category X'0002' on the DR TS7700 cluster, the TS7700 will recall volume ABC123 either from cache or tape.

   For more information, refer to the White Paper *IBM Virtualization Engine TS7700 Series Best Practices - Return-to-Scratch Considerations for Disaster Recovery Testing with a TS7700 Grid,* which you can find at:

   http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101281

   > **Note:** A logical volume that has been returned cannot normally be requested as a specific mount. When a specific mount request is issued, DFSMS checks its TCDB for the status of the volume. If it is in scratch status, the mount request is failed. If the volume was still in the PRIVATE state when the TCDB snapshot for the test environment was taken, the specific mount request on the test system will be passed to the TS7700.
   >
   > On the TS7700, you need to make sure that the fast-ready attribute for the host production scratch category is not set on the DR TS7700. If the fast-ready attribute is set for the production host scratch category on the DR TS7700, the TS7700 presents a newly initailized tape image. It does not present the existing data. The fast-ready attribute can be set on the production host TS7700.

3. Suspend only the return-to-scratch processing for the production volume needed for the test. For RMM, this can be done using policy management through VRSs. A volume VRS can be set up that covers each production volume such that this will override any existing

policies for data sets. For example, say the production logical volumes to be used in the test are in a volser range of 990000-990999. To prevent them from being returned to scratch, the following subsommand would be run on the production system:

```
RMM AS VOLUME(990*) COUNT(99999) OWNER(VTSTEST) LOCATION(CURRENT) PRIORITY(1)
```

Then EDGHSKP EXPROC can be run and not expire the data required for test.

After the test completes, you have a set of tapes in TS7700 that belong to the test activities. You need to decide what to do with these tapes. As the test ends, probably RMM database and TCDB are destaged (as all the data that is used in the Test), but in the LM database, the tapes remain defined. Some will be *private* and others *scratch*.

Whether you want to delete those logical volumes or want to keep them, we recommend that you perform the cleanup on the DR test system first. If you want to delete the logical volumes, perform these steps:

1. Stop the RMM address space and subsystem and, using Interactive Storage Management Facility (ISMF) 2.3, return all private cartridges to scratch.
2. After the volumes are in scratch status, using again ISMF 2.3, eject all the cartridges. Keep in mind, that the LM only accepts 1000 eject commands at one time. If you have to eject a high number of cartridges, this process might be time consuming.

If you want to keep the logical volumes defined in the LM database, just return the volume to scratch as described in Step 1. Cartridges remain in scratch status and ready to use for the next test. We recommend that you set the “Delete Expired Volume Data” value to a very low value so that these logical volumes do not use physical cartridge space. See 4.3.6, “Define logical volume Expiration Time” on page 156 for more information.

Note: Cartridges defined in the LM remain ready to use. Thus, when you create the test environment the next time, you must ensure that these cartridges are defined to RMM as well as in the TCDB before you can use them.

9.6.7 Copies flush or kept as LRU

The default management of the data in the cache is different for host-created versus copied data. By default, data copied to a TS7700 from another TS7700 is preferred to be removed from the cache, largest first, where data created by a host is preferred to keep in the cache by most recently modified/accessed.

The default behavior for cache preference is done so that where a host is connected to both TS7700s in a grid configuration, the effective cache size is the combination of both TS7700 TVCs. This way, more mount requests can be satisfied from the cache. These “cache hits” result in faster mount times because no physical tape mount is required. It does have the disadvantage in that in the case of a disaster, most of the recently copied logical volumes are not going to be resident in the cache at the recovery site, because they were copies and would have been managed to be removed from the cache.

You can ask the IBM SSR to modify this behavior. You can change the default cache management behavior to manage the copies to be kept in cache by most recently modified/accessed order, just as though the host directly created the data there.

9.6.8 Ownership takeover

If you are going to perform the test with the links broken between sites, you will need to enable Read Ownership Takeover so that the test site can access the data on the production volumes owned by the production site. Because the production volumes are created by mounting them on the production site's TS7700, that TS7700 will have volume ownership.

If you attempt to mount one of those volumes from the test system, without ownership takeover enabled, the mount will fail because the test site's TS7700 will not be able to request ownership transfer from the production site's TS7700. By enabling Read Ownership Takeover, the test host will now be able to mount the production logical volumes and read their contents.

The test host will not be able to modify the production site-owned volumes or change their attributes. The volume looks to the test host as a write protected volume. Because the volumes that are going to be used by the test system for writing data were inserted through the management interface associated with the TS7700 at the test site, that TS7700 will already have ownership of those volumes and the test host will have complete read and write control of them.

Note: Never enable Write Ownership Takeover mode for a test! Write ownership mode should only be enabled in the event of a loss or failure of the production TS7700.

If you are not going to break the links between the sites, then normal ownership transfer will occur whenever the test system requests a mount of a production volume

9.6.9 DR testing with Copy Export

The part of Copy Export Recovery that relates to a real disaster recovery is described in 9.4.1, “Execute Copy Export Recovery” on page 504. Read that chapter as well before starting with DR testing as described in the following pages.

Planning and considerations for testing Copy Export Recovery

There are several areas to consider before executing the Copy Export Recovery:

1. First, Copy Export Recovery requires an empty TS7700 that is not part of a Multi Cluster Grid. The TS7700 could be used for DR testing by different customers as long as only one DR test is done at a time. If data exists on the TS7700, an option is provided to erase that data. Refer to “Cleanup after testing Copy Export Recovery” on page 517 for a description of the erase process.
2. The hardware on the DR site must be compatible with the original setup where the physical tapes were created.
   - The recovery TS7700 must have physical tape drives that are capable of reading the physical volumes from a source TS7700. If a source TS7700 writes using the native E05 format, then the recovery TS7700 must also have 3592 E05 drives running in native format mode.
   - If the exporting pool on the source TS7700 is set up to encrypt the data, then the recovery TS7700 must also be able to handle encrypted volumes and have access to the encryption key manager with replicated keys from the production site.

– If the source TS7700 writes the volumes in J1A or emulated J1A mode, then any 3592 model drive in the recovery TS7700 can read the data.

– Bandwidth, cache size and number of drives in the Single Cluster Grid do not need to be the same as for production. It could be more financially attractive to choose a small configuration for the DR test, although you have to accept a degrade in performance.

Cleanup after testing Copy Export Recovery

After DR testing is done, the last part that is needed is to remove the physical tapes and do an erase of the TS7700 database to be prepared for the next DR testing or for a real disaster.

1. Remove the physical tapes manually by:
   – Switching the library into pause mode
   – Opening a door and removing the physical tapes
   – Closing the door and switching the library into auto mode again
   – Returning the tapes to their offsite location
2. Make sure that the TS7700 is online but with all drives offline. If a DR test is finished and the DR host stopped, then this is already done.
3. With the TS7700 online, use the MI and select **Copy Export Recovery** as in Figure 9-15 on page 517.

Figure 9-15 Menu with the Copy Export Recovery possibility

4. If there is data on the TS7700 and database entries, Figure 9-16 will appear telling you to take action. If you are sure that all data must be erased, then check the “Erase all existing volumes before the recovery” check box and click **Submit**. If the TS7700 is empty, the next page will not appear. At this point you must be sure that you are running on the right TS7700.

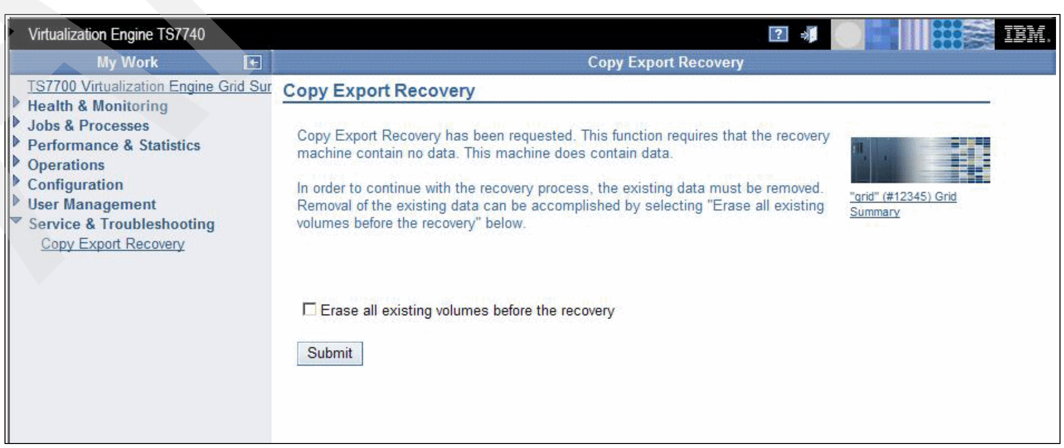

Figure 9-16 Menu with Copy Export Recovery initiation

5. Type the password and click **Submit** on the menu in Figure 9-17.

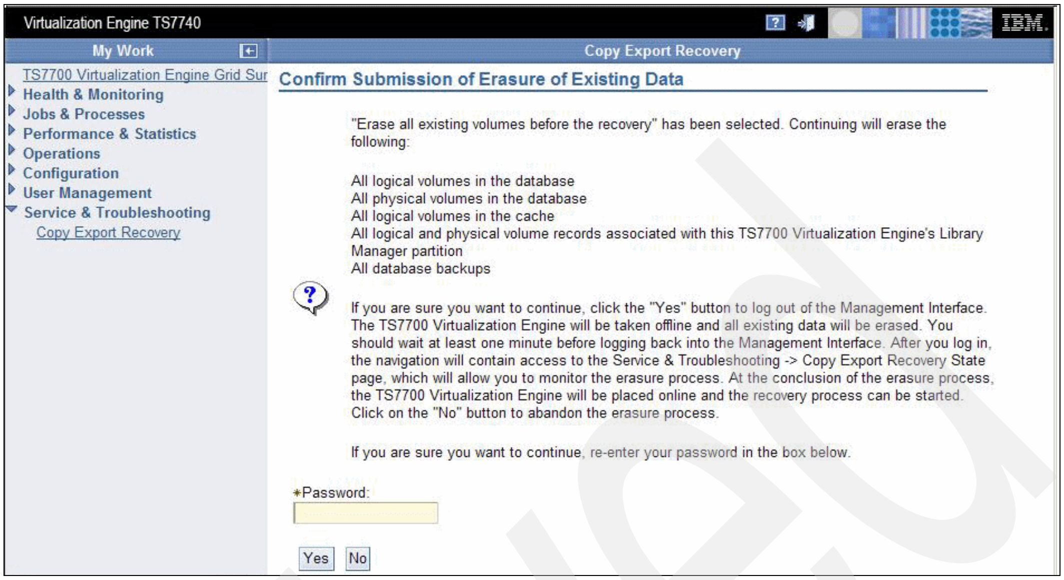

Figure 9-17 Menu with Confirm submission of Copy Export Recovery

6. The TS7700 will initiate the erase. Data and all database records will be erased and you will be logged off the MI.
7. Wait a while (one minute) and log on to the MI. The only available selection right now is the Service & Troubleshooting menu shown in Figure 9-18. From there you can select **Copy Export Recovery** where you can follow the progress of the erase. The menu shows the total number of steps to complete and the current state.

Figure 9-18 Menu with Copy Export Recovery State - Erase

Use the Refresh button to refresh the status in the menu in Figure 9-18. The following tasks will show:

– The TS7700 will be taken offline.
– The existing data in the TS7700 cache is being removed.
– The existing data in the library partition associated with this TS7700 is being removed.
– Cleanup (removal) of existing data.
– Requesting the TS7700 to go online.

8. If errors occur during the erasure, error messages will appear. Common for all of these is that you should call for IBM support.

After completion, the TS7700 will be online and you can continue with the Copy Export Recovery as described in 9.4.1, “Execute Copy Export Recovery” on page 504, even though you are doing a DR test.

A

Availability configurations

In this appendix, we describe how you can use a TS7700 Multi Cluster Grid configuration to improve the data availability in your data center. Data availability is a key issue for any organization. Your current information and communication technology (ICT) infrastructure already has been implemented with data availability as an important design objective.

Data availability has many aspects, ranging from outages for software maintenance to interrupts caused by hardware failures. Before we describe the new opportunities that you have with the TS7700 Multi Cluster Grid configuration, we will discuss and clarify the continuous availability terminology in use in the first part of this chapter.

Disaster recovery adds another dimension to continuous availability. There are several levels of disaster recovery distinguished by the level of data loss and service loss. The relationship between disaster recovery and continuous availability is also discussed. For an in-depth discussion of these topics, see *Continuous Availability - Systems Design Guide*, SG24-2085, and *Continuous Availability S/390 Technology Guide*, SG24-2086. We hope that these discussions will help you to decide what level of availability you want to seek and how you might begin to achieve your availability objectives for tape by implementing a TS7700 Multi Cluster Grid configuration.

The TS7700 Multi Cluster Grid configuration uses Gigabit Ethernet connections between its components. It is important that the technical details of Gigabit Ethernet and its consequences are well understood before you plan a configuration.

We describe three single-site configurations and four dual-site configurations in more detail. For each configuration, we document the availability targets, the configuration details and the configuration considerations in detail. The configurations might assist you in determining the preferred configuration for your environment.

A TS7700 Multi Cluster Grid configuration helps you to improve your data availability for tape data. As a consequence, it can result in changing your current implementation in terms of infrastructure and daily procedures.

We also describe the FIbre CONnection channel architecture. It provides information about distance issues and channel configuration alternatives using fiber optic technology that can be used for a TS7700 Multi Cluster Grid configuration attachment.

More information about using FICON is available in the following publications:

- *Introduction to IBM S/390 FICON*, SG24-5176
- *Continuous Availability S/390 Technology Guide*, SG24-20866
- *IBM System z Connectivity Handbook*, SG24-5444
- *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-44099
- *Fiber Saver (2029) Implementation Guide*, SG24-5608
- *Introduction to IBM S/390 FICON*, SG24-5176

Continuous availability concepts

Computer outages are less and less accepted in today's data centers. More and more business applications have to be online 24 hours a day and 365 days a year. Outages can have a serious impact on the performance of an enterprise and should be kept to a minimum.

Sometimes outages are caused by hardware, software, and data maintenance work, which needs to be performed at regular intervals. Such planned outages are usually scheduled at a period of low activity.

Sometimes outages are caused by faulty hardware, software, and procedures. Such *unplanned* outages have become quite rare in recent years with advanced multi-user systems.

Today, an increasing number of enterprises require that both types of computer outages, planned and unplanned, be eliminated or substantially reduced. An advanced computer system is not a single box, but a complex assembly of such components as processors, peripherals, networks, operating systems, and application software. Consequently, an outage-free system cannot be bought off the shelf. It requires an appropriate systems design, using elements such as redundant and fault-tolerant components and appropriate systems management.

System availability terms

We use three terms to describe the concept of system availability: *high availability*, *continuous operation*, and *continuous availability*. These terms are now explained.

- High availability

  This is the ability of a system to provide service to its users during defined service periods, at an acceptable or agreed level.

  These service periods, as well as the definition of an "acceptable" service level, are either stated in a service level commitment by the service provider, or in a service level agreement between end users and the service provider.

  Typically, a service level above 99.7 percent is accepted as high availability. High availability is maintained by avoiding or reducing any unplanned outages.

- Continuous operation

  This is the ability of a system to provide service to its users at all times, day and night, without scheduled outages to perform system and data maintenance activities.

  It is obviously difficult to perform change and maintenance work on a system that is supposed to be in continuous operation. However, without preventative maintenance, a system can be in continuous operation, but its availability might not be as high as it should be because it might suffer unscheduled outages more frequently.

Even if true continuous operation turns out to be impossible to implement for some applications, a realistic goal might be to increase the defined service period, for example, from 14 to 18 hours per day.

- Continuous availability

  This is the property of a system that provides both high availability and continuous operation at the same time. The system must be designed so that users do not experience scheduled or unscheduled outages.

  This goal seems difficult to achieve, as hardware and software components are usually not entirely error-free and maintenance-free, and large computer systems undergo frequent component additions and changes. The solution is to employ hardware components, software, and operational procedures that mask outages from the user. This solution usually requires that recovery from an outage must be performed so quickly that the user does not perceive it as an outage. It also frequently requires the use of redundant components, so that an alternate component can be used in case of a permanent component failure, or while a component is in maintenance.

An additional term, which you often hear used with high availability and is sometimes erroneously used instead of high availability, is *fault tolerant*.

Many major components of today's computer systems are fault tolerant to some degree, which means they will tolerate some faults. These components might have:

- Redundant sub-components
- Error checking and correction for data
- Retry capabilities for basic operations
- Alternate paths for I/O requests
- Duplexed data facilities on DASD and tape

However, they might also have a single point of failure that, despite the fault tolerance, can cause these components to fail. Similarly, if one important component in a system is not fault-tolerant, then the system is not fault-tolerant even though all other components are.

Note: A set of fault-tolerant components or products does not necessarily make a fault-tolerant system. A fault-tolerant system design will result in a fault-tolerant system only if it is implemented and managed properly.

Figure A-1 illustrates the relationship between the components of continuous availability.

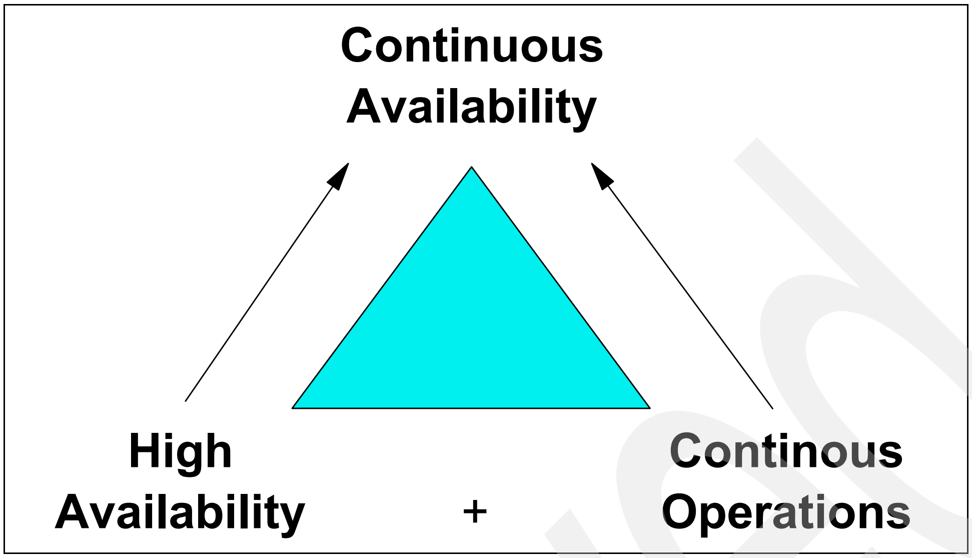

Figure A-1 Components of continuous availability

Disaster recovery

Disaster recovery is the process of reacting to a disaster by being able to provide computing services from another location. In most cases, the countermeasures you employ to be able to recover from a disaster are entirely different from the solution you use to achieve continuous availability.

In a disaster situation, users normally are aware that an outage has happened to the central computer facility, and the duration of the outage is mainly dependent on the recovery solution.

Usually this duration is measured in two different components:

- Data loss

  This represents the loss of data you have, that is, how much work you must re-execute when your system is recovered.

- Service loss

  This represents the loss of computing you experienced from the moment of disaster up to the moment when your system has been recovered.

At SHARE 78 held in Anaheim in 1992, session M028, the Automated Remote Site Recovery Task Force presented seven tiers of recoverability, which were ranked based on the recovery method and recovery time. The following sections describe these tiers.

Tier 0: No off-site data

This tier provides no preparation in saving information, determining requirements, establishing a backup hardware platform, or developing a contingency plan (Figure A-2).

The typical recovery time is unpredictable. In fact, you cannot guarantee recovery at all.

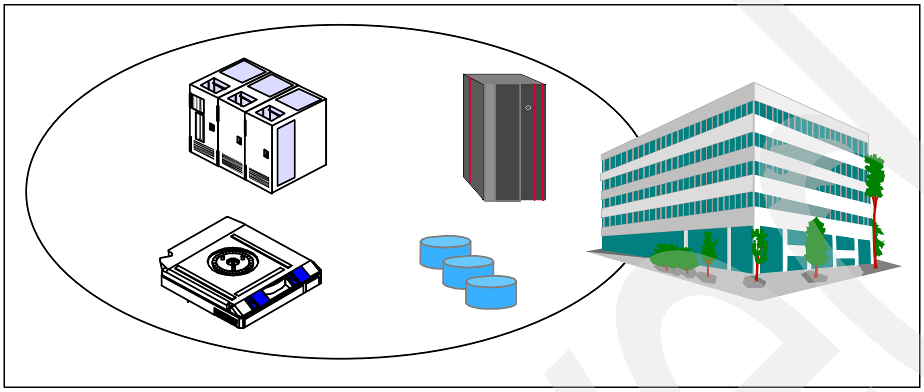

Figure A-2 Tier 0: No recovery solution

Tier 1: Pickup truck access method (PTAM)

To be at Tier 1, an installation would need to develop a contingency plan, back up required information, and store it in contingency storage (at an off-site location), determine recovery requirements, and optionally establish a backup platform supporting a conditioned facility without processing hardware (Figure A-3).

The typical recovery time is usually more than a week.

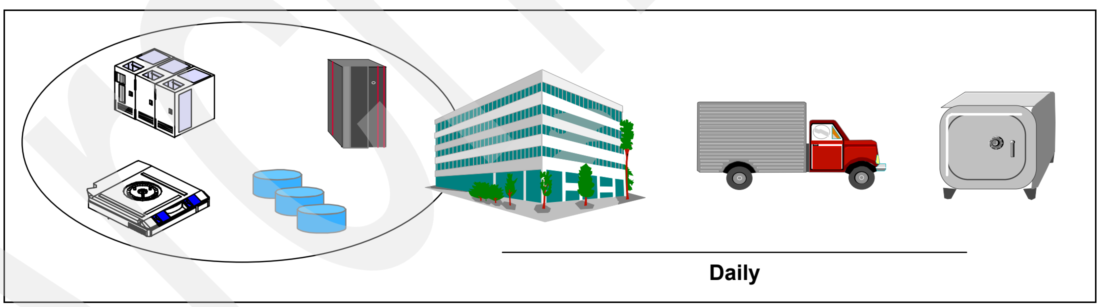

Figure A-3 Tier 1 recovery solution

Tier 2: PTAM with hot site

Tier 2 encompasses all requirements of Tier 1 and also requires a backup platform to have sufficient hardware and network to support the installation's critical processing requirements. Processing is considered critical if it must be supported on hardware that exists at the time of the disaster (Figure A-4).

The typical recovery time is usually more than one day.

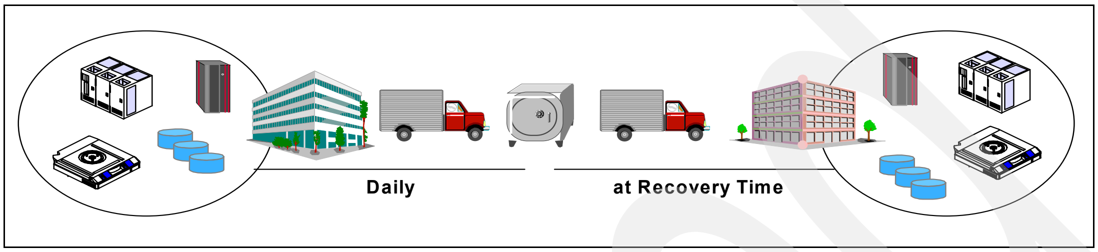

Figure A-4 Tier 2 recovery solution

Tier 3: Electronic vaulting

Tier 3 encompasses all the requirements of Tier 2 and, in addition, supports electronic vaulting of some subset of the information. The receiving hardware must be physically separated from the primary platform and the data stored for recovery after the disaster (Figure A-5).

The typical recovery time is usually about one day.

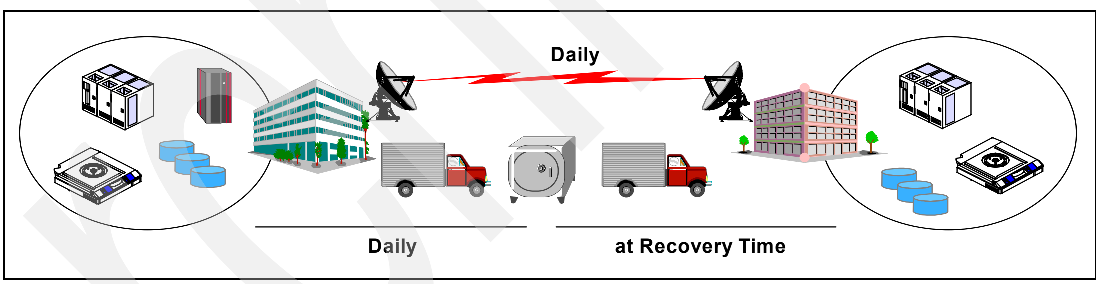

Figure A-5 Tier 3 recovery solution

Tier 4: Active secondary site

Tier 4 introduces the requirements of active management of the recovery data by a processor at the recovery site, and bi-directional recovery. The receiving hardware must be physically separated from the primary platform (Figure A-6).

The typical recovery time is usually up to one day.

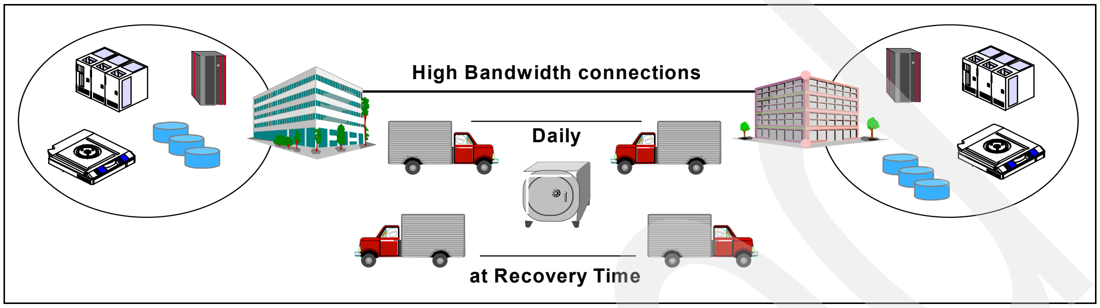

Figure A-6 Tier 4 recovery solution

Tier 5: Two-site two-phase commit

Tier 5 encompasses all the requirements of Tier 4 and, in addition, will maintain selected data in image status (updates will be applied to both the local and remote copies of the databases within a single commit scope).

Tier 5 requires both the primary and secondary platforms' data to be updated before the update request is considered satisfied. Tier 5 requires partially- or fully-dedicated hardware on the secondary platform, with the capability to automatically transfer the workload to the secondary platform (Figure A-7).

The typical recovery time is usually less than 12 hours.

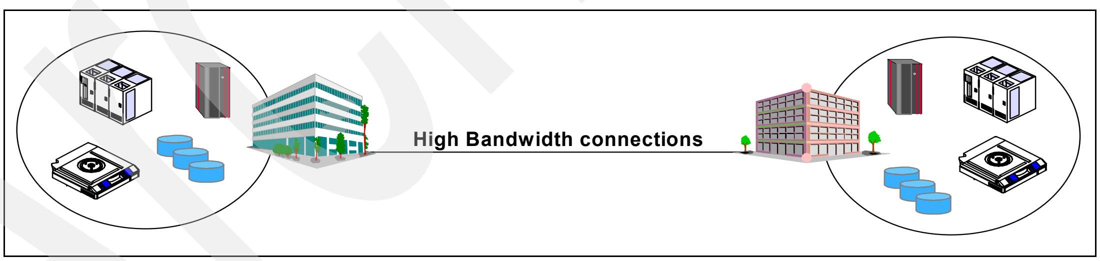

Figure A-7 Tier 5 recovery solution

Tier 6: Zero data loss

Tier 6 encompasses zero loss of data and immediate and automatic transfer to the secondary platform. Data is considered lost if ENTER has been accepted (at the terminal), but the request has not been satisfied (Figure A-8).

The typical recovery time is usually a few minutes.

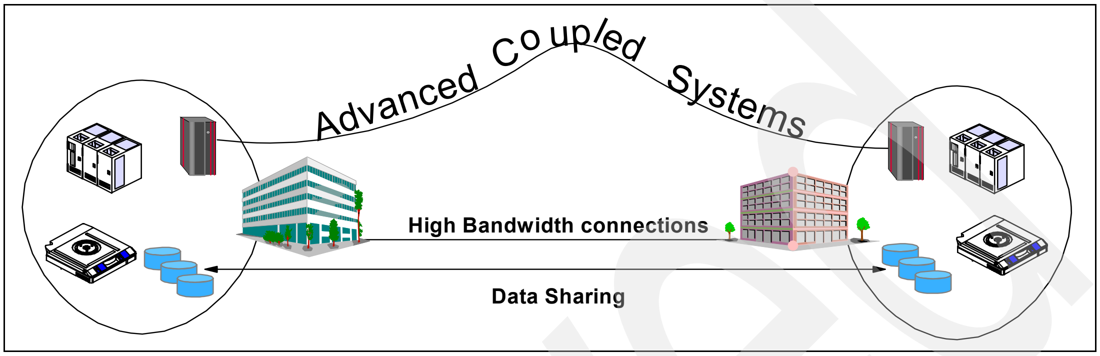

Figure A-8 Tier 6 recovery solution

TS7700 Dual Cluster Grid as a Tier 3 configuration

Figure A-9 shows a TS7700 Dual Cluster Grid configuration as a Tier 3 implementation. This kind of installation focuses on high availability of the stored data.

The two TS7700 Virtualization Engines are typically located at one site and are interconnected through a Local Area Network. The one or more local hosts are attached to both TS7700 Virtualization Engines. Therefore, data is available through either TS7700 cluster in case of a failure.

A ownership takeover has to be enabled before all data can be accessed. This is done through the TS7700 Management Interface.

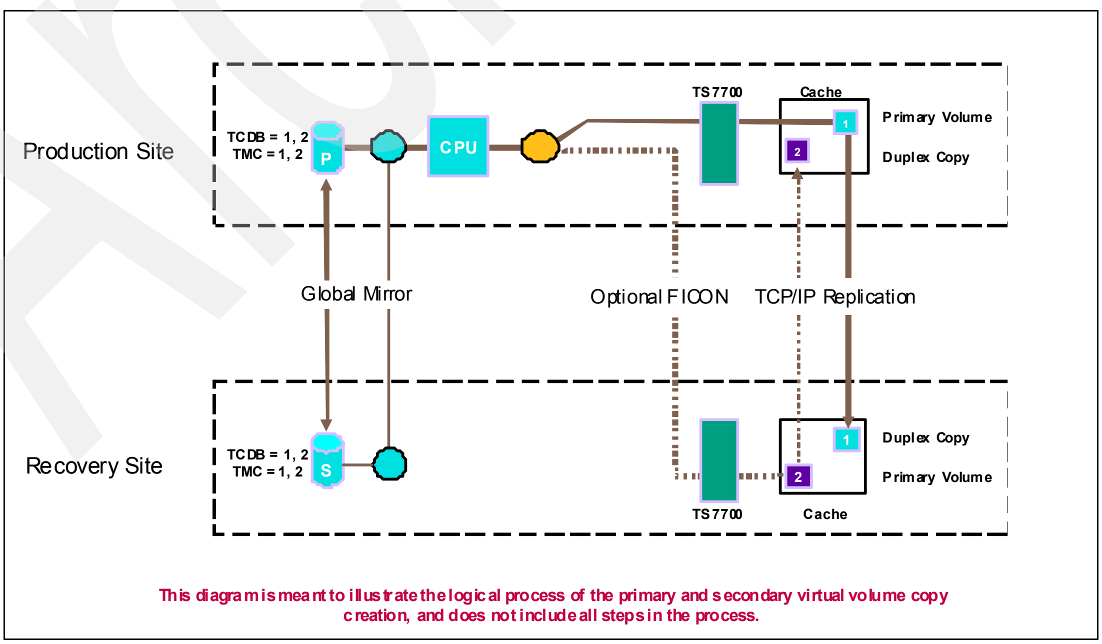

Figure A-9 TS7700 as a Tier 3 configuration

TS7700 Dual Cluster Grid as a Tier 5 configuration

In a Tier 5 configuration, the two TS7700 Virtualizations Engines are located at two sites geographically separated. The interconnection is accomplished through a Wide Area Network. This tier always requires a second host in the remote location.

There can be two flavors of a Tier 5 implementation:

- The first implementation form is for disaster recovery, where only the second host for disaster recovery is connected to the remote TS7700 cluster. If the local TS7700 cluster is unavailable, the data is only accessible at the remote TS7700 cluster. An ownership takeover needs to be enabled beforehand here as well.
- The other Tier 5 setup focuses on disaster recovery and availability. Therefore, both hosts connect to each of the remote TS7700 clusters through channel extended FICON interfaces making data available through either TS7700. You need to vary online the remote devices only when needed. Ownership Takeover has to be enabled before as well in this scenario.

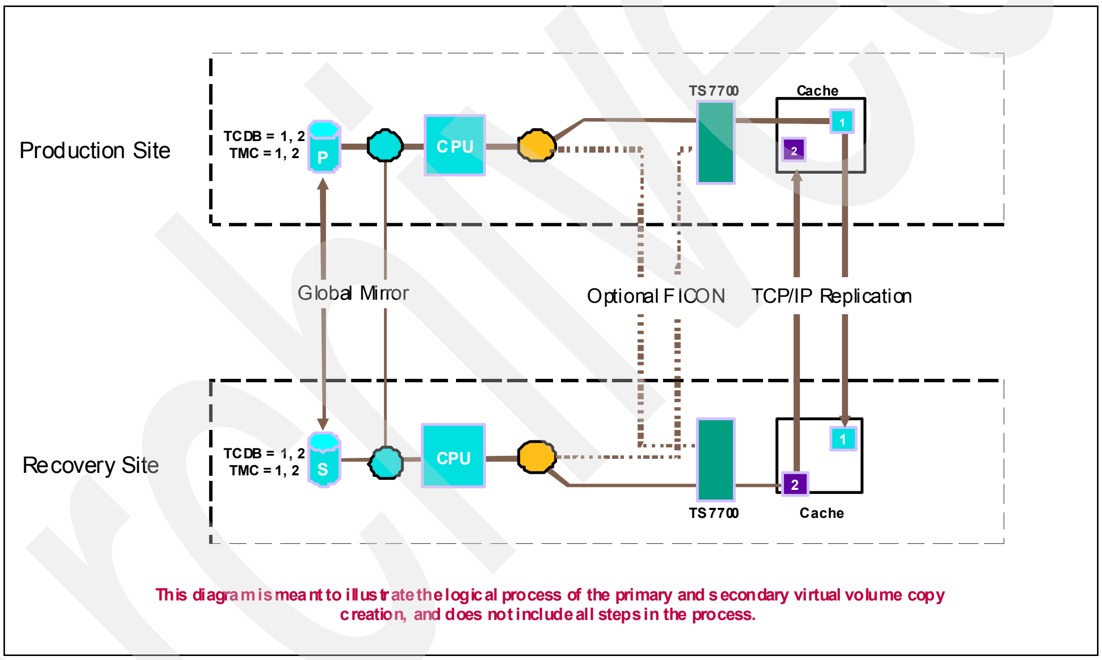

Figure A-10 TS7700 as a Tier 5 configuration

Data loss and service loss

The typical recovery time associated with each tier is just a rough indication of the time that an installation usually needs to restore its computing services. However, in a disaster situation, there are many other points to consider.

For example, some installations can tolerate resuming their services after longer periods of time, but with maximum data currency. Other installations must resume their services as soon as possible, regardless of the currency of their data. Still others need both a short recovery time and maximum data currency (Figure A-11).

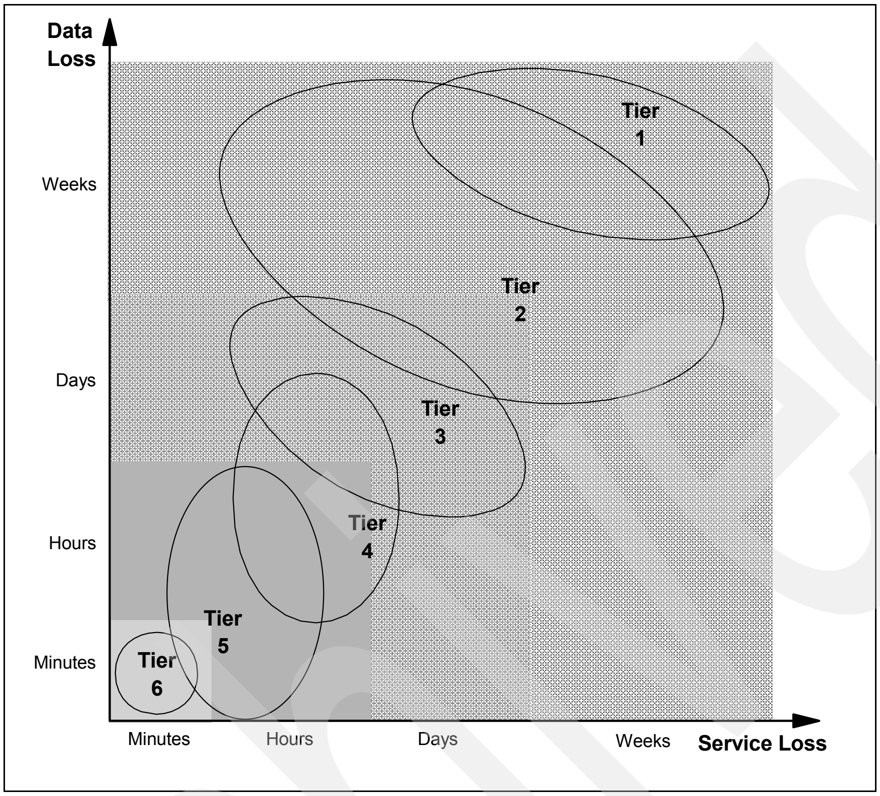

Figure A-11 Data loss and service loss

Relating continuous availability and disaster recovery

All the components that make up continuous availability in a computer system are usually situated in the same building. Therefore, the building itself can represent a single point of failure. In this case, and despite all the continuous availability you have designed into the system, a disaster could cause you to lose all your computing services.

While you must of course be prepared to react to a disaster, the solution you would apply in this case might be more of a recovery solution than a continuous availability solution. A recovery solution can be defined by a trade-off between implementation costs, maintenance costs, and financial impacts resulting from a business impact analysis of your business.

Furthermore, as you can see from the tier definitions, only a disaster recovery Tier 6 solution can be compared to a continuous availability solution, although currently there is limited technology available that fully fulfills that definition. However, with the current z/OS technologies, it is possible to achieve a Tier 6 workload transfer if you have enough time to do it in a controlled fashion before all computing capabilities are lost in a disaster at your production site.

B

Library Manager volume categories

Table B-1 lists all default library manager volume categories, the platforms on which they are used, and their definitions.

> **Note:** z/OS users might define any category up to X'FEFF' with the DEVSUPxx member SYS1.PARMLIB. The appropriate member must be pointed to by IEASYSxx.

