IBM

IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide

Find out about the IBM TS3500 with the IBM 3953 Tape System

Optimize your VTS management with APM

Improve performance with IBM 3592 tape drives

Babette Haeusser
Martin Berger
Jodi Buechler
Michael Dede
Wolfgang Kessel

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide

June 2006

SG24-6115-03

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

Fourth Edition (June 2006)

This edition applies to IBM TotalStorage Virtual Tape Server, IBM TotalStorage 3494 and TS3500 Tape Library, and IBM System Storage 3953 Tape System models current at the time of publishing.

Contents

Figures

Tables

Notices

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS 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 may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may 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 may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

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 illustrates programming techniques on various operating platforms. You may 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. You may 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 IBM's application programming interfaces.

Trademarks

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

AIX®
AS/400®
CICS®
CUA®
DB2®
DFSMSdfp™
DFSMSdss™
DFSMShsm™
DFSMSrmm™
DFSORT™
Enterprise Storage Server®
ESCON®
eServer™
FICON®
FlashCopy®
GDPS®
Geographically Dispersed Parallel Sysplex™
IBM®
IMS™
Magstar®
MVS™
Netfinity®
NetView®
OS/390®
OS/400®
Parallel Sysplex®
RACF®
RAMAC®
Redbooks™
Redbooks (logo) ™
RETAIN®
RMF™
RS/6000®
S/390®
System i™
System p™
System x™
System z™
Tivoli®
TotalStorage®
VM/ESA®
VSE/ESA™
VTAM®
z/OS®
z/VM®
z/VSE™

The following terms are trademarks of other companies:

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

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

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Preface

This IBM® Redbook discusses enhancements to the IBM TotalStorage® Virtual Tape Server (VTS) in the area of automated tape copying for z/OS® systems.

These enhancements broaden the usage of the VTS extensively and make the VTS a possible replacement product for any existing z/OS tape solution. However, they also add another level of complexity in terms of planning for and implementing the Peer-to-Peer VTS solution.

Starting with a detailed description of the architecture and the design of the Peer-to-Peer VTS, we give you a comprehensive overview of all the planning tasks, provide guidance on a wide variety of possible configuration choices, and discuss several possible recovery scenarios as well as considerations for day-to-day operations.

A good knowledge of IBM virtual tape technology is assumed.

The team that wrote this redbook

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

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 Tape. 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 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 with herself focussing on Enterprise Tape including tape libraries and Virtual Tape Servers.

Martin Berger is an Advisory IT Specialist for IBM Global Services, Integrated Technology Services, Germany, providing Software Support in z/OS and Storage Backoffice for IBM EMEA. He has 23 years of practical experience in IBM storage management software (DFSMShsm™, DFSMS OAM, DFSMSdss™, DFSMSdfp™, and DFSMSrmm™) and IBM disk and tape storage products (IBM ESS, RVA, XRC, PPRC and FlashCopy®, IBM 3494, and VTS).

Jodi Buechler is an Advisory Programmer from Tucson Arizona working in DFSMS Software Support. She has 16 years of experience in IBM storage management software (DFDSS, DFSMS OAM, and most recently AOM, TAPE/DASD ERP, BTLS, and ICKDSF.). She has specialized in IBM ATL, VTS, MTL, and PtP service support over the last several years.

Michael Dede is an ITAP System Services Representative for the Greater Toronto Area in Canada. He has 17 years of experience in servicing Enterprise systems. His areas of expertise include the 3494 Tape Library, VTS and Peer-to-Peer VTS. This includes assignment to one of IBM SO's accounts, which utilize the 3494 Peer-to-Peer VTS libraries in a GDPS® environment.

Wolfgang Kessel is a Senior System Engineer within the FTSS organization in Germany with more than 30 years of experience in the Storage Management area which includes Hardware and Software solutions. Currently, he is responsible for technical pre-sales support of RMSS products at specific large z/OS accounts in the Central Region.

Martin, Jodi, Wolfgang, Babette, Michael

The authors of the previous three editions of this redbook were:
Marco Ceresoli, Babette Haeusser, Tom Koudstaal, Nigel Morton, Taisei Takai, Paul Rizzio, David McFarlane, Barry Kadleck, Christina Coutts, Wolfgang Kessel

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

Emma Jacobs, Yvonne Lyon, Leslie Parham, Deanna Polm, Sangam Racherla
International Technical Support Organization, San Jose Center

Joe Bacco, Jennifer Bappert, Carl Bauske, Thomas Bish, Erika Dawson, Jim Fisher, Kevin Gibble, Dennis Haight, J.D. Metzger, Collette Morton, Corie Neri, Christina Orosco, Jon Peake, Jeffrey Pilch, Craig Schultz, David Shields, Jesse Thrall
IBM Tucson

Immo Compas, Antoinette Kaschner, Joachim Mueller, Patrick Wolf
IBM Germany

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll 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 Redbooks™ to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

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

  `ibm.com/redbooks`

- Send your comments in an e-mail to:

  `redbook@us.ibm.com`

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. QXXE Building 80-E2
  650 Harry Road
  San Jose, California 95120-6099

Summary of changes

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

Summary of Changes
for SG24-6115-03
for IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide
as created or updated on June 23, 2006.