Table B-1 Library manager volume categories

Category (in hex)	Used by	Definition
0000	Null Category	This pseudo category is used in certain library commands to specify that the category that is already associated with the volume is to be used by default, or that no category is specified. Use of the null category does not affect the volume's order within the category to which it is assigned. No volumes are associated with this category.
0001	DFSMS/MVS	Indicates scratch MEDIA1. MEDIA1 is a standard-capacity cartridge system tape.
0002	DFSMS/MVS	Indicates scratch MEDIA2. MEDIA2 is an enhanced-capacity cartridge system tape.
0003	DFSMS/MVS	Indicates scratch MEDIA3. MEDIA3 is the IBM TotalStorage Enterprise 3590 High Performance Tape Cartridge.
0004	DFSMS/MVS	Indicates scratch MEDIA4. MEDIA4 is the IBM TotalStorage Enterprise 3590 Extended High Performance Tape Cartridge.
0005	DFSMS/MVS	Indicates scratch MEDIA5. MEDIA5 is the IBM TotalStorage Enterprise Tape Cartridge 3592 DATA.
0006	DFSMS/MVS	Indicates scratch MEDIA6. MEDIA6 is the IBM TotalStorage Enterprise Tape Cartridge 3592 WORM.

Category (in hex)	Used by	Definition
0007	DFSMS/MVS	Indicates scratch MEDIA7. MEDIA7 is the IBM TotalStorage Enterprise Tape Cartridge 3592 ECONOMY.
0008	DFSMS/MVS	Indicates scratch MEDIA8. MEDIA8 is the IBM TotalStorage Enterprise Tape Cartridge 3592 ECONOMY WORM.
0009	DFSMS/MVS	Indicates scratch MEDIA9. MEDIA8 is the IBM System Storage Extended Enterprise Tape Cartridge 3592.
000A	DFSMS/MVS	Indicates scratch MEDIA10. MEDIA 10 is the IBM System Storage Extended Enterprise Tape Cartridge 3592 Economy.
000B to 000D	DFSMS/MVS	Reserved
000E	DFSMS/MVS	Indicates an error volume. Volumes in this category are scratch volumes for which the software detected an error during processing.
000F	DFSMS/MVS	Indicates a private volume. Volumes in this category contain user data or are assigned to a user.
0010 to 007F	DFSMS/MVS	Reserved. These volume categories can be used for library partitioning.
0080	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH0.
0081	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH1.
0082	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH2.
0083	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH3.
0084	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH4.
0085	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH5.
0086	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH6.
0087	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH7.
0088	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH8.
0089	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCH9.
008A	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHA.
008B	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHB.
008C	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHC.
008D	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHD.

Category (in hex)	Used by	Definition
008E	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHE.
008F	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHF.
0090 to 009F	-	Currently not assigned.
00A0	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH00.
00A1	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH01.
00A2	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH02.
00A3	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH03.
00A4	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH04.
00A5	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH05.
00A6	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH06.
00A7	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH07.
00A8	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH08.
00A9	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH09.
00AA	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH10.
00AB	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH11.
00AC	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH12.
00AD	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH13.
00AE	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH14.
00AF	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH15.
00B0	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH16.
00B1	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH17.
00B2	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH18.
00B3	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH19.

Category (in hex)	Used by	Definition
00B4	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH20.
00B5	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH21.
00B6	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH22.
00B7	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH23.
00B8	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH24.
00B9	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH25.
00BA	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH26.
00BB	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH27.
00BC	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH28.
00BD	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH29.
00BE	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH30.
00BF	Native z/VSE	Indicates that the volume belongs to the VSE category SCRATCH31.
00C0 to 00FF	-	Currently not used.
0100	OS/400® (MLDD)	Indicates that the volume has been assigned to category *SHARE400. Volumes in this category can be shared with all attached iSeries® and AS/400® systems.
0101	OS/400 (MLDD)	Indicates that the volume has been assigned to category *NOSHARE. Volumes in this category can be accessed only by the OS/400 system that assigned it to the category.
0102 to 012B	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager on request of a host.
012C	Tivoli Storage Manager for AIX®	Indicates a private volume. Volumes in this category are managed by Tivoli Storage Manager.
012D	Tivoli Storage Manager for AIX	Indicates an IBM 3490 scratch volume. Volumes in this category are managed by Tivoli Storage Manager.
012E	Tivoli Storage Manager for AIX	Indicates an IBM 3590 scratch volume. Volumes in this category are managed by Tivoli Storage Manager.
012F to 0FF1	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager on request of a host.
0FF2	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH2.

Category (in hex)	Used by	Definition
0FF3	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH3.
0FF4	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH4.
0FF5	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH5.
0FF6	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH6.
0FF7	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH7.
0FF8	BTLS	Indicates a scratch volume. Volumes in this category belong to the optional scratch pool SCRTCH8.
0FF9 to 0FFE	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager on request of a host.
0FFF	BTLS	Indicates a scratch volume. Volumes in this category belong to the default scratch pool used by BTLS. **Note**: If you are planning to migrate to DFSMS/MVS, you should use this default scratch category only.
1000 to F00D	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager on request of a host.
F00E	BTLS	Indicates a volume in error. Volumes are assigned to the error category during demount if the volume serial specified for demount does not match the external label of the volume being demounted.
F00F to FEFF	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager on request of a host.
FF00	All	Insert category. When a tape volume is added to an automated tape library, the library reads the external label on the volume, creates an inventory entry for the volume and assigns the volume to the insert category. This category might be updated by operator interaction through Librarian Workstation Support.
FF01	Virtual Tape Server & TS7700 Virtualization Engine	Stacked Volume Insert category for a Virtual Tape Server & TS7700 Virtualization Engine. A volume is set to this category when its volume serial number is in the range specified for stacked volumes for any VTS library partition.
FF02	Virtual Tape Server	Stacked Volume Scratch category 0 for a Virtual Tape Server & TS7700 Virtualization Engine. This category is reserved for future use for scratch stacked volumes.
FF03	Virtual Tape Server	Stacked Volume Scratch category 1 for a Virtual Tape Server & TS7700 Virtualization Engine. This category is used by the VTS for its scratch stacked volumes. This category is not used if LIC is 527 or higher.

Category (in hex)	Used by	Definition
FF04	Virtual Tape Server & TS7700 Virtualization Engine	Stacked Volume Private category for a Virtual Tape Server & TS7700 Virtualization Engine. This category includes both scratch and private volumes (since VTS LIC level 527).
FF05	Virtual Tape Server & TS7700 Virtualization Engine	Stacked Volume Disaster Recovery category for a Virtual Tape Server & TS7700 Virtualization Engine. A volume is set to this category when its volume serial number is in the range specified for stacked volumes for any VTS library partition *and* the Library Manager is in disaster recovery mode.
FF06	Virtual Tape Server & TS7700 Virtualization Engine	This category is used by the VTS as a temporary category for disaster recovery. When a stacked volume in category FF05 is processed, it is put into this category. This category is also used be the PFE tool called “movedata” as a temporary category.
FF07	Virtual Tape Server & TS7700 Virtualization Engine	This category is reserved for future hardware functions.
FF08	Virtual Tape Server	This category is used by the VTS to indicate a Read-Only-Recovery Stacked Volume with active data cannot be recovered.
FF09 to FF0F	-	Reserved for future hardware functions
FF10	Library manager	Convenience-Eject category. When a tape volume is assigned to the convenience-eject category, it becomes eject pending and the Library Manager queues the tape volume to be moved to a convenience output station. When the volume is delivered to an output station, it is deleted from the Library Manager’s inventory. **Note:** Logical volumes cannot be ejected from the library. They can be deleted or exported.
FF11	Library manager	Bulk-Eject category. Set when the Library Manager accepts an eject request. The volume becomes eject pending and is queued to be moved to the high capacity output station. When the cartridge accessor delivers the volume to the output rack, it is deleted from the Library Manager’s inventory. **Note:** Logical volumes cannot be ejected from the library. They can be deleted or exported.
FF12	Virtual Tape Server	Export-Pending category. A logical volume to be exported is assigned to this category at the beginning of a Virtual Tape Server export operation. Logical volumes in this category are considered in use. Any attempt by a host to mount, audit, or change the category of a volume fails.

Category (in hex)	Used by	Definition
FF13	Virtual Tape Server	Exported category. Set when the Virtual Tape Server has exported the logical volume. The attached hosts are notified when volumes are assigned to this category. Any attempt by a host to mount, audit, or change the category of a volume fails, except a Library Set Volume Category order assigning the volume to the purge-volume category.
FF14	Virtual Tape Server	Import category. Stacked volumes that contain logical volumes to import into the Virtual Tape Server are assigned to this category by an operator at the Library Manager, after they are entered into the library through the convenience I/O station and placed in the Unassigned category.
FF15	Virtual Tape Server	Import-Pending category. Logical volumes to be imported from a stacked volume are added to the Library Manager inventory and assigned to this category when the Virtual Tape Server starts importing them. At completion, successfully imported volumes are assigned to the insert category (FF00). The attached hosts are then notified of volumes assigned to the insert category. Any host attempt to use a volume assigned to this category will be failed.
FF16	Virtual Tape Server & TS7700 Virtualization Engine	Unassigned Category. Volumes are assigned to this category by the Library Manager whenever volumes are added to the library through the convenience I/O station and the library contains one or more VTS subsystems that have the Import/Export functions installed and enabled. Manual intervention is required to assign the cartridges to the proper category. For exported stacked volumes, this is the import category (FF14).
FF17	Virtual Tape Server	Export-Hold category. Physical volumes are assigned to this category on completion of processing for an export stacked volume.
FF18 & FF19	-	Reserved for library. These categories are reserved for future hardware functions.
FF20	PtP Virtual Tape Server & TS7700 Virtualization Engine	Corrupted-Token Volume Category. In a Peer to Peer VTS, volumes are assigned to this category when it has been determined that the tokens associated with the volume have been corrupted. This is to prevent the volume from being selected by a category mount request.
FF21 to FFF5	-	Reserved for library. These categories are reserved for future hardware functions.
FFF4	Library manager	3592 Cleaner Volume. Cleaner volumes for 359 type devices in the library are assigned to this category automatically.
FFF5	Library manager	3592 Service Volume. Volumes are assigned to this category by the Library Manager when it detects that a volume has a unique service cartridge volser and a media type compatible with a 3592 device.

Category (in hex)	Used by	Definition
FFF6	Library manager	3590-Service-Volume Category. Volumes are assigned to this category by the Library Manager when it detects that a volume has a unique service cartridge VOLSER and a media type compatible with a 3590 device.
FFF7 & FFF8	-	Reserved for library. These categories are reserved for internal library functions.
FFF9	Library manager	3490-Service-Volume Category. Volumes are assigned to this category by the Library Manager when it detects that a volume has a unique service cartridge VOLSER and a media type compatible with a 3490 device.
FFFA	Library manager	Manually-Ejected Category. Volumes are assigned to this category when they are removed from the library under the control of an operator, not the control program. Volumes in this category are no longer available for any other operations expect purge-volume category assignment.
FFFB	Library manager	Purge-Volume Category. When this category is specified in a Perform Library Function command with the Library Set Volume Category order and the volume is either in the misplaced state, is assigned to the exported category or is assigned to the manually-ejected category, the specified VOLSER's record is deleted from the inventory. No volumes are associated with this category.
FFFC	Library manager	Unexpected-Volume Category. This category is reserved for future use.
FFFD	Library manager	3590-Cleaner-Volume Category. Cleaner volumes for 3590 type devices in the library are assigned to this category automatically.
FFFE	Library manager	3490-Cleaner-Volume Category. Cleaner volumes for 3490 type devices in the library are assigned to this category automatically.
FFFF	Library manager	Volser-Specific Category. This category is for general use by programming except that any Library Mount request to this category must be for a specific VOLSER and not based on the category only.

C

Feature codes

In this Appendix, we list all the feature codes (FC) related to the installation of a TS7700 Virtualization Engine. We provide a quick reference list as well as more detailed explanations.

Quick reference list

This is a quick reference list of all feature codes for the TS7700 Virtualization Engine by component machine type and model:

FC0201	9 micron LC/LC 31 meter - TS7740 Node
FC0202	9 micron LC/SC 31 meter - TS7740 Node
FC0203	50 micron LC/LC 31 meter - TS7740 Node
FC0204	50 micron LC/SC 31 meter - TS7740 Node
FC0205	62.5 micron LC/LC 31 meter - TS7740 Node
FC0206	62.5 micron LC/SC 31 meter - TS7740 Node
FC0521	Functional enhancement field - TS7740 Node
FC0522	Move VTS database (available until 5 December 2008)
FC0523	Move VTS database and cartridges (available until 5 December 2008)
FC1030	1 Gb Grid copper connection - TS7740 Node
FC1031	1 Gb Grid optical SW connection - TS7740 Node
FC1903	Dual AC power - IBM 3952 tape frame
FC2714	Console expansion - TS7740 Node
FC2715	Console attachment - TS7740 Node
FC2719	Console upgrade - TS7740 Node
FC2720	TS3000 System Console - TS7740 Node
FC3441	FICON shortwave attachment - TS7740 Node
FC3442	FICON longwave attachment - TS7740 Node
FC3443	FICON 10 km longwave attachment - TS7740 Node
FC3488	4 Gb Fibre Channel Switch
FC4015	Grid enablement
FC4016	Remove Cluster from Grid
FC4017	Cluster cleanup
FC5240	Attach 3592 Tape Drives - TS7740 Node
FC5267	1 TB cache enablement - TS7740 Node
FC5268	100 MB/s increment - TS7740 Node
FC5628	Plant install 3957 V06 - IBM 3952 tape frame
FC5638	Plant install 3956 CC6 - IBM 3952 tape frame
FC5648	Plant install 3956 CX6 - IBM 3952 tape frame
FC5649	Field install 3956 CX6 - IBM 3952 tape frame
FC5759	Integrated control path - IBM 3952 tape frame
FC6000	Intraframe fibre cable to 3956-CC6
FC6003	Intraframe fibre cable to 3957 V06 - TS7740 cache controller
FC7120	1.7 TB fibre storage, for the TS7740 cache controller
FC7120	1.7 TB fibre storage, for the TS7740 cache drawer
FC7312	TS7700 base frame - IBM 3952 tape frame
FC9000	Mainframe attachment - TS7740 Node
FC9010	Attach to a VTS
FC9013	Attach to TS7700
FC9217	Attach to 3953 LM - TS7740 Node
FC9218	Attach to 3494 LM - TS7740 Node
FC9230	Attach to 3957 V06 - TS7740 cache controller
FC9350	Plant install V06 in F05 - TS7740 Node
FC9352	Plant install CC6 in 3952 F05 - TS7740 cache controller
FC9354	Plant install 3956-CX6 in 3952 F05
FC9355	Field merge CX6 in 3952 F05
FC9700	No factory cables - TS7740 Node
FC9900	Encryption configuration - TS7740 Node
FC9954	NEMA L6-30 power cord - IBM 3952 tape frame
FC9955	RS 9750 DP power cord - IBM 3952 tape frame

FC9956	IEC 309 power cord - IBM 3952 tape frame
FC9957	PDL 4.3 power cord - IBM 3952 tape frame
FC9958	Korean 4.3-m power cord - IBM 3952 tape frame

Feature code details

In this section, we list the feature codes under the machine type and model together with a short description.

3952 Frame features

This section describes each feature code for the 3952 tape frame in detail.

FC1903, Dual AC power

FC1903, Dual AC power, provides a second power distribution unit (PDU) to allow connection to independent power sources. This feature supplies two power cords when ordered. See FC1903, Dual AC power, in Table C-1 on page 544 for requirements and other information.

FC5628, Plant install 3957 V06

FC5628, Plant install 3957 V06, allows plant installation of a TS7740 3957 V06 into a new 3952 Tape Frame. This feature must appear on the 3952 Tape Frame order. You must also order FC9350, Plant install V06 in F05, against the TS7740 3957 V06 when ordering this feature. See FC5628, Plant install 3957 V06, in Table C-1 on page 544 for requirements and other information.

FC5638, Plant install 3956 CC6

FC5638, Plant install 3956 CC6, allows plant installation of a TS7740 Cache Controller model CC6 into a new 3952 Tape Frame. This feature must appear on the 3952 Tape Frame order. You must also order FC9352, Plant install CC6 in 3952 F05, against the TS7740 Cache Controller model CC6 when ordering this feature. See FC5638, Plant install 3956 CC6, in Table C-1 on page 544 for requirements and other information.

FC5648, Plant install 3956 CX6

FC5648, Plant install 3956 CX6, allows plant installation of aTS7740 Cache Drawer model CX6 into a new 3952 Tape Frame. This feature must appear on the 3952 Tape Frame order. You must also order FC9354, Plant install 3956 CX6 in 3952 F05, against the TS7740 Cache Drawer model CX6 when ordering this feature. See FC5648, Plant install 3956 CX6, in Table C-1 on page 544 for requirements and other information.

FC5649, Field install 3956 CX6

FC5649, Field install 3956 CX6, allows field installation of a TS7740 Cache Drawer model CX6 into an installed 3952 Tape Frame. This feature must be ordered on the 3592 Tape Frame and the Field Merge a 3956 Model CX6 feature FC9355 must appear on the TS7740 Cache Drawer order.

For Release 1.3, when the 3952 Model F05 Frame contains the TS7740 Base Frame feature FC7312, valid quantities of FC5648 plus FC5649 must be 1 or 3. For Release 1.4, when the 3953 Model F05 Frame contains the TS7740 Base Unit feature FC7312, valid quantities of FC5648 plus 5649 are 0, 1, or 3.

See FC5649, Field install 3956 CX6, in Table C-1 on page 544 for requirements and other information.

FC5759, Integrated control path

FC5759, Integrated control path, provides the router, switches, and cables necessary to create the control path between the TS7700 Virtualization Engine and the Library Manager.

See FC5759, Integrated control path, in Table C-1 on page 544 for requirements and other information.

FC7312, Base unit TS7740

FC7312, Base unit TS7740, identifies this 3952 tape frame as the base unit for the TS7700 Virtualization Engine. See FC7312, Base unit TS7740, in Table C-1 on page 544 for requirements and other information.

FC9954, NEMA L6-30 power cord

FC9954, NEMA L6-30 power cord, provides a National Electrical Manufacturers Association (NEMA) L6-30 non-watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac or 240 V ac and 24 A. This power cord is suggested for use in the USA, Canada, Latin America, and Japan. See FC9954, NEMA L6-30 power cord, in Table C-1 on page 544 for requirements and other information.

FC9955, RS 9750 DP power cord

FC9955, RS 9750 DP power cord, provides a Russellstoll 3750 DP, watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac or 240 V ac and 24 A. This power cord is suggested for use in the USA (highly recommended in Chicago, Illinois, to conform with local requirements), Canada, Latin America, and Japan. See FC9955, RS 9750 DP power cord, in Table C-1 on page 544 for requirements and other information.

FC9956, IEC 309 power cord

FC9956, IEC 309 power cord, provides an International Electrotechnical Commission (IEC) 309 p+n+g 4.3-m (14-ft.) power cord, rated for 230 V ac and 24 A. This power cord is suggested for use in Europe, the Middle East, and Africa. See FC9956, IEC 309 power cord, in Table C-1 on page 544 for requirements and other information.

FC9957, PDL 4.3 power cord

FC9957, PDL 4.3 power cord, provides a PDL 4.3-m (14-ft.) power cord, rated for 230 V ac to 240 V ac and 24 A. This power cord is suggested for use in Australia and New Zealand. See FC9957, PDL 4.3 power cord, in Table C-1 on page 544 for requirements and other information.

FC9958, Korean 4.3-m power cord

FC9958, Korean 4.3-m power cord, provides a NEMA L6-30 non-watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac or 240 V ac and 24 A, with a South Korean plug. This power cord is suggested for use in South Korea. See FC9958, Korean 4.3-m power cord, in Table C-1 on page 544 for requirements and other information.

FC 9959, Unterminated power cord

FC 9959, Unterminated power cord, provides an unterminated, non-watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac and 24 Amp, with IRAM and BSMI agency certifications. This is the recommended cord for Argentina and Taiwan. See FC 9959, Unterminated power cord, in Table C-1 on page 544 for requirements and other information.

FC 9966, Unterminated power cord (China)

FC 9966, Unterminated power cord (China), provides an unterminated, non-watertight, 4.3-m (14-ft.) power cord, rated for 200 V ac to 208 V ac and 24 Amp, with CCC agency certifications. This is the recommended cord for China. See FC 9966, Unterminated power cord (China), in Table C-1 on page 544 for requirements and other information.

Table C-1 3952 tape frame feature codes

Feature	Minimum	Maximum	Requirements	Installation	Removable?
FC1903, Dual AC power	0	1	This feature is required if you are ordering FC7312.	Plant	No
FC5628, Plant install 3957 V06	0	1	Corequisite: FC9350, Plant install V06 in F05, against the IBM Virtualization Engine TS7700	Plant only	No
FC5638, Plant install 3956 CC6	0	1	Corequisite: FC9352, Plant install CC6 in 3952 F05, against the TS7700 cache controller	Plant only	No
FC5648, Plant install 3956 CX6	0	3	Corequisite: FC9354, Plant install 3956 CX6 in 3952 F05, against the TS7700 cache drawer	Plant only	No
FC5649, Field Install 3956 CX6	1	2	Prerequisite: FC5648, Plant install 3956 CX6 quantity one.	Field only	No
FC5759, Integrated control path	1	1	Prerequisite: FC7312	Plant	No
FC7312, Base unit TS7740	0	1	Mutually exclusive with FC7310, FC7311, 7315, and 7316. Prerequisite: FC1903. Corequisite with one FC5628, one FC5638, and three of FC5648.	Plant	No
FC9954, NEMA L6-30 power cord	0	1	One of FC9954, 9955, 9956, 9957, or 9958	Plant or Field	Not applicable
FC9955, RS 9750 DP power cord	0	1	One of FC9954, 9955, 9956, 9957, or 9958	Plant or Field	Not applicable
FC9956, IEC 309 power cord	0	1	One of FC9954, 9955, 9956, 9957, or 9958	Plant or Field	Not applicable
FC9957, PDL 4.3 power cord	0	1	One of FC9954, 9955, 9956, 9957, or 9958	Plant or Field	Not applicable
FC9958, Korean 4.3-m power cord	0	1	One of FC9954, 9955, 9956, 9957, or 9958	Plant or Field	Not applicable
FC 9959, Unterminated power cord	0	1	One of FC 9954, 9955, 9956, 9957, 9958, 9959, or 9966	Plant or Field	Not applicable
FC 9966, Unterminated power cord (China)	0	1	One of FC 9954, 9955, 9956, 9957, 9958, 9959, or 9966	Plant or Field	Not applicable

TS7740 Node features

This section describes each feature code for the TS7740 Node, machine type and model 3957 V06, in detail.

FC0201, 9 micron LC/LC 31 meter

FC0201, 9 micron LC/LC 31 meter, provides a 31 meter, 9 micron FICON cable with LC/LC connectors. See FC0201, 9 micron LC/LC 31 meter, in Table C-2 on page 550 for requirements and other information.

FC0202, 9 micron LC/SC 31 meter

FC0202, 9 micron LC/SC 31 meter, provides a 31 meter, 9 micron FICON cable with LC/SC connectors. See FC0202, 9 micron LC/SC 31 meter, in Table C-2 on page 550 for requirements and other information.

FC0203, 50 micron LC/LC 31 meter

FC0203, 50 micron LC/LC 31 meter, provides a 31 meter, 50 micron FICON cable with LC/LC connectors. See FC0203, 50 micron LC/LC 31 meter, in Table C-2 on page 550 for requirements and other information.

FC0204, 50 micron LC/SC 31 meter

FC0204, 50 micron LC/SC 31 meter, provides a 31 meter, 50 micron FICON cable with LC/SC connectors. See FC0204, 50 micron LC/SC 31 meter, in Table C-2 on page 550 for requirements and other information.

FC0205, 62.5 micron LC/LC 31 meter

FC0205, 62.5 micron LC/LC 31 meter, provides a 31 meter, 62.5 micron FICON cable with LC/LC connectors. See FC0205, 62.5 micron LC/LC 31 meter, in Table C-2 on page 550 for requirements and other information.

FC0206, 62.5 micron LC/SC 31 meter

FC0206, 62.5 micron LC/SC 31 meter, provides a 31 meter, 62.5 micron FICON cable with LC/SC connectors. See FC0206, 62.5 micron LC/SC 31 meter, in Table C-2 on page 550 for requirements and other information.

FC0521, Functional enhancement field

FC0521, Functional enhancement field, provides an update to the microcode of an installed TS7700 to provide enhanced functions contained in the latest level of functional microcode firmware support. Newer microcode levels might be required when adding new functions.

> **Note:** Installation of this enhanced firmware requires an update to the microcode firmware of an installed library manager in a 3494 or 3953 Tape Frame, which also requires ordering FC 0520, LM Functional enhancement field, on the library manager frame.

FC 0522, Move VTS database

FC 0522, Move VTS database, provides the utilities and services to move a VTS database to a TS7700 Virtualization Engine when the VTS database and TS7700 Virtualization Engine are attached to the same tape library. If FC 4010, Peer-to-peer VTS, is installed on the VTS, the database from each VTS must be moved to a separate TS7700 Virtualization Engine and FC 0522, Move VTS database, or FC 0523, Move VTS database & cartridges, must be ordered for each TS7700 Virtualization Engine.

FC 0523, Move VTS database & cartridges

FC 0523, Move VTS database & cartridges, provides the utilities and services to move a VTS database to a TS7700 Virtualization Engine when the VTS and TS7700 Virtualization Engine are attached to different tape libraries. You are responsible for moving the tape cartridges between the tape libraries. If FC 4010, Peer-to-peer VTS, is installed then each VTS database must be moved to a separate TS7700 Virtualization Engine, and FC 0523, Move VTS database & cartridges, or FC 0522, Move VTS database, must be ordered for each TS7700 Virtualization Engine.

FC1030, 1 Gb Grid copper connection

FC1030, 1 Gb Grid copper connection, provides an Ethernet 1000 BaseT adapter for communication between TS7700 Virtualization Engines in a Grid configuration. You must supply your own Ethernet cables when FC4015, Grid enablement, is installed. The Ethernet cables used with the adapter must be of Category 5e or 6, though Category 6 is preferred. Category 5 Ethernet cables are not allowed. See FC1030, 1 Gb Grid copper connection, in Table C-2 on page 550 for requirements and other information.

FC1031, 1 Gb Grid optical SW connection

FC 1031, 1 Gb Optical SW connection, provides a 1 Gb shortwave adapter for Grid communication between TS7700 Virtualization Engines. It replaces RPQ #Q8B3409. This adapter has an LC Duplex connector for attaching a 50.0- or 62.5 micron multi-mode fibre cable. This is a standard shortwave (850 nm) adapter that conforms to the IEEE 802.3z standards. It supports distances of 2 to 260 meters for 62.5 micron cable and 2 to 550 meters for 50.0 micron cable. Multi-mode fibre cables must be used with the adapter when FC 4015, Grid enablement, is installed.See FC1031, 1 Gb Grid optical SW connection, in Table C-2 on page 550 for requirements and other information.

FC2714, Console expansion

FC2714, Console expansion, provides an Ethernet hub and product attachment cable for expanding the number of components that can attach to the TSSC that is provided by FC2713, System console for service. See FC2714, Console expansion, in Table C-2 on page 550 for requirements and other information.

> **Note:** If you have FC2713, System console for service, you must order FC2720 against the TS7740, or FC2721 against the 3953 F05 Frame.

FC2715, Console attachment

FC2715, Console attachment, provides a cable for attaching a TS7740 Node to the Ethernet hub provided by FC2718, FC2720, FC2721, or by FC2714, Console expansion. A single TSSC facility can support a maximum of forty instances of this feature. See FC2715, Console attachment, in Table C-2 on page 550 for requirements and other information.

> **Note:** If you have FC2713, System console for service, you must order FC2720 against the TS7740, or FC2721 against the 3953 F05 Frame.

FC2719, Console upgrade

FC2719, Console upgrade, provides a second Ethernet network interface card for the TSSC to allow redundant attachment to the service network. This feature provides a memory upgrade to 2 GB total RAM and a second Ethernet card for the Service Console to allow redundant connections into the service network. This feature only applies to consoles shipped with features FC2718, FC2720 and FC2721. It is required on any console used by TS7700 features FC2714, FC2715 or FC2720. This feature is only available on TSSCs shipped before 26 might 2006; TSSCs shipped after this date will include the second Ethernet network interface card. See FC2719, Console upgrade in Table C-2 on page 550 for requirements and other information.

FC2720, TS3000 System Console

FC2720, TS3000 System Console, provides the IBM TotalStorage Master Console, an Ethernet hub, and a cable and connectors to enable remote enhanced service connection of a TS7740 to an IBM-supplied modem. The Ethernet hub provides 14 additional connections for cables supplied with FC2714, Console expansion, or FC2715, Console attachment. You should specify FC2720, TS3000 System Console on the first unit in an installation connected to a master console facility. See FC2720, TS3000 System Console, in Table C-2 on page 550 for requirements and other information.

FC3441, FICON shortwave attachment

FC3441, FICON shortwave attachment, provides a shortwave 4 Gbps FICON adapter with an LC duplex connector for attachment to a FICON host system shortwave channel using a 50 micron or 62.5 micron multimode fibre cable. The maximum fibre cable length allowed by 50 micron cable is 150 m (492 ft.). Each 4 Gbps FICON attachment can support up to 128 logical channels (for TS7700 R1.3 and below) or up to 256 logical channels (for TS7700 R1.4). See FC3441, FICON shortwave attachment, in Table C-2 on page 550 for requirements and other information.

FC3442, FICON longwave attachment

FC3442, FICON longwave attachment, provides a longwave 4 Gbps FICON adapter with an LC duplex connector for attachment to a FICON host system longwave channel using a 9 micron single-mode fibre cable. The maximum fibre cable length is 4 km (2.48 mi.). Each 4 Gbps FICON attachment can support up to 128 logical channels (for TS7700 R1.3 and below) or up to 256 logical channels (for TS7700 R1.4). See FC3442, FICON longwave attachment, in Table C-2 on page 550 for requirements and other information.

FC3443, FICON 10-km longwave attachment

FC3443, FICON 10-km longwave attachment, provides a longwave 4 Gbps FICON adapter with an LC duplex connector for attachment to a FICON host system longwave channel using a 9 micron single-mode fibre cable. The maximum fibre cable length is 10 km (6.21 mi.). Each 4 Gbps FICON attachment can support up to 128 logical channels (for TS7700 R1.3 and below) or up to 256 logical channels (for TS7700 R1.4). See FC3443, FICON 10-km longwave attachment, in Table C-2 on page 550 for requirements and other information.LC

FC4015, Grid enablement

FC 4015, Grid enablement, provides a key to enable the communication function that allows a TS7700 Virtualization Engine to communicate with other TS7700 Virtualization Engines in a grid. Each TS7700 Virtualization Engine must have this feature to be able to participate in a Grid configuration. See FC 4015, Grid enablement, in Table C-2 on page 550 for requirements and other information.

Note: More than one situation has occurred where a TS7700 was ordered with FC4015 (Grid enablement) and the intention was to initially use the TS7700 in a standalone test environment. This TS7700 should be installed completely (service code 20 CIA 1) before being used in this test configuration. This use of a grid-enabled TS7700 in a standalone environment is an unintended use and requires that the TS7700 be "Reset to Factory Settings," also sometimes referred to as *manufacturing cleanup*, before using the TS7700 in a grid configuration.

If all clusters in a Grid configuration already have FC 4015, Grid enablement installed, contact your IBM Service Representative to properly setup, connect, and configure the Grid environment.

FC 4016, Remove Cluster from Grid

FC 4016, Remove Cluster from Grid, delivers instructions for a one-time process to remove a TS7740 Cluster from a TS7700 Grid. You must order this feature before you can perform a cluster cleanup (FC 4017, Cluster cleanup) on any TS7740 Cluster configured to participate in a TS7700 Grid. If aTS7740 Cluster is removed from a TS7700 Grid, cleaned up using FC 4016, Remove Cluster from Grid, and then joined to a new TS7700 Grid, another instance of FC 4016, Remove Cluster from Grid, is required to remove the cluster from the new Grid.

FC 4017, Cluster cleanup

FC 4017, Cluster cleanup, facilitates a one-time cluster cleanup to clean the database, delete logical volumes from the Tape Volume Cache, and remove configuration data for host connections from a TS7740 Cluster. If the cluster is a member of a TS7700 Grid, the cluster must first be removed from the Grid using FC 4016, Remove Cluster from Grid. The target TS7740 Cluster must have FC 4015, Grid enablement, installed. If another cleanup is required on this TS7740 Cluster, another instance of FC 4017, Cluster cleanup, is required.

FC5240, Attach 3592 Tape Drive

FC5240, Attach 3592 Tape Drive, installs two Fibre Channel interface cards in the TS7740 Node and provides two Fibre Channel cables to connect the TS7740 Node to the Fibre Channel switch. See FC5240, Attach 3592 Tape Drive, in Table C-2 on page 550 for requirements and other information.

FC5267, 1 TB cache enablement

FC5267, 1 TB cache enablement, enables a 1 TB increment of disk cache to store logical volumes. Enabling a cache increment does not guarantee that amount of additional cache capacity. The amount of additional cache capacity provided will be limited by the capacity of the underlying physical cache installed. This is a key-controlled feature. See FC5267, 1 TB cache enablement, in Table C-2 on page 550 for requirements and other information.

FC5268, 100 MB/s increment

FC5268, 100 MB/s increment, delivers a key to enable an additional 100 MB/s increment of potential Peak Data Throughput. Enabling a data throughput increment does not guarantee that the overall VTS will perform at that data throughput level. See FC5268, 100 MB/s increment, in Table C-2 on page 550 for requirements and other information.

FC9000, Mainframe attachment

FC 9000, Mainframe attachment, is a specify feature that indicates attachment to one of the following host platforms:

- z/OS 1.4 and higher
- z/VM, 4.4.0 and higher
- z/VSE, 3.1.2 and higher
- zTPF, V1.1 and higher
- TPF, 4.1 and higher

See FC9000, Mainframe attachment, in Table C-2 on page 550 for requirements and other information.

FC9217, Attach to 3953 LM

FC9217, Attach to 3953 LM, is a specify feature that indicates that the TS7700 Virtualization Engine will attach to a 3953 L05 Library Manager. See FC9217, Attach to 3953 LM, in Table C-2 for requirements and other information.

FC9218, Attach to 3494 LM

FC9218, Attach to 3494 LM, is a specify feature that indicates the TS7700 Virtualization Engine will attach to a 3494 Library Manager. See FC9218, Attach to 3494 LM, in Table C-2 for requirements and other information.

FC9350, Plant install V06 in F05

FC9350, Plant install V06 in F05, allows the plant installation of a TS7700 Virtualization Engine into a new 3952 Tape Frame. This feature must appear on the TS7700 Virtualization Engine order. You must also order FC5628, Plant install 3957 V06, against the 3952 Tape Frame order. See FC9350, Plant install V06 in F05, in Table C-2 for requirements and other information.

FC9700, No factory cables

FC9700, No factory cables, instructs the plant not to ship any FICON cables with the TS7740 Node. See FC9700, No factory cables, in Table C-2 for requirements and other information.

FC9900, Encryption configuration

FC9900, Encryption configuration, should be ordered when encryption will be used in a 3494 Tape Library, TS3500 Tape Library, or TS7700 Cluster. It includes publication updates with information about enabling and configuring the library (virtual or real) to support encryption. This feature also provides an Encryption Key Manager (EKM) publication. Customer-initiated procedures need to be completed for enabling and configuring the 3494 Tape Library, TS3500 Tape Library, or TS7700 Cluster to support encryption with the TS1120 encryption-capable tape drive.

> **Note:** FC 9900, Encryption configuration is only supported with encryption-capable TS1120 3592 E05 Tape Drives. FC 9900, Encryption configuration, is not supported by 3592 J1A or 3590 Tape Drives.

Table C-2 TS7740 Node features

Feature	Minimum	Maximum	Requirements	Installation	Removable?
FC0201, 9 micron LC/LC 31 meter	0	4	None	Plant or field	Yes
FC0202, 9 micron LC/SC 31 meter	0	4	None	Plant or field	Yes
FC0203, 50 micron LC/LC 31 meter	0	4	None	Plant or field	Yes
FC0204, 50 micron LC/SC 31 meter	0	4	None	Plant or field	Yes
FC0205, 62.5 micron LC/LC 31 meter	0	4	None	Plant or field	Yes
FC0206, 62.5 micron LC/SC 31 meter	0	4	None	Plant or field	Yes
FC0521, Functional enhancement field	0	No Maximum	This feature is a one-time code update. Installation of this feature might also require an update to the microcode firmware of an installed library manager in a 3494 or 3953 Tape Frame	Field only	No
FC1030, 1 Gb Grid copper connection	2 of FC1030 or FC1031	2	This feature is mutually exclusive with FC 1031. You must order two of either FC 1030 *or* FC 1031. This feature is available as a concurrent MES.	Plant or field	Yes
FC1031, 1 Gb Grid optical SW connection	2 of FC1030 or FC1031	2	This feature is mutually exclusive with FC 1030. You must order two of either FC 1031 **OR** FC 1030. This feature is available as a concurrent MES.	Plant or field	Yes
FC2714, Console expansion	0	1	You must specify one instance of FC2720, FC2714, or FC2715. FC2714, Console expansion, must attach to FC2718, FC2720, or FC2721. Attachment to FC2713 is not allowed.	Plant or field	Yes

Feature	Minimum	Maximum	Requirements	Installation	Removable?
FC2715, Console attachment	0	1	You must specify one instance of FC2720, FC2714, or FC2715. Console attachment, must attach to FC2718, FC2720, or FC2721. Attachment to FC2713 is not allowed.	Plant or field	Yes
FC 2719, Console upgrade	0	1	Required on TSSC (FC 2720 or 2721) used by the TS7700 Virtualization Engine, whether TS7700 Virtualization Engine is attached directly to the TSSC using FC 2720, or indirectly using FC 2714 or FC 2715.	Field	No
FC2720, TS3000 System Console	0	1	You must order one of FC2713, FC2714, or FC2715.	Plant only	No
FC3441, FICON shortwave attachment	0[a]	4	You must order this feature in quantities of zero, two, or four. You can order two instances of this feature and two instances of FC 3442 or FC 3443. When taken together, the total minimum quantity of FC 3441, FC 3442, and FC 3443 is two; the total maximum quantity is four. This feature is available as a concurrent MES.	Plant or field	Yes

Feature	Minimum	Maximum	Requirements	Installation	Removable?
FC3442, FICON longwave attachment	0[a]	4	You must order this feature in quantities of zero, two, or four. You can order two instances of this feature and two instances of FC 3441 or FC 3443. When taken together, the total minimum quantity of FC 3441, FC 3442, and FC 3443 is two; the total maximum quantity is four. This feature is available as a concurrent MES.	Plant or field	Yes
FC3443, FICON 10-km longwave attachment	0[a]	4	You must order this feature in quantities of zero, two, or four. You can order two instances of this feature and two instances of FC 3441 or FC 3442. When taken together, the total minimum quantity of FC 3441, FC 3442, and FC 3443 is two; the total maximum quantity is four. This feature is available as a concurrent MES.	Plant or field	Yes
FC4015, Grid enablement	0	1	You must order two instances of either FC 1030 or FC 1031	Plant or field	No
FC 4016, Remove Cluster from Grid	0	99	Prerequisite: FC 4015, Grid enablement	Field only	No
FC 4017, Cluster cleanup	0	99	None	Field only	No
FC5240, Attach 3592 Tape Drive	1	1	You must order one instance of this feature.	Plant	No

Feature	Minimum	Maximum	Requirements	Installation	Removable?
FC5267, 1 TB cache enablement	1	6	Total hardware capacity might not be enabled unless appropriate quantities of FC5267 are ordered. This feature is available as a concurrent MES.	Plant or field	No
FC5268, 100 MB/s increment	1	6	None This feature is available as a concurrent MES.	Plant or field	No
FC9000, Mainframe attachment	1	1	Corequisite: FC3441, FC3442, or FC3443	Plant only	No
FC9217, Attach to 3953 LM	0	1	This feature is mutually exclusive with FC 9218. You must order one of either FC 9217 **OR** FC 9218.	Plant only	No
FC9218, Attach to 3494 LM	0	1	This feature is mutually exclusive with FC 9217. You must order one of either FC 9218 OR FC 9217. Prerequisite: FC 9013 is required on the 3494 D22 Frame containing the 3592 Tape Drives for this TS7700 Virtualization Engine.	Plant only	No
FC9350, Plant install V06 in F05	1	1	Corequisite: FC5628 against 3952 F05	Plant only	No
FC9700, No factory cables	0	1	If you do not order FC0201, FC0202, FC0203, FC0204, FC0205, or FC0206, you must order one FC9700	Plant only	Yes
FC9900, Encryption configuration	0	1	This feature is available as a concurrent MES if the TS1120 drives were running in native mode prior to the TS7700's current online process.	Plant or field	No

a. Two of either FC3441, FC3442 or FC3443 must be ordered.

TS7740 cache controller features

This section describes each feature code for the TS7740 cache controller, machine type and model 3956 CC6, in detail.

FC6003, Intraframe fibre cable to 3957 V06

FC6003, Intraframe fibre cable to 3957 V06, provides one 2-m fibre cable. See FC6003, Intraframe fibre cable to 3957 V06, in Table C-3 for requirements and other information.

FC7120, 1.7 TB fiber storage

FC7120, 1.7 TB fiber storage, installs a complete set of sixteen Fibre Channel-capable disk drive modules in the TS7740 cache controller, providing a total usable storage capacity of 1.7 terabytes. See FC7120, 1.7 TB fiber storage, in Table C-3 for requirements and other information.

FC9230, Attach to 3957 V06

FC9230, Attach to 3957 V06, is a specify feature that indicates a TS7740 cache controller is attached to a IBM Virtualization Engine TS7700. You must order FC9352, Plant install CC6 in 3952 F05, against the IBM Virtualization Engine TS7700 and FC5638, Plant install 3956 CC6, against the 3952 Tape Frame. See FC9230, Attach to 3957 V06 in Table C-3 for requirements and other information.

FC9352, Plant install CC6 in 3952 F05

FC9352, Plant install CC6 in 3952 F05, allows the plant installation of a TS7740 Cache Controller (3956 CC6) into a new 3952 Tape Frame. You must also order FC9230, Attach to 3957 V06 against the 3956 CC6 and FC5638, Plant install 3956 CC6 against the 3952 Tape Frame. This feature must appear on the Model CC6 order. You must also order FC5638, Plant install 3956 CC6FC5738, Plant install 3956 CC6, against the 3952 Tape Frame. See FC9352, Plant install CC6 in 3952 F05, in Table C-3 for requirements and other information.

Table C-3 TS7740 cache controller features

Feature	Minimum	Maximum	Requirements	Installation	Removable?
FC6003, Intraframe fibre cable to 3957 V06	4	4	2-m cable	Plant only	No
FC7120, 1.7 TB fibre storage, for the TS7740 cache controller	1	1	1.7 TB is usable capacity	Plant only	No
FC9230, Attach to 3957 V06	1	1	Corequisite: Four instances of FC6003	Plant only	No
FC9352, Plant install CC6 in 3952 F05	1	1	Corequisite: FC5638 against the 3952 F05	Plant only	No

TS7740 cache drawer features

This section describes each feature code for the TS7740 cache drawer, machine type and model 3956 CX6, in detail.

FC6000, Intraframe fibre cable to 3956 CC6

FC6000, Intraframe fibre cable to 3956 CC6, provides one 1-m fibre cable and is also used to connect cache expansion drawers. See FC6000, Intraframe fibre cable to 3956 CC6, in Table C-4 for requirements and other information.

FC7120, 1.7 TB fibre storage

FC7120, 1.7 TB fibre storage, installs a complete set of 16 Fibre Channel-capable disk drive modules in the TS7740 cache drawer, providing a total usable storage capacity of 1.7 terabytes. See FC7120, 1.7 TB fibre storage, in Table C-4 for requirements and other information.

FC9354, Plant install 3956 CX6 in 3952 F05

FC9354, Plant install 3956 CX6 in 3952 F05 is a specified feature that allows the plant installation of a TS7740 cache expansion (3956 CX6) into a new 3952 Tape Frame. FC9354, Plant install 3956 CX6 in 3952 F05, must appear on the 3956 CX6 order. You must also order FC5648, Plant install 3956 CX6, against the 3952 Tape Frame. See FC9354, Plant install 3956 CX6 in 3952 F05, in Table C-4 for requirements and other information.

FC9355, Field merge 3956 CX6 in 3952 F05

FC3955, Field merge 3956 CX6 in 3952 F05, allows the field merge of a TS7740 cache expansion (3956 CX6) into an installed 3952 Tape Frame. This feature must be ordered on the Model CX6 and the Field Install a 3956 Model CX6 feature FC5649 must be ordered on the 3952 Tape Frame.

Table C-4 TS7740 cache drawer features

Feature	Minimum	Maximum	Requirements	Installation	Removable?
FC6000, Intraframe fibre cable to 3956 CC6	4	4	1-m Cable. One cable will be shipped per feature code.	Plant only	No
FC7120, 1.7 TB fibre storage, for the TS7740 cache drawer	1	1	1.7 TB is usable capacity.	Plant only	No
FC9354, Plant install 3956 CX6 in 3952 F05	0	1	Corequisite: FC5648 against the 3952 F05. You must order one of FC 9354 or FC 9355.	Plant only	No
FC9355, Field merge 3956 CX6 in 3952 F05	0	1	Prerequisite: FC5649 against the 3952 F05. You must order one of FC 9354 or FC 9355.	Field only	No

Feature conversions

This section lists supported TS7700 Virtualization Engine feature conversions

Table C-5 displays supported TS7700 Virtualization Engine feature code conversions.

Table C-5 Supported TS7700 Virtualization Engine feature code conversions

Feature code	might be converted to
FC 3441, FICON short-wavelength attachment	FC 3442, FICON long-wavelength attachment
	FC 3443, FICON 10-km long-wavelength attachment
FC 3442, FICON long-wavelength attachment	FC 3441, FICON short-wavelength attachment
	FC 3443, FICON 10-km long-wavelength attachment
FC 3443, FICON 10-km long-wavelength attachment	FC 3441, FICON short-wavelength attachment

Feature code	might be converted to
FC1030, 1 Gb Grid copper connection	FC103, 1 Gb Grid optical SW connection1
FC1031, 1 Gb Grid optical SW connection	FC1030, 1 Gb Grid copper connection

Table C-6 on page 556 displays supported 3494 D22 Frame conversions specifically for use with the TS7700 Virtualization Engine

Table C-6 Supported feature code conversions for the 3494 D22 Frame

Feature code	might be converted to
FC 9010, Attach to a VTS	FC 9013, Attach to TS7700

The following features are available until December 5, 2008 on the 3494 Model B10 or B20 VTS

Table C-7 VTS conversion features

Feature	Minimum	Maximum	Requirements	Installation	Removable?
FC 0522, Move VTS database	0	1	VTS and TS7700 Virtualization Engine must be attached to the same tape library.	Field only	No
FC 0523, Move VTS database & cartridges	0	1	VTS and TS7700 Virtualization Engine must be attached to different tape libraries.	Field only	No

D

TS3500 checklists

You use the checklists that we include in this appendix when attaching a TS7740 to a TS3500 tape library. For the most recent version, refer to the *TS7740 TS3500 Worksheets* that are part of the Solution Assurance documents for the TS7740. The document is accessible by your IBM representative.

TS7740 Installation Worksheets

The following customer-supplied information (as applicable) is required when you install the TS7740. You and the local IBM account team need to complete these worksheets and give them to the installation team.

> **Note:**
>
> These worksheets are the same as the worksheets that are contained in the service document *IBM Virtualization Engine TS7740 Installation Roadmap for use with IBM Systems Storage TS3500, IBM 3584 Tape Library*, PN 23R7608.
>
> There can be additional Tape Control Units (3592-J70/C06), or Virtual Tape Subsystems (3494-B10/B20/B18) attached to the same 3584 Tape Library or 3953. The installation or configuration of these devices is not within the scope of this document. In most cases, you need to install these devices using documentation that is provided with the devices.
>
> The terms *Cluster 0*, *Cluster 1*, and *Cluster 2* refer to the separate TS7740 Virtualization Engines participating in a GRID (peer-to-peer) configuration.
>
> If you are installing a standalone (non-GRID) TS7740 Virtualization Engine, complete the Cluster 0 information but leave the Cluster 1 and 2 fields blank.
>
> If you are installing a Two-Cluster TS7740 Virtualization Engine Grid, complete the Cluster 0 and 1 information but leave the Cluster 2 fields blank.
>
> Provide the SSR with a serial number and a description of which TS7740 Virtualization Engine will be Cluster 0, which TS7740 will be Cluster 1, and which TS7740 will be Cluster 2. The information in Table D-2 on page 560 tells the SSR which TS7740 is designated as Cluster 0, which is Cluster 1, and which is Cluster 2.
>
> Complete the *Description of Location* column with geographical information that is unique to each cluster. Complete as much information as possible. Suggestions for descriptions are the x, y coordinates within a data center, room number, floor number, building numbers, city, and so forth.

Grid cluster descriptions

Use Table D-1 to define information about the grid clusters.

Table D-1 Grid cluster descriptions

Cluster	Machine Type-Model	Serial Number	Description of Location
Cluster 0	3957-V06		
	3953-F05		
	3584-Lxx		
Cluster 1	3957-V06		
	3953-F05		
	3584-Lxx		
Cluster 2	3957-V06		
	3953-F05		
	3584-Lxx		

3584 Tape Library configuration information

Use Table D-2 to define information about the Tape Library configuration.

Table D-2 3584 Tape Library configuration information

Field	Value	Notes
3584 Tape Library Ethernet Network Configuration Method (customer network).	Cluster 0: DHCP [ ] Fixed IP [ ]	The network configuration method is specified by the customer's LAN administrator. It is either Fixed IP or Dynamic Host Configuration Protocol (DHCP). **Note:** Fixed IP is recommended.
	Cluster 1: DHCP [ ] Fixed IP [ ]	
	Cluster 2: DHCP [ ] Fixed IP [ ]	
3584 Tape Library Ethernet IP Address (customer network) (Used for 3584 Tape Library web specialist access).	Cluster 0: ____.____.____.____	If Network Configuration Method is DHCP then this field is not used. **Note:** The Ethernet ports are 10/100 Mb only.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
3584 Tape Library Ethernet Hostname	Cluster 0:	The entry in this field is used to identify the machine in remote support logs. **Note:** Even if host names are not typically used by the customer, this host name is still required. A typical host name is ATL1.
	Cluster 1:	
	Cluster 2:	
3584 Tape Library Ethernet Subnet Mask (customer network)	Cluster 0: ___.___.___.___	If Network Configuration Method is DHCP then this field is not used.
	Cluster 1: ___.___.___.___	
	Cluster 2: ___.___.___.___	
3584 Tape Library Ethernet Gateway (customer network)	Cluster 0: ___.___.___.___	If Network Configuration Method is DHCP then this field is not used.
	Cluster 1: ___.___.___.___	
	Cluster 2: ___.___.___.___	

Field	Value	Notes
3584 Tape Library Logical Library Name (for 3953 attachment)	Cluster 0:	Each 3953 (or dual 3953 for redundancy) must be connected to a single 3584 Tape Library logical library. This logical library must have a name, which should have been assigned when the logical library was created. Record the logical library name (assign it if necessary). The logical library name will be needed when performing the following tasks: ▶ Configuring the logical library. ▶ Obtaining the Starting Element Address for the logical library. ▶ Obtaining the physical position of tape drives within the logical library. ▶ Obtaining the WWNNs of those tape drives. ▶ Setting the Cartridge Assignment Policy. ▶ Configuring ALMS
	Cluster 1:	
	Cluster 2:	
Max Cartridges	Cluster 0:	This value defaults to the number of physical cartridge slots currently installed in the TS3500 Tape Library. You might want to set a different value in the following cases the customer wants to restrict the number of cartridges in the logical library to manage the cost of application licensing. The TS3500 Tape Library will be expanded at a later date and you want to avoid reconfiguring the host.
	Cluster 1	
	Cluster 2:	
Has FC 0500 been installed on the 3584 Tape Library?	Cluster 0: Yes [ ] No [ ]	▶ FC 0500 is the **Library and Drive Code Update in a 3584 (TS3500) Tape Library**. It provides a customer the ability to have library and drive code updated by a CE, and a way for the CE to write off their time updating library or drive code for the customer. ▶ The Order Detail indicates if FC 0500 has been installed. ▶ If FC 0500 has been installed, the CE will update microcode, if not, the customer must perform the microcode update.
	Cluster 1: Yes [ ] No [ ]	
	Cluster 2: Yes [ ] No [ ]	

3584 Tape Library drive information

Use Table D-3 to enter the Tape Library drive information.

Table D-3 3584 Tape Library drive information

Field	Value	Notes
► Tape Drive Physical Positions (Fn, Rnn) in the 3584 Tape Library for the drives assigned to this TS7740. ► Record the physical position and the last two digits of the World Wide Node Name (WWNN) for each tape drive. ► Circle the drives that will be control paths. ► Read the notes in the right column of this table for guidance in completing the values.	Cluster 0: ► F = Frame ► R = Row ► WWNN = World Wide Node Name 1. F____, R____, WWNN_____ 2. F____, R____, WWNN_____ 3. F____, R____, WWNN_____ 4. F____, R____, WWNN_____ 5. F____, R____, WWNN_____ 6. F____, R____, WWNN_____ 7. F____, R____, WWNN_____ 8. F____, R____, WWNN_____ 9. F____, R____, WWNN_____ 10. F____, R____, WWNN_____ 11. F____, R____, WWNN_____ 12. F____, R____, WWNN_____ 13. F____, R____, WWNN_____ 14. F____, R____, WWNN_____ 15. F____, R____, WWNN_____ 16. F____, R____, WWNN_____	► A 3953 can be connected to one or two TS7740 subsystems. The first is known as VTS-1, the second (if installed) is known as VTS-2. VTS-1 and VTS-2 will share the same logical library, but will have different tape drives. A minimum of 4 and a maximum of 16 drives can be connected to each TS7740 Virtualization Engine ► It is recommended that the drives for a single cluster are listed in order using the frame and row. The lowest numbered drive should be in the lowest numbered frame and row for its assign drive slots. This is not a requirement, but it might help avoid confusion when identifying drives during future troubleshooting. ► Distributing the tape drives across two 3584 frames is recommended to improve availability. Placing eight drives (two control path drives) in one frame and either drives (two control path drives) in a second frame is recommended. ► To determin the WWNN, press **Menu → Vital Product Data → World Wide Node Names.**

Field	Value	Notes
	Cluster 1: ► F = Frame ► R = Row ► WWNN = World Wide Node Name 1. F____, R____, WWNN_____ 2. F____, R____, WWNN_____ 3. F____, R____, WWNN_____ 4. F____, R____, WWNN_____ 5. F____, R____, WWNN_____ 6. F____, R____, WWNN_____ 7. F____, R____, WWNN_____ 8. F____, R____, WWNN_____ 9. F____, R____, WWNN_____ 10. F____, R____, WWNN_____ 11. F____, R____, WWNN_____ 12. F____, R____, WWNN_____ 13. F____, R____, WWNN_____ 14. F____, R____, WWNN_____ 15. F____, R____, WWNN_____ 16. F____, R____, WWNN_____	
	Cluster 2: ► F = Frame ► R = Row ► WWNN = World Wide Node Name 1. F____, R____, WWNN_____ 2. F____, R____, WWNN_____ 3. F____, R____, WWNN_____ 4. F____, R____, WWNN_____ 5. F____, R____, WWNN_____ 6. F____, R____, WWNN_____ 7. F____, R____, WWNN_____ 8. F____, R____, WWNN_____ 9. F____, R____, WWNN_____ 10. F____, R____, WWNN_____ 11. F____, R____, WWNN_____ 12. F____, R____, WWNN_____ 13. F____, R____, WWNN_____ 14. F____, R____, WWNN_____ 15. F____, R____, WWNN_____ 16. F____, R____, WWNN_____	

Media volume serial range

Complete Table D-4 on page 565 with the requested information. There might be one or many **Media Volume Serial Ranges**, so complete as many rows as apply to your system.

Refer to the following list for a description of the values needed for Table D-4 on page 565:

- **From, To:** A range that contains the bar code label volsers of all the cartridges that are assigned to a single TS7740 Virtualization Engine. As an example, if cartridges assigned to a TS7740 have barcode labels in the range from A00000JA - A00500JA, then you record the following information:
  - From: A00000
  - To: A00500
- **Media Type:** Indicated by the last two characters of the 8 character bar code label volser on the cartridges. As an example, if the cartridges are labeled 123456JA, then the media type is JA. See the rules that are defined in the following notes:
  - JA and JJ tape cartridges are supported and can be mixed in a TS7740 Virtualization Engine.
  - JB tape cartridges are also supported (and can be mixed with JA and JJ tape cartridges) if all of the tape drives associated with the TS7740 are 3592 model E05 drives and none of them are in J1A emulation mode.
  - No other tape types are currently supported for use with the TS7740.
  - If there is at least one 3592 model J1A tape drive associated with the TS7740, then *all* 3592 model E05 tape drives that are associated with the TS7740 must be in J1A emulation mode.
  - If at least one tape cartridge that is associated with a TS7740 has been written by a 3592 model E05 tape drive that is *not* in J1A emulation mode, then the TS7740 will no longer support any 3592 model J1A tape drives or any 3592 model E05 tape drive that is in J1A emulation mode. After you go to E05 Native mode, you cannot go back to J1A mode because the J1A tape drive cannot read or write a tape cartridge written in E05 mode and the TS7740 does not currently support mixed J1A and E05 tape drives (unless all E05 tape drives are in J1A emulation mode).
  - The capacity of a JJ tape cartridge is 60 GB if written by a J1A drive (or an E05 drive that is in J1A emulation mode) or 100 GB if written by an E05 drive that is *not* in J1A emulation mode. The capacity of a JA tape cartridge is 300 GB if written by a J1A drive (or an E05 drive that is in J1A emulation mode) or 500 GB if written by an E05 drive that is *not* in J1A emulation mode. The capacity of a JB cartridge is 700 GB.
  - TS7740 Feature Code 9900 (Encryption) requires that all tape drives are in E05 Native mode. The Encryption feature is *not* compatible with J1A emulation mode.

  Using the rules defined in “3584 Tape Library drive information” on page 562, determine whether you want the tape drives attached to the TS7740 to be in J1A emulation mode or in E05 native mode. Record this in the Tape Drive Format entry in Table D-6 on page 568.

- **(Distributed) Library Sequence Number**: Typically assigned to the system administrator. This is a 5 character name that is used as an identifier for a specific cluster and the associated LM and Library in a Grid configuration. This identifier is specified in the 3953 L05 Library Manager configuration. *It is required even if the TS7740 is not in a Grid configuration*.

> **Note:** *Each VTS partition must have a single, unique value for the Distributed Library Sequence Number*. For the TS7740, a typical value is the last 5 digits of the 3952-F05 frame serial number.

- **Home Pool** (also called *Scratch Pool*): You might have assigned a Home Pool value. If one has not been set, the default value is 00.

> **Note:** A *pool* is a group of physical tape cartridges. A *scratch pool* is a group of cartridges that are considered to be scratch, meaning that they are ready for use by any write job.
>
> If there is at least one J70 or C06 Tape Control Unit connected to the 3953 L05, you can choose to record the **Native Lib Seq Num** in the column labeled **(Distributed) Library Sequence Number**. This information will not be used for the TS7740.

Table D-4 Media volume serial range

Cluster (0, 1, 2)	From	To	Media Type	Distributed Library Sequence Number	Home Pool	VTS Number (1 or 2)	Fast Ready Category

3953 configuration information

Use Table D-5 to record 3953 configuration information.

Table D-5 3953 configuration information

Field	Value	Notes
Logical Device IDs	Cluster 0:	► This is a 3-digit ID chosen by the customer's administrator. The value is usually the last 3 digits of the host HCD device address. It is used to identify a range of Logical Device IDs (128 or 256). ► The Logical Device ID is the identifier for this device in error log output. Specifying an ID that is already assigned will not cause errors, however it will be harder to distinguish this device during trouble shooting
	Cluster 1:	
	Cluster 2:	
Customer IP Address	Cluster 0: LM-A ____.____.____.____ LM-B ____.____.____.____	► Used for 3953 Library Manager Web Specialist. ► The Ethernet ports are 10/100 Mb only.
	Cluster 1: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 2: LM-A ____.____.____.____ LM-B ____.____.____.____	
Customer Subnet mask	Cluster 0: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 1: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 2: LM-A ____.____.____.____ LM-B ____.____.____.____	
Customer Gateway	Cluster 0: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 1: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 2: LM-A ____.____.____.____ LM-B ____.____.____.____	

Field	Value	Notes
Hostname	Cluster 0: LM-A ____.____.____.____ LM-B ____.____.____.____	Used for 3953 Web Specialist If you have only one 3953-F05 then typical values would be LMA and LMB.
	Cluster 1: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 2: LM-A ____.____.____.____ LM-B ____.____.____.____	

TS7740 Virtualization Engine configuration information

Use Table D-6 to record TS7740 Virtualization Engine configuration information.

Table D-6 TS7740 Virtualization Engine configuration information

Field	Value	Notes
Composite Library Sequence Number		► This 5 character name must be the same on all clusters (peers) within the same grid. This identifier is specified in the TS7740 configuration. It is required even if the machine is not in a grid configuration. ► The Composite Library Sequence Number must be different than the Distributed Library Sequence number specified in the 3953-L05 configuration.
VTS Partition	Cluster 0: Cluster 1: Cluster 2:	► 1 or 2. Each 3953 Library Manager (or pair of LMs if they are in a dual LM configuration for redundancy) can be attached to one or two TS7740 subsystems. The first TS7740 connected to a 3953 is called VTS-1 and the second (if present) is called VTS-2. ► It is possible to have more than one 3953 connected to the same 3584 Tape Library, but each 3953 must be connected to a different logical library.
Customer IP 1 (Virtual)	Cluster 0: ____.____.____.____ Cluster 1: ____.____.____.____ Cluster 2: ____.____.____.____	► Used for TS7740 web Management Interface ► This number is a virtual IP that is not associated with a physical cable. It will communicate through the Primary IP, and will automatically fail over to the Alternate IP when required.
Customer IP 2 (Primary)	Cluster 0: ____.____.____.____ Cluster 1: ____.____.____.____ Cluster 2: ____.____.____.____	► This is the IP address used to connect to the TS7740 through the internal primary network. This IP address should not be used by the customer unless the Virtual IP is inaccessible. ► The Ethernet ports are 10/100 Mb only.
Customer IP 3 (Alternate)	Cluster 0: ____.____.____.____ Cluster 1: ____.____.____.____ Cluster 2: ____.____.____.____	► This is the IP address used to connect to the TS7740 through the internal alternate network. This IP address should not be used by the customer unless the Virtual IP is inaccessible. ► The Ethernet ports are 10/100 Mb only.

Field	Value	Notes
Customer Gateway	Cluster 0: ____.____.____.____	This is used with the virtual, primary, and alternate customer IP addresses.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Customer Subnet Mask	Cluster 0: ____.____.____.____	This is used with the virtual, primary, and alternate customer IP addresses.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
NTP server IP address (if used) *Using the NTP server is strongly recommended to ensure that all components have consistent time settings.*	____.____.____.____	The TCP/IP address you obtain from the customer is either the NTP server at their site (if they maintain one locally), or an internet server. Use of an internet server assumes that the customer allows access to the internet on the NTP services port. (TCP/IP port 123).
Tape Drive Format	Cluster 0: J1A Emulation Mode [ ] E05 Native Mode [ ]	To determine whether to use J1A emulation mode or E05 native mode, refer to the rules listed above under **Media Type.**
	Cluster 1: J1A Emulation Mode [ ] E05 Native Mode [ ]	
	Cluster 2: J1A Emulation Mode [ ] E05 Native Mode [ ]	

TSSC remote support information

Use Table D-7 to record TSSC remote support information.

Table D-7 TSSC remote support information

Field	Value	Notes
Dial Out Prefix (if required)	Cluster 0:	If you need to dial a prefix to get an outside line, specify it here. As an example, if you need to dial 9 to get an outside line, record a 9 in this field.
	Cluster 1:	
	Cluster 2:	
Call-in Modem Number	Cluster 0:	This is the telephone number assigned to the analog phone line where the call-home modem will be connected.
	Cluster 1:	
	Cluster 2:	
Customer Company Name	Cluster 0:	This is the value that is sent out with the call home records.
	Cluster 1:	
	Cluster 2:	
Customer Contact Phone Number	Cluster 0:	This is the phone number that IBM contacts when call-home indicates a failure.
	Cluster 1:	
	Cluster 2:	
Off Shift Customer Contact Phone Number	Cluster 0:	This is the alternate phone number that IBM contacts when call-home indicates a failure.
	Cluster 1:	
	Cluster 2:	

Grid local address

The following notes apply to Table D-8:

- If the TS7740 that you are installing is a standalone machine (not part of a grid configuration) then leave Table D-8 blank.
- The grid interfaces are the 1 Gb Internet connections between clusters, allowing them to automatically remain in sync.
- IBM strongly recommends that the primary and alternate grid interfaces be on separate sub-nets. As an example, 10.10.1.n for the primary interface and 10.11.1.n for the alternate interface. Notice that the second set of octets must be different. If the grid interfaces are direct connected (without using Ethernet switches) then using separate sub-nets is required.

The Grid interfaces require connections using the following TCP/IP ports:

- 7 (Ping)
- 9 (Discard Service for bandwidth measuring tools)
- 1415 (WebSphere® message queues Grid - Grid)
- 1416 (WebSphere message queue HDM - HDM)
- 123 (Network Time Protocol)
- 350 (Distributed Library - Distributed Library File Transfer)

The following TCP/IP ports are also useful in service scenarios if allowed:

- 23 (Telnet)
- 20 and 21 (FTP)

Table D-8 Grid local address

Field	Value	Notes
Primary Grid Interface IP address	Cluster 0: ____.____.____.____	The Primary Grid Interface is the 1 Gb Ethernet Adapter located in slot C4 of the 3957-V06. The Primary Grid Interface at each cluster will connect to the Primary Grid Interface at each of the other clusters in the same Grid.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Primary Grid Interface Network Mask	Cluster 0: ____.____.____.____	
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Primary Grid Interface Gateway	Cluster 0: ____.____.____.____	If a Gateway is not used, leave this field blank. If using crossover cables you must not specify a Gateway.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Alternate Grid Interface IP address	Cluster 0: ____.____.____.____	The Alternate Grid Interface is the 1 Gb Ethernet Adapter located in slot C5 of the 3957-V06. The Alternate Grid Interface at each cluster will connect to the Alternate Grid Interface at each of the other clusters in the same Grid.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	

Field	Value	Notes
Alternate Grid Interface Network Mask	Cluster 0: ____.____.____.____	
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Alternate Grid Interface Gateway	Cluster 0: ____.____.____.____	If a Gateway is not used, leave this field blank. If using crossover cables you must not specify a Gateway.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	

TSSC Grid configuration information

The following notes apply to Table D-9:

1. If the TS7740 you are installing is a standalone machine (not part of a grid configuration) then leave Table D-9 blank.
2. If you will be using the Grid in a Cascade Deferred style then you do not need the Autonomic Ownership Takeover Manager (AOTM) and should leave Table D-9 blank.
3. Refer to *The Autonomic Ownership Takeover Manager (AOTM)* of the *IBM Virtualization Engine TS7740 Installation Roadmap* for more information about AOTM before you continue. Do *not* attempt to configure AOTM, but use the information to make an informed decision on whether to use AOTM.
4. If you do not want to use the Autonomic Ownership Takeover Manager (AOTM), then leave Table D-9 blank.
5. The TSSC grid interface is used only for the Autonomic Ownership Takeover Manager (AOTM).
6. Each cluster can be configured to use AOTM to provide ownership takeover for one cluster.

The AOTM requires the following TCP/IP ports:

- 7 (Ping)
- 80 (HTTP)

Table D-9 TSSC Grid configuration information

Field	Value	Notes
TSSC Grid Interface IP Address	Cluster 0: ____.____.____.____	The TSSC Grid Interface is used to allow the TSSC at one cluster to communicate with the TSSC at one other cluster. This is required if the Autonomic Ownership Takeover Manager (AOTM) will be used.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
TSSC Grid Interface Subnet Mask	Cluster 0: ____.____.____.____	
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
TSSC Grid Interface Gateway	Cluster 0: ____.____.____.____	
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Takeover Mode	Cluster 0: ROT [ ] WOT [ ]	The Takeover Mode must be either Read Only Takeover (ROT) or Write Only Takeover (WOT).
	Cluster 1: ROT [ ] WOT [ ]	
	Cluster 2: ROT [ ] WOT [ ]	

Field	Value	Notes
Grace Period (Minutes)	Cluster 0:	When a cluster detects that another cluster in the same Grid has failed, it will wait for the number of minutes specified as the Grace Period before it attempts to take over the volumes of the failed cluster. A typical Grace Period is 25 minutes.
	Cluster 1:	
	Cluster 2:	
Retry Period (Minutes)	Cluster 0:	The retry period is the number of minutes between attempts to take over ownership of the volumes associated with a failed cluster. A typical Retry Period is 5 minutes. In most cases this will be the same for both clusters.
	Cluster 1:	
	Cluster 2:	

E

3494 Checklists

You use these checklists when attaching a TS7700 to a 3494 tape library. For the most recent version, refer to the *TS7700 3494 Worksheets* that are part of the Solution Assurance documents for the TS7700. The document is accessible by your IBM representative.

TS7700 Installation Worksheets

The following customer-supplied information (as applicable) is required when you install the TS7700. You and the local IBM account team need to complete these worksheets and give them to the installation team.

Note:

These worksheets are the same as those contained in the service document *Roadmap: Install Virtualization Engine TS7700 for use with IBM 3494 Tape Library*, PN 95P6747.

There might be additional Tape Control Units (3592-J70/C06) or Virtual Tape Subsystems (3494-B10/B20/B18) attached to the same 3494 Tape Library. The installation or configuration of these devices is not within the scope of this document. In most cases, you need to install these devices using the documentation that is provided with the devices.

The terms *Cluster 0*, *Cluster 1*, and *Cluster 2* refer to two separate TS7700 Virtualization Engines participating in a GRID (peer-to-peer) configuration.

If you are installing a standalone (non-GRID) TS7700 Virtualization Engine, then complete the Cluster 0 information but leave the Cluster 1 and 2 fields blank.

If you are installing a two cluster TS7700 Virtualization Engine grid, complete the Cluster 0 and 1 information but leave the cluster 2 fields blank.

Provide the serial number and a description of which TS7700 Virtualization Engine will be Cluster 0, which TS7740 will be Cluster 1, and which will be Cluster 2. The information in Table E-1 on page 577 helps the installer know which TS7740 is designated as Cluster 0, which is Cluster 1, and which is Cluster 2.

Complete the *Description of Location* column with geographical information that is unique to each cluster. Complete as much information as possible. Suggestions for descriptions are the x, y coordinates within a data center, room number, floor number, building numbers, city, and so forth.

Cluster descriptions

Use Table E-1 to enter information about cluster descriptions.

Table E-1 Cluster descriptions

Cluster	Machine Type-Model	Serial Number	Description of Location
Cluster 0	3957-V06		
Cluster 0	3494-Lxx		
Cluster 1	3957-V06		
Cluster 1	3494-Lxx		
Cluster 2	3957-V06		
Cluster 2	3494-Lxx		

3494 Tape Library drive information

Use Table E-2 to enter Tape Library drive information.

Table E-2 3494 Tape Library drive information

Field	Value	Notes
Physical Frame Number where Tape Drives are physically installed for this TS7700 Virtualization Engine	FRAME NUMBER:	► A 3494 can be connected to one or two TS7700 subsystems. The first is known as VTS-1, the second (if installed) is known as VTS-2. Each TS7700 will have a separate set of tape drives. A minimum of 4 and a maximum of 12 drives can be connected to each TS7700 Virtualization Engine attached to a 3494 Tape Library. All the tape drives for a TS7700 must be installed in the same 3494 D22 frame. The frame number value that is to be entered into this table is based on the 3494 D22 frame position in relation to the 3494 Lxx frame. The 3494 Lxx frame will always be frame 1. If the 3494 D22 frame is the next frame connected to the 3494 Lxx frame then the frame number value will be 2. If the 3494 D22 frame is the 5th frame attached to the 3494 Lxx frame, then the frame number value will be 6.
Number of Tape Drives within the 3494 D22 that will be attached to the TS7700 Virtualization Engine.	NUMBER OF TAPE DRIVES:	

Media volume serial range

Complete Table E-3 on page 579 with the requested information. There might be one or many **Media Volume Serial Ranges**, so complete as many rows as apply to your system.

Refer to the following list for a description of the values needed for Table E-3 on page 579.

- **From, To:** Contains the bar code label volsers of all the cartridges that are assigned to a single TS7740 Virtualization Engine. As an example, if cartridges assigned to a TS7700 have bar code labels in the range from A00000JA - A00500JA, then record the following information:
  - From: A00000J
  - To: A00500
- **Media Type:** Indicated by the last two characters of the 8 character bar code label volser on the cartridges. As an example, if the cartridges are labeled 123456JA, then the media type is JA. See the rules defined in the following notes:
  - JA and JJ tape cartridges are supported and can be mixed in a TS7700 Virtualization Engine.
  - JB tape cartridges are also supported (and can be mixed with JA and JJ tape cartridges) if all of the tape drives that are associated with the TS7700 are 3592 model E05 drives and none of them are in J1A emulation mode.
  - No other tape types are currently supported for use with the TS7700.
  - If there is at least one 3592 model J1A tape drive associated with the TS7700, then *all* 3592 model E05 tape drives associated with the TS7700 must be in J1A emulation mode.
  - If at least one tape cartridge associated with a TS7700 has been written by a 3592 model E05 tape drive that is *not* in J1A emulation mode then the TS7700 will no longer support any 3592 model J1A tape drives or any 3592 model E05 tape drive that is in J1A emulation mode. After you go to E05 Native mode, you cannot go back to J1A mode because the J1A tape drive cannot read or write a tape cartridge written in E05 mode and the TS7700 does *not* currently support mixed J1A and E05 tape drives (unless all E05 tape drives are in J1A emulation mode).
  - The capacity of a JJ tape cartridge is 60 GB if written by a J1A drive (or an E05 drive that is in J1A emulation mode) or 100 GB if written by an E05 drive that is *not* in J1A emulation mode. The capacity of a JA tape cartridge is 300 GB if written by a J1A drive (or an E05 drive that is in J1A emulation mode) or 500 GB if written by an E05 drive that is *not* in J1A emulation mode. The capacity of a JB cartridge is 700 GB.
  - TS7700 Feature Code 9900 (Encryption) requires that all tape drives are in E05 Native mode. The Encryption feature is *not* compatible with J1A emulation mode.

  Using the rules defined in "3494 Tape Library drive information" on page 577, determine whether you want the tape drives attached to the TS7700 to be in J1A emulation mode or in E05 native mode. Record this in the Tape Drive Format entry in Table E-4 on page 580.

- **(Distributed) Library Sequence Number**: Typically assigned by the system administrator. This is a 5 character name that is used as an identifier for a specific cluster and the associated LM and Library in a Grid configuration. This identifier is specified in the 3494 Library Manager configuration. *It is required even if the TS7700 is not in a Grid configuration*.

> **Note:** Each VTS partition must have a single, unique value for the Distributed Library Sequence Number. For the TS7700, a typical value is the last 5 digits of the 3494-Lxx frame serial number.

- **Home Pool** (also called *Scratch Pool*): You might have assigned a Home Pool value. If one has not been set, the default value is 00.

> **Note:**
>
> A *pool* is a group of physical tape cartridges. A *scratch pool* is a group of cartridges that are considered to be scratch, meaning they are ready for use by any write job.
>
> If there is at least one J70 or C06 Tape Control Unit connected to the 3494 tape Library, you can choose to record the **Native Lib Seq Num** in the column labeled **(Distributed) Library Sequence Number**. This info will not be used for the TS7700.

Table E-3 Media volume serial range

Cluster (0, 1, or 2)	From	To	Media Type	Distributed Library Sequence Number	Home Pool	VTS Number (1 or 2)	Fast Ready Category

3494 configuration information

Use Table E-4 to record 3494 configuration information.

Table E-4 3494 configuration information

Field	Value	Notes
Logical Device IDs	Cluster 0:	► This is a 3-digit ID chosen by the customer's administrator. The value is usually the last 3 digits of the host HCD device address. It is used to identify a range of Logical Device IDs (128 or 256). ► The Logical Device ID is the identifier for this device in error log output. Specifying an ID that is already assigned will not cause errors, however it will be harder to distinguish this device during trouble shooting
	Cluster 1:	
	Cluster 2:	
Customer IP Address	Cluster 0: LM-A ____.____.____.____ LM-B ____.____.____.____	Used for 3494 Library Manager Web Specialist The Ethernet ports are 10/100 Mb only.
	Cluster 1: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 2: LM-A ____.____.____.____ LM-B ____.____.____.____	
Customer Subnet mask	Cluster 0: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 1: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 2: LM-A ____.____.____.____ LM-B ____.____.____.____	
Customer Gateway	Cluster 0: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 1: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 2: LM-A ____.____.____.____ LM-B ____.____.____.____	

Field	Value	Notes
Hostname	Cluster 0: LM-A ____.____.____.____ LM-B ____.____.____.____	Used for 3494 Web Specialist.
	Cluster 1: LM-A ____.____.____.____ LM-B ____.____.____.____	
	Cluster 2: LM-A ____.____.____.____ LM-B ____.____.____.____	

TS7700 Virtualization Engine configuration information

Use Table E-5 to enter TS7700 Virtuatlization Engine configuration information.

Table E-5 TS7700 Virtualization Engine configuration information

Field	Value	Notes
Composite Library Sequence Number		► This 5 character name must be the same on all clusters (peers) within the same grid. This identifier is specified in the TS7700 configuration. It is required even if the machine is not in a grid configuration. ► The Composite Library Sequence Number must be different than the Distributed Library Sequence number specified in the 3494 configuration.
VTS Partition	Cluster 0:	1 or 2. Each 3494 Library Manager (or pair of LMs if they are in a dual LM configuration for redundancy) can be attached to one or two TS7700s. The first TS7700 connected to a 3494 is called VTS-1 and the second (if present) is called VTS-2.
	Cluster 1:	
	Cluster 2:	
Customer IP 1 (Virtual)	Cluster 0: ____.____.____.____	► Used for TS7700 web Management Interface ► This number is a virtual IP that is not associated with a physical cable. It will communicate through the Primary IP, and will automatically fail over to the Alternate IP when required.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Customer IP 2 (Primary)	Cluster 0: ____.____.____.____	► This is the IP address used to connect to the TS7700 through the internal primary network. This IP address should not be used by the customer unless the Virtual IP is inaccessible. ► The Ethernet ports are 10/100 Mb only.
	Cluster 1: ____.____.____.____	
	Cluster2: ____.____.____.____	
Customer IP 3 (Alternate)	Cluster 0: ____.____.____.____	► This is the IP address used to connect to the TS7700 through the internal alternate network. This IP address should not be used by the customer unless the Virtual IP is inaccessible. ► The Ethernet ports are 10/100 Mb only.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Customer Gateway	Cluster 0: ____.____.____.____	This is used with the virtual, primary, and alternate customer IP addresses.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	

Field	Value	Notes
Customer Subnet Mask	Cluster 0: ____.____.____.____	This is used with the virtual, primary, and alternate customer IP addresses.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
NTP server IP address (if used) Using the NTP server is strongly recommended to ensure that all components have consistent time settings.	____.____.____.____	The TCP/IP address you obtain from the customer is either the NTP server at their site (if they maintain one locally), or an internet server. Use of an internet server assumes that the customer allows access to the internet on the NTP services port. (TCP/IP port 123).
Tape Drive Format	Cluster 0: J1A Emulation Mode [ ] E05 Native Mode [ ]	To determine whether to use J1A emulation mode or E05 native mode, refer to the rules listed above under **Media Type**.
	Cluster 1: J1A Emulation Mode [ ] E05 Native Mode [ ]	
	Cluster 2: J1A Emulation Mode [ ] E05 Native Mode [ ]	

TSSC remote support information

Use Table E-6 to enter TSSC remote support information.

Table E-6 TSSC remote support information

Field	Value	Notes
Dial Out Prefix (if required)	Cluster 0:	If you need to dial a prefix to get an outside line, specify it here. As an example, if you need to dial 9 to get an outside line, record a 9 in this field.
	Cluster 1:	
	Cluster 2:	
Call-in Modem Number	Cluster 0:	This is the telephone number assigned to the analog phone line where the call-home modem will be connected.
	Cluster 2:	
	Cluster 2:	
Customer Company Name	Cluster 0:	This is the value that will be sent out with the call home records.
	Cluster 1:	
	Cluster 2:	
Customer Contact Phone Number	Cluster 0:	This is the phone number that IBM contacts when call-home indicates a failure.
	Cluster 1:	
	Cluster 2:	
Off Shift Customer Contact Phone Number	Cluster 0:	This is the alternate phone number that IBM contacts when call-home indicates a failure.
	Cluster 1:	
	Cluster 2:	
IP Address (for optional Ethernet Call Home)	Cluster 0:	► This is an optional IP address on the customer network that will allow the TSSC to Call Home using a high speed connection through the customer network to the internet. Ports 443 (https) and 80 (http) on this IP address must be free for outbound traffic through the customer network to the internet. ► Call Home through an Ethernet connection greatly improves the speed of the data transfer, making the data available to support personnel more quickly. At this time the Call Home through an Ethernet connection is only used for Call Home (outbound). IBM recommends enabling Call Home to improve data collection when a problem occurs. Call Home through an Ethernet connection provides better throughput and lower cost than Call Home through a modem connection ► The modem is still needed for Call In (remote support).
	Cluster 1:	
	Cluster 2:	

Field	Value	Notes
Subnet Mask (for optional Ethernet Call Home)	Cluster 0:	This is the subnet mask associated with the IP Address (for optional Ethernet Call Home).
	Cluster 1:	
	Cluster 2:	
Gateway (for optional Ethernet Call Home)	Cluster 0:	This is the gateway associated with the IP Address (for optional Ethernet Call Home).
	Cluster 1:	
	Cluster 2:	

Grid local addresses

The following notes apply to Table E-7:

- If the TS7700 that you are installing is a standalone machine (not part of a grid configuration) then leave Table E-7 blank.
- The grid interfaces are the 1 Gb internet connections between clusters, allowing them to automatically remain in sync.
- IBM strongly recommends that the primary and alternate grid interfaces be on separate sub-nets. As an example, 10.10.1.n for the primary interface and 10.11.1.n for the alternate interface. Notice that the second set of octets must be different. If the grid interfaces are direct connected (without using Ethernet switches), then using separate sub-nets is required.

The Grid interfaces require connections using the following TCP/IP ports:

- 7 (Ping)
- 9 (Discard Service for bandwidth measuring tools)
- 1415 (WebSphere message queues Grid - Grid)
- 1416 (WebSphere message queue HDM - HDM)
- 123 (Network Time Protocol)
- 350 (Distributed Library - Distributed Library File Transfer).

The following TCP/IP ports are also useful in service scenarios if allowed:

- 23 (Telnet)
- 20 and 21 (FTP)

Table E-7 Grid local addresses

Field	Value	Notes
Primary Grid Interface IP address	Cluster 0: ____.____.____.____	The Primary Grid Interface is the 1 Gb Ethernet Adapter located in slot C4 of the TS7700. The cluster 0 Primary Grid Interface will connect to the cluster 1 Primary Grid Interface.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Primary Grid Interface Network Mask	Cluster 0: ____.____.____.____	
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Primary Grid Interface Gateway	Cluster 0: ____.____.____.____	If a Gateway is not used, leave this field blank. If using crossover cables you must not specify a Gateway.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Alternate Grid Interface IP address	Cluster 0: ____.____.____.____	The Alternate Grid Interface is the 1 Gb Ethernet Adapter located in slot C5 of the 3957-V06. The cluster 0 Alternate Grid Interface will connect to the cluster 1 Alternate Grid Interface.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Alternate Grid Interface Network Mask	Cluster 0: ____.____.____.____	If a Gateway is not used, leave this field blank. If using crossover cables you must not specify a Gateway.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	

Field	Value	Notes
Alternate Grid Interface Gateway	Cluster 0: ____.____.____.____	
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	

TSSC Grid configuration information

The following notes apply to Table E-8:

- If the TS7700 that you are installing is a standalone machine (not part of a grid configuration), then leave Table E-8 blank.
- If you will be using the Grid in a Cascade Deferred style, then you do not need the Autonomic Ownership Takeover Manager (AOTM) and should leave Table E-8 blank.
- Refer to *The Autonomic Ownership Takeover Manager (AOTM)* chapter of the *IBM Virtualization Engine TS7740 Installation Roadmap* for more information about AOTM before you continue. Do *not* attempt to configure AOTM, but use the information to make an informed decision on whether to use AOTM.
- If you do not want to use the Autonomic Ownership Takeover Manager (AOTM) then leave Table E-8 0 blank.
- The TSSC grid interface is only used for the Autonomic Ownership Takeover Manager (AOTM).
- Each cluster can be configured to use AOTM to provide ownership takeover for one cluster.

The AOTM requires the following TCP/IP ports:

- 7 (Ping)
- 80 (HTTP)

Table E-8 TSSC Grid configuration information

Field	Value	Notes
TSSC Grid Interface IP Address	Cluster 0: ____.____.____.____	The TSSC Grid Interface is used to allow the TSSC at one cluster to communicate with the TSSC at another cluster. This is required if the Autonomic Ownership Takeover Manager (AOTM) will be used.
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
TSSC Grid Interface Subnet Mask	Cluster 0: ____.____.____.____	
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
TSSC Grid Interface Gateway	Cluster 0: ____.____.____.____	
	Cluster 1: ____.____.____.____	
	Cluster 2: ____.____.____.____	
Takeover Mode	Cluster 0: ROT [ ] WOT [ ]	The Takeover Mode must be either Read Only Takeover (ROT) or Write Only Takeover (WOT).
	Cluster 1: ROT [ ] WOT [ ]	
	Cluster 2: ROT [ ] WOT [ ]	

Field	Value	Notes
Grace Period (Minutes)	Cluster 0:	When a cluster detects that another cluster in the same Grid has failed, it will wait for the number of minutes specified as the Grace Period before it attempts to take over the volumes of the failed cluster. A typical Grace Period is 25 minutes.
	Cluster 1:	
	Cluster 2:	
Retry Period (Minutes)	Cluster 0:	The retry period is the number of minutes between attempts to take over ownership of the volumes associated with a failed cluster. A typical Retry Period is 5 minutes. In most cases this will be the same for both clusters.
	Cluster 1:	
	Cluster 2:	

F

TS7700 Virtualization Engine implementation, step-by-step

In this appendix, we document the typical implementation scenario for the TS7700. The steps that we discuss here are a guide only. You need to tailor the steps specifically for your site.

We include two examples in this appendix, and both examples assume that the host system is DFSMS:

- First-time TS7700 library installation *without* use of outboard policy management.
- First-time TS7700 library installation *with* use of outboard policy management.

First-time TS7700 library installation

This example shows details of a sample definition procedure for a first-time tape library installation without use of outboard policy management:

1. Analyze TS7700 Virtualization Engine targeted workloads to determine the appropriate TS7700 configuration for your environment. The BatchMagic tool can be helpful in making this determination.

   > **Important:** This is the crucial step in configuring your TS7700 Virtualization Engine correctly. You should perform this step in conjunction with your IBM Storage Specialist.

2. Check the latest TS7700, TS3500 and 3953 Systems Assurance Product Review (SAPR) Guide.
3. Order stacked volume media or labels.
4. Set up the HCD (see 5.2.1, “Defining devices through HCD” on page 200).
5. Vary CHPIDs, paths, and TS7700 virtual drives online.
6. JES3 system considerations (see Appendix G, “JES3 implementation” on page 617).
7. For DFSMSrmm:
   a. Create VLPOOLS for logical volumes.
   b. Add logical volumes/location to DFRMM.
   c. Check REJECT PARMs.
   d. Refresh DFRMM.
8. For CA1:
   a. If your planned logical volume range has not been defined to your Tape Management System (TMS), you need to extend the TMC. Also check the amount of free DSNBs available and add any, if required, concurrent with this extension.
   b. Check that the SCRTCH PARM is set to YES in PPOPTION. When a tape becomes scratch or a DSNB record gets deleted, control is given to the security exit TMSUX2S. This is the scratch notification interface between CA1 and the 3953 Library Manager.
   c. Depending on the parameters specified in PPOPTION, you might need to initialize the logical volumes after they have been inserted.

   > **Note:** Check with the vendor of your tape management product to ensure that the level of software installed on your system will support the TS35000/3953/TS7700.

9. Check and, if necessary, modify current SYS1.PARMLIB members:
   – IEFSSNxx

     Add or update the OAM1 entry with the name of the initialization module (CBRINIT) executed at IPL.

   – IGDSMSxx

     If you want OAM to start automatically as part of SMS initialization, add the OAMPROC and OAMTASK optional parameters.

> **Note:** If you use an Independent Storage Vendor (ISV) tape management system, it might require that the OAM address space be started after the tape management initialization. In this case, do not start OAM automatically. Check with the vendor of the tape management product.

– CONSOLxx

  If you want to receive library messages at a specific console, update the CONSOLxx member referenced by IEASYSxx. This console name must also be defined to SMS and is set when you define the library through Interactive Storage Management Facility (ISMF).

– DEVSUPxx

  This member is used to hard partition your TS7700 by specifying category codes to provide a unique range of tapes for each system. If no category codes are specified, defaults are assigned.

> **Important:** We recommend that you do not use the DEVSUPxx default categories. Avoiding the use of the defaults will make future partitioning of the library with other systems easier and more secure. See *Guide to Sharing and Partitioning IBM Tape Libraries*, SG24-4409 for additional information.

– COMMNDxx

  If you want your library automatically brought online after each IPL, add the VARY SMS,LIBRARY command. Another area where you can bring your library automatically online is using ISMF when you define your library. Set the Initial Online Status to YES.

– GRSCNFxx

  If you are going to share the tape library among two or more systems in a DFSMS complex, a global resource serialization ring can be created to include all sharing systems. This enables OAM to serialize the cartridge entry process.

– LOADxx (optional)

  The default data set name for the TCDB is SYS1.VOLCAT.VGENERAL. If you want to use this name, no change is required to the LOADxx member. If you want to use a different HLQ than the defined default, update columns 64 through 71 of the SYSCAT statement with the preferred HLQ.

– COFVLFxx

  Add the volume catalogs to the IGGCAS class definitions where you have other ICF catalogs.

– ALLOCxx

  Add tape automation policies.

– IECIOSxx

  Set Missing Interrupt Handler values. This value should be 45 minutes for the TS7700 logical drives.

> **Note:** For coding examples and more information regarding SYS1.PARMLIB updates, refer to *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-46322.

10. If your IBM System Service Representative (SSR) has installed your TS7700, perform the following Library Manager definitions:

   – Define stacked volume ranges (see 4.3.4, “Define VOLSER ranges for physical (stacked) volumes” on page 151).

   – Define Fast Ready categories (see 4.3.5, “Define Fast Ready categories” on page 154). The Fast Ready category number you define is the same value as those defined for MEDIA1 or MEDIA2 as specified in the DEVSUPxx of your PARMLIB.

   – Define the expired logical volume data policy (see 4.3.6, “Define logical volume Expiration Time” on page 156).

   – Define TS7700 management policies (see 4.3.7, “Define TS7700 management policies” on page 157):

     - Inhibit reclaim schedule
     - Reclaim threshold percentage
     - Free storage threshold

11. If stacked volume media is available, insert the physical TS7700 stacked volumes (see “Inserting physical volumes into the IBM TS3500 Tape Library” on page 142.

12. Create TCDB as shown in Example F-1.

Example F-1 Define TCDB

```
//DEFINECAT JOB
//STEP1    EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
   DEFINE USERCATALOG -
   (NAME(SYS1.VOLCAT.VGENERAL) -
   VOLCATALOG -
   VOLUME(volser) -
   CYLINDERS(30,30))
/*
```

When sizing the TCDB, calculate 275 bytes for each volume entry. For 10 000 logical volumes, 2.75 MB or about 3 cylinders of 3390 disk space is required. Allow sufficient room for growth.

> **Note:** If you are sharing the tape library (non-hard partitioned), use IDCAMS to IMPORT CONNECT the VOLCAT to the other sharing systems.

13. Define security profiles:

   – ISMF
   – DFSMS constructs
   – STGADMIN
   – DFSMSrmm
   – z/OS Operator Commands

> **Note:** The following manuals are helpful in setting up your security environment:
>
> - *z/OS DFSMSdfp Storage Administrator Reference*, SC35-0422
> - *z/OS DFSMSrmm Implementation and Customization Guide*, SC26-7405
> - *z/OS MVS Planning: Operation*, SC22-7601

14. Create the cataloged OAM procedure.

15. Start OAM. Upon startup, you might receive the following messages:

– `CBR1115I - No libraries defined to OAM.`
– `CBR0094E - OAM has initialized without tape or object support.`

These messages are informational and expected. They will be resolved when the library has been defined and the DFSMS SCDS activated.

> **Tip:** Code the RESTART=YES PARM in your OAM procedure. This PARM, upon an activation of the SCDS, automatically restarts OAM. This is important when you have changed tape library related constructs. Without this PARM, you must restart OAM manually after an SCDS activation if you have changed tape library related constructs. This can be accomplished by the following command:
>
> F OAM,RESTART

16. Define the following DFSMS constructs using ISMF:

– Tape library
  - Figure F-12 on page 605
  - Figure F-13 on page 605
– Data class
  - Figure F-1 on page 596
  - Figure F-2 on page 596
  - Figure F-3 on page 597
  - Figure F-4 on page 597
  - Figure F-5 on page 598
– Storage class
  - Figure F-6 on page 600
  - Figure F-7 on page 600
– Management class
  - Figure F-8 on page 602
  - Figure F-9 on page 602
  - Figure F-10 on page 603
– Storage group
  - Figure F-11 on page 604

17. Create the following ACS routines for DFSMS-managed tape:

– Data class
  - Example F-2 on page 599
– Storage class
  - Example F-3 on page 601
– Management class
  - Example F-4 on page 603
– Storage group
  - Example F-5 on page 604

```
                              DATA CLASS DEFINE                Page 1 of 5
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVE

To DEFINE Data Class, Specify:
  Description ==> Default TS7700 Data Class
              ==>
Recfm . . . . . . . . . .                (any valid RECFM combination or blank)
Lrecl . . . . . . . . . .                (1 to 32761 or blank)
Space Avgrec  . . . . . .                (U, K, M or blank)
        Avg Value . . . . .              (0 to 65535 or blank)
        Primary . . . . . .              (0 to 999999 or blank)
        Secondary . . . . .              (0 to 999999 or blank)
        Directory . . . . .              (0 to 999999 or blank)
Retpd or Expdt  . . . . .                (0 to 9999, YYYY/MM/DD or blank)
Volume Count  . . . . . . 1              (1 to 59 or blank)
Add'l Volume Amount . . .                (P=Primary, S=Secondary or blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-1 Data Class Define ISMF panel 1 of 5 - first-time setup

```
                              DATA CLASS DEFINE                Page 2 of 5
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVE
To DEFINE Data Class, Specify:

  Data Set Name Type  . . . . . .       (EXT, HFS, LIB, PDS or blank)
    If Ext  . . . . . . . . . . .       (P=Preferred, R=Required or blank)
    Extended Addressability . . . N     (Y or N)
    Record Access Bias  . . . . .       (S=System, U=User or blank)
  Space Constraint Relief . . . . N     (Y or N)
    Reduce Space Up To (%)  . . .       (0 to 99 or blank)
    Dynamic Volume Count  . . . .       (1 to 59 or blank)
  Compaction  . . . . . . . . . . Y     (Y, N, T, G or blank)

  Spanned / Nonspanned  . . . . .       (S=Spanned, N=Nonspanned or blank)



Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-2 Data Class Define ISMF panel 2 of 5 - first-time setup

```
                              DATA CLASS DEFINE                Page 3 of 5
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVE

To DEFINE Data Class, Specify:

  Media Interchange
    Media Type . . . . . . . . 2         (1, 2, 3, 4, 5, 6, 7, 8, 9, 10 or
blank)
   Recording Technology . . . 36       (18,36,128,256,384,E1,E2,EE2 or blank)
    Performance Scaling  . . .  N        (Y, N or blank)
    Performance Segmentation. .          (Y, N or blank)
  Block Size Limit . . . . .             (32760 to 2GB or blank)
  Recorg . . . . . . . . . .             (KS, ES, RR, LS or blank)
  Keylen . . . . . . . . . .             (0 to 255 or blank)
  Keyoff . . . . . . . . . .             (0 to 32760 or blank)
  CIsize Data  . . . . . . .             (1 to 32768 or blank)
  % Freespace CI . . . . . .             (0 to 100 or blank)
              CA . . . . . .             (0 to 100 or blank)
Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-3 Data Class Define ISMF panel 3 of 5 - first-time setup

```
                              DATA CLASS DEFINE                Page 4 of 5

Command ===>

 SCDS Name . . . : CATIA.SCDS
 Data Class Name : DCVE

 To DEFINE Data Class, Specify:

   Encryption Management
     Key Label 1 . . .     (1 to 64 characters or blank)

     Key Label 2 . . .


   Encoding for Key Label 1  . . . . .           (L, H or blank)
   Encoding for Key Label 2  . . . . .           (L, H or blank)




 Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
 Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-4 Data Class Define ISMF panel 4 of 5 - first-time setup

```
                                                  DATA CLASS DEFINE
             Page 5 of 5
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVE

To DEFINE Data Class, Specify:

  Shareoptions Xregion  . . .          (1 to 4 or blank)
               Xsystem  . . .          (3, 4 or blank)
  Reuse . . . . . . . . . . . N        (Y or N)
  Initial Load  . . . . . . . R        (S=Speed, R=Recovery or blank)
  BWO . . . . . . . . . . . .          (TC=TYPECICS, TI=TYPEIMS, NO or blank)
  Log . . . . . . . . . . . .          (N=NONE, U=UNDO, A=ALL or blank)
  Logstream Id  . . . . . . .
  FRlog . . . . . . . . . . .          (A=ALL, N=NONE, R=REDO, U=UNDO or blank)
  RLS CF Cache Value  . . . . A        (A=ALL, N=NONE, U=UPDATESONLY)
  RLS Above the 2-GB Bar  . . N        (Y or N)
  Extent Constraint Removal   N        (Y or N)

Use ENTER to perform Verification; Use UP Command to View previous Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit
```

Figure F-5 Data Class Define ISMF panel 5 of 5 - first-time setup

Example F-2 Data Class ACS routine

```
PROC 0 DATACLAS

 /*------------------------------------------------------------------*/
 /* DEFINE VALID TAPE UNIT ESOTERICS                                  */
 /*------------------------------------------------------------------*/
        FILTLIST TAPE_UNIT INCLUDE(348*,TAPE*,'CART','AFF=SMST','3420',
                                   '3590','3590-1','3592-J1A','3592-E05','3490')

        FILTLIST DBASE_BKP   INCLUDE(PROD.DBASE.BKP.**)

        FILTLIST PROD_SMF    INCLUDE(PROD.SMF.**)

          SELECT
             WHEN ((&UNIT = &TAPE_UNIT) OR (&ANYVOL = 'REF=ST'))
                DO
                  SELECT
                     WHEN (&DATACLAS NE '')                /* Allows users to
specify  */
                     DO                                /* data class for tape
*/
                         SET &DATACLAS = &DATACLAS
                         EXIT
                       END
                                WHEN (&DSN = &DBASE_BKP)
                       DO
                         SET &DATACLAS = 'DC3590K'
                         EXIT
                       END
                     WHEN (&DSN = &PROD_SMF)
                       DO
                         SET &DATACLAS = 'DCVE'
                         EXIT
                       END
                  END
                END
                   OTHERWISE
                    DO
                      SET &DATACLAS = ''
                      EXIT
                    END
           END /* END SELECT */
      END      /* END PROC   */
```

```
                              STORAGE CLASS DEFINE              Page 1 of 2
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Class Name  : SCTAPE
To DEFINE Storage Class, Specify:
  Description ==> Default Storage Class
              ==>
  Performance Objectives
   Direct Millisecond Response . . . .               (1 to 999 or blank)
   Direct Bias . . . . . . . . . . . .               (R, W or blank)
   Sequential Millisecond Response . .               (1 to 999 or blank)
   Sequential Bias . . . . . . . . . .               (R, W or blank)
   Initial Access Response Seconds . .               (0 to 9999 or blank)
   Sustained Data Rate (MB/sec)  . . .               (0 to 999 or blank)
  Availability . . . . . . . . . . . . N             (C, P ,S or N)
  Accessibility  . . . . . . . . . . . N             (C, P ,S or N)
   Backup  . . . . . . . . . . . . . .               (Y, N or Blank)
   Versioning  . . . . . . . . . . . .               (Y, N or Blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Page;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-6   Storage Class Define ISMF panel 1 of 2 - first-time setup

```
                            STORAGE CLASS DEFINE               Page 2 of 2
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Class Name  : SCTAPE

To DEFINE Storage Class, Specify:

  Guaranteed Space  . . . . . . . . . N           (Y or N)
  Guaranteed Synchronous Write  . . . N           (Y or N)
  Multi-Tiered SG . . . . . . . . . .             (Y, N, or blank)
  Parallel Access Volume Capability   N           (R, P, S, or N)
  CF Cache Set Name . . . . . . . . .             (up to 8 chars or blank)
  CF Direct Weight  . . . . . . . . .             (1 to 11 or blank)
  CF Sequential Weight  . . . . . . .             (1 to 11 or blank)




Use ENTER to Perform Verification; Use UP Command to View previous Page;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-7   Storage Class Define ISMF panel 2 of 2 - first-time setup

Example F-3 Storage Class ACS routine

```
PROC 0 STORCLAS

 /*-----------------------------------------------------------------*/
 /* DEFINE VALID TAPE UNIT ESOTERICS                                */
 /*-----------------------------------------------------------------*/
    FILTLIST TAPE_UNIT INCLUDE(348*,TAPE*,'CART','AFF=SMST','3420',
                           '3590','3590-1','3592-J1A','3592-E05','3490')

    FILTLIST SMS_TAPE  INCLUDE('DCVE','DC3590K')

     SELECT

       WHEN ((&UNIT = &TAPE_UNIT) OR (&ANYVOL = 'REF=ST'))
         DO
           SELECT
              WHEN (&DATACLAS = &SMS_TAPE)
                DO
                   SET &STORCLAS = 'SCTAPE'
                   EXIT
                END
           END
         END
           OTHERWISE
              DO
                SET &STORCLAS = ''
                EXIT
              END
     END /* END SELECT */
  END  /* END PROC */
```

```
                            MANAGEMENT CLASS DEFINE          Page 1 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCTAPE

To DEFINE Management Class, Specify:

  Description ==> Default Management Class
              ==>

  Expiration Attributes
    Expire after Days Non-usage  . . NOLIMIT     (1 to 9999 or NOLIMIT)
    Expire after Date/Days . . . . . NOLIMIT     (0 to 9999, yyyy/mm/dd or
                                                 NOLIMIT)

  Retention Limit  . . . . . . . . . NOLIMIT     (0 to 9999 or NOLIMIT)



Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-8 Management Class Define ISMF panel 1 of 3 - first-time setup

```
                            MANAGEMENT CLASS DEFINE          Page 2 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCTAPE

To DEFINE Management Class, Specify:

  Partial Release . . . . . . . . . N          (Y, C, YI, CI or N)

  Migration Attributes
    Primary Days Non-usage  . . . .            (0 to 9999 or blank)
    Level 1 Days Non-usage  . . . .            (0 to 9999, NOLIMIT or blank)
    Command or Auto Migrate . . . . NONE       (BOTH, COMMAND or NONE)


  GDG Management Attributes
    # GDG Elements on Primary . . .            (0 to 255 or blank)
    Rolled-off GDS Action . . . . .            (MIGRATE, EXPIRE or blank)

Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-9 Management Class Define ISMF panel 2 of 3 - first-time setup

```
                            MANAGEMENT CLASS DEFINE            Page 3 of 5
Command ===>

SCDS Name . . . . . . . : CATIA.SCDS
Management Class Name : MCTAPE

To DEFINE Management Class, Specify:
  Backup Attributes
    Backup Frequency  . . . . . . . .          (0 to 9999 or blank)
    Number of Backup Vers . . . . . .          (1 to 100 or blank)
      (Data Set Exists)
    Number of Backup Vers . . . . . .          (0 to 100 or blank)
      (Data Set Deleted)
    Retain days only Backup Ver . . .          (1 to 9999, NOLIMIT or blank)
      (Data Set Deleted)
    Retain days extra Backup Vers . .          (1 to 9999, NOLIMIT or blank)
    Admin or User command Backup  . . NONE     (BOTH, ADMIN or NONE)
    Auto Backup . . . . . . . . . . . N        (Y or N)
    Backup Copy Technique . . . . . . S        (P=Conc Preferred, R=Conc
                                               Required or S=Standard)
Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; Cancel to Exit.
```

Figure F-10 Management Class Define ISMF panel 3 of 3 - first-time setup

> **Note:** For subsequent Management Class panels, provide null values.

Example F-4 Management Class ACS routine

```
PROC 0 MGMTCLAS

/*******************************************************************/
/*            DEFINE TAPE DATASET FILTERING CRITERIA               */
/*******************************************************************/

 FILTLIST SMS_TAPE  INCLUDE('DCVE','DC3590K')

 IF (&DATACLAS = &SMS_TAPE) THEN
   DO
     SET &MGMTCLAS = 'MCTAPE'
   EXIT
 END
END /* END PROC *
```

```
                              TAPE STORAGE GROUP DEFINE
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Group Name  : VE1

To DEFINE Storage Group, Specify:

  Description ==> Test Storage Group
              ==>

  Library Names  (1 to 8 characters each):
    ===> BBBVE1   ===>          ===>          ===>
    ===>           ===>          ===>          ===>



  DEFINE    SMS Storage Group Status . . Y (Y or N)


Use ENTER to Perform Verification and Selection;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-11 Tape Storage Group Define ISMF panel - first-time setup

Notes:

- Select Storage Group Type=TAPE to define a tape Storage Group.
- Select Y in the SMS Storage Group Status to alter the Storage Group status of your connected systems or system groups.
- The Library Name you assign in the Define Storage Group panel is matched with the name you use when you define the library (Figure F-12).

Example F-5 Storage group ACS routine

```
PROC 0 STORGRP

  SELECT
     WHEN (&ANYVOL = 'REF=ST')
        DO
           SET &STORGRP = &STORGRP
           EXIT
        END
     WHEN (&DATACLAS = 'DCVE')
        DO
           SET &STORGRP = 'VE1'
           EXIT
        END
     WHEN (&DATACLAS = 'DC3590K')
        DO
           SET &STORGRP = 'ATL1'
           EXIT
        END
 END /* END SELECT */
END  /* END PROC   */
```

18. Define the Tape Library using ISMF, as shown in Figure F-13.

```
                                TAPE LIBRARY DEFINE              Page 1 of 2
Command ===>

SCDS Name  . : CATIA.SCDS
Library Name : BBBVE1

To Define Library, Specify:
  Description ==> BBBVE1 Library Define
              ==>
  Library ID . . . . . . . . . . . 12345     (00001 to FFFFF)
  Console Name . . . . . . . . . .
  Entry Default Data Class . . . . DCVE
  Entry Default Use Attribute  . . S         (P=PRIVATE or S=SCRATCH)
  Eject Default  . . . . . . . . . P         (P=PURGE or K=KEEP)

  Media Type:  Scratch Threshold        Media1  . . . . 0        (0 to 999999)
    Media2  . . . 400                   Media3  . . . . 0        (0 to 999999)
    Media4  . . . 0                     Media5  . . . . 0        (0 to 999999)
    Media6  . . . 0                     Media7  . . . . 0        (0 to 999999)
    Media8  . . . 0                     Media9  . . . . 0        (0 to 999999)
    Media10 . . . 0

Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-12 Tape Library Define ISMF panel 1 of 2 - first-time setup

```
                                TAPE LIBRARY DEFINE              Page 2 of 2
Command ===>

SCDS Name  . : CATIA.SCDS
Library Name : BBBVE1

  Initial Online Status (Yes, No, or Blank):
    SYS1     ===>  YES  SYS2     ===> NO   SYS3     ===>      SYS4     ===>
    SYS5     ===>






Warning:
  When you connect a tape library to a system group rather than a system,
  you lose the ability to vary that library online or offline to the
  individual systems in the system group.  It is strongly recommended that
  the tape library be connected to individual systems only.
Use ENTER to Perform Verification; Use UP Command to View previous Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit
```

Figure F-13 Tape Library Define ISMF panel 2 of 2 - first-time setup

19. Translate the DFSMS ACS routines through ISMF.
20. Validate the DFSMS SCDS through ISMF.
21. Test the ACS routines through ISMF.
22. Activate the DFSMS SCDS.

> **Note:** By varying the TS7700 drives online previously, upon activation of the SCDS, the library comes online automatically. You must have a TS7700 drive online before the library can come online. To bring a library online manually, issue the following command:
>
> VARY SMS,LIB(library-name),ONLINE

After OAM is started and the library is online, if the host has even a single path initialized for communication with the library, that host is capable of doing cartridge entry processing. The library drives do not have to currently be online for the path to be usable. If this library is being partitioned with other host systems, it is important that each system's tape management CBRUXENT first be coded and linked to prevent the host from entering volumes that do not belong to it.

23. Insert logical volumes (see 4.5.8, "Insert logical volumes using the TS7700 Management Interface" on page 182. TS7700 is now ready for use.

> **Tip:** For help in navigating the ISMF panels, as well as detailed information pertaining to the defining of the DFSMS constructs, refer to the following manuals:
>
> - *IBM TotalStorage 3494 Tape Library: A Practical Guide to Tape Drives and Tape Automation*, SG24-4632
> - *z/OS DFSMSdfp Storage Administrator Reference*, SC35-0422

If you do not want to take advantage of the outboard policy management functions, no further actions are required. The next section describes how to set up these functions.

Set up Volume Pooling in your environment

The following example shows how to use Outboard Policy Management in a TS7700 Single Cluster Grid environment. This example assumes use of logical volumes that are currently assigned to the system. It uses the following criteria:

- Segregate production and testing logical volume data into pools
- Set up logical volume dual copy
- Set up early cache removal of logical volumes

Define Outboard Policy Management constructs

Through the ETL Specialist or Library Manager console, set up the Outboard Policy Management constructs using these steps:

1. Add two Storage Groups (see 4.4.2, “Creating Storage Groups” on page 167). For the purpose of this exercise, the Storage Groups are named SGPROD and SGTEST.
   - SGPROD is assigned Primary Pool 02. This pool contains logical volume production application data (Figure F-14).
   - SGTEST is assigned Primary Pool 01. This pool contains logical volume test application data.

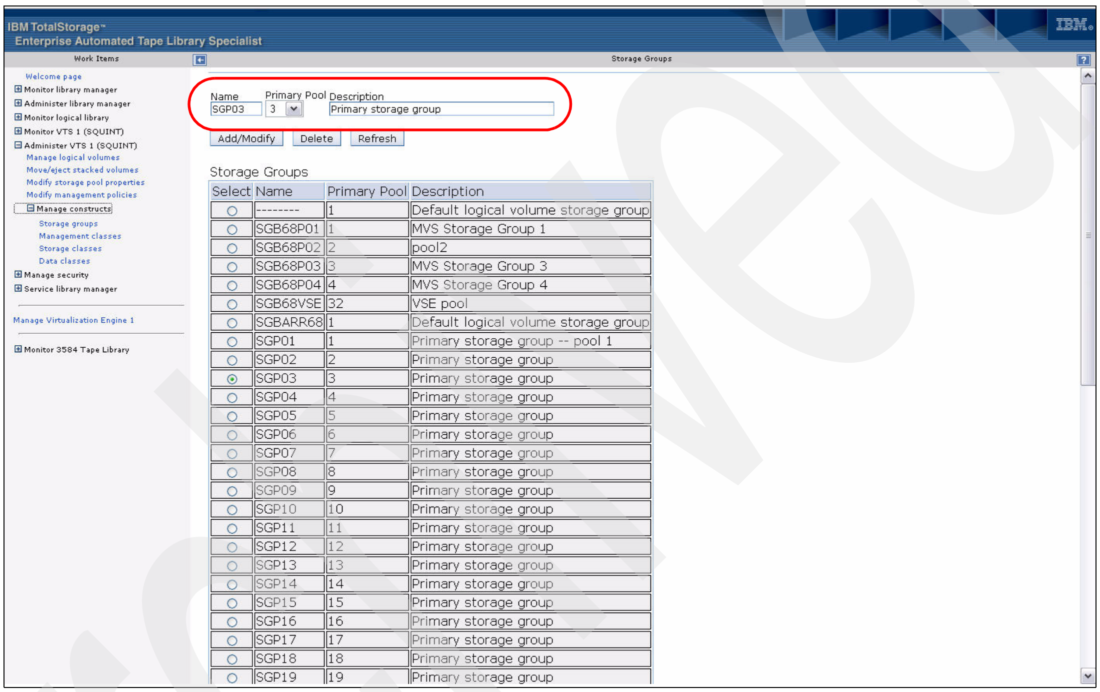

Figure F-14 ETL Specialist Manage Storage Groups - OPM setup

2. Repeat the process to create the Library Manager construct for Storage Group SGTEST.

3. Create the DFSMS Storage Group construct for SGPROD through ISMF (as shown in Figure F-15).

```
                         TAPE STORAGE GROUP DEFINE
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Group Name  : SGPROD

To DEFINE Storage Group, Specify:

  Description ==> Production Application Storage Group
              ==>

  Library Names  (1 to 8 characters each):
    ===> BBBVE1   ===>          ===>          ===>
    ===>          ===>          ===>          ===>



  DEFINE    SMS Storage Group Status . . N  (Y or N)


Use ENTER to Perform Verification and Selection;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-15 Storage Group Define ISMF panel - OPM setup

4. Repeat the process to create the DFSMS Storage Group construct SGTEST.
5. Create or modify your Storage Group DFSMS ACS routine (as shown in Example F-6).

Example F-6 Storage group routine - OPM setup

```
PROC 0 STORGRP

  SELECT

     WHEN (&ANYVOL = 'REF=ST')
        DO
           SET &STORGRP = &STORGRP
           EXIT
        END
     WHEN (&STORCLAS = 'SCCACHE')
        DO
           SET &STORGRP = 'SGPROD'
           EXIT
        END
     WHEN (&STORCLAS = 'SCNCACHE')
        DO
           SET &STORGRP = 'SGTEST'
           EXIT
        END
END /* END SELECT */
END /* END PROC */
```

6. Modify the pool properties for Pools 01, 02, and 03 (see 4.4.1, “Defining stacked volume pool properties” on page 164). Here, you need to define the parameters for your pools as described in 4.4.1, “Defining stacked volume pool properties” on page 164.
7. Repeat the same process for Pools 02 and 03.

> **Note:** For the purpose of this example, Pool 03 is the pool for the logical volume dual copies. Therefore, we do not need to create the DFSMS Storage Group construct or to make any references to it in the DFSMS Storage Group ACS routine.

8. Add or modify stacked volser ranges for Pools 01, 02, and 03 (see 4.3.4, “Define VOLSER ranges for physical (stacked) volumes” on page 151). Assign the following values to the pools:
   - Home pool
   - Media type
9. Insert stacked volumes.
10. Add two Storage Classes (see 4.4.4, “Creating Storage Classes” on page 170). For the purpose of this exercise, the Storage Classes are named SCCACHE and SCNCACHE. Here, you need to define the TVC preference using the Storage Class panel as described in 4.4.4, “Creating Storage Classes” on page 170.
    - SCCACHE is assigned the TVC Preference of 1. Logical Volumes assigned to this Storage Class are removed from the TVC by a Least Recently Used (LRU) algorithm. The logical volumes are copied to tape and deleted when space needs to be made available in the TVC.
    - SCNCACHE is assigned the TVC Preference of 0. Logical Volumes assigned to this Storage Class are removed from the TVC as soon as they are copied to tape. This case uses the *largest first* algorithm.

> **Note:** Using the above Preference Levels of 0 and 1 overrides any IART values that you currently have in effect. If you want to use current assigned IART values, specify IART as the Preference Level.

11. Repeat the process to create the Library Manager construct for Storage Class SCNCACHE.

12. Create the DFSMS Storage Class construct through ISMF (as shown in Figure F-16 and Figure F-17).

```
                              STORAGE CLASS DEFINE              Page 1 of 2
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Class Name  : SCCACHE
To DEFINE Storage Class, Specify:
  Description ==> KEEP IN CACHE STORAGE CLASS
              ==>
  Performance Objectives
   Direct Millisecond Response . . . .             (1 to 999 or blank)
   Direct Bias . . . . . . . . . . . .             (R, W or blank)
   Sequential Millisecond Response . .             (1 to 999 or blank)
   Sequential Bias . . . . . . . . . .             (R, W or blank)
   Initial Access Response Seconds . .             (0 to 9999 or blank)
   Sustained Data Rate (MB/sec)  . . .             (0 to 999 or blank)
  Availability . . . . . . . . . . . . N           (C, P ,S or N)
  Accessibility  . . . . . . . . . . . N           (C, P ,S or N)
   Backup  . . . . . . . . . . . . . .             (Y, N or Blank)
   Versioning  . . . . . . . . . . . .             (Y, N or Blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Page;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-16 Storage Class Define ISMF panel 1 of 2 - OPM setup

```
                            STORAGE CLASS DEFINE               Page 2 of 2
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Class Name  : SCCACHE

To DEFINE Storage Class, Specify:

  Guaranteed Space . . . . . . . . . N           (Y or N)
  Guaranteed Synchronous Write . . . N           (Y or N)
  CF Cache Set Name  . . . . . . . .             (up to 8 chars or blank)
  CF Direct Weight . . . . . . . . .             (1 to 11 or blank)
  CF Sequential Weight . . . . . . .             (1 to 11 or blank)






Use ENTER to Perform Verification; Use UP Command to View previous Page;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-17 Storage Class Define ISMF panel 2 of 2 - OPM setup

13. Repeat the process to create the DFSMS Storage Class construct SCNCACHE.
14. Create or modify your Storage Class DFSMS ACS routine as shown in Example F-7.

Example F-7 Storage Class routine - APM setup

```
PROC 0 STORCLAS

 /*-----------------------------------------------------------------*/
 /* DEFINE VALID TAPE UNIT ESOTERICS                                */
 /*-----------------------------------------------------------------*/
   FILTLIST TAPE_UNIT INCLUDE(348*,TAPE*,'CART','AFF=SMST','3420',
                       '3590','3590-1','3592-J1A','3592-E05','3490')

   FILTLIST PROD_DATA    INCLUDE(PROD.**)

   FILTLIST TEST_DATA    INCLUDE(TEST.**)

   SELECT
     WHEN ((&UNIT = &TAPE_UNIT) OR (&ANYVOL = 'REF=ST'))
       DO
         SELECT
             WHEN (&DSN = &PROD_DATA)
                 DO
                   SET &STORCLAS = 'SCCACHE'
                   EXIT
                 END
             WHEN (&DSN = &TEST_DATA)
                 DO
                   SET &STORCLAS = 'SCNCACHE'
                   EXIT
                 END
          END
       END
         OTHERWISE
           DO
             SET &STORCLAS = ''
               EXIT
           END
    END /* END SELECT */
 END  /* END PROC  */
```

15. Add two Management Classes. For the purpose of this exercise, the Management Classes is named MCTEST and MCPROD. The ETL Specialist panel for this definition as described in 4.4.3, “Creating Management Classes” on page 168.
    - MCTEST is assigned a Secondary Pool of 00. No dual copy.
    - MCPROD is assigned a Secondary Pool of 03. Logical volumes assigned to this Management Class are dual copied. Stacked volumes defined to Pool 03 are used.

> **Note:** For a TS7700 Virtualization Engine, the PTP Copy Control Mode values have no effect.

16. Repeat the process to create the Library Manager construct for Management Class MCTEST. Specify a Secondary Pool value of 0 for no dual copy.

17. Create the DFSMS Management Class construct through ISMF (see Figure F-18 through Figure F-20).

```
                         MANAGEMENT CLASS DEFINE               Page 1 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCPROD

To DEFINE Management Class, Specify:

  Description ==> PRODUCTION MANAGEMENT CLASS - DUAL COPY
              ==>

  Expiration Attributes
    Expire after Days Non-usage  . . NOLIMIT     (1 to 9999 or NOLIMIT)
    Expire after Date/Days . . . . . NOLIMIT     (0 to 9999, yyyy/mm/dd or
                                                 NOLIMIT)

  Retention Limit  . . . . . . . . . NOLIMIT     (0 to 9999 or NOLIMIT)



Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-18   Management Class Define ISMF panel 1 of 3 - OPM setup

```
                         MANAGEMENT CLASS DEFINE               Page 2 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCPROD

To DEFINE Management Class, Specify:

  Partial Release . . . . . . . . . N           (Y, C, YI, CI or N)

  Migration Attributes
    Primary Days Non-usage  . . . .             (0 to 9999 or blank)
    Level 1 Days Non-usage  . . . .             (0 to 9999, NOLIMIT or blank)
    Command or Auto Migrate . . . . NONE        (BOTH, COMMAND or NONE)


  GDG Management Attributes
    # GDG Elements on Primary . . .             (0 to 255 or blank)
    Rolled-off GDS Action . . . . .             (MIGRATE, EXPIRE or blank)

Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-19   Management Class Define ISMF panel 2 of 3 - OPM setup

```
                         MANAGEMENT CLASS DEFINE            Page 3 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCPROD

To DEFINE Management Class, Specify:
  Backup Attributes
    Backup Frequency  . . . . . . . .         (0 to 9999 or blank)
    Number of Backup Vers . . . . . .         (1 to 100 or blank)
      (Data Set Exists)
    Number of Backup Vers . . . . . .         (0 to 100 or blank)
      (Data Set Deleted)
    Retain days only Backup Ver . . .         (1 to 9999, NOLIMIT or blank)
      (Data Set Deleted)
    Retain days extra Backup Vers . .         (1 to 9999, NOLIMIT or blank)
    Admin or User command Backup  . . NONE    (BOTH, ADMIN or NONE)
    Auto Backup . . . . . . . . . . . N       (Y or N)
    Backup Copy Technique . . . . . . S       (P=Conc Preferred, R=Conc
                                              Required or S=Standard)
Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; Cancel to Exit.
```

Figure F-20 Management Class Define ISMF panel 3 of 3 - OPM setup

> **Note:** For subsequent Management Class panels, provide null values.

18. Repeat the process to create the DFSMS Management Class construct MCTEST.
19. Create or modify your Management Class DFSMS ACS routine similar to that shown in Example F-8.

Example F-8 Management Class routine - OPM setup

```
PROC 0 MGMTCLAS

/*******************************************************************/
/*            DEFINE TAPE DATASET FILTERING CRITERIA               */
/*******************************************************************/
    SELECT
       WHEN (&STORCLAS = 'SCCACHE')
           DO
              SET &MGMTCLAS = 'MCPROD'
              EXIT
           END
       WHEN (&STORCLAS = 'SCNCACHE')
           DO
              SET &MGMTCLAS = 'MCTEST'
              EXIT
           END
    END  /* END SELECT */
 END  /* END PROC */
```

20. Define a Data Class (see 4.4.5, “Creating Data Classes” on page 172). For the purpose of this exercise, the Data Class is named DCVE1. You can define it using the panel shown in Figure 4-26 on page 173.

21. Create the DFSMS Data Class construct through ISMF as shown in Figure F-21, Figure F-22, and Figure F-23.

```
                              DATA CLASS DEFINE                 Page 1 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVE1

To DEFINE Data Class, Specify:
  Description ==>
              ==>
  Recorg  . . . . . . . .               (KS, ES, RR, LS or blank)
  Recfm . . . . . . . . .               (any valid RECFM combination or blank)
  Lrecl . . . . . . . . .               (1 to 32761 or blank)
  Keylen  . . . . . . . .               (0 to 255 or blank)
  Keyoff  . . . . . . . .               (0 to 32760 or blank)
  Space Avgrec  . . . . .               (U, K, M or blank)
        Avg Value . . . .               (0 to 65535 or blank)
        Primary . . . . .               (0 to 999999 or blank)
        Secondary . . . .               (0 to 999999 or blank)
        Directory . . . .               (0 to 999999 or blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure F-21 Data Class Define ISMF panel 1 of 3 - OPM setup

```
                              DATA CLASS DEFINE                 Page 2 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVE1
To DEFINE Data Class, Specify:
  Retpd or Expdt  . . . . . .            (0 to 9999, YYYY/MM/DD or blank)
  Volume Count  . . . . . . .            (1 to 59 or blank)
    Add'l Volume Amount . . .            (P=Primary, S=Secondary or blank)
  Imbed . . . . . . . . . . .            (Y, N or blank)
  Replicate . . . . . . . . .            (Y, N or blank)
  CIsize Data . . . . . . . .            (1 to 32768 or blank)
  % Freespace CI  . . . . . .            (0 to 100 or blank)
              CA  . . . . . .            (0 to 100 or blank)
  Shareoptions Xregion  . . .            (1 to 4 or blank)
               Xsystem  . . .            (3, 4 or blank)
  Compaction  . . . . . . . . Y          (Y, N, T, G or blank)
  Media Interchange
    Media Type  . . . . . . . 2          (1, 2, 3, 4 or blank)
    Recording Technology  . . 36         (18, 36, 128, 256 or blank)
  Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
  Use HELP Command for Help; Use END Command to Save and Exit; Cancel to Exit.
```

Figure F-22 Data Class Define ISMF panel 2 of 3 - OPM setup

```
                               CLASS DEFINE                    Page 3 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVE1

To DEFINE Data Class, Specify:
  Data Set Name Type  . . . . . .      (EXT, HFS, LIB, PDS or blank)
    If Ext  . . . . . . . . . . .      (P=Preferred, R=Required or blank)
    Extended Addressability . . . N    (Y or N)
    Record Access Bias  . . . . .      (S=System, U=User or blank)
  Reuse . . . . . . . . . . . . . N    (Y or N)
  Initial Load  . . . . . . . . . R    (S=Speed, R=Recovery or blank)
  Spanned / Nonspanned  . . . . .      (S=Spanned, N=Nonspanned or blank)
  BWO . . . . . . . . . . . . . .      (TC=TYPECICS, TI=TYPEIMS, NO or blank)
  Log . . . . . . . . . . . . . .      (N=NONE, U=UNDO, A=ALL or blank)
  Logstream Id  . . . . . . . . .
  Space Constraint Relief . . . . N    (Y or N)
    Reduce Space Up To (%)  . . .      (0 to 99 or blank)
  Block Size Limit  . . . . . . .              (32760 to 2GB or blank)
Use ENTER to perform Verification; Use UP Command to View previous Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit
```

Figure F-23 Data Class Define ISMF panel 3 of 3 - OPM setup

22. Create or modify your Data Class DFSMS ACS routine as shown in Example F-9.

Example F-9 Data Class routine - OPM setup

```
PROC 0 DATACLAS

 /*----------------------------------------------------------------*/
 /* DEFINE VALID TAPE UNIT ESOTERICS                               */
 /*----------------------------------------------------------------*/
     FILTLIST TAPE_UNIT INCLUDE(348*,TAPE*,'CART','AFF=SMST','3420',
                               '3590','3590-1','3592-J1A','3592-E05','3490')

     FILTLIST VE_DATA    INCLUDE(PROD.**,TEST.**)

     SELECT
        WHEN ((&UNIT = &TAPE_UNIT) OR (&ANYVOL = 'REF=ST'))
          DO
            SELECT
               WHEN (&DATACLAS NE '')              /* Allows users to specify
*/
                 DO                                /* data class for tape
*/
                   SET &DATACLAS = &DATACLAS
                   EXIT
                 END
               WHEN (&DSN = &VE_DATA)
                 DO
                   SET &DATACLAS = 'DCVE1'
                   EXIT
                 END
            END
          END
            OTHERWISE
              DO
                SET &DATACLAS = ''
                EXIT
              END
     END /* END SELECT */
END  /* END PROC */
```

23. If the following DFSMS ACS routines have been changed, translate them through ISMF:

- Data class
- Storage class
- Management class
- Storage group

24. Validate the DFSMS SCDS through ISMF.

25. Test the ACS routines through ISMF.

26. Activate the DFSMS SCDS.

G

JES3 implementation

This appendix describes JES3 IBM TS3500 Tape Library support with DFSMS. The primary purpose of this support is to maintain JES3 resource allocation and share tape allocations. For detailed information, see *z/OS JES3 Initialization and Tuning Reference*, SA22-7550.

DFSMS has support that provides JES3 allocation with the appropriate information to select an IBM TS3500 Tape Library device. This is done by referencing device strings with a common name among systems in a JES3 complex.

Implementation overview

The following steps are necessary to set up an IBM TS3500 Tape Library in a JES3 environment:

1. Define library device groups (LDG). Prepare the naming conventions in advance. Clarify all the names for the library device groups you need.
2. Include the esoteric names from step1 in HCD and activate the new EDT.
3. Update the JES3 INISH deck:
   a. Define all devices in the IBM TS3500 Tape Library through DEVICE statements.
   b. Set JES3 device names through the SETNAME statement.
   c. Define which device names are subsets of other device names through the HWSNAME statement.

All IBM TS3500 Tape Library units can be shared between processors in a JES3 complex. They must also be shared among systems in the same SMSplex.

Note: Tape drives in the IBM TS3500 Tape Library cannot be used by JES3 dynamic support programs (DSPs).

Define all devices in the libraries through DEVICE statements. All IBM TS3500 Tape Library tape drives within a complex should be either JES3-managed or non-JES3-managed. Do not mix managed and nominated devices. Mixing might prevent the non-managed devices from being used for new data set allocations and reduce device eligibility for existing data sets. Allocation failures or delays in job setup result.

Neither JES3 or DFSMS verifies that a complete and accurate set of initialization statements is defined to the system. Incomplete or inaccurate IBM TS3500 Tape Library definitions can result in jobs failing to be allocated.

Library device groups

Library device groups (LDGs) isolate the IBM TS3500 Tape Library drives from other tape drives in the complex. They allow JES3 main device scheduler (MDS) allocation to select an appropriate set of library-resident tape drives.

The DFSMS JES3 support requires LDGs to be defined to JES3 for SETNAME groups and HWSNAME names in the JES3 initialization statements. During converter/interpreter (C/I) processing for a job, the LDG names are passed to JES3 by DFSMS for use by MDS in selecting library tape drives for the job. Unlike a JES2 environment, a JES3 operating environment requires the specification of esoteric unit names for the devices in a library. These unit names are used in the required JES3 initialization statements.

Important: Even if the LDG definitions are defined as esoterics in the HCD, they are not used in the JCL. There is no need for any UNIT= Parameter in JES3 JCL for libraries. The allocation goes through the ACS routines. Coding a UNIT= Parameter might cause problems.

The only need for coding the LDG definition in HCD as an esoteric name is the HWSNAME definitions in the JES3 INISH deck.

Each device in a library must have exactly four special esoteric names associated with it. These are:

- The *complex-wide name* is always LDGW3495. It allows you to address every device and device type in every library.
- The *library-specific name* is an 8 character string composed of LDG prefixing the 5-digit library identification number. It allows you to address every device and device type in that specific library.
- The *complex-wide device type*, shown in Table G-1, defines the used device types. It contains a prefix of LDG and a device type identifier. It allows you to address a specific device type in every tape library.

Table G-1 Library device groups - complex-wide device type specifications

Device type	Complex-wide device type definition
3490E	LDG3490E
3592-J1A	LDG359J
3592-E05	LDG359K
3592-E05 with Encryption enabled	LDG359L

- A *library-specific device type name*, an 8 character string, starts with a different prefix for each device type, followed by the 5-digit library device number. See Table G-2.

Table G-2 Library device groups - library-specific device types

Device type	Library-specific device type	Content
3490E	LDE + library number	All 3490E in lib *xx*
3592-J1A	LDJ + library number	All 3592 Model J1A in lib *xx*
3592-E05	LDK + library number	All 3592 Model E05 in lib *xx*
3592-E05 with Encryption enabled	LDL + library number	All 3592 Model E05 in lib *xx* which are Encryption-enabled

It also allows you to address a specific device type in a specific tape library. In a Multi Cluster Grid environment installed in two physical libraries, there is still only one library-specific device name. The library ID of the Composite Library is used.

Updating the JES3 INISH deck

To allow JES3 to allocate the appropriate device, you must code some definitions:

- Device statements
- Setname statements
- High water mark setup statements

These statements are described in detail in the following sections.

Device statement: Defining I/O devices for IBM TS3500 tape libraries

Use the DEVICE format to define a device so that JES3 can use it. A device statement (see Figure G-1) must be defined for each string of IBM TS3500 Tape Library drives in the complex. XTYPE specifies a 1- to 8 character name, given by the user. There is no default or specific naming convention for this statement. This name is used in other JES3 initialization statements to group the devices together for some JES3 processes (for example, allocation). Therefore, it is necessary that all the devices with the same XTYPE belong to:

- The same library
- The same device type

The letters CA in the XTYPE definition tell us that this is a CARTRIDGE device.

```
*/  Devices 3592-J1A and 3592-E05 in Libary 1 ............................/*

DEVICE,XTYPE=(LB13592J,CA),XUNIT=(1000,*ALL,,OFF),numdev=4
DEVICE,XTYPE=(LB13592K,CA),XUNIT=(1104,*ALL,,OFF),numdev=4
```

Figure G-1 DEVICE statement sample

> **Note:** TS3500 Tape Library tape drives cannot be used as support units by JES3 DSPs. Therefore, do not specify DTYPE, JUNIT, or JNAME parameters on the DEVICE statements. No check is made during initialization to prevent TS3500 Tape Library drives from being defined as support units, and no check is made to prevent the drives from being allocated to a DSP if they are defined. Any attempt to call a tape DSP by requesting a TS3500 Tape Library fails because the DSP is unable to allocate a TS3500 Tape Library drive.

SETNAME statement

The SETNAME statement is used for proper allocation in a JES3 environment. For tape devices, it tells JES3 which tape device belongs to which library. This is done by specifying the relationships between the XTYPE values (coded in the DEVICE Statement) and the LDG names (see Figure G-2). A SETNAME statement must be defined for each unique XTYPE in the device statements.

The rules for this SETNAME statement are:

- Each SETNAME statement has one entry from each LDG category.
- The complex-wide library name must be included in all statements.
- A library-specific name must be included for XTYPEs in the referenced library.
- The complex device type name must be included for all XTYPEs of the corresponding device type in the complex.
- A library-specific device type name must be included for the XTYPE associated with the devices in the library.

```
SETNAME,XTYPE=LB1359K,
        NAMES=(LDGW3495,LDGF4001,LDG359K,LDKF4001)
              Complex   Library   Complex   Library
               Wide      Specific  Wide      Specific
               Library   Library   Device    Device
               Name      Name      Name      Name
```

Figure G-2 SETNAME rules

Note: Do not specify esoteric and generic unit names such as 3480, 3480X, 3490, SYS3480R, and SYS348XR. Also, never use esoteric names such as TAPE and CART.

High watermark setup names

Use the HWSNAME statement to define which device names are subsets of other device names. You need to specify all applicable varieties. Therefore, HWSNAME coding is not easily understood. To clarify the subject, we introduce only the rules and coding and discuss the details in the configuration examples.

```
HWSNAME, TYPE=(groupname,{altname})
```

Note the following explanation:

groupname Specifies a device type valid for a high watermark setup.

altname Specifies a list of valid user-supplied or IBM-supplied device names. These are alternate units to be used in device selection.

Consider the following example:

```
HWSNAME,TYPE=(LDGW3495,LDGF4001,LDGF4006,LDG359J,LDG359K,
      LDJF4001,LDKF4001,LDKF4006)
```

The rules for the LDG HWSNAME statement are:

- The complex-wide library name, LDGW3495, must include all other LDG names as alternates.
- The library-specific name must include all LDG names for the corresponding library as alternates. When all tape devices of a type in the complex are within a single IBM 3494 Tape Library, the complex device type name must also be included as an alternate name.
- The complex device type name must include all library-specific device type names. When all devices of one type in the complex are within a single IBM TS3500 Tape Library, the complex device type name is equivalent to that library name. In this case, the library name should also be specified as an alternate.
- The library-specific device type name must be included. Alternate names can be specified as follows:
  - When all drives within the IBM TS3500 Tape Library have the same device type, the library-specific device type name is equivalent to the library name. In this case, the library-specific name should be specified as an alternate.
  - When these are the only drives of this type in the complex, the complex device type name is equivalent to the library-specific device type name.

Make sure that all valid alternate names are specified.

H

DEVSERV QLIB command

This appendix lists the syntax and parameter explanations, which are copied from the cover letter of the PTFs for APAR OA07505.

Tip: Use DEVSERV QLIB,? to get the SYNTAX of the command.

The DEVSER QLIB command can be used to:

- Request a list of tape library subsystems that are defined to the host. Libraries are listed by serial number (library-id).
- Request a list of devices within a library. Devices are listed by device number and the library port for each device is displayed.
- Validate the connection status of devices in a library. For example, devices that are connected to the host.
- Delete an improperly defined library control block in preparation for an IODF activate.
- Issue a diagnostic state save order to a library when requested by the IBM Service Center. For the state save you also need to apply the PTF for APAR OA09599.
- When using the DEVSERV QLIB command to display the subsystems (port-ids) and drives associated with the specified library-id, if the library-id specified is for a composite library, the command will now display the distributed library IDs associated with the composite library.

Important: Do not use this state save command for just testing purposes. It will impact the performance of your VTS/ATL, because it consumes time to take the dump in the hardware.

Example H-1 DEVSERV QLIB Command (APAR OA07505)

```
DOCUMENTATION:
        This new function APAR adds support to the DEVSERV command for
        a new Query Library option.

        Use the Query Library option of the DEVSERV command to:
```



```
* Request a list of tape library subsystems that are defined to
  the host. Libraries are listed by serial number (library-id).
* Request a list of devices within a library. Devices are
  listed by device number and the library port for each
  device is displayed.
* Validate the connection status of devices in a library. For
  example, devices that are connected to the host.
* Delete an improperly defined library control block in
  preparation for an IODF activate.
* Issue a diagnostic state save order to a library when
  requested by the IBM Service Center.

Query Library can be abbreviated QLIB or QL and supports the
following parameters:

DS QL,LIST(,filter)
DS QL,LISTALL(,filter)
DS QL,libid(,filter)
DS QL,dddd,SS

Parameters:

LIST-     Indicates that QLIB should display a list of the
          ACTIVE library-ids (ACTIVE is the default).
          You can optionally generate a list of INACTIVE
          library-ids or QUEUE'd library orders.
          LIST uses the sub-parameters ACTIVE, INACTIVE,
          and QUEUE.

LISTALL-  Produces a detailed list of all libraries,
          including the devices and port-ids within each
          library. LISTALL uses the sub-parameters ACTIVE
          and INACTIVE (ACTIVE is the default).

libid-    List information for the library with serial number
          'libid'. The 'libid' parameter uses sub-parameters
          ACTIVE, INACTIVE, VALIDATE, QUEUE and DELETE.
          ACTIVE is the default.

dddd-     Indicates that the request is either for the library
          that contains device dddd, or is for the device dddd
          itself. A sub-parameter is required when dddd is
          specified. dddd uses the sub-parameter SS.

?-        Causes QLIB to display the command syntax.

Sub-Parameters:

ACTIVE-   Displays information about the library configuration
          that is currently in use by the system.

INACTIVE- Displays information about the library configuration
          that will become active following the next IODF
          activate. The INACTIVE configuration is similar to
```

```
          ACTIVE, but may contain additional devices or
          libraries.

VALIDATE- Displays the INACTIVE configuration. However,
          before the configuration is displayed, I/O is issued
          to each device in the configuration to validate the
          devices connectivity to the host.

DELETE-   Indicates that QLIB should delete the INACTIVE control
          blocks for library LIBID, but not affect the existing
          ACTIVE library definition. The DELETE command is used
          to remove incorrectly defined library control blocks
          so they can be rebuilt. DEVSERV DELETE provides an
          alternative to the method described in information
          APAR II09065, which requires two IODF activates.

          The DEVSERV QLIB method is as follows:

          1. Use QLIB DELETE to delete all of the devices
             from the incorrect control blocks.
          2. Use QLIB LIST to display that the INACTIVE
             control blocks have been deleted.
          3. Use ACTIVATE IODF to redefine the devices.
          4. Use QLIB LIST to display that the ACTIVE control
             blocks are properly defined.

          Note: the steps above assume that library devices
                are HCD defined with LIBID and PORTID. Using
                LIBID and PORTID enables the activate in step
                3 (above) to build library control blocks. If
                LIBID and PORTID are not defined, then the
                following alternate method must be used:

          1.Use QLIB DELETE to delete all of the devices from
            the incorrect control blocks.
          2.Attempt to vary ONLINE each device in the library.
            Each VARY should fail with message:

            IEA437I TAPE LIBRARY DEVICE(dddd), ACTIVATE IODF
                    IS REQUIRED

          3.Each VARY attempt in the previous step should add a
            device to the INACTIVE configuration. Use QLIB LIST
            to list the INACTIVE configuration and verify that
            devices are configured correctly. If there are
            configuration errors, correct them and begin at
            step 1.
          4.Use ACTIVATE IODF to rebuild the ACTIVE configuration.
             This step replaces the currently ACTIVE
            configuration with the INACTIVE configuration.
            This step also rebuilds the allocation EDT's.
          5.Use QLIB LIST to display that the ACTIVE control
            blocks are properly defined.

QUEUE-    Lists the library orders that are waiting to be
```

```
              completed. Such orders include:

              MOUNT,DEMOUNT,EJECT and AUDIT

              When an order completes, the library notifies the
              host and the order is removed from the queue. The
              QL display can list orders for all libraries, or can
              be limited to a single library.

   SS-        Indicates that QLIB should issue a diagnostic state
              save to the library containing device dddd. This
              This command is intended to be used at the request of
              IBM Support Center. For example, SS can be used to
              diagnose a hardware error that results in a mount
              failure message. Automated Operator code can extract
              the failing device number from the failure message,
              then insert the device in a QLIB SS command.
Examples-

   DS QL,LIST
   IEE459I 13.59.01 DEVSERV QLIB 478
   The following are defined in the   ACTIVE configuration:
    10382  15393

   DS QL,10382
   IEE459I 13.59.09 DEVSERV QLIB 481
   The following are defined in the   ACTIVE configuration:
   LIBID  PORTID   DEVICES
   10382  04       0940  0941  0942  0943  0944  0945  0946  0947
                   0948  0949  094A  094B  094C  094D  094E  094F
          03       09A0  09A1  09A2  09A3  09A4  09A5  09A6  09A7
                   09A8  09A9  09AA  09AB  09AC  09AD  09AE  09AF
          02       09D0  09D1  09D2  09D3  09D4  09D5  09D6  09D7
                   09D8  09D9  09DA  09DB  09DC  09DD  09DE  09DF
          01       F990  F991  F992  F993  F994  F995  F996  F997
                   F998  F999  F99A  F99B  F99C  F99D  F99E  F99F
  DISTRIBUTED LIBID(S)
  AAAAA BBBBB

   DS QL,10382,DELETE
   *04 REPLY 'YES' TO DELETE THE INACTIVE CONFIGURATION FOR
       LIBRARY 10382, ANY OTHER REPLY TO QUIT.
   IEF196I Reply 'YES' to delete the INACTIVE configuration for
       library 10382, any other reply to quit.
   R 4,YES
   IEE459I 14.01.19 DEVSERV QLIB 490
   Inactive configuration for library 10382 successfully deleted

 COMMENTS:
   CROSS REFERENCE-MODULE/MACRO NAMES TO APARS
    IGUDSL01  OA07505

   CROSS REFERENCE-APARS TO MODULE/MACRO NAMES
    OA07505   IGUDSL01
```

```
     THE FOLLOWING MODULES AND/OR MACROS ARE AFFECTED BY THIS PTF:
MODULES
      IGUDSL01
LISTEND
 */.
++ HOLD(UA17546) SYS FMID(HDZ11G0) REASON(DOC) DATE(05098)
   COMMENT
    (This new function APAR adds support to the DEVSERV command for
     a new Query Library option.

     Use the Query Library option of the DEVSERV command to:

     * Request a list of tape library subsystems that are defined to
       the host. Libraries are listed by serial number (library-id).
     * Request a list of devices within a library. Devices are
       listed by device number and the library port for each
       device is displayed.
     * Validate the connection status of devices in a library. For
       example, devices that are connected to the host.
     * Delete an improperly defined library control block in
       preparation for an IODF activate.
     * Issue a diagnostic state save order to a library when
       requested by the IBM Service Center.

     Query Library can be abbreviated QLIB or QL and supports the
     following parameters:

     DS QL,LIST(,filter)
     DS QL,LISTALL(,filter)
     DS QL,libid(,filter)
     DS QL,dddd,SS

     Parameters:

     LIST-     Indicates that QLIB should display a list of the
               ACTIVE library-ids (ACTIVE is the default).
               You can optionally generate a list of INACTIVE
               library-ids or QUEUE'd library orders.
               LIST uses the sub-parameters ACTIVE, INACTIVE,
               and QUEUE.

     LISTALL-  Produces a detailed list of all libraries,
               including the devices and port-ids within each
               library. LISTALL uses the sub-parameters ACTIVE
               and INACTIVE (ACTIVE is the default).

     libid-    List information for the library with serial number
               'libid'. The 'libid' parameter uses sub-parameters
               ACTIVE, INACTIVE, VALIDATE, QUEUE and DELETE.
               ACTIVE is the default.

     dddd-     Indicates that the request is either for the library
               that contains device dddd, or is for the device dddd
               itself. A sub-parameter is required when dddd is
               specified. dddd uses the sub-parameter SS.
```

```
?-       Causes QLIB to display the command syntax.


Sub-Parameters:

ACTIVE-   Displays information about the library configuration
          that is currently in use by the system.

INACTIVE- Displays information about the library configuration
          that will become active following the next IODF
          activate. The INACTIVE configuration is similar to
          ACTIVE, but may contain additional devices or
          libraries.

VALIDATE- Displays the INACTIVE configuration. However, before
          before the configuration is displayed, I/O is issued
          to each device in the configuration to validate the
          devices connectivity to the host.

DELETE-   Indicates that QLIB should delete the INACTIVE control
          blocks for library LIBID, but not affect the existing
          ACTIVE library definition. The DELETE command is used
          to remove incorrectly defined library control blocks
          so they can be rebuilt. DEVSERV DELETE provides an
          alternative to the method described in information
          APAR II09065, which requires two IODF activates.

          The DEVSERV QLIB method is as follows:

          1. Use QLIB DELETE to delete all of the devices
             from the incorrect control blocks.
          2. Use QLIB LIST to display that the INACTIVE
             control blocks have been deleted.
          3. Use ACTIVATE IODF to redefine the devices.
          4. Use QLIB LIST to display that the ACTIVE control
             blocks are properly defined.

          Note: the steps above assume that library devices
                are HCD defined with LIBID and PORTID. Using
                LIBID and PORTID enables the activate in step
                3 (above) to build library control blocks. If
                LIBID and PORTID are not defined, then the
                following alternate method must be used:

          1.Use QLIB DELETE to delete all of the devices from
            the incorrect control blocks.
          2.Attempt to vary ONLINE each device in the library.
            Each VARY should fail with message:

            IEA437I TAPE LIBRARY DEVICE(dddd), ACTIVATE IODF
                    IS REQUIRED

          3.Each VARY attempt in the previous step should add a
            device to the INACTIVE configuration. Use QLIB LIST
```

```
      to list the INACTIVE configuration and verify that
      devices are configured correctly. If there are
      configuration errors, correct them and begin at
      step 1.
    4.Use ACTIVATE IODF to rebuild the ACTIVE configuration
      This step replaces the currently ACTIVE
      configuration with the INACTIVE configuration.
      This step also rebuilds the allocation EDT's.
    5.Use QLIB LIST to display that the ACTIVE control
      blocks are properly defined.

QUEUE-    Lists the library orders that are waiting to be
          completed. Such orders include:

          MOUNT,DEMOUNT,EJECT and AUDIT

          When an order completes, the library notifies the
          host and the order is removed from the queue. The
          QL display can list orders for all libraries, or can
          be limited to a single library.

SS-       Indicates that QLIB should issue a diagnostic state
          save to the library containing device dddd. This
          This command is intended to be used at the request of
          IBM Support Center. For example, SS can be used to
          diagnose a hardware error that results in a mount
          failure message. Automated Operator code can extract
          the failing device number from the failure message,
          then insert the device in a QLIB SS command.

Examples-

DS QL,LIST
IEE459I 13.59.01 DEVSERV QLIB 478
The following are defined in the   ACTIVE configuration:
 10382  15393

DS QL,10382
IEE459I 13.59.09 DEVSERV QLIB 481
The following are defined in the   ACTIVE configuration:
LIBID  PORTID   DEVICES
10382  04       0940  0941  0942  0943  0944  0945  0946  0947
                0948  0949  094A  094B  094C  094D  094E  094F
       03       09A0  09A1  09A2  09A3  09A4  09A5  09A6  09A7
                09A8  09A9  09AA  09AB  09AC  09AD  09AE  09AF
       02       09D0  09D1  09D2  09D3  09D4  09D5  09D6  09D7
                09D8  09D9  09DA  09DB  09DC  09DD  09DE  09DF
       01       F990  F991  F992  F993  F994  F995  F996  F997
                F998  F999  F99A  F99B  F99C  F99D  F99E  F99F
 DISTRIBUTED LIBID(S)
 AAAAA BBBBB

DS QL,10382,DELETE
*04 REPLY 'YES' TO DELETE THE INACTIVE CONFIGURATION FOR
    LIBRARY 10382, ANY OTHER REPLY TO QUIT.
```

```
IEF196I Reply 'YES' to delete the INACTIVE configuration for
    library 10382, any other reply to quit.
R 4,YES
IEE459I 14.01.19 DEVSERV QLIB 490
Inactive configuration for library 10382 successfully deleted).
```

I

TS7700 Statistics record format

This appendix is based on the *Virtualization Engine TS7700 Series Statistical Data Format* White Paper and provides the record layout of the statistics data that the TS7700 provides.

> **Note:** The content of this appendix is current at the time of writing. We recommend that you refer to the latest version of the White Paper to check for the latest updates to the statistics records. You can find the White Paper using the following links:
>
> `http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100829`

Record types and sizes

Table I-1 shows the record types, a description of the record type, the length of each record, and the number of records produced.

Table I-1 Record types and sizes

Data type (hexadecimal)	Description	Record length (in bytes)	Number of records
x01	Vnode Virtual Device Point in Time Record	96 + (number-of-virtual-devices-in-this-cluster x 32)	1 per Vnode
x02	Vnode Adapter Point in Time Record	384	1 per Vnode
x20	Vnode Virtual Device Historical Record	192	1 per Vnode
x21	Vnode Adapter Historical Record	384	1 per Vnode
x10	Hnode HSM Point in Time Record	96 + (number-of-physical-libraries-attached-to-this-cluster x 1568)	1 per Hnode
x11	Hnode Grid Point-in-Time Record	96 + (number-of-clusters-in-the-grid x 129)	1 per Hnode
x30	Hnode HSM Historical Record	1152	1 per Hnode
x31	Reserved		
x32	Hnode Library Historical Record	8256	1 per physical library attached to a cluster
x33	Hnode Grid Historical Record	96 + (number-of-clusters-in-the-grid x 256)	1 per Hnode

Vnode Virtual Device Point-In-Time record

This Vnode Point-In-Time (PIT) record, as listed in Table I-2, has the following nested structure:

- Header
- General Information Container
- Virtual Device Container
  - Virtual Device 0 info
  - Virtual Device 1 info
  - ...
  - Virtual Device xxx info

Table I-2 Vnode Virtual Device PIT record

Bytes	Name	Description	When Data is Sampled/ Updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record the value is set to x01 indicating this is a Vnode Virtual Device Point-In-Time record.	
4	Node ID	This 1 byte hexadecimal field indicates the Vnode ID which this interval's data represents. Valid values are x00 – x0F.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID which this Vnode is a part of. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the Epoch (00:00:00 UTC, 01 January 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to “3957”.	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to “V06”.	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine (VE). The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification. Fix in a decimal form. For example the code level of 8.0.0.104 would be represented in the 8 bytes as: x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the grid (Composite) library.	
40-63	Reserved	All bytes set to x00	

Bytes	Name	Description	When Data is Sampled/ Updated
Vnode PIT General Information Container **Bytes 64-95** **The fields below provide information concerning the configuration of the Vnode whose data is being reported.**			
64	Node State	This 1 byte hexadecimal field indicates the state of the Vnode at the end of this interval. Possible values are: **Value** – **Description** x00 – Offline x01 – Online x02 – Going Offline x03 – Going Online x04-xFE – Reserved xFF – Node is not working at all	
65-66	Configured maximum throughput	This 2 byte hexadecimal field contains the maximum throughput for this Vnode. The value is expressed in MB/s (1 MB = 1024x1024 bytes). This field is set to x0000 if there is no restriction for the maximum throughput. This is the value at the end of the interval.	Set to x0000 for first release.
67-68	Installed Virtual Devices	This 2 byte hexadecimal field indicates the number of installed virtual devices in this Vnode. This field can be used to determine how many Virtual Device containers will be attached to this record. This is the value at the end of the interval.	
69-95	Reserved	All bytes set to x00	
Virtual Device Container Bytes 96 and up (Number of Installed Virtual Devices x 32 bytes/device). For example, if there are 128 installed virtual devices there will be 128 sets of data with 32 bytes each which totals 4096 bytes. There is a maximum of 256 virtual devices per Vnode. This next segment of the record contains one set of data for each virtual device installed in the Vnode as defined in bytes 67-68 above. Each set of data contains 32 bytes. The following fields define the 32 bytes of data and are numbered starting with byte 0. The first virtual device's data can be found in bytes 96-127, the second device's data can be found in bytes 128-159, and so forth.			

<table>
<tr><th>Bytes</th><th>Name</th><th>Description</th><th>When Data is Sampled/ Updated</th></tr>
<tr><td>0-9</td><td>Mounted Volume</td><td>This 10 byte EBCDIC field contains the volser of the logical volume, if any, that is currently mounted in the virtual device or was most recently mounted in the device.

For device mount states that indicate a mount in progress or is mounted (x01, x02, x03, x06 – see byte 11) this field will contain the volser of the volume that is in the process of being mounted.

For device mount states that indicate a device is not mounted or in the process of being mounted (x00, x04, x05 – see byte 11) this field will contain the volser of the last successfully mounted volume, if any, or will be filled with EBCDIC blanks.

This field is left justified and padded with EBCDIC blanks.

This is the value at the end of the interval.</td><td>Updated whenever the logical volume changes in the virtual device.</td></tr>
<tr><td>10</td><td>Cluster Access Point</td><td>This 1 byte hexadecimal field indicates the Cluster ID which is sourcing or has most recently sourced the logical volume for a mount. In the case of the most recently sourced cluster, the current device mount state will indicate “Device unloaded, failed or cancelled”. (see byte 11 below) Valid values for this field are x00 – x07.

The subsystem has the ability to access a logical volume in any cluster from any cluster. Where a logical volume is sourced from is based on a set of criteria including volume consistency, access policies, and so forth.

This is the value at the end of the interval.</td><td>Updated for each virtual mount.</td></tr>
<tr><td>11</td><td>Device Mount State</td><td>This 1 byte hexadecimal field indicates the mount state of the virtual device. Valid values are as follows:
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x00</td><td>Device unloaded. This indicates the previous mount completed successfully and the virtual device is now unloaded. This is also reported when the virtual device has not mounted any volume yet and is not in the process of performing a mount.</td></tr>
<tr><td>x01</td><td>Mount request being processed (DE owed, PLF received)</td></tr>
<tr><td>x02</td><td>Mount accepted and DE given (initial status received from Library Manager)</td></tr>
<tr><td>x03</td><td>Mount in progress (device mount command received from Library Manager)</td></tr>
<tr><td>x04</td><td>Mount failed. This indicates the previous mount failed (error reported by device) and the virtual device is unloaded.</td></tr>
<tr><td>x05</td><td>Mount cancelled. This indicates the previous mount request was cancelled before the mount occurred. (Demount received before Unload.) The virtual device is now unloaded.</td></tr>
<tr><td>x06</td><td>Mounted</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
</table></td><td>Updated whenever the virtual device mount state changes</td></tr>
</table>

Bytes	Name	Description	When Data is Sampled/ Updated
		For a non-configured or non-installed device, this field will indicate Device unloaded. This is the value at the end of the interval.	

Bytes	Name	Mask Value	Description	When Data is Sampled/ Updated
12-13	Device Flags		This 2 byte field contains 16 single bit flags related to the current state of the virtual device. Several bits can be set to 1 at the same time. The bit mask values are defined as follows:	Updated whenever the virtual device mount state changes.
		Mask Value	**Description**	
		x8000	Device is ready. Provides information as to the mount state of the virtual drive. If the bit is set, then a virtual volume is loaded into the drive and is ready for host I/O activity.	
		x4000	Device is write-protected. This flag indicates the logical write protect state of the virtual drive/virtual volume. This prevents any modifications to a virtual volume loaded into this drive.	
		x2000	Write data is in the buffer. This flag indicates that the control unit has write data for this virtual volume held in the controller's memory, and not yet committed to the virtual device. This is normal for streaming write data to the device.	
		x1000	Write mode - The last IO to the device was a write operation. Indicates that the virtual device is in write mode. The device enters write mode upon the first write I/O operation, and stays in this mode until a non-write command (position change, read, etc.) is encountered.	
		x0800	Volume is in the LEOP (Logical End of Partition). Indicates that the current block position of the virtual tape is in the LEOP region of the virtual volume, and the host should be performing end of volume processing on the virtual volume.	
		x0400	Volume is at BOT (Beginning of Tape). Indicates the virtual volume loaded is currently at the beginning of the virtual volume. This would be the tape position after a mount request, or if the host repositioned the volume back to block 0.	
		x0200	Device is fenced. The control unit presents unit check status with associated sense data indicating ERA 47, Volume Fenced, if a condition has occurred which has resulted in the loss of volume integrity due to lost positioning or assignment. The control unit prevents further access to the tape volume by generating deferred unit checks with associated sense data indicating ERA 47, Volume Fenced, for all eligible commands until the condition is reset or until the cartridge is unloaded. The original condition which subsequently caused ERA 47, Volume Fenced, to be presented has already been indicated by a previous unit check and associated sense data.	
		x0100	Device is in standalone mount mode. This indicates that the virtual volume in this virtual drive was loaded as part of an operator request through the Management Interface, and not through a connected host. Standalone mounts are required to be able to IPL the operating system image from a virtual tape device, without the operating system having to have tape mount capabilities.	
			All other bits are reserved.	
			For a non-configured or non-installed device, this field will be set to x0000. This is the value at the end of the interval.	

Bytes	Name	Description	When Data is Sampled/ Updated
14	Buffer CCR Conditions	This 1 byte hexadecimal field contains the number of times the virtual device had to CCR (Channel Command Retry) the channel due to a buffer condition during this interval. If more than 255 buffer CCRs occur during the interval this field will indicate 255 buffers CCRs (xFF). This field can be used in conjunction with the Device Flags field, Write mode bit, to determine if the device was in write mode (buffer full condition) or read mode (buffer empty condition) during the interval. This value is reset to 0 at the beginning of the interval. **Note:** This field is not accurate at this time.	Count is incremented whenever a Buffer CCR condition occurs.
15-16	Channel Bytes Read	This 2 byte hexadecimal field contains the number of bytes transferred from the virtual device to the channel for a read from this device. The value is reported in increments of 100K bytes (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This value is reset to 0 at the beginning of the interval	Count is incremented for each read of a logical volume's data.
17-18	Channel Bytes Written	This 2 byte hexadecimal field contains the number of bytes transferred from the channel to the virtual device for a write to this device. The value is reported in increments of 100K bytes (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This value is reset to 0 at the beginning of the interval.	Count is incremented for each write of data to a logical volume.
19-31	Reserved	All bytes set to x00.	

Vnode Adapter Point-In-Time record

This Vnode PIT record, as listed in Table I-3, has the following nested structure:

- Header
- Adapter Container
  - Adapter 0 general information
    - Port 0 information
    - Port 1 information
  - Adapter 1 general information
    - Port 0 information
    - Port 1 information
  - Adapter 2 general information
    - Port 0 information
    - Port 1 information
  - Adapter 3 general information
    - Port 0 information
    - Port 1 information

Table I-3 Vnode Adapter PIT record

Bytes	Name	Description	When Data is Sampled/Updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record the value is set to x02 indicating this is a Vnode Adapter Point-In-Time record.	
4	Node ID	This 1 byte hexadecimal field indicates the Vnode ID which this interval's data represents. Valid values are x00 – x0F.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID which this Vnode is a part of. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the Epoch (00:00:00 UTC, 01 January 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine (VE). The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example the code level of 8.0.0.104 would be represented in the 8 bytes as: x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the grid (Composite) library.	
40-63	Reserved	All bytes set to x00	

<table>
<tr><th>Bytes</th><th>Name</th><th>Description</th><th>When Data is Sampled/Updated</th></tr>
<tr><td colspan="3">Adapter Container

Bytes 64-383 (4 sets of data x 80 bytes/set = 320 bytes)

This next set of bytes contains information for up to 4 host bus adapters (HBA). For each adapter there is data for up to 2 ports on the adapter. Each set of data contains 80 bytes. The following fields define the 80 bytes of data and are numbered starting with byte 0. The first set of data is for adapter 0, the second for adapter 1, and so forth. The first HBA’s data can be found in bytes 64-143, the second adapter’s data can be found in bytes 144-223, and so forth.</td><td></td></tr>
<tr><td>0</td><td>Adapter Type</td><td>This 1 byte hexadecimal field identifies the type of Host Bus Adapter (HBA). The possible values are:
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x00</td><td>No adapter installed</td></tr>
<tr><td>x01-x08</td><td>Reserved</td></tr>
<tr><td>x09</td><td>FICON – 1 Port (Arctic Circle)</td></tr>
<tr><td>x0A</td><td>FICON – 2 Port (Yukon)</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
</table>
This is the value at the end of the interval.</td><td></td></tr>
<tr><td>1</td><td>Adapter State</td><td>This 1 byte hexadecimal field identifies the current state of the adapter. The possible values are:
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x00</td><td>No adapter installed</td></tr>
<tr><td>x01</td><td>The adapter is online</td></tr>
<tr><td>x02</td><td>The adapter is offline</td></tr>
<tr><td>x03</td><td>The adapter is not working at all</td></tr>
<tr><td>x04</td><td>The adapter is reloading itself</td></tr>
<tr><td>x05</td><td>The adapter is in a Check1 condition</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
</table>
This is the value at the end of the interval.</td><td></td></tr>
</table>

Bytes	Name	Description	When Data is Sampled/Updated
2	HBA Drawer	This 1 byte hexadecimal field indicates which drawer the HBA is located in. The possible values are: **Value** – **Description** x00 – The HBA is in the left drawer when looking at the drawers from the back which is the side the cables plug into. x01 – The HBA is in the right drawer when looking at the drawers from the back which is the side the cables plug into. All other values are reserved. This is the value at the end of the interval.	
3	HBA Slot Number	This 1 byte hexadecimal field indicates the physical slot number of the HBA within its drawer. This is the value at the end of the interval.	
4-15	Reserved	All bytes set to x00.	
Adapter-Port Container Relative bytes 16 – 79 (2 sets of data x 32 bytes/set = 64 bytes) This next set of bytes contains information for up to 2 ports on the HBA. Each set of data contains 32 bytes. The following fields define the 32 bytes and are numbered starting with byte 0. The first port’s data can be found in relative bytes 16-47 and the second port’s data can be found in relative bytes 48-79.			
0-1	RCD Interface ID	This 2 byte hexadecimal field contains the internal ID of the HBA port that is reported in the RCD (Read Configuration Data), General NEQ (Node Element Qualifier) record. This is the value at the end of the interval.	
2-4	Reserved	All bytes set to x00	
5-8	Bytes Read by the Channel	This 4 byte hexadecimal field contains the number of bytes transferred to the channel from this HBA port as part of a read operation. This is the value after the data has been decompressed by the HBA. The value is reported in increments of 4K bytes (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure I-1, Bytes Read by the Channel is indicated by the **A** label. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data read.

Bytes	Name	Description	When Data is Sampled/Updated
9-12	Bytes Written by the Channel	This 4 byte hexadecimal field contains the number of bytes transferred from the channel to this HBA port as part of a write operation. This is the value before the effect of the HBA compression. The value is reported in increments of 4K bytes (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure I-1, Bytes Written by the Channel is indicated by the **B** label. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written.
13-16	Bytes Read from the Virtual Devices	This 4 byte hexadecimal field contains the number of bytes transferred from virtual devices to this HBA port as part of a read operation. The value is for data previously compressed by the HBA. The value is reported in increments of 4K bytes (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure I-1, Bytes Read from Virtual Devices is indicated by the **C** label. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data read.
17-20	Bytes Written to Virtual Devices	This 4 byte hexadecimal field contains the number of bytes transferred to virtual devices from this HBA port as part of a write operation. The value is for data compressed by the HBA. The value is reported in increments of 4K bytes (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure I-1, Bytes Written to Virtual Devices is indicated by the **D** label. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written.
21-31	Reserved	All bytes set to x00.	

Figure I-1 shows the data flow and the labels used in Table I-3 on page 639.

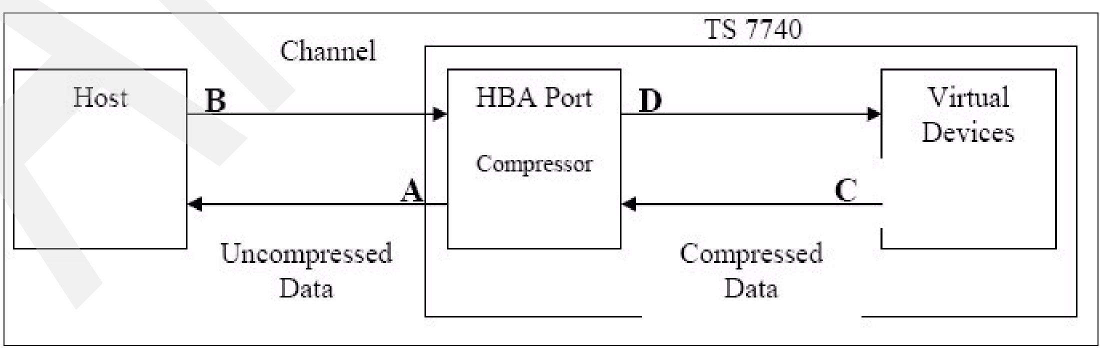

Figure I-1   Data Flow and labels for Table I-3

Vnode Virtual Device Historical record

This Vnode Historical record, as listed in Table I-4, has the following nested structure:

- Header Virtual
- Device Container

Table I-4 Vnode Virtual Device Historical record

Bytes	Name	Description	When Data is Sampled/Updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record the value is set to x20 indicating this is a Vnode Virtual Device Historical record.	
4	Node ID	This 1 byte hexadecimal field indicates the Vnode ID which this interval's data represents. Valid values are x00 – x0F.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID which this Vnode is a part of. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the Epoch (00:00:00 UTC, 01 January 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to 3957.	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to V06.	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine (VE). The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example the code level of 8.0.0.104 would be represented in the 8 bytes as: x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the grid (Composite) library.	

Bytes	Name	Description	When Data is Sampled/Updated
40-63	Reserved	All bytes set to x00	
Vnode Virtual Device Container **Bytes 64-191**			
64-65	Installed Virtual Devices	This 2 byte hexadecimal field indicates the number of installed virtual devices in this Vnode. This is the value at the end of the interval.	
66-69	Virtual Device Type	This 4 byte EBCDIC field indicates the device type emulated by the virtual devices. Initially this is set to 3490. This field is left justified and padded with blanks.	
70-72	Virtual Device Model	This 3 byte EBCDIC field indicates the device model emulated by the virtual devices. Initially this is set to C2A. This field is left justified and padded with blanks. This is the value at the end of the interval.	
73-80	Channel Blocks Written 1-2048 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this Vnode that had a size of between 1 and 2048 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written that fits the size.
81-88	Channel Blocks Written 2049-4096 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this Vnode that had a size of between 2049 and 4096 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written that fits the size.
89-96	Channel Blocks Written 4097-8192 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this Vnode that had a size of between 4097 and 8192 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written that fits the size.
97-104	Channel Blocks Written 8193-16384 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this Vnode that had a size of between 8193 and 16384 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written that fits the size.
105-112	Channel Blocks Written 16385-32768 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this Vnode that had a size of between 16385 and 32768 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written that fits the size.
113-120	Channel Blocks Written 32769-65536 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this Vnode that had a size of between 32769 and 65536 bytes inclusive for this interval. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written that fits the size.

Bytes	Name	Description	When Data is Sampled/Updated
121-128	Channel Blocks Written above 65536 byte range	This 8 byte hexadecimal field indicates the number of channel blocks written to all the virtual devices in this Vnode that had a size of 65537 or higher for this interval. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written that fits the size.
129-130	Configured Maximum Throughput	This 2 byte hexadecimal field contains the current maximum throughput for this Vnode. The value is expressed in MB/s (1 MB = 1024x1024 bytes). This field is set to x0000 if there is no restriction for the maximum throughput. This is the value at the end of the interval.	This is set to x0000 for the first release.
131-132	Minimum Virtual Devices Mounted	This 2 byte hexadecimal field indicates the minimum number of virtual devices that were mounted at the same time over the interval.	The count of mounted virtual devices is sampled every 15 seconds. The min/max/avg is updated over the interval.
133-134	Maximum Virtual Devices Mounted	This 2 byte hexadecimal field indicates the maximum number of virtual devices that were mounted at the same time over the interval.	The count of mounted virtual devices is sampled every 15 seconds. The min/max/avg is updated over the interval.
135-136	Average Virtual Devices Mounted	This 2 byte hexadecimal field indicates the average number of virtual devices that were mounted at the same time over the interval. The average is calculated by recording the number of mounted devices on a periodic basis then averaging it over the interval.	The count of mounted virtual devices is sampled every 15 seconds. The min/max/avg is updated over the interval.
137-191	Reserved	All bytes set to x00.	

Vnode Adapter Historical record

This Vnode Historical record, as listed in Table I-5 on page 646, has the following nested structure:

- Header
- Adapter 0 Container
  - Adapter-Port 0 Container
  - Adapter-Port 1 Container
- Adapter 1 Container
  - Adapter-Port 0 Container
  - Adapter-Port 1 Container
- Adapter 2 Container
  - Adapter-Port 0 Container
  - Adapter-Port 1 Container
- Adapter 3 Container
  - Adapter-Port 0 Container
  - Adapter-Port 1 Container

Table I-5 Vnode Adapter Historical record

Bytes	Name	Description	When Data is Sampled/Updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record the value is set to x21 indicating this is a Vnode Adapter Historical record.	
4	Node ID	This 1 byte hexadecimal field indicates the Vnode ID which this interval's data represents. Valid values are x00 – x0F.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID which this Vnode is a part of. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the Epoch (00:00:00 UTC, 01 January 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to 3957.	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to V06.	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine (VE). The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example the code level of 8.0.0.104 would be represented in the 8 bytes as: x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the grid (Composite) library.	
40-63	Reserved	All bytes set to x00	

<table>
<tr><th>Bytes</th><th>Name</th><th colspan="2">Description</th><th>When Data is Sampled/Updated</th></tr>
<tr><td colspan="4">Vnode Adapter Container

Bytes 64-383 (4 sets x 80 bytes/set = 320 bytes)

This next set of bytes contains information for up to 4 host bus adapters (HBA). For each adapter there is data for up to 2 ports. Each set of data contains 80 bytes. The following fields define the bytes of data and are numbered starting with byte 0. The first set of data is for the adapter 0, the second for adapter 1, and so forth. The first adapter's data can be found in bytes 64-143, the second adapter's data can be found in bytes 144-223, and so forth.</td><td></td></tr>
<tr><td rowspan="9">0</td><td rowspan="9">Adapter Type</td><td colspan="2">This 1 byte hexadecimal field identifies the type of Host Bus Adapter (HBA) this data is for. The possible values are:</td><td rowspan="9"></td></tr>
<tr><td>Value</td><td>Description</td></tr>
<tr><td>x00</td><td>No adapter installed</td></tr>
<tr><td>x02-x08</td><td>Reserved</td></tr>
<tr><td>x09</td><td>FICON – 1 Port (Arctic Circle)</td></tr>
<tr><td>x0A</td><td>FICON – 2 Port (Yukon)</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
<tr><td colspan="2">This is the value at the end of the interval.</td></tr>
<tr style="display:none"></tr>
<tr><td rowspan="10">1</td><td rowspan="10">Adapter State</td><td colspan="2">This 1 byte hexadecimal field identifies the current state of the adapter. The possible values are:</td><td rowspan="10"></td></tr>
<tr><td>Value</td><td>Description</td></tr>
<tr><td>x00</td><td>No adapter installed</td></tr>
<tr><td>x01</td><td>The adapter is online</td></tr>
<tr><td>x02</td><td>The adapter is offline</td></tr>
<tr><td>x03</td><td>The adapter is not working at all</td></tr>
<tr><td>x04</td><td>The adapter is reloading itself</td></tr>
<tr><td>x05</td><td>The adapter is in a Check1 condition</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
<tr><td colspan="2">This is the value at the end of the interval.</td></tr>
</table>

Bytes	Name	Description	When Data is Sampled/Updated
2	HBA Drawer	This 1 byte hexadecimal field indicates which drawer the HBA is located in. The possible values are: **Value** – **Description** x00 – The HBA is in the left drawer when looking at the drawers from the back which is the side the cables plug into. x01 – The HBA is in the right drawer when looking at the drawers from the back which is the side the cables plug into. All other values are reserved. This is the value at the end of the interval.	
3	HBS Slot Number	This 1 byte hexadecimal field indicates the physical slot number of the HBA within its drawer. This is the value at the end of the interval.	
4-15	Reserved	All bytes set to x00.	
Vnode Adapter-Port Container Relative bytes 16-79 (2 sets x 32 bytes/set = 64 bytes) This next set of bytes contains information for up to 2 ports on the HBA. Each set of data contains 32 bytes. The following fields define the 32 bytes and are numbered starting with byte 0. The first port’s data can be found in relative bytes 16-47 and the second port’s data can be found in relative bytes 48-79.			
0-1	RCD Interface ID	This 2 byte hexadecimal field contains the internal ID of the HBA port that is reported in the RCD (Read Configuration Data), General NEQ (Node Element Qualifier) record. This is the value at the end of the interval.	
2	Maximum Data Rate	This 1 byte hexadecimal field indicates the maximum data rate the FICON port is capable of at the end of the interval. The value is reported in Giga-bits (Gb) per second. For ESCON ports this value is reported as x00 because ESCON does not have a variable data rate.	
3	Actual Data Rate	This 1 byte hexadecimal field indicates the actual data rate of the FICON port at the end of the interval. The value is reported in Giga-bits (Gb) per second. A value of x00 in this field indicates that auto-negotiate is enabled. For ESCON ports this value is reported as x00 because ESCON does not have a variable data rate.	

Bytes	Name	Description	When Data is Sampled/Updated
4-7	Bytes Read by the Channel	This 4 byte hexadecimal field contains the number of bytes transferred to the channel from this HBA port as part of a read operation. This is the value after the data has been decompressed by the HBA. The value is reported in increments of 4K bytes (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure I-2, Bytes Read by the Channel is indicated by the **A** label. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data read.
8-11	Bytes Written by the Channel	This 4 byte hexadecimal field contains the number of bytes transferred from the channel to this HBA port as part of a write operation. This is the value before the effect of the HBA compression. The value is reported in increments of 4K bytes (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure I-2, Bytes Written by the Channel is indicated by the **B** label. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written
12-15	Bytes Read from Virtual Devices	This 4 byte hexadecimal field contains the number of bytes transferred from virtual devices to this HBA port as part of a read operation. The value is for data previously compressed by the HBA. The value is reported in increments of 4K bytes (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure I-2, Bytes Read from Virtual Devices is indicated by the **C** label. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data read.
16-19	Bytes Written to Virtual Devices	This 4 byte hexadecimal field contains the number of bytes transferred to virtual devices from this HBA port as part of a write operation. The value is for data compressed by the HBA. The value is reported in increments of 4K bytes (4 x 1024). Any residual data will cause the value to be rounded up to the next higher value. In Figure I-2, Bytes Written to Virtual Devices is indicated by the **D** label. This value is reset to 0 at the beginning of the interval.	Count is incremented for each block of data written.
20	Selective Resets	This 1 byte hexadecimal field indicates the number of selective resets this port received during the interval. This field is set to xFF when the number of selective resets is greater than 255 for the interval. This value is reset to 0 at the beginning of the interval.	Count is incremented for each selective reset.

Bytes	Name	Description	When Data is Sampled/Updated
21	System Resets	This 1 byte hexadecimal field indicates the number of system resets this port received during the interval. This field is set to xFF when the number of system resets is greater than 255 for the interval. This value is reset to 0 at the beginning of the interval.	Count is incremented for each system reset.
22-31	Reserved	All bytes set to x00.	

Figure I-2 describes the data flow and shows the labels used in Table I-5 on page 646.

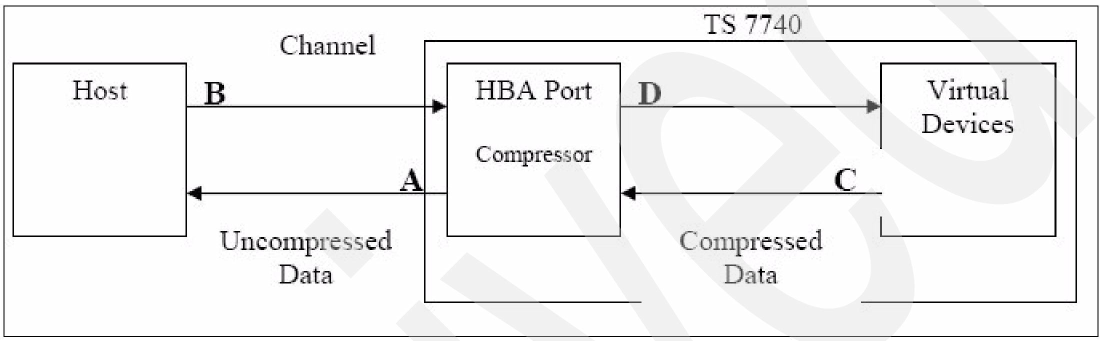

Figure I-2 Data Flow and labels for Table I-5

Hnode HSM Point-In-Time record

This Hnode Point-In-Time record, as listed in Table I-6, has the following nested structure:

- Header
- HSM Container
  - HSM-Library 0 Container
    - HSM-Library-Physical Device 0 Container
    - HSM-Library-Physical Device 1 Container
    - ...
    - HSM-Library-Physical Device 31 Container
  - HSM-Library 1 Container (If installed)
    - Same sub-containers as HSM-Library 0 Container
  - HSM-Library 2 Container (If installed)
    - Same sub-containers as HSM-Library 0 Container
  - HSM-Library 3 Container (If installed)
    - Same sub-containers as HSM-Library 0 Container

Table I-6 Hnode HSM Point-In-Time record

Bytes	Name	Description	When Data is Sampled/Updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record the value is set to x10 indicating this is an Hnode HSM Point-In-Time record.	
4	Node ID	This 1 byte hexadecimal field indicates the Hnode ID which this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID which this Hnode is a part of. Valid values are x00 – x07.	
6-7	Interval Duration	This 1 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 1 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the Epoch (00:00:00 UTC, January 1, 1970)	
12-15	Machine Type	This 1 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to 3957.	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to V06.	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine (VE). The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example the code level of 8.0.0.104 would be represented in the 8 bytes as: x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the grid (Composite) library.	
40-63	Reserved	All bytes set to x00	
HSM Container Bytes 64 and up. This container provides information concerning HSM related items. The total length depends upon the number of physical libraries attached to this Hnode.			

Bytes	Name	Description	When Data is Sampled/Updated
64-65	Recalls in Queue	This 2 byte hexadecimal field contains the current number of queued recall operations at the end of the interval.	This count is updated as recalls added or removed from the queue.
66-67	Pre-migrates in Queue	This 2 byte hexadecimal field contains the current number of queued pre-migrate operations at the end of the interval.	This count is updated with the current value every 30 seconds.
68-71	Write Overrun Throttle	This 4 byte hexadecimal field contains the write overrun throttling value over the interval. This is for throttling where write overrun was the predominant reason for throttling. The value is reported in thousandths of a second. This is the value at the end of the interval.	This count is updated with the current value every 30 seconds.
72-75	Copy Throttle	This 4 byte hexadecimal field contains the copy throttling value over the interval. This is for throttling where copy was the predominant reason for throttling. The value is reported in thousandths of a second. This is the value at the end of the interval.	This count is updated with the current value every 30 seconds.
76	Number of Physical Libraries	This 1 byte hexadecimal field indicates the number of physical libraries this Hnode is attached to. This field can be used to determine how many HSM-Library containers will be attached to this record. This is the value at the end of the interval.	
77-95	Reserved	All bytes set to x00.	
HSM-Library Container Bytes 96-and up (Number of Physical Libraries x 1568 bytes/set) For example, if there is just one physical library attached to this node there will be 1 set of data with 1568 bytes. There is a maximum of 4 physical libraries attached to a single cluster. This next segment of the record contains one set of data for each physical library attached to this Hnode as defined in byte 76 above. Each set of data contains 1568 bytes. The data for the first library is found in bytes 96-1663; the second library's data (if the library exists) is found in bytes 1664-3231, and so forth. These fields contain information concerning the underlying automation.			
0-15	Library Sequence Number	This 16 byte EBCDIC field indicates the Library Sequence Number of the underlying automation. This field is left justified and padded with EBCDIC blanks. This is the value at the end of the interval.	
16-31	Reserved	All bytes set to x00.	

<table>
<tr><th>Bytes</th><th>Name</th><th>Description</th><th>When Data is Sampled/Updated</th></tr>
<tr><td colspan="3">HSM-Library-Physical Device Container

Relative bytes 32-415 (32 sets of device data x 48 bytes/set = 1536 bytes)

For a system with just 1 class of physical device (homogeneous), the data for the devices is found in the first 16 sets of data. For a system with 2 classes of physical devices (heterogeneous), the data for one class of drives is found in the first 16 sets of data, and the data for the second class of devices is found in the second 16 sets of data.

There are 32 sets of data per possible library, one for each of 32 possible physical devices per physical library. The data for the first device of the first device class is in relative bytes 0-47; the second device’s data of the first device class is in bytes 48-95, and so forth. For the second device class, the first device’s data is found in relative bytes 768-815, the second device’s data is in bytes 816-863, and so forth.</td><td></td></tr>
<tr><td>0</td><td>Device Class ID</td><td>This 1 byte hexadecimal field contains this device’s device class identifier.
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x00</td><td>No device installed</td></tr>
<tr><td>x11</td><td>3590 Model B1</td></tr>
<tr><td>x13</td><td>3590 Model E1A</td></tr>
<tr><td>x14</td><td>3590 Model H1A</td></tr>
<tr><td>x20</td><td>3592 Model J1A. This also includes a 3592-E05 that is emulating a 3592-J1A device.</td></tr>
<tr><td>x22</td><td>3592 Model E05. This is for a 3592-E05 that is behaving as a 3590-E05.</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
</table>
This is the value at the end of the interval.</td><td></td></tr>
<tr><td>1-10</td><td>Physical Volume</td><td>This 10 byte EBCDIC field contains the volser of the volume that is loaded in this device at the end of the interval, if any. This field is left justified and padded with EBCDIC blanks. This field is filled with EBCDIC blanks when there is no volume in the device.</td><td>Volser is updated whenever there is a change.</td></tr>
<tr><td>11</td><td>Volume Pool</td><td>This 1 byte hexadecimal field indicates the pool associated with the volume, if any, loaded in the device at the end of the interval. This field is set to x00 when there is not a volume loaded in the device. Values 1-32 are also valid in this field.</td><td>Pool is updated whenever there is a change.</td></tr>
<tr><td>12</td><td>Device State</td><td>This 1 byte hexadecimal field indicates the state of the device at the end of the interval. Possible values are:
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x00</td><td>Device is online</td></tr>
<tr><td>x01</td><td>Device is offline</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
</table></td><td>State is updated whenever there is a change.</td></tr>
</table>

<table>
<tr><th>Bytes</th><th>Name</th><th colspan="2">Description</th><th>When Data is Sampled/Updated</th></tr>
<tr><td rowspan="10">13</td><td rowspan="10">Device Role</td><td colspan="2">This 1 byte hexadecimal field contains the role of the device at the end of the interval. Possible values are:</td><td rowspan="10">Role is updated whenever there is a change.</td></tr>
<tr><td>Value</td><td>Description</td></tr>
<tr><td>x00</td><td>Idle</td></tr>
<tr><td>x01</td><td>Recalling</td></tr>
<tr><td>x02</td><td>Migrating</td></tr>
<tr><td>x03</td><td>Reclaiming</td></tr>
<tr><td>x04</td><td>Importing</td></tr>
<tr><td>x05</td><td>Exporting</td></tr>
<tr><td>x06</td><td>Data Security Erase</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
<tr><td>14-23</td><td>Logical Volume</td><td colspan="2">This 10 byte EBCDIC field contains the volser of the logical volume, if any, that is being processed by the physical device at the end of the interval. This field is left justified and padded with EBCDIC blanks. This field contains all EBCDIC blanks when there is not a logical volume being processed at the end of the interval.</td><td>Volser updated whenever there is a change.</td></tr>
<tr><td>24-27</td><td>Data Read From Device</td><td colspan="2">This 4 byte hexadecimal field indicates the number of bytes transferred from the physical device during this interval. The value is reported in increments of 100K bytes (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value.

This value is reset to 0 at the beginning of the interval.</td><td>Count is incremented after every read.</td></tr>
<tr><td>28-31</td><td>Data Written to Device</td><td colspan="2">This 4 byte hexadecimal field indicates the number of bytes transferred to the physical device during this interval. The value is reported in increments of 100K bytes (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value.

This value is reset to 0 at the beginning of the interval.</td><td>Count is incremented after each write.</td></tr>
<tr><td>32-47</td><td>Reserved</td><td colspan="2">All bytes set to x00.</td><td></td></tr>
</table>

Hnode Grid Point-In-Time record

This Hnode Point-In-Time record, as listed in Table I-7, has the following nested structure:

- Header
- Grid Container
  - Grid-Cluster 0 Container
  - Grid-Cluster 1 Container (If installed)
  - Grid-Cluster 2 Container (If installed)
  - Grid-Cluster 3 Container (If installed)
  - Grid-Cluster 4 Container (If installed)
  - Grid-Cluster 5 Container (If installed)
  - Grid-Cluster 6 Container (If installed)
  - Grid-Cluster 7 Container (If installed)

Table I-7 Hnode Grid Point-In-Time (PIT) record

Bytes	Name	Description	When Data is Sampled/Updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record the value is set to x11 indicating this is an Hnode Grid Point-In-Time record.	
4	Node ID	This 1 byte hexadecimal field indicates the Hnode ID which this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID which this Hnode is a part of. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the Epoch (00:00:00 UTC, January 1, 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine (VE). The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example the code level of 8.0.0.104 would be represented in the 8 bytes as: x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the grid (Composite) library.	
40-63	Reserved	All bytes set to x00.	

Bytes	Name	Description	When Data is Sampled/Updated
Grid Container Bytes 64 and up This container provides information concerning the grid aspects of this cluster. The length of the message depends on the number of clusters in the grid. There are 32 fixed bytes plus 32 bytes for each cluster in the grid.			
64-67	Immediate Copy Queue	This 4 byte hexadecimal field indicates the number of logical volumes in the immediate copy queue targeted for this cluster at the end of the interval.	Count is updated every 5 minutes.
68-71	Deferred Copy Queue	This 4 byte hexadecimal field indicates the number of logical volumes in the deferred copy queue targeted for this cluster at the end of the interval.	Count is updated every 5 minutes.
72-75	Active Copies	This 4 byte hexadecimal field indicates the number of active copies that are targeting this cluster at the end of the interval.	Count is updated every 5 minutes.
76-77	Network CRC Errors	This 2 byte hexadecimal field indicates the number of CRC (Cyclic Redundancy Check) errors that were detected by this cluster during this interval. This field is reset to 0 at the beginning of the interval.	Count is incremented for each error.
78	Link Usage	This 1 byte hexadecimal field indicates the percentage of the estimated network throughput was used for grid functions during this interval. This field is reset to 0 at the beginning of the interval.	Value is calculated at the end of the interval.
79	Number of clusters	This 1 byte hexadecimal field indicates the number of clusters in the grid. This field can be used to determine how many grid-cluster containers will be attached to this record. There is a maximum of 8 clusters. This is the value at the end of the interval.	
80-95	Reserved	All bytes set to x00	
Grid-Cluster Container Bytes 96 and up (Number of clusters x 129 bytes/set) For example, if there are 3 clusters in the grid there will be 3 sets of data. There is a maximum of 8 clusters in a grid. This next segment of the record contains one set of data for each cluster in the grid as defined in byte 79 above. Each set of data contains 129 bytes. The data for the first cluster (Cluster 0) can be found in bytes 96-223; the second cluster's (Cluster 1) data can be found in bytes 224-351, and so forth.			

<table>
<tr><th>Bytes</th><th>Name</th><th colspan="2">Description</th><th>When Data is Sampled/Updated</th></tr>
<tr><td rowspan="7">0</td><td rowspan="7">Cluster Link State</td><td colspan="2">This 1 byte hexadecimal field indicates the state of the link between this cluster and the others in the grid at the end of the interval.</td><td rowspan="7">Updated whenever the link state changes.</td></tr>
<tr><td>Value</td><td>Description</td></tr>
<tr><td>x00</td><td>All links are fully operational. There are no detected error conditions between the clusters in the grid. For a grid with only a single cluster, this is the value always reported in this field.</td></tr>
<tr><td>x01</td><td>Degraded. One or more of the logical communication paths between the clusters is not operational.</td></tr>
<tr><td>x02</td><td>Failed. All of the logical communication paths between this cluster and the other clusters in the grid are not working.</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
<tr><td colspan="2">This field represents the overall link state as determined by the cluster. The actual connection can be comprised of several physical connections and underlying infrastructure which is transparent to the system.</td></tr>
<tr><td>1-32</td><td>Data transferred into a cluster’s cache from other clusters</td><td colspan="2">This field contains eight 4 byte fields indicating the number of bytes transferred to this cluster’s cache from other clusters as part of a remote file access during this interval.

The first 4 byte field indicates the number of bytes transferred into this cluster’s cache from Cluster 0. The second 4 byte field indicates the number of bytes transferred from Cluster 1, and so forth. A value of x00000000 is indicated for a cluster’s own value and for non-existent clusters.

The value is reported in increments of 100K bytes (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value.

This value is reset to 0 at the beginning of the interval.</td><td>Count is incremented every time a block of data is written to this cluster.</td></tr>
<tr><td>33-64</td><td>Data transferred from a cluster’s cache to other clusters</td><td colspan="2">This field contains eight 4 byte fields indicating the number of bytes transferred from this cluster’s cache from other clusters as part of a copy and remote file access during this interval.

The first 4 byte field indicates the number of bytes transferred from this cluster’s cache to Cluster 0. The second 4 byte field indicates the number of bytes transferred to Cluster 1, and so forth. A value of x00000000 is indicated for a cluster’s own value and for non-existent clusters.

The value is reported in increments of 100K bytes (100 x 1024). Any residual data will cause the value to be rounded up to the next higher value.

This value is reset to 0 at the beginning of the interval.</td><td>Count is incremented every time a block of data is read from this cluster</td></tr>
<tr><td>65-128</td><td>Reserved</td><td colspan="2">All bytes set to x00.</td><td></td></tr>
</table>

Hnode HSM Historical record

This Hnode HSM Historical record, as listed in Table I-8 on page 658, has the following nested structure:

- Header
- HSM Container

The HSM Container has the following structure:

- HSM-Disaster Recovery Container
- HSM-Cache Container
  - HSM-Cache-Partition 0 Container
    - HSM-Cache-Partition-Preference Group 0 Container
    - HSM-Cache-Partition-Preference Group 1 Container
  - HSM-Cache-Partition 1 Container
    - HSM-Cache-Partition-Preference Group 0 Container
    - HSM-Cache-Partition-Preference Group 1 Container
  - ...
  - HSM-Cache-Partition 7 Container
    - HSM-Cache-Partition-Preference Group 0 Container
    - HSM-Cache-Partition-Preference Group 1 Container

Table I-8 Hnode HSM Historical record

Bytes	Name	Description	When Data is Sampled/Updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record the value is set to x30 indicating this is an Hnode HSM Historical record.	
4	Node ID	This 1 byte hexadecimal field indicates the Hnode ID which this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID which this Hnode is a part of. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the Epoch (00:00:00 UTC, January 1, 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	

Bytes	Name	Description	When Data is Sampled/Updated
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine (VE). The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example the code level of 8.0.0.104 would be represented in the 8 bytes as: x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the grid (Composite) library.	
40-63	Reserved	All bytes set to x00	
HSM Container Bytes 64-1151 This record contains information pertinent to the Hydra Storage Manager.			
HSM – Disaster Recovery Container Bytes 64 – 95 This set of 32 bytes contains information concerning HSM Disaster Recovery (DR).			
64-73	Disaster Recovery Volser	This 10 byte EBCDIC field contains the volser of the physical volume that contains the latest subsystem backup information. This field is set to all EBCDIC blanks if a disaster recovery volume does not exist. This field is left justified and padded with EBCDIC blanks. This is the value at the end of the interval.	Volser is updated whenever it changes.
74-95	Reserved	All bytes set to x00.	
HSM – Cache Container Bytes 96 – 1151 This set of bytes contains information concerning the Tape Volume Cache (TVC) and this Hnode.			

<table>
<tr><th>Bytes</th><th>Name</th><th colspan="2">Description</th><th>When Data is Sampled/Updated</th></tr>
<tr><td rowspan="5">96</td><td rowspan="5">Active Hnode</td><td colspan="2">This 1 byte hexadecimal field indicates if this Hnode was the active node in charge of Tape Volume Cache (TVC) management at the end of the interval being reported.</td><td rowspan="5">This value is set to x00 for the first release</td></tr>
<tr><td>Value</td><td>Description</td></tr>
<tr><td>x00</td><td>This is not the active Hnode</td></tr>
<tr><td>x01</td><td>This is the active Hnode</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
<tr><td>97-100</td><td>TVC Size</td><td colspan="2">This 4 byte hexadecimal field indicates the current size of the cluster’s TVC in increments of 1 GB (1024 x 1024 x 1000). A TVC that is 1.7 TB in size will be reported as x000006A4 (1700 decimal).

This field is filled in by the Hnode that is currently in charge of TVC management. The Hnode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active Hnode field.

This is the value at the end of the interval.</td><td></td></tr>
<tr><td>101</td><td>Percent Write Overrun Throttle</td><td colspan="2">This 1 byte hexadecimal field indicates the percentage of 30-second periods where there was at least one throttling value greater than zero and that write overrun was the predominant reason for throttling.

This field is filled in by the Hnode that is currently in charge of TVC management. The Hnode that is not in charge of the TVC management reports x00 in this field. Refer to the Active Hnode field.

This value is recomputed at the end of the interval based on the data from the interval.</td><td>The current throttle value is sampled every 30 seconds. The percentage is computed over the interval.</td></tr>
<tr><td>102-105</td><td>Average Write Overrun Throttle</td><td colspan="2">This 4 byte hexadecimal field indicates the average write overrun throttle value during the interval. The value presented is the average of the non-zero throttling values where write overrun was the predominant reason for throttling. The value is reported in thousandths of a second.

This field is filled in by the Hnode that is currently in charge of TVC management. The Hnode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active Hnode field.

This value is recomputed at the end of the interval based on the data from the interval.</td><td>The current throttle value is sampled every 30 seconds. The average is computed over the interval.</td></tr>
</table>

Bytes	Name	Description	When Data is Sampled/Updated
106	Percent Copy Throttle	This 1 byte hexadecimal field indicates the percentage of 30-second periods where there was at least one throttling value greater than zero and that copy was the predominant reason for throttling. This field is filled in by the Hnode that is currently in charge of TVC management. The Hnode that is not in charge of the TVC management reports x00 in this field. Refer to the Active Hnode field. This value is recomputed at the end of the interval based on the data from the interval.	The current throttle value is sampled every 30 seconds. The percentage is computed over the interval.
107-110	Average Copy Throttle	This 4 byte hexadecimal field indicates the average copy throttle value during the interval. The value presented is the average of the non-zero throttling values where copy was the predominant reason for throttling. The value is reported in thousandths of a second. This field is filled in by the Hnode that is currently in charge of TVC management. The Hnode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active Hnode field. This value is recomputed at the end of the interval based on the data from the interval.	The current throttle value is sampled every 30 seconds. The average is computed over the interval.
111-114	Average Overall Throttle	This 4 byte hexadecimal field indicates the average of all throttling values during the interval. The calculation includes samples for periods where throttling was both zero and non-zero. The value is reported in thousandths of a second. This field is filled in by the Hnode that is currently in charge of TVC management. The Hnode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active Hnode field. This value is recomputed at the end of the interval based on the data from the interval.	The current throttle values are sampled every 30 seconds. The average is computed over the interval.
115-127	Reserved	All bytes set to x00.	
HSM – Cache – Partition Container Bytes 128 – 1151 This next set of bytes contains information for up to 8 cache partitions for the cluster. Each set of data contains 128 bytes. The following fields define the 128 bytes and are numbered starting with byte 0. The first cache partition's data can be found in relative bytes 0-127 and the second cache partition's data can be found in relative bytes 128-255, and so forth. These fields are filled in by the Hnode that is currently in charge of TVC management. The Hnode that is not in charge of the TVC management reports x00 in these fields. Note: At this time, the TS7700 only supports a single cache partition.			

Bytes	Name	Description	When Data is Sampled/Updated
0-3	Partition Size	This 4 byte hexadecimal field indicates the amount of cache assigned to this partition. The value is reported in increments of 1 GB (1024 x 1024 x 1000). This field is filled in by the Hnode that is currently in charge of TVC management. The Hnode that is not in charge of the TVC management reports x00000000 in this field. Refer to the Active Hnode field. This is the value at the end of the interval.	The size is updated when it changes.
4-5	Fast Ready Mounts	This 2 byte hexadecimal field indicates the number of mount requests completed using the Fast Ready method during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The count is incremented when the mount-complete message is received.
6-9	Average Fast Ready Mount Time	This 4 byte hexadecimal field indicates the average time, in milliseconds, taken to complete Fast-Ready mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval's time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The time is incremented for each mount and averaged at the end of the interval.
10-11	Cache Hit Mounts	This 2 byte hexadecimal field indicates the number of mount requests completed that the data was resident in the Tape Volume Cache (TVC) during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager.	The count is incremented when the mount-complete message is received.
		This value is reset to 0 at the beginning of the interval.	
12-15	Average Cache Hit Mount Time	This 4 byte hexadecimal field indicates the average time, in milliseconds, taken to complete Cache Hit mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval's time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The time is incremented for each mount and averaged at the end of the interval.
16-17	Cache Miss Mounts	This 2 byte hexadecimal field indicates the number of mount requests completed that required recall from a stacked volume during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The count is incremented when the mount-complete message is received.

Bytes	Name	Description	When Data is Sampled/Updated
18-21	Average Cache Miss Mount Time	This 4 byte hexadecimal field indicates the average time, in milliseconds, taken to complete Cache Miss mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval's time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	The time is incremented for each mount and averaged at the end of the interval.
22-31	Reserved	All bytes set to x00.	
HSM – Cache – Partition – Preference Group Container Relative bytes 32 – 127 (2 x 48 bytes = 96 bytes) These bytes are relative to the HSM – Cache - Partition Container. This next set of bytes contains information for 2 preference groups for the cache partition. Each set of data contains 48 bytes. The following fields define the 48 bytes and are numbered starting with byte 0. The first preference group's (PG0) data can be found in bytes 0-47 and the second preference group's (PG1) data can be found in relative bytes 48-95. Because these fields are part of the HSM-Cache-Partition container, these fields are filled in by the Hnode that is currently in charge of TVC management. The Hnode that is not in charge of the TVC management reports x00 in these fields.			
0-3	Virtual Volumes in Cache	This 4 byte hexadecimal field contains the number of virtual volumes in the Tape Volume Cache (TVC) partition that are assigned to the preference group this data is for. This is the value at the end of the interval.	The count is updated every 10 minutes.
4-7	Data Resident in Cache	This 4 byte hexadecimal field contains the amount of data in the TVC partition whose volumes are assigned to the preference this data is for. The value is reported in increments of 1 MB (1024 x 1024). Any residual data will cause the value to be rounded up to the next higher value. This is the value at the end of the interval.	The count is updated every 10 minutes.
8-11	4 Hour Average Cache Age	This 4 byte hexadecimal field contains a 4 hour rolling average of cache age, in minutes, of the virtual volumes migrated out of the cache partition that were assigned to the preference group this data is for. Cache age is measured from when a volume is created or recalled into cache until it has been migrated from cache. Each volume's cache age is rounded up to the nearest minute. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The age is updated once an hour, on the hour.
12-15	Volumes Migrated Last 4 Hours	This 4 byte hexadecimal field contains the number of virtual volumes migrated from the cache partition over the past 4 hours that are assigned to the preference group this data is for. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The count is updated once an hour, on the hour.

Bytes	Name	Description	When Data is Sampled/Updated
16-19	48 Hour Average Cache Age	This 4 byte hexadecimal field contains a 48 hour rolling average of cache age, in minutes, of the virtual volumes migrated out of the cache partition that were assigned to the preference group this data is for. Cache age is measured from when a volume is created or recalled into cache until it has been migrated from cache. Each volume's cache age is rounded up to the nearest minute. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The age is updated once an hour, on the hour.
20-23	Volumes Migrated Last 48 Hours	This 4 byte hexadecimal field contains the number of virtual volumes migrated from the cache partition over the past 48 hours that are assigned to the preference group this data is for. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The count is updated once an hour, on the hour.
24-27	35 Day Average Cache Age	This 4 byte hexadecimal field contains a 35 day rolling average of cache age, in minutes, of the virtual volumes migrated out of the cache partition that were assigned to the preference group this data is for. Cache age is measured from when a volume is created or recalled into cache until it has been migrated from cache. Each volume's cache age is rounded up to the nearest minute. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The age is updated once an hour, on the hour.
28-31	Volumes Migrated Last 35 Days	This 4 byte hexadecimal field contains the number of virtual volumes migrated from the cache partition over the past 35 days that are assigned to the preference group this data is for. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.	The count is updated once an hour, on the hour.
32-47	Reserved	All bytes set to x00.	

Hnode Library Historical record

This Hnode Library Historical record, as listed in Table I-9 on page 665, has the following nested structure:

- Header
- Library Container

The Library Container has the following structure:

- Library Device Type 0 Usage Container
- Library Device Type 1 Usage Container
- Library Device Type 2 Usage Container
- Library Device Type 3 Usage Container
- Library-Pooling Container
  - Library-Pooling-CSP Media Type 0 Container
  - ...
  - Library-Pooling-CSP Media Type 7 Container

- Library-Pooling-GUP 1 Container
  - Library-Pooling-GUP-Media Type 0 Container
  - ...
  - Library-Pooling-GUP-Media Type 7 Container
  - Library-Pooling-GUP-Reclaim Container
  - Library-Pooling-GUP-Properties Container
- Library-Pooling-GUP 2 Container
  - Same sub-containers as GUP 1
- ...
- Library-Pooling-GUP 32 Container
  - Same sub-containers as GUP 1

Table I-9 Hnode Library Historical record

Bytes	Name	Description	When Data is Sample/Updated
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record the value is set to x32 indicating this is an Hnode Library Historical record.	
4	Node ID	This 1 byte hexadecimal field indicates the Hnode ID which this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID which this Hnode is a part of. Valid values are x00 – x07.	
6-7	Interval Duration	This 2 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the Epoch (00:00:00 UTC, 01 January 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	

Bytes	Name	Description	When Data is Sample/Updated
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine (VE). The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example the code level of 8.0.0.104 would be represented in the 8 bytes as: x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the grid (Composite) library.	
40-63	Reserved	All bytes set to x00	
Library Container Bytes 64-255 (192 bytes) This record contains information pertinent to the Library operations associated with this Hnode.			
64-69	Library Machine Type	This 6 character EBCDIC field contains the machine type of the underlying automation. The field is left justified and padded with EBCDIC blanks. Initially this field will be set to 3494 or 3584. This is the value at the end of the interval.	
70-72	Library Model Number	This 3 character EBCDIC field contains the model number of the underlying automation. The field is left justified and padded with EBCDIC blanks. Initially this field will be set to L10 when attached to a 3494 or L22 when attached to a 3584. This is the value at the end of the interval.	
73-75	Library Manufacturer	This 3 byte EBCDIC field contains an abbreviation of the manufacturer of the underlying automation. The field is left justified and padded with EBCDIC blanks. Initially this will be set to IBM. This is the value at the end of the interval.	
76-77	Library Plant of Manufacture	This 2 byte EBCDIC field contains the library plant of manufacture of the underlying automation. The field is left justified and padded with EBCDIC blanks. This is the value at the end of the interval.	
78-93	Library Sequence Number	This 16 byte EBCDIC field contains the sequence number of the underlying automation. The field is left justified and padded with EBCDIC blanks. This is the value at the end of the interval.	
94-127	Reserved	All bytes set to x00.	
Library – Tape Device Usage (TDU) Container Bytes 128-255 (4 sets of data x 32 bytes/set = 128 bytes) This container contains 4 sets of data. The 4 sets of data allow up to 4 device types/models to be attached to the Hnode. Each set of data contains 32 bytes. The following fields define these bytes and are numbered starting with 0. Data for the first device type is found in bytes 128-159, the data for the second device type is found in bytes 160-191, and so forth. Each Hnode reports this data from its perspective.			

<table>
<tr><th>Bytes</th><th>Name</th><th>Description</th><th>When Data is Sample/Updated</th></tr>
<tr><td>0</td><td>Device Class ID</td><td>This 1 byte hexadecimal field contains this device’s device class identifier.<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x00</td><td>No device installed</td></tr>
<tr><td>x11</td><td>3590 Model B1A</td></tr>
<tr><td>x13</td><td>3590 Model E1A</td></tr>
<tr><td>x14</td><td>3590 Model H1A</td></tr>
<tr><td>x20</td><td>3592 Model J1A. This also includes a 3592-E05 that is emulating a 3592-J1A device.</td></tr>
<tr><td>x22</td><td>3592 Model E05. This is for a 3592-E05 that is behaving as a 3590-E05.</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
</table>This is the value at the end of the interval.</td><td></td></tr>
<tr><td>1</td><td>Installed Physical Devices</td><td>This 1 byte hexadecimal field contains the number of physical devices, of the device class indicated, that are installed at the end of the interval.</td><td></td></tr>
<tr><td>2</td><td>Available Physical Devices</td><td>This 1 byte hexadecimal field contains the number of physical devices, of the device class indicated, that are available for use at the end of the interval.</td><td>Count is updated when the number of available physical devices changes.</td></tr>
<tr><td>3</td><td>Maximum Physical Devices Mounted</td><td>This 1 byte field contains the maximum number of physical devices, of the device class indicated, that were concurrently mounted during the interval.

his value is reset to 0 at the beginning of the interval.</td><td>Count is updated whenever the number of mounted physical devices changes.</td></tr>
<tr><td>4</td><td>Minimum Physical Devices Mounted</td><td>This 1 byte field contains the minimum number of physical devices, of the device class indicated, that were concurrently mounted during the interval.

This value is reset to 0 at the beginning of the interval.</td><td>Count is updated whenever the number of mounted physical devices changes.</td></tr>
<tr><td>5</td><td>Average Physical Devices Mounted</td><td>This 1 byte field contains the average number of physical devices, of the device class indicated, that were concurrently mounted during the interval. The average is calculated by recording the number of mounted devices on a periodic basis then averaging it over the interval.

This value is reset to 0 at the beginning of the interval.</td><td>The count of mounted physical devices is sampled every 15 seconds. The average is computed over the interval.</td></tr>
<tr><td>6-7</td><td>Maximum Physical Mount Time</td><td>This 2 byte field contains the maximum time, in seconds, that it took to complete the execution of a mount request for a physical device, of the device class indicated, over the interval. Mount time is accrued from the time the mount request is sent to the Library Manager (LM) until the mount complete is received from the LM. The mount time is accredited to the interval it was completed.

This value is reset to 0 at the beginning of the interval.</td><td>The time for each mount is examined when the mount completes.</td></tr>
</table>

Bytes	Name	Description	When Data is Sample/Updated
8-9	Minimum Physical Mount Time	This 2 byte field contains the minimum time, in seconds, that it took to complete the execution of a mount request for a physical device, of the device class indicated, over the interval. Mount time is accrued from the time the mount request is sent to the Library Manager (LM) until the mount complete is received from the LM. The mount time is accredited to the interval it was completed. This value is reset to 0 at the beginning of the interval.	The time for each mount is examined when the mount completes.
10-11	Average Physical Mount Time	This 2 byte field contains the average time, in seconds, that it took to complete the execution of a mount request for a physical device, of the device class indicated, over the interval. Mount time is accrued from the time the mount request is sent to the Library Manager (LM) until the mount complete is received from the LM. The mount time is accredited to the interval it was completed. This value is reset to 0 at the beginning of the interval.	The time for each mount is examined when the mount completes.
12-13	Physical Recall Mounts	This 2 byte hexadecimal field contains the number of physical mount requests completed by the library during the interval to satisfy recall mounts for the device class indicated. A mount is accredited to the interval the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	Count is incremented each time a recall mount is completed.
14-15	Physical Pre-Migrate Mounts	This 2 byte hexadecimal field contains the number of physical mount requests completed by the library during the interval to satisfy pre-migrate mounts for the device class indicated. A mount is accredited to the interval the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	Count is incremented each time a pre-migrate mount is completed.
16-17	Physical Reclaim Mounts	This 2 byte hexadecimal field contains the number of physical mount requests completed by the library during the interval to satisfy reclaim mounts for the device class indicated. A mount is accredited to the interval the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	Count is incremented each time a reclaim mount is completed.
18-19	Physical Security Data Erase Mounts	This 2 byte hexadecimal field contains the number of physical mount requests completed by the library during the interval to satisfy Security Data Erase mounts for the device class indicated. A mount is accredited to the interval the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	Count is incremented each time a Security Data Erase mount is completed.
20-31	Reserved	All bytes set to x00.	
Library - Pooling Container Bytes 256-8255 The data from each Hnode will be the same within a cluster.			

<table>
<tr><th>Bytes</th><th>Name</th><th colspan="2">Description</th><th>When Data is Sample/Updated</th></tr>
<tr><td colspan="4">Library - Pooling – Common Scratch Pool (CSP) Media Container

Bytes 256-319 (8 sets of data x 8 bytes/set = 64 bytes)

These fields contain 8 sets of data describing the physical media types in the CSP. The first media type is referred to as Media Type 0; the second is referred to as Media Type 1 and so forth. Each set of data contains 8 bytes. The following fields define these bytes and are numbered starting with 0. The data for media type 0 can be found in bytes 256-263; the data for media type 1 can be found in bytes 264-271, and so forth.</td><td></td></tr>
<tr><td rowspan="13">0</td><td rowspan="13">Physical Media Type</td><td colspan="2">This 1 byte hexadecimal field contains the identifier for the media type associated with following common scratch pool volume counts. The value is recorded at the end of the interval. The following are values for this field:</td><td rowspan="13">Type is updated every 2 minutes.</td></tr>
<tr><td>Value</td><td>Description</td></tr>
<tr><td>x00</td><td>No media type for this set of data</td></tr>
<tr><td>x10</td><td>3590 J media</td></tr>
<tr><td>x11</td><td>3590 K media</td></tr>
<tr><td>x20</td><td>3592 JA media</td></tr>
<tr><td>x21</td><td>3592 JW media (reserved for future use)</td></tr>
<tr><td>x22</td><td>3592 JJ media</td></tr>
<tr><td>x23</td><td>3592 JR media (reserved for future use)</td></tr>
<tr><td>x24</td><td>3592 JB Media</td></tr>
<tr><td>x25</td><td>3592 JX Media (Reserved for future use)</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
<tr><td colspan="2">This is the value at the end of the interval.</td></tr>
<tr><td>1-2</td><td>Physical Media Count</td><td colspan="2">This 2 byte hexadecimal field contains the number of scratch stacked volumes, of the type identified, assigned to the common scratch pool.

This is the value at the end of the interval.</td><td>Count is updated every 2 minutes.</td></tr>
<tr><td>3-7</td><td>Reserved</td><td colspan="2">All bytes set to x00.</td><td></td></tr>
<tr><td colspan="4">Library - Pooling – General Use Pool (GUP) Container

Bytes 320-8255 (32 sets of data x 248 bytes/set = 7936 bytes)

These fields contain 32 sets of data describing each of the 32 General Use Pools. Each set of data contains 248 bytes. The following fields define these bytes and are numbered starting with 0. The data for GUP 1 can be found in bytes 320-567; the data for GUP 2 can be found in bytes 568-815, and so forth.</td><td></td></tr>
</table>

Bytes	Name	Description	When Data is Sample/Updated
0-3	Active Logical Volumes	This 4 byte hexadecimal field contains the number of logical volume images resident in the volume pool at the end of the interval. The number is updated and the reported value is the snapshot of that count when statistics are calculated at the end of the interval. To be considered resident in a pool, the logical volume must be on one of the physical volumes assigned to the pool. Cache resident only volumes, although assigned to the pool, are not included.	Count is updated every 60 minutes.
4-7	Active Data	This 4 byte hexadecimal field contains the number of MBs of logical volume image data managed in the volume pool. The number is updated dynamically and the reported value is the snapshot of that count when statistics are calculated at the end of the interval. To be considered resident in a pool, the logical volume must be on one of the physical volumes assigned to the pool. Cache resident only volumes, although assigned to the pool, are not included. The value is reported in increments of 1 MB (1024x1024). Any volume with a count of less than 1 MB is rounded up to 1 MB.	Count is updated every 60 minutes.
8-11	Data Written to Pool	This 4 byte hexadecimal field represents the number bytes written to the media associated with this pool during the last interval. This is data pre-migrated from the Tape Volume Cache (TVC), and does not include data moved as part of reclamation. The value is reported in increments of 1 MB (1024x1024). This value is reset to 0 at the beginning of the interval.	Count is incremented when a logical volume's data is written to a pool.
12-15	Data Read from Pool	This 4 byte hexadecimal field represents the number bytes read from the media associated with this pool during the last interval. This is recall data written to the Tape Volume Cache (TVC), and does not include data moved as part of reclamation. The value is reported in increments of 1 MB (1024x1024). This value is reset to 0 at the beginning of the interval.	Count is incremented when a logical volume's data is read from a pool.

<table>
<tr><th>Bytes</th><th>Name</th><th>Description</th><th>When Data is Sample/Updated</th></tr>
<tr><td>16</td><td>Device Class</td><td>This 1 byte hexadecimal field indicates the device class identifier for the pool.
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x00</td><td>No device</td></tr>
<tr><td>x11</td><td>3590 Model B1</td></tr>
<tr><td>x13</td><td>3590 Model E1A</td></tr>
<tr><td>x14</td><td>3590 Model H1A</td></tr>
<tr><td>x20</td><td>3592 Model J1A. This also includes a 3592-E05 that is emulating a 3592-J1A device.</td></tr>
<tr><td>x22</td><td>3592 Model E05. This is for a 3592-E05 that is behaving as a 3590-E05.</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
</table>
This is the value at the end of the interval.</td><td></td></tr>
<tr><td>17-18</td><td>Average time since data began expiring on volumes.</td><td>This 2 byte hexadecimal field contains the average age, in days, of the residual data that resides on the stacked volumes assigned to the pool. Statically assigned and borrowed volumes are included in this calculation. A physical volume has residual data on it if it is not full. This value is calculated based on the date a volume transitions to not full and the current date.

This data does not include stacked volumes that are in Read-Only-Recovery or are Unavailable.

The data for this field is reported at the end of the interval.</td><td>Average age is updated every 360 minutes.</td></tr>
<tr><td>19-20</td><td>Maximum time since data began expiring on volumes.</td><td>This 2 byte hexadecimal field contains the maximum age, in days, of the residual data that resides on the stacked volumes assigned to the pool. Statically assigned and borrowed volumes are included in this calculation. A physical volume has residual data on it if it is not full. This value is calculated based on the date a volume transitions to not full and the current date.

This data does not include stacked volumes that are in Read-Only-Recovery or are Unavailable.

The data for this field is reported at the end of the interval.</td><td>Maximum age is updated every 360 minutes.</td></tr>
<tr><td>21-22</td><td>Average Age of Full Private Volumes</td><td>This 2 byte hexadecimal field contains the average age, in days, of private stacked volumes in the pool. Statically assigned and borrowed volumes are included in this calculation. This value is calculated based when the volume is marked as full until it is reclaimed.

This data does not include stacked volumes that are in Read-Only-Recovery or are Unavailable.

The data for this field is reported at the end of the interval.</td><td>Average age is updated every 360 minutes.</td></tr>
</table>

<table>
<tr><th>Bytes</th><th>Name</th><th colspan="2">Description</th><th>When Data is Sample/Updated</th></tr>
<tr><td>23-24</td><td>Maximum Age of Full Private Volumes</td><td colspan="2">This 2 byte hexadecimal field contains the maximum age, in days, of private stacked volumes in the pool. Statically assigned and borrowed volumes are included in this calculation. This value is calculated based when the volume is marked as full until it is reclaimed.

This data does not include stacked volumes that are in Read-Only-Recovery or are Unavailable.

The data for this field is reported at the end of the interval.</td><td>Maximum age is updated every 360 minutes.</td></tr>
<tr><td>25-31</td><td>Reserved</td><td colspan="2">All bytes set to x00.</td><td></td></tr>
<tr><td colspan="4">Library - Pooling – GUP - Media Container

Relative bytes 32-223 (8 sets of data x 24 bytes/set = 192 bytes)
These bytes are relative to the Library - Pooling – GUP Container.

These fields contain 8 sets of data describing up to 8 physical media types in the pool. The first media type is referred to as Media Type 0; the second is referred to as Media Type 1 and so forth. Each set of data contains 24 bytes. The following fields define these bytes and are numbered starting with 0. The data for media type 0 can be found in relative bytes 32-55; the data for media type 1 can be found in relative bytes 56-79, and so forth.</td><td></td></tr>
<tr><td rowspan="13">0</td><td rowspan="13">Physical Media Identifiers</td><td colspan="2">This 1 byte hexadecimal field contains the identifier for the media type associated with following general use pool (GUP) volume counts. The following are values for this field:</td><td rowspan="13"></td></tr>
<tr><td>Value</td><td>Description</td></tr>
<tr><td>x00</td><td>No media type for this set of data</td></tr>
<tr><td>x10</td><td>3590 J media</td></tr>
<tr><td>x11</td><td>3590 K media</td></tr>
<tr><td>x20</td><td>3592 JA media</td></tr>
<tr><td>x21</td><td>3592 JW media (reserved for future use)</td></tr>
<tr><td>x22</td><td>3592 JJ media</td></tr>
<tr><td>x23</td><td>3592 JR media (reserved for future use)</td></tr>
<tr><td>x24</td><td>3592 JB media</td></tr>
<tr><td>x25</td><td>3592 JX media (reserved for future use)</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
<tr><td colspan="2">This is the value at the end of the interval.</td></tr>
</table>

Bytes	Name	Description	When Data is Sample/Updated
1-2	Scratch Volume Count	This 2 byte hexadecimal field indicates the number of scratch stacked volumes associated with this pool of the media type indicated. The count includes volumes statically assigned to the pool at the end of the interval. It does not include any volumes that are borrowed from the CSP or are waiting to be erased due to a Security Data Erase. This data does not include stacked volumes that are in Read-Only-Recovery or are Unavailable. This is the value at the end of the interval.	Count is updated every 60 minutes.
3-4	Private Volume Count	This 2 byte hexadecimal field indicates the number of private stacked volumes associated with this pool of the media type indicated. The count includes volumes statically assigned to the pool at the end of the interval. It does not include any volumes that are borrowed from the CSP or are waiting to be erased due to a Security Data Erase. This data does not include stacked volumes that are in Read-Only-Recovery or are Unavailable. This is the value at the end of the interval.	Count is updated every 60 minutes.
5-6	Waiting for Security Data Erase	This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are waiting for Security Data Erase. The count includes volumes statically assigned to the pool and any volumes borrowed from the CSP. This data does not include stacked volumes that are in Read-Only-Recovery or are Unavailable. This is the value at the end of the interval.	Count is updated every 60 minutes.
7-8	Borrowed Scratch Volume Count	This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are in scratch status that have been borrowed from the CSP. This data does not include stacked volumes that are in Read-Only-Recovery or are Unavailable. This is the value at the end of the interval.	Count is updated every 60 minutes.
9-10	Borrowed Private Volume Count	This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are in private status that have been borrowed from the CSP. This data does not include stacked volumes that are in Read-Only-Recovery or are Unavailable. This is the value at the end of the interval.	Count is updated every 60 minutes.

Bytes	Name	Description	When Data is Sample/Updated
11-12	Read Only Recovery Volume Count	This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are in read only recovery status. The count includes stacked volumes that are both statically assigned to the pool and are borrowed stacked volumes. This is the value at the end of the interval.	Count is updated every 60 minutes.
13-14	Unavailable Volume Count	This 2 byte hexadecimal field indicates the number of stacked volumes associated with this pool of the media type indicated that are in unavailable status. The count includes stacked volumes that are both statically assigned to the pool and are borrowed stacked volumes. This is the value at the end of the interval.	Count is updated every 60 minutes.
15-23	Reserved	All bytes set to x00.	
Pooling – GUP - Reclaim Container Relative bytes 224-239 (1 set of data x 16 bytes/set = 16 bytes) These bytes are relative to the Library - Pooling – GUP Container. These fields contain Reclaim information for this GUP.			
224	Reclaim Threshold	This 1 byte hexadecimal field contains the reclaim threshold percentage for the pool as defined at the end of the interval.	Threshold is updated when its value changes.
225	Reclaim Pool	This 1 byte hexadecimal field contains the reclaim pool for the pool as defined at the end of the interval.	Pool is updated when its value changes.
226-227	Last Access Policy	This 2 byte hexadecimal field indicates, in days, when a physical volume is eligible for reclaim based on last access. A volume is eligible for reclaim when the number of days specified has elapsed since any data on the volume has been accessed because of a recall. If this field contains a value of 0, it is not used as criteria for reclaim. This is the value at the end of the interval.	Policy is updated when its value changes.
228-229	Last Written Policy	This 2 byte hexadecimal field indicates, in days, when a physical volume is eligible for reclaim based on when it was last written to. A volume is eligible for reclaim when the number of days specified has elapsed since any data has been written to the volume. If this field contains a value of 0, it is not used as criteria for reclaim. This is the value at the end of the interval.	Policy is updated when its value changes.
230-231	Last Data Invalidation Policy	This 2 byte hexadecimal field indicates, in days, when a physical volume is eligible for reclaim based on when data was last invalidated on it. A volume is eligible for reclaim when the number of days specified has elapsed since any data has invalidated on the volume. If this field contains a value of 0, it is not used as criteria for reclaim. This is the value at the end of the interval.	Policy is updated when its value changes.
232	Minimum Active Data Percentage Policy	This 1 byte hexadecimal field indicates the minimum active data percentage a physical volume's active data must fall below before it can be reclaimed using the days since last data invalidation reclamation policy. This is the value at the end of the interval.	Policy is updated when its value changes.

<table>
<tr><th>Bytes</th><th>Name</th><th>Description</th><th>When Data is Sample/Updated</th></tr>
<tr><td>233-234</td><td>Force Erasure Policy</td><td>This 2 byte hexadecimal field indicates the number of days before a physical volume must complete the erase process. The time starts when the first data is invalidated on the volume. Supported values are 0 to 365. If this field contains a value of 0, it is not used as criteria for reclaim. When this field contains a non-zero value, the pool is operating in the secure data erasure mode. Any volume that is reclaimed in the pool is physically erased before being returned to scratch. This is the value at the end of the interval.</td><td>Policy is updated when its value changes.</td></tr>
<tr><td>235-239</td><td>Reserved</td><td>All bytes set to x00.</td><td></td></tr>
<tr><td colspan="3">Pooling – GUP - Properties Container

Relative bytes 240-247 (1 set of data x 8 bytes/set = 8 bytes)
These bytes are relative to the Library - Pooling – GUP Container.

These fields contain pool property information for this GUP.</td><td></td></tr>
<tr><td>240</td><td>Pool Type</td><td>This 1- byte hexadecimal field indicates the type of the pool. The valid bit masks are as follows:
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x80</td><td>Reserved</td></tr>
<tr><td>x40</td><td>Reserved</td></tr>
<tr><td>x20</td><td>Export Pool</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
</table>
This is the value at the end of the interval.</td><td>Pool type is updated when its value changes</td></tr>
<tr><td>241</td><td>Return Borrowed Volumes</td><td>This 1 byte hexadecimal field indicates whether volumes borrowed from the CSP should be returned to the CSP after they are no longer needed by the pool. The valid values are:
<table>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x00</td><td>Don’t return borrowed volumes</td></tr>
<tr><td>x01</td><td>Return borrowed volumes</td></tr>
<tr><td></td><td>All other values are reserved.</td></tr>
</table>
This is the value at the end of the interval.</td><td>Policy is updated when its value changes.</td></tr>
</table>

<table>
<tr><th>Bytes</th><th>Name</th><th colspan="2">Description</th><th>When Data is Sample/Updated</th></tr>
<tr><td rowspan="10">242</td><td rowspan="10">First Media Types to Borrow</td><td colspan="2">This 1 byte bit-mapped field indicates the first media types to be borrowed from the CSP if additional scratch physical volumes are needed by the pool. When a bit is set to one (1) then the media type defined by the corresponding Library-Pooling-GUP-Media - Physical Media Identifier set can be used as one of the first media types to borrow from the CSP. When a bit is set to zero (0) the corresponding physical media type, if any, is not one of the first media types to borrow.
For example, if the first Physical Media Identifier (first set of 24 bytes) of the Library-Pooling-GUP-Media container indicates J media, the second (second set of 24 bytes) indicates K media, the third (third set of 24 bytes) indicates JA media, and the fourth (fourth set of 24 bytes) indicates JJ media, a value of x20 in this field would indicate the media type identified in the third set of 24 bytes should be borrowed first. In this example JA media will be borrowed first.
A value of x00 indicates borrowing is turned off for this pool.</td><td rowspan="10">Media types is updated when its value changes.</td></tr>
<tr><th>Value</th><th>Description</th></tr>
<tr><td>x80</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the first set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice.</td></tr>
<tr><td>x40</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the second set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice.</td></tr>
<tr><td>x20</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the third set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice.</td></tr>
<tr><td>x10</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the fourth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice.</td></tr>
<tr><td>x08</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the fifth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice.</td></tr>
<tr><td>x04</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the sixth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice.</td></tr>
<tr><td>x02</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the seventh set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice.</td></tr>
<tr><td>x01</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the eighth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a first choice.</td></tr>
</table>

<table>
<tr><th>Bytes</th><th>Name</th><th colspan="2">Description</th><th>When Data is Sample/Updated</th></tr>
<tr><td rowspan="3">242</td><td rowspan="3">First Media Types to Borrow</td><td colspan="2"></td><td rowspan="3">Media types is updated when its value changes.</td></tr>
<tr><td>Value</td><td>Description</td></tr>
<tr><td colspan="2">This is the value at the end of the interval.</td></tr>
<tr><td rowspan="8">243</td><td rowspan="8">Second Media Types to Borrow</td><td colspan="2">This 1 byte bit-mapped field indicates the second media types to be borrowed from the CSP if additional scratch physical volumes are needed by the pool. When a bit is set to one (1) then the media type defined by the corresponding Library-Pooling-GUP-Media – Physical Media Identifier set can be used as one of the second media types to borrow from the CSP. When a bit is set to zero (0) the corresponding physical media type, if any, is not one of the second media types to borrow.
For example, if the first Physical Media Identifier (first set of 24 bytes) of the Library-Pooling-GUP-Media container indicates J media, the second (second set of 24 bytes) indicates K media, the third (third set of 24 bytes) indicates JA media, and the fourth (fourth set of 24 bytes) indicates JJ media, a value of x10 in this field would indicate the media type identified in the fourth set of 24 bytes should be borrowed second. In this example JJ media will be borrowed second.
A value of x00 indicates a second media type to borrow is not specified</td><td rowspan="8">Media types is updated when its value changes.</td></tr>
<tr><td>Value</td><td>Description</td></tr>
<tr><td>x80</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the first set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice.</td></tr>
<tr><td>x40</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the second set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice.</td></tr>
<tr><td>x20</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the third set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice.</td></tr>
<tr><td>x10</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the fourth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice.</td></tr>
<tr><td>x08</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the fifth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice.</td></tr>
<tr><td>x04</td><td>The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the sixth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice.</td></tr>
</table>

Bytes	Name	Description	When Data is Sample/Updated
243	Second Media Types to Borrow	**Value** – **Description** x02 – The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the seventh set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. x01 – The media type defined by the Physical Media Identifier (Relative-Relative byte 0) in the eighth set of eight sets of data in the Library-Pooling-GUP-Media container can be borrowed as a second choice. This is the value at the end of the interval.	Media types is updated when its value changes.
244-247	Reserved	All bytes set to x00.	

Hnode Grid Historical record

This Hnode Historical record has the following nested structure:

- Header
- Grid Container
  - Grid-Cluster 0 Container
  - Grid-Cluster 1 Container (If installed)
  - Grid-Cluster 2 Container (If installed)
  - Grid-Cluster 3 Container (If installed)
  - ...
  - Grid-Cluster 7 Container (If installed)

Table I-10 Hnode Grid Historical record

Bytes	Name	Description	When Data is Sampled/Update
0-1	Length	This 2 byte hexadecimal field contains the length of this record. The length includes these 2 bytes.	
2	Version	This 1 byte hexadecimal field contains the version of the data presented in this record. Initially the version is set to x01.	
3	Data Type	This 1 byte hexadecimal field indicates the type of data contained in this record. For this record the value is set to x33 indicating this is an Hnode Grid Historical record.	
4	Node ID	This 1 byte hexadecimal field indicates the Hnode ID which this interval's data represents. Valid values are x00 – x01.	
5	Cluster ID	This 1 byte hexadecimal field indicates the Cluster ID which this Hnode is a part of. Valid values are x00 – x07.	
6-7	Interval Duration	This 1 byte hexadecimal field indicates the interval in seconds that this interval's data was taken over.	

Bytes	Name	Description	When Data is Sampled/Update
8-11	Time Stamp	This 4 byte hexadecimal field indicates the end time of the interval this data was taken over. This value is the time in seconds since the Epoch (00:00:00 UTC, 01 January 1970)	
12-15	Machine Type	This 4 byte EBCDIC field contains this node's machine type. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "3957".	
16-18	Machine Model	This 3 byte EBCDIC field contains this node's machine model. The field is left justified padded with EBCDIC blanks. Initially this field will be set to "V06".	
19-26	Machine Serial Number	This 8 character EBCDIC field contains the serial number of this node. This field is left justified and padded with EBCDIC blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed.	
27-34	VE Code Level	This 8 byte hexadecimal field contains the code level of the TS7700 Virtualization Engine (VE). The 8 bytes are actually four 2 byte fields. Each 2 byte field represents a portion of the code level. The VE code level is expressed as Version.Release.Modification.Fix in a decimal form. For example the code level of 8.0.0.104 would be represented in the 8 bytes as: x0008000000000068.	
35-39	Grid Library Sequence Number	This 5 character EBCDIC field contains the Library Sequence Number of the grid (Composite) library.	
40-63	Reserved	All bytes set to x00	
Grid Container Bytes 64-95 The data from each Hnode will be the same within a cluster.			
64-67	Logical Volumes for Copy	This 4 byte hexadecimal field indicates the number of logical volumes that are scheduled to be copied to this cluster. This is the value at the end of the interval.	Count is updated every 5 minutes.
68-71	Data to Copy	This 4 byte hexadecimal field indicates the amount of data that is scheduled to be copied to this cluster. This represents the amount of data contained in the logical volumes that are scheduled to be copied. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This is the value at the end of the interval.	Count is updated every 5 minutes.
72-75	Average Deferred Queue Age	This 4 byte hexadecimal field indicates the average age, in seconds, of the logical volumes in the deferred copy queue destined to be copied to this cluster. This is the value at the end of the interval.	Age is updated every 5 minutes.
76-79	Average Immediate Queue Age	This 4 byte hexadecimal field indicates the average age, in seconds, of the logical volumes in the immediate copy queue destined to be copied to this cluster. This is the value at the end of the interval.	Age is updated every 5 minutes.

Bytes	Name	Description	When Data is Sampled/Update
80	Number of clusters	This 1 byte hexadecimal field indicates the number of clusters in the grid. This field can be used to determine how many grid-cluster containers will be attached to this record. There is a maximum of 8 clusters. This is the value at the end of the interval.	
81-95	Reserved	All bytes set to x00.	
Grid-Cluster Container Bytes 96 and up (Number of clusters x 256 bytes/set) For example, if there are 3 clusters in the grid there will be 3 sets of data. There is a maximum of 8 clusters. This next segment of the record contains one set of data for each cluster in the grid as defined in byte 80 above. Each set of data contains 256 bytes. The data for the first cluster (Cluster 0) can be found in bytes 96-351; the second cluster's (Cluster 1) data can be found in bytes 352-447, and so forth.			
0-31	Data transferred into a cluster's cache from other clusters as part of a copy operation	This field contains eight 4 byte fields indicating the number of bytes transferred to this cluster's cache from other clusters as part of a copy operation during this interval. The first 4 byte field indicates the number of bytes transferred into this cluster's cache from Cluster 0. The second 4 byte field indicates the number of bytes transferred from Cluster 1, and so forth. A value of x00000000 is indicated for a cluster's own value and for non-existent clusters. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This value is not currently supported. The fields will be set to zeroes. This value is reset to 0 at the beginning of the interval.	Count is incremented every time a block of data is transferred to this cluster.
32-63	Data transferred from a cluster's cache to other clusters as part of a copy operation	This field contains eight 4 byte fields indicating the number of bytes transferred from this cluster's cache to other clusters as part of a copy operation during this interval. The first 4 byte field indicates the number of bytes transferred from this cluster's cache to Cluster 0. The second 4 byte field indicates the number of bytes transferred to Cluster 1, and so forth. A value of x00000000 is indicated for a cluster's own value and for non-existent clusters. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This value is reset to 0 at the beginning of the interval.	Count is incremented every time a block of data is transferred.

Bytes	Name	Description	When Data is Sampled/Update
64-95	Logical mounts directed to other clusters	This field contains eight 4 byte fields indicating the number of logical mounts from all Vnodes in this cluster which were satisfied by accessing another cluster. The first 4 byte field indicates the number of logical mounts directed to Cluster 0 from this cluster. The second 4 byte field indicates the number of logical mounts directed to Cluster 1, and so forth. A value of x00000000 is indicated for a cluster's own value and for non-existent clusters. A logical mount is counted when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.	Count is incremented when a logical mount is completed.
96-127	Data transferred into a cluster's cache from other clusters as part of a remote write operation	This field contains eight 4 byte fields indicating the number of bytes transferred to this cluster's cache from other clusters as part of a remote write operation during this interval. The first 4 byte field indicates the number of bytes transferred into this cluster's cache from Cluster 0. The second 4 byte field indicates the number of bytes transferred from Cluster 1, and so forth. A value of x00000000 is indicated for a cluster's own value and for non-existent clusters. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This value is reset to 0 at the beginning of the interval.	Count is incremented every time a block of data is transferred.
128-159	Data transferred from a cluster's cache to other clusters as part of a remote read operation	This field contains eight 4 byte fields indicating the number of bytes transferred from of this cluster's cache to other clusters as part of a Remote read operation during this interval. The first 4 byte field indicates the number of bytes transferred from this cluster's cache to Cluster 0. The second 4 byte field indicates the number of bytes transferred to Cluster 1, and so forth. A value of x00000000 is indicated for a cluster's own value and for non-existent clusters. The value is reported in increments of 1 MB (1024x1024). Any residual data is rounded up to 1 MB. This value is reset to 0 at the beginning of the interval.	Count is incremented every time a block of data is transferred.
160-255	Reserved	All bytes set to x00.	

J

Sample JCL

This appendix provides the sample JCL to run:

- Bulk Volume Information Retrieval (BVIR) jobs to obtain statistical data from the TS7700 Virtualization Engine.
- JCL to run the VEHSTATS tool to analyze Point-in-time and Historical statistics records obtained through BVIR.
- JCL used for the Copy Export function. How to create an Export List Volume.

> **Note:** You can also find tailored JCL to run BVIR jobs and to analyze the data using VEHSTATS in the IBMTOOLS libraries. To access the IBM Tape Tools, go to:
>
> `ftp://ftp.software.ibm.com/storage/tapetool/`

BVIR JCL examples

The following are fragments of JCL that show how to use a standard IBM utility, IEBGENER, to request and process the requested data for the BVIR function.

The sample JCL shown in Example J-1 obtains a scratch volume to perform the BVIR request.

Example J-1 Obtain a scratch volume for BVIR request

```
//VEQUERY JOB ...
//********************************************************************
//* DO NOT USE COMPACTION WHEN WRITING THE REQUEST FILE
//********************************************************************
//* SUBSTITUTE YOUR OWN DATA SET NAME, JOB NAME, ETC.
//* DATA SET IS CATALOGED
//********************************************************************
//* USING A LOGICAL SCRATCH VOLUME, CREATE THE REQUEST FILE WITH
//* THE 2 REQUIRED RECORDS. IN ORDER TO ENSURE THAT A SCRATCH VOLUME
//* IS ALLOCATED IN THE TARGET LIBRARY FOR THE QUERY OPERATION,
//* THE ACS ROUTINES NEED TO HAVE LOGIC TO ALLOCATE A TAPE
//* DRIVE IN THE TARGET LIBRARY. ONE WAY TO ACCOMPLISH THIS IS TO
//* HAVE A STORAGE GROUP UNIQUE TO EACH VTS LIBRARY PROVIDING A
//* 1 TO 1 RELATIONSHIP BETWEEN STORAGE GROUP AND LIBRARY.
//* THE ACS ROUTINES WOULD THEN NEED TO KEY OFF OF SOMETHING
//* UNIQUE IN THE DD STATEMENT (DATA SET NAME, DATA CLASS SPECIFICATION,
//* UNIT SPECIFICATION, ETC ...) TO GET THE CORRECT STORAGE GROUP
//* AND THE RIGHT TARGET LIBRARY SELECTED.
//********************************************************************
//* FILE SEQUENCE 1: REQUEST/RESPONSE FILE
//* RECORDS MUST BE SPECIFIED AS ILLUSTRATED BELOW, STARTING IN
//* THE FIRST COLUMN:
//* SPECIFY THE SEQUENCE NUMBER OF THE TS7700 THE REQUEST IS TO GO TO AS A
//* CHECK THAT IT IS GOING TO THE CORRECT TS7700
//* //********************************************************************
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=A
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=SYSBADM.CQUERY,
// UNIT=3490,LABEL=(,SL),
// DISP=(NEW,CATLG),
// DCB=(RECFM=F,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1   DD *
VTS BULK VOLUME DATA REQUEST
CACHE CONTENTS
/*
...
```

The JCL shown in Example J-2 uses a specific volume to perform the BVIR request.

Example J-2 Obtain a specific volume for BVIR request

```
//VEQUERY JOB ...
//*********************************************************************
//* DO NOT USE COMPACTION WHEN WRITING THE REQUEST FILE
//*********************************************************************
//* SUBSTITUTE YOUR OWN DATA SET NAME, JOB NAME, ETC.
//* DATA SET IS CATALOGED
//*********************************************************************
//* USING A LOGICAL SPECIFIC VOLUME, CREATE THE REQUEST FILE WITH
//* THE 2 REQUIRED RECORDS.
//*********************************************************************
//* FILE SEQUENCE 1: REQUEST/RESPONSE FILE
//* RECORDS MUST BE SPECIFIED AS ILLUSTRATED BELOW, STARTING IN
//* THE FIRST COLUMN:
//* SPECIFY THE SEQUENCE NUMBER OF THE TS7700 THE REQUEST IS TO GO TO AS A
//* CHECK THAT IT IS GOING TO THE CORRECT TS7700

//* //*********************************************************************
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=SYSBADM.CQUERY,
// UNIT=3490,LABEL=(,SL),
// DISP=(NEW,CATLG),VOL=SER=VTA001,
// DCB=(RECFM=F,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1   DD *
VTS BULK VOLUME DATA REQUEST
CACHE CONTENTS
/*
...
```

For the request types that return 80 byte fixed block records (grid volume status, cache contents, volume map and physical media pools), the following sample JCL (Example J-3) reads the BVIR response data and sends it to a printer.

Example J-3 Reads the BVIR response data and sends it to a printer

```
//VERESP JOB ...
//*********************************************************************
//* THE RESPONSE DATA IS NOT COMPACTED
//*********************************************************************
//* SUBSTITUTE YOUR OWN DATA SET NAME, JOB NAME, ETC.
//*********************************************************************
//* USING THE DATA SET CATALOGED IN THE REQUEST JOB
//*********************************************************************
//* FILE SEQUENCE 1: REQUEST/RESPONSE FILE
//*********************************************************************
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=A
//SYSIN    DD DUMMY
//SYSUT1   DD DSN=SYSBADM.CQUERY, DISP=OLD

//SYSUT2   DD  SYSOUT=A,
//         DCB=(DSORG=PS,RECFM=F,LRECL=80,BLKSIZE=80)
...
```

For the request types that return variable length unformatted records (Point-in-time and Historical statistics), the JCL shown in Example J-4 reads the BVIR response data to a DASD data set.

Example J-4 Reads the BVIR response data to a DASD data set

```
//VERESP JOB ...
//*******************************************************************
//* THE RESPONSE DATA IS NOT COMPACTED
//*******************************************************************
//* SUBSTITUTE YOUR OWN DATA SET NAME, JOB NAME, ETC.
//*******************************************************************
//* USING THE DATA SET CATALOGED IN THE REQUEST JOB
//*******************************************************************
//* FILE SEQUENCE 1: REQUEST/RESPONSE FILE
//*******************************************************************
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=A
//SYSIN    DD DUMMY
//SYSUT1   DD DSN=SYSBADM.PITSTAT, DISP=OLD
//         DCB=(RECFM=U,BLKSIZE=24000)

//SYSUT2   DD  DSN=SYSBADM.PITSTAT.CAPTURE,
//         UNIT=SYSDA,SPACE=(CYL,(2,1)),
//         DISP=(NEW,CATLG),
//         DCB=(RECFM=U,BLKSIZE=24000)
...
```

The following sample JCL in Example J-5 for JES2 (this will not work for JES3 because it does not demount/mount the volume between steps) combines the BVIR request and read steps into a single job with REF=* and includes the response data in the job output:

Example J-5 Combine the BVIR request and read steps into a single job for job output

```
//BVIRINFO JOB ...
//STEP1    EXEC PGM=IEBGENER
//SYSPRINT  DD  SYSOUT=*
//SYSUT2  DD DSN=TAPE.BVIR.LIB5MPLS,
//          UNIT=B63M2N36,LABEL=(,SL),
//          DISP=(NEW,KEEP),
//          DCB=(RECFM=F,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1   DD *
VTS BULK VOLUME DATA REQUEST
PHYSICAL MEDIA POOLS
/*
//SYSIN   DD DUMMY
//*
//STEP2    EXEC PGM=IEBGENER
//SYSPRINT DD  SYSOUT=*
//SYSUT1   DD DSN=TAPE.BVIR.LIB5MPLS,
//         VOLUME=(,,REF=*.STEP1.SYSUT2),
//         DCB=*.STEP1.SYSUT2,
//         DISP=(OLD),LABEL=(1,SL)
//SYSUT2   DD  SYSOUT=A,
//         DCB=(DSORG=PS,RECFM=F,LRECL=80,BLKSIZE=80)
//SYSIN    DD DUMMY
/*
```

The sample JCL shown in Example J-6 for JES2 (this will not work for JES3 because it does not demount/mount the volume between steps) combines the BVIR request and read steps into a single job where the request tape is cataloged and the response is written to an output file on DASD.

Example J-6 Combine the BVIR request and read steps into a single job for DASD output

```
//BVIRINFO JOB ...
//STEP1    EXEC PGM=IEBGENER
//SYSPRINT  DD  SYSOUT=*
//SYSUT2  DD DSN=TAPE.BVIR.B63M2N36,
//           UNIT=B63M2N36,LABEL=(,SL),
//           DISP=(NEW,CATLG),
//           DCB=(RECFM=F,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1   DD *
VTS BULK VOLUME DATA REQUEST
CACHE CONTENTS
/*
//SYSIN   DD DUMMY
//*
//STEP2    EXEC PGM=IEBGENER
//SYSPRINT DD  SYSOUT=*
//SYSUT1   DD DSN=TAPE.BVIR.B63M2N36,
//         DISP=(OLD.UNCATLG),LABEL=(1,SL)
//SYSUT2   DD  DSN=TAPE.BVIR.B63M2N36.OUTPUT,
//         DISP=(NEW,CATLG),SPACE=(CYL,(5,5)),
//         UNIT=3390,
//         DCB=(DSORG=PS,RECFM=F,LRECL=80,BLKSIZE=80)
//SYSIN    DD DUMMY
/*
```

The final sample JCL in Example J-7 for JES3 separates the create and read steps into two separate jobs.

Example J-7 JES3 separate create and read steps

```
JOB1:
//JS010   EXEC PGM=IEBGENER
//SYSPRINT  DD SYSOUT=*
//SYSIN     DD DUMMY
//SYSUT1    DD *
VTS BULK VOLUME DATA REQUEST
VOLUME MAP
/*
//SYSUT2   DD  DSN=OUTPUT.DATASET.NAME,
//             DISP=(NEW,CATLG,DELETE),
//             UNIT=CTAPE,
//             RETPD=14,
//             DCB=(LRECL=80,BLKSIZE=80,TRTCH=NOCOMP)
//*
JOB2:
//JS020   EXEC PGM=IEBGENER
//SYSPRINT DD  SYSOUT=*
//SYSIN    DD  DUMMY
//SYSUT1   DD  DSN=OUTPUT.DATASET.NAME,
//             DISP=OLD
//SYSUT2   DD  SYSOUT=U,LRECL=80,RECFM=F
```

VEHSTATS sample JCL

In this section, we provide some sample JCL which you can use for VEHSTATS. The following jobs are part of the IBMTOOLS.EXE package:

BVIRHSTS Requests the statistics for one or more days and writes them to the SMF log file. See Example J-8.

BVIRHSTU Requests the statistics for one or more days and writes them to a disk data set as RECFM=U. See Example J-9 on page 689.

BVIRHSTV Requests the statistics for one or more days and writes them to a disk data set as RECVM=VB.

VEHSTATS Processes the BVIRHIST file to produce reports. See Example J-10 on page 690.

BVIRPOOL Is used to get Point-in-time pooling information from any cluster in a grid.

VEHSCAN Is more of a debugging tool which simply lists the contents of every field in the Historical statistics for a user-selected interval.

VEPSCAN Is also a debugging tool that simply lists the contents of every field in the Point-in-time statistics.

The first JCL (Example J-8) is called BVIRHSTS, requests the statistics for one or more days, and writes them to a disk data set which can be kept as long as you want. As documented before, the TS7700 Virtualization Engine keeps the last 90 days of Historical statistics.

Example J-8 Request Historical statistics for one or more days and write them to the SMF log file

```
//*JOB1
//*JOB2
//*JOB3
//*JOB4
/*JOBPARM SYSAFF=*
//*
//*  THIS JOB ISSUES THE BULK VOLUME INFORMATION (BVIR) REQUEST FOR
//*  HISTORICAL STATISTICS FROM THE VTS ASSOCIATED WITH THE VIRTUAL
//*  DRIVE ADDRESS USED.  THE BVIR FEATURE MUST BE ACTIVATED ON THE
//*  VTS RECEIVING THE REQUEST.  THE FINAL OUTPUT IS DIRECTED TO SMF.
//*  NEXT, RUN VEHSTATS TO GET REPORTS.
//*
//BVIRHIST PROC USERHLQ=USERID,      HI-LEVEL FOR USER DATA FILES
//        TOOLHLQ=TOOLID,        HLQ FOR LOAD AND CNTL
//        VTSID=VTSID,      CLUSTER SERIAL NUMBER TO BE PART OF DSN
//        UNIT=VTAPE             UNITNAME ON THIS VTS
//*
//STEP1    EXEC PGM=IEFBR14
//DEL1     DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
//             DSN=&USERHLQ..#&VTSID..BVIRTAPE
//*
//STEP2    EXEC PGM=GETHIST         ISSUE HISTORICAL STATS REQUEST
//STEPLIB  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.LOAD
//SYSLIST  DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//BVIRREQ  DD  DSN=&USERHLQ..#&VTSID..BVIRTAPE,
//             UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
//             DCB=(RECFM=F,LRECL=80,BLKSIZE=80,TRTCH=NOCOMP)
//*
```

```
//STEP3    EXEC PGM=CPYHIST   APF LIBRARY NEEDED IF WRITING TO SMF
//STEPLIB  DD  DISP=SHR,DSN=USER.AUTH.LIB
//SYSLIST  DD  SYSOUT=*
//SYSPRINT DD   SYSOUT=*
//RECLIST  DD   SYSOUT=*
//SYSUT1   DD   DSN=&USERHLQ..#&VTSID..BVIRTAPE,
//             DCB=(RECFM=U,BLKSIZE=22000),
//             DISP=(OLD,DELETE)
//SYSUT2 DD  DSN=NULLFILE,DCB=(RECFM=U,BLKSIZE=22000),
//         DISP=(NEW,CATLG),SPACE=(CYL,(40,25),RLSE),UNIT=SYSDA
//   PEND
//*
//GETHIST  EXEC BVIRHIST
//STEP2.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
//     DD  *
 SDATE= 01FEB2007;         USE HERE AS DDMONYEAR
 EDATE= 01FEB2007;         USE HERE AS DDMONYEAR
*SDATE= TODAY- 1;         THIS FORMAT PULLS STATS FROM PREVIOUS DAY
*EDATE= TODAY- 1;
//STEP3.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
//     DD  *
 SMFNUM = 194;   USER SELECTABLE SMF # FOR QUARTER HOUR STATISTICS
*   THE SMF TIME STAMP WILL BE THE QUARTER HOUR DATE/TIME ORIGINALLY
*    WRITTEN EVEN IF PULLED SEVERAL DAYS LATER.
//*
```

The next example (Example J-9) shows the BVIRHSTU job to write the Historical statistics to a disk data set with RECFM = U.

Example J-9   Write the Historical statistics to a disk data set with RECFM = U

```
//*JOB1
//*JOB2
//*JOB3
//*JOB4
/*JOBPARM SYSAFF=*
//*
//*  THIS JOB ISSUES THE BULK VOLUME INFORMATION (BVIR) REQUEST FOR
//*  HISTORICAL STATISTICS FROM THE VTS ASSOCIATED WITH THE VIRTUAL
//*  DRIVE ADDRESS USED.  THE BVIR FEATURE MUST BE ACTIVATED ON THE
//*  VTS RECEIVING THE REQUEST.  THE FINAL OUTPUT IS WRITTEN TO A
//*  DISK DATA SET AS RECFM=U.
//*  NEXT, RUN VEHSTATS TO GET REPORTS.
//*
//BVIRHIST PROC USERHLQ=USERID,      HI-LEVEL FOR USER DATA FILES
//        TOOLHLQ=TOOLID,         HLQ FOR LOAD AND CNTL
//        VTSID=VTSID,       CLUSTER SERIAL NUMBER TO BE PART OF DSN
//        SDATE=,                 BEGINNING DATE
//        EDATE=,                 ENDING DATE
//        UNIT=VTAPE              UNITNAME ON THIS VTS
//*
//STEP1    EXEC PGM=IEFBR14
//DEL1     DD  UNIT=(&UNIT,,DEFER),DISP=(MOD,DELETE),
//             DSN=&USERHLQ..#&VTSID..BVIRTAPE
//DEL2     DD  UNIT=SYSDA,DISP=(MOD,DELETE),SPACE=(TRK,1),
```

```
//              DSN=&USERHLQ..#&VTSID..BVIRHIST.D&SDATE..D&EDATE
//*
//STEP2    EXEC PGM=GETHIST         ISSUE HISTORICAL STATS REQUEST
//STEPLIB  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.LOAD
//SYSLIST  DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//BVIRREQ  DD  DSN=&USERHLQ..#&VTSID..BVIRTAPE,
//             UNIT=&UNIT,LABEL=(,SL),DISP=(NEW,CATLG),
//             DCB=(RECFM=F,LRECL=80,BLKSIZE=80,TRTCH=NOCOMP)
//*
//STEP3    EXEC PGM=CPYHIST
//STEPLIB  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.LOAD
//SYSLIST  DD  SYSOUT=*
//SYSPRINT DD   SYSOUT=*
//RECLIST  DD   SYSOUT=*
//SYSUT1   DD   DSN=&USERHLQ..#&VTSID..BVIRTAPE,
//             DCB=(RECFM=U,BLKSIZE=22000),
//             DISP=(OLD,DELETE)
//SYSUT2 DD  DSN=&USERHLQ..#&VTSID..BVIRHIST.D&SDATE..D&EDATE.,
//           DCB=(RECFM=U,BLKSIZE=22000),UNIT=SYSDA,
//           DISP=(NEW,CATLG),SPACE=(CYL,(40,25),RLSE)
//   PEND
//*
//GETHIST  EXEC BVIRHIST,SDATE=070201,EDATE=070201    HERE AS YYMMDD
//STEP2.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
//     DD  *
 SDATE= 01FEB2007;         USE HERE AS DDMONYEAR
 EDATE= 01FEB2007;         USE HERE AS DDMONYEAR
*SDATE= TODAY- 1;         THIS FORMAT PULLS STATS FROM PREVIOUS DAY
*EDATE= TODAY- 1;
//STEP3.SYSCNTL DD DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
//*
```

The next JCL VEHSTATS (Example J-10) processes the BVIRHIST file to produce the VEHSTATS reports. Do not forget to activate one of the input DD statements at the end of the job before submitting it.

Example J-10 Producing the VEHSTATS reports from the BVIRHIST file

```
//*JOB1
//*JOB2
//*JOB3
//*JOB4
//*
//*    VIRTUALIZATION ENGINE HISTORICAL STATISTICS REPORTING
//*  RUN THE BVIRHIST JOB FIRST TO GET THE STATISTICS FILE(S)
//*
//VEHSTATS PROC TOOLHLQ=TOOLID,     HLQ FOR LIBRARIES
//           USERHLQ=USERID,        FOR THE INPUT BVIR FILE
//           ORDER=ORDER1,          DEFAULT ORDER STATEMENTS
//           VTSID=12345            ID FOR REPORTING SYSTEM
//*
//DELETE  EXEC  PGM=IEFBR14
//HOURFLAT  DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..#&VTSID..HOURFLAT.TXT,
```

```
//          DCB=(RECFM=FB,LRECL=6000,BLKSIZE=0)
//DAYHSMRY  DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..#&VTSID..DAYHSMRY.TXT,
//          DCB=(RECFM=FB,LRECL=6000,BLKSIZE=0)
//WEKHSMRY  DD UNIT=SYSDA,SPACE=(CYL,1),DISP=(MOD,DELETE),
//          DSN=&USERHLQ..#&VTSID..WEKHSMRY.TXT,
//          DCB=(RECFM=FB,LRECL=6000,BLKSIZE=0)
//*
//RPTSTEP  EXEC  PGM=VEHSTATS,REGION=0M
//STEPLIB  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.LOAD
//SYSLIST  DD  SYSOUT=*        CONTROL PARAMETERS USED
//RECLIST  DD  SYSOUT=*        DETAIL LIST OF BVIR COMMANDS
//*                            RPTS ARE ALL LRECL=140
//H20VIRT  DD  SYSOUT=*        VIRTUAL DRIVE ACTIVITY
//H21ADP00 DD  SYSOUT=*        VNODE ADAPTER 0 ACTIVITY
//H21ADP01 DD  SYSOUT=*        VNODE ADAPTER 1 ACTIVITY
//H21ADP02 DD  SYSOUT=*        VNODE ADAPTER 2 ACTIVITY
//H21ADP03 DD  SYSOUT=*        VNODE ADAPTER 3 ACTIVITY
//H21ADPXX DD  SYSOUT=*        ADAPTER COMPARISON
//H21ADPSU DD  SYSOUT=*        ADAPTER SUMMARY
//H30TVC1  DD  SYSOUT=*        CACHE PARTITION 1 ACTIVITY
//H30TVC2  DD  SYSOUT=*        CACHE PARTITION 2 ACTIVITY
//*   ADD DD STATEMENTS FOR TVC3, 4, ETC TO CORRESPOND WITH THE
//*     NUMBER OF PARTITIONS RUNNING.  IF A DD STATEMENT
//*     IS NEEDED, BUT MISSING, YOU WILL GET AN S0C1 ABEND.
//H31IMEX  DD  SYSOUT=*        IMPORT/EXPORT ACTIVITY
//H32TDU12 DD  SYSOUT=*        PHYS DRIVE TYPES 1/2 ACTIVITY
//H32TDU34 DD  SYSOUT=*        PHYS DRIVE TYPES 3/4 ACTIVITY
//H32CSP   DD  SYSOUT=*        COMMON SCRATCH POOL
//H32GUP01 DD  SYSOUT=*        GENERAL USE POOLS 01/02
//H32GUP03 DD  SYSOUT=*        GENERAL USE POOLS 03/04
//H32GUP05 DD  SYSOUT=*        GENERAL USE POOLS 05/06
//H32GUP07 DD  SYSOUT=*        GENERAL USE POOLS 07/08
//H32GUP09 DD  SYSOUT=*        GENERAL USE POOLS 09/10
//H32GUP11 DD  SYSOUT=*        GENERAL USE POOLS 11/12
//H32GUP13 DD  SYSOUT=*        GENERAL USE POOLS 13/14
//H32GUP15 DD  SYSOUT=*        GENERAL USE POOLS 15/16
//H32GUP17 DD  SYSOUT=*        GENERAL USE POOLS 17/18
//H32GUP19 DD  SYSOUT=*        GENERAL USE POOLS 19/20
//H32GUP21 DD  SYSOUT=*        GENERAL USE POOLS 21/22
//H32GUP23 DD  SYSOUT=*        GENERAL USE POOLS 23/24
//H32GUP25 DD  SYSOUT=*        GENERAL USE POOLS 25/26
//H32GUP27 DD  SYSOUT=*        GENERAL USE POOLS 27/28
//H32GUP29 DD  SYSOUT=*        GENERAL USE POOLS 29/30
//H32GUP31 DD  SYSOUT=*        GENERAL USE POOLS 31/32
//H33GRID  DD  SYSOUT=*    DISTRIBUTED LIBRARY STATS
//HOURXFER DD  SYSOUT=*    HOURLY XFER FOR ON DEMAND REPORTING
//DAYXFER  DD  SYSOUT=*    DAILY XFER FOR ON DEMAND REPORTING
//DAYSMRY  DD  SYSOUT=*    DAILY SUMMARY - VERTICAL 7 DAYS ACROSS
//MONSMRY  DD  SYSOUT=*    MONTHLY SUMMARY - VERTICAL 10 MONTHS ACROSS
//WEKHSMRY  DD UNIT=SYSDA,SPACE=(CYL,(2,3),RLSE),DISP=(,CATLG),
//          DSN=&USERHLQ..#&VTSID..WEKHSMRY.TXT,
//          DCB=(RECFM=FB,LRECL=6000,BLKSIZE=0)
//DAYHSMRY  DD UNIT=SYSDA,SPACE=(CYL,(2,3),RLSE),DISP=(,CATLG),
//          DSN=&USERHLQ..#&VTSID..DAYHSMRY.TXT,
```

```
//          DCB=(RECFM=FB,LRECL=6000,BLKSIZE=0)
//HOURFLAT  DD UNIT=SYSDA,SPACE=(CYL,(2,3),RLSE),DISP=(,CATLG),
//          DSN=&USERHLQ..#&VTSID..HOURFLAT.TXT,
//          DCB=(RECFM=FB,LRECL=6000,BLKSIZE=0)
//SORTIN    DD UNIT=(SYSDA,1),SPACE=(CYL,(300,100)),DISP=(,CATLG),
//             DSN=&&SORTIN,DCB=(RECFM=VB,LRECL=12000,BLKSIZE=0)
//SORTOUT   DD UNIT=(SYSDA,1),SPACE=(CYL,(300,100)),DISP=(,CATLG),
//             DSN=&&SORTED,DCB=(RECFM=VB,LRECL=12000,BLKSIZE=0)
//SORTWK01  DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK02  DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK03  DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK04  DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK05  DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK06  DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SORTWK07  DD UNIT=SYSDA,SPACE=(CYL,(200,100))
//SYSOUT   DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//  PEND
//*
//RUNRPTS  EXEC  VEHSTATS
//*********************************************************************
//RPTSTEP.SYSCNTL  DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.JCL(EXPIRE)
//       DD  DISP=SHR,DSN=&TOOLHLQ..IBMTOOLS.CNTL(&ORDER)
//       DD *
*********************************************************************
*        FILL IN THE FOLLOWING RECORDS AS APPROPRIATE:              *
*********************************************************************
 CUSTOMER= TITLENAME VEHSTATS; 1-50 CHAR
 REPORT= QTR VDSUM;   HOURLY ROLL-UP REPORTING AND VERTICAL SUMMARY
*      = QTR;     REQUEST 15 MINUTE REPORTING AS GENERATED BY TS7740
*      = HRS;     REQUEST HOURLY ROLL-UP REPORTING
*      = VDSUM;   DAILY SUMMARY - VERTICAL 7 DAYS ACROSS
*               PLUS MONTHLY SUMMARY - VERTICAL 10 MONTHS ACROSS
*               PLUS WEEKLY FLAT FILE - HORIZONTAL 1 WEEK/LINE
*      = HDSUM;   DAILY SUMMARY - HORIZONTAL 1 DAY/LINE
*      = DXFR;    FOR DAILY ON DEMAND TRANSFER REPORTING
*UTCMINUS= 07;     ADJUST UTC TO LOCAL TIME WEST OF GREENWICH
*UTCPLUS=  02;     ADJUST UTC TO LOCAL TIME EAST OF GREENWICH
*SDATE=    14JAN2007;     START DATE FOR OUTPUT REPORTING
*SDATE=    TODAY-  1;     REPORT JUST YESTERDAY'S DATA
*STIME=    10:10;         START TIME FOR OUTPUT REPORTING
*EDATE=    14JAN2007;     END DATE FOR OUTPUT REPORTING
*EDATE=    TODAY-  1;     REPORT JUST YESTERDAY'S DATA
*ETIME=    10:20;         END TIME FOR OUTPUT REPORTING
 LINES= 58;
*SMFNUM = 194;   USER SELECTABLE SMF #
*
* THE ORDER STATEMENTS DETERMINE WHICH FIELDS WILL BE REPORTED IN THE
*  DAYSMRY, MONSMRY, HOURFLAT, DAYHSMRY, AND WEKHSMRY REPORTS AND WHAT
*  ORDER THEY WILL APPEAR IN.
*  PICK AND CHOOSE FROM THIS LIST AND RE-ARRANGE TO FIT YOUR NEEDS.
*
*  IBMTOOLS.CNTL(ORDER1) IS THE DEFAULT MEMBER OR YOU CAN CREATE YOUR
*   OWN MEMBER WITH YOUR FIELDS AND SEQUENCE.
*
```

```
//*
//* ACTIVATE ONE OR MORE OF THE FOLLOWING DD STATEMENTS FOR YOUR DATA
//*STATSU  DD DISP=SHR,
//*           DSN=&USERHLQ..#&VTSID..BVIRHIST.D070205.D070205
//*STATSVB DD DISP=SHR,
//*           DSN=&USERHLQ..#&VTSID..BVIRHIST.D070206.D070206
//*STATSMF DD DISP=SHR,      RECORDS WILL BE SELECTED BASED ON SMFNUM
//*           DSN=&USERHLQ..#&VTSID..SMF194
```

Export List Volume sample JCL

In this section, we provide some sample JCL to create an Export List Volume. Example J-11 shows how to create the three needed files. In 5.3.2, "Implementing Copy Export" on page 212 there is an example of how the use Export List Volume as part of the Export Copy function.

Example J-11 Creation of Export List Volume

```
//****************************************

//* FILE 1: EXPORT LIST
//****************************************
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=HILEVELQ.EXPLIST,MGMTCLAS=MCNOCOPY,
// UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(1,SL),
// VOL=(,RETAIN),
// DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1 DD *
EXPORT LIST 03
EXPORT PARAMETERS PHYSICAL POOL TO EXPORT:09
OPTIONS1,COPY,EJECT
/*
//****************************************
//* FILE 2: RESERVED FILE
//****************************************
//STEP2 EXEC PGM=IEBGENER,COND=(4,LT)
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=HILEVELQ.RESERVED,MGMTCLAS=MCNOCOPY,
// UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(2,SL),
// VOL=(,RETAIN,REF=*.STEP1.SYSUT2),
// DCB=*.STEP1.SYSUT2
//SYSUT1 DD *
RESERVED FILE
/*
//****************************************
//* FILE 3: EXPORT STATUS FILE
//****************************************
//STEP3 EXEC PGM=IEBGENER,COND=(4,LT)
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
```

```
//SYSUT2 DD DSN=HILEVELQ.EXPSTATS,MGMTCLAS=MCNOCOPY,
// UNIT=VTS1,DISP=(NEW,CATLG),LABEL=(3,SL),
// VOL=(,,REF=*.STEP1.SYSUT2),
// DCB=*.STEP1.SYSUT2
//SYSUT1 DD *
EXPORT STATUS 01
/*
```

Example J-11 serves as input to do an Export Copy function, which enables you to do offsite vaulting of physical volumes from a TS7700.

Related publications

We consider the publications that we list in this section particularly suitable for a more detailed discussion of the topics that we cover in this book.

IBM Redbooks publications

For information about ordering these publications, see “How to get IBM Redbooks publications” on page 697. Note that some of the documents referenced here might be available in softcopy only.

- *IBM Virtualization Engine TS7740 R1.5 and TS7720: New Virtualization Options for Mainframe Servers,* SG24-7712 (available in early 2009)
- *IBM TS3500 Tape Library with System z Attachment: A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789
- *Introduction to SAN Distance Solutions*, SG24-6408
- *Guide to Sharing and Partitioning IBM Tape Library Data*, SG24-4409
- *IBM TotalStorage 3494 Tape Library: A Practical Guide to IBM Tape Drives and Tape Automation*, SG24-4632
- *DFSMShsm Primer*, SG24-5272
- *Continuous Availability - Systems Design Guide*, SG24-2085
- *Continuous Availability S/390 Technology Guide*, SG24-2086
- *IBM System Storage Tape Library Guide for Open Systems*, SG24-5946
- *Introduction to IBM S/390 FICON*, SG24-5176
- *FICON Native Implementation and Reference Guide*, SG24-6266
- *IBM System z Connectivity Handbook*, SG24-5444
- *Fiber Saver (2029) Implementation Guide*, SG24-5608

Other publications

These publications are also relevant as further information sources:

- *IBM Virtualization Engine TS7700 Series, Introduction and Planning Guide,* GA32-0567
- *IBM System Storage 3953 Library Manager Model L05 Operator Guide*, GA32-0558
- *IBM System Storage 3953 Library Manager Model L05 Operator Guide*, GA32-0448
- *IBM System Storage TS3500 Tape Library Introduction and Planning Guide*, GA32-0559
- *IBM System Storage TS3500 Tape Library Operator Guide,* GA32-0560
- *z/OS V1R3.0-V1R4.0 DFSMS Using Magnetic Tapes*, SC26-7412-01
- *Implementing System Managed Storage*, SC26-3123
- *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427

- *Lights Out! Advanced Tape Automation Using VM/ESA*, GG24-4347
- *DFSMSdfp Utilities*, SC26-7414-02
- *z/OS DFSMSdss Storage Administration Reference,* SC35-0424
- *DFSMShsm Storage Administration Guide*, SC35-0421
- *z/OS DFSMSdfp Storage Administrator Reference*, SC35-0422
- *z/OS DFSMSrmm Guide and Reference*, SC26-7404
- *z/OS DFSMSrmm Implementation and Customization Guide*, SC26-7405
- *IBM System Storage Tape System 3592 Introduction and Planning Guide,* GA32-0464
- *IBM TotalStorage Enterprise Tape System 3592 Operators Guide,* GA32-0465
- *IBM TotalStorage TS3500 Operator Guide,* GA32-0468
- *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473
- *IBM System Storage 3953 Library Manager Model L05 Operator Guide*, GA32-0448
- *IBM Encryption Key Manager component for the Java platform Introduction, Planning, and User's Gui*de, GA76-0418
- *IBM System Storage TS1120 Tape Drive and Controller Introduction and Planning Guide,* GA32-0555
- *IBM System Storage TS1120 Tape Drive and Controller Operator Guide,* GA32-0556
- *IBM System Storage 3953 Tape System Introduction and Planning Guide,* GA32-0557
- *IBM System Storage 3953 L05 Library Manager Operator Guide,* GA32-0558
- *IBM System Storage TS3500 Tape Library Introduction and Planning Guide,* GA32-0559
- *IBM System Storage TS3500 Tape Library Operator Guide,* GA32-0560
- *IBM Virtualization Engine TS7700 Series Introduction and Planning Guide,* GA32-0567
- *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141
- *z/OS V1R7.0 MVS System Commands*, SA22-7627
- *z/OS JES3 Initialization and Tuning Reference*, SA22-7550
- *z/OS MVS Planning: Operation*, SC22-7601
- *MVS Initialization and Tuning Reference*, SA22-7592-08
- *z/VSE System Administration Guide*, SC33-8224
- *z/VSE System Macros Reference*, SC33-8230
- *z/VSE V4R1.0 Administration,* SC33-8304
- *VM/ESA DFSMS/VM Removable Media Services User's Guide and Reference*
- *z/VM V5R2.0 Running Guest Operating Systems*, SC24-6115
- *z/VM V5R1.0 DFSMS/VM Planning Guide*, SC24-6089
- *z/VM V5R2.0 DFSMS/VM Removable Media Services,* SC24-6090
- *z/VM V5R2.0 DFSMS/VM Storage Administration,* SC24-6091

Online resources

These Web sites are also relevant as further information sources:

- TS7700 Infocenter

  http://publib.boulder.ibm.com/infocenter/ts7700ic/v1r0/index.jsp
- Web-Based Enterprise Management (WBEM)

  http://www.dmtf.org/standards/wbem/
- Common Information Model (CIM)

  http://www.dmtf.org/standards/cim/
- FICON Director support matrix

  http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ116133
- Technote Performance Considerations for a Cascaded FICON Director Environment

  http://www-1.ibm.com/servers/eserver/zseries/library/techpapers/pdf/gm130237.pdf
- IBM Global Services home page

  http://www.ibm.com/services
- IBM Business Continuity and Recovery Services

  http://www.ibm.com/services/continuity

How to get IBM Redbooks publications

You can search for, view, or download IBM Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

D

Z

IBM Virtualization Engine TS7700 Release 1.4a: Tape Virtualization for System z Servers

(1.0" spine)
0.875"<->1.498"
460 <-> 788 pages

IBM Virtualization Engine TS7700 Release 1.4a:

Tape Virtualization for System z Servers

Consolidate and protect your data with a Three-Cluster Grid configuration

Prepare for Disaster Recovery and implement Copy Export

Protect your data using TS1120 Tape Encryption

The IBM Virtualization Engine TS7700 provides tape virtualization for the System z environment. As the follow-on product to the highly successful IBM TotalStorage Virtual Tape Server, the TS7700 Virtualization Engine is designed to provide improved performance and capacity to help lower the total cost of ownership for tape processing. It introduces a new modular, scalable, high-performing architecture for mainframe tape virtualization.

The TS7700 Virtualization Engine integrates the advanced performance, capacity, and data integrity design of the 3592 Tape Drives, industry-leading tape technology from IBM, with high-performance disk and a new advanced System p server to form a storage hierarchy managed by robust storage management firmware with extensive self management capability. It includes functions such as advanced policy management to control physical volume pooling, cache management, dual copy, dual copy across a grid network, and copy mode control. The TS7700 offers a new standards-based management interface and enhanced statistical reporting, compared to the VTS.

This IBM Redbooks publication provides a detailed description of the TS7700 Virtualization Engine. In addition, it discusses the tasks involved with planning, implementing, and operating the TS7700 Virtualization Engine and describes migration considerations.

This book covers the TS7700 Virtualization Engine releases up to and including TS7700 R1.4a and includes information about supported tape drives and libraries current and supported at the time of writing. For detailed information about the IBM Virtualization Engine TS7740 R1.5 and the disk-only TS7720 Virtualization, refer to *IBM Virtualization Engine TS7740 R1.5 and TS7720: New Virtualization Options for Mainframe Servers*, SG24-7712, which will be available in early 2009.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7312-02 ISBN 0738432237