June 2006, Fourth Edition

The latest edition of this redbook includes a complete rework of the book, which now includes detailed information about:

- VTS Release 7.2:
  - Advanced Policy Management (APM) enhancements
  - Bulk Volume Information Retrieval (BVIR)
- VTS Release 7.3:
  - IBM 3592 tape drive support
  - Reclamation enhancements
  - APM enhancements
- VTS Release 7.4:
  - IBM TS3500 / 3953 support[1]
  - APM enhancements

[1] The IBM System Storage TS3500 Tape Library was previously know as the IBM TotalStorage 3584 Tape Library.

1

Introducing the IBM TotalStorage Peer-to-Peer VTS

The IBM TotalStorage Peer-to-Peer Virtual Tape Server, an extension of the IBM TotalStorage Virtual Tape Server, builds on a proven base configuration to provide even greater benefits for tape processing and operations. By enhancing data backup and recovery capabilities, the IBM TotalStorage Peer-to-Peer Virtual Tape Server is specifically designed to enhance data availability. It accomplishes this by providing dual volume copy, remote functionality, and automatic recovery and switchover capabilities. With a design that reduces single points of failure (including the physical media where logical volumes are stored) the IBM TotalStorage Peer-to-Peer Virtual Tape Server improves system reliability and availability, as well as data access. To help protect current hardware investments, existing IBM TotalStorage Virtual Tape Servers can be upgraded for use in this new configuration (see Figure 1-1).

Figure 1-1 Peer-to-Peer configuration

The IBM TotalStorage Peer-to-Peer Virtual Tape Server (PtP VTS) consists of new models and features of the IBM 3494 that are used to join two separate Virtual Tape Servers into a single interconnected system. Supported libraries are the IBM TotalStorage 3494 Tape Library and the IBM System Storage TS3500 Tape Library. The two virtual tape systems can be located at the same site or at different sites that are geographically remote. This provides a remote copy capability for remote vaulting applications.

The Peer-to-Peer VTS appears to the z/OS host as a single automated tape library with 64, 128, or 256 virtual tape drives and up to 500,000 virtual volumes. The configuration of this system has up to 3.5 TB of Tape Volume Cache native (10.4 TB with 3:1 compression), up to 24 IBM TotalStorage 3592 tape drives, and up to 12 IBM TotalStorage 3590 tape drives models B1A, E1A, or H1A, and up to 16 host ESCON® or FICON® channels.

The components introduced to build a Peer-to-Peer VTS complex require no additional host processor resources, and provide the dual copy function transparently to the host processor.

Advantages from this new configuration include:

- Automated outboard data copying of virtual volumes between two Virtual Tape Servers
- Remote data copying between two ESCON/FICON connected Virtual Tape Servers
- Improved data availability by eliminating single points of failure
- Concurrent maintenance capabilities
- Workload balancing between the two connected Virtual Tape Servers

1.1 Design

The Peer-to-Peer VTS automatically creates a copy of any newly created or updated tape volume in both Virtual Tape Servers, unless with advanced policy management support (APM), the assigned Management Class does not request a copy. This process is performed transparently to the customer application and with no host processor resources required.

Either volume copy can then be used to satisfy a specific customer mount. This copy of all new or updated virtual volumes can be created using one of two possible modes of operation:

- Immediate copy: Creates a copy of the logical volume in the companion connected Virtual Tape Server prior to completion of a rewind/unload command. This mode provides the highest level of data protection.
- Deferred copy: Creates a copy of the logical volume in the companion connected Virtual Tape Server as activity permits after receiving a rewind/unload command. This mode provides protection that is superior to most currently available backup schemes.

A Peer-to-Peer VTS configuration is built from the following components:

- The 3494 Virtual Tape Controller (VTC)
- The 3494 Auxiliary Tape Frame Model CX1
- The 3494 Model B10 and B20 Virtual Tape Server attached to a IBM tape library
- Peer-to-Peer Copy features

The VTS model B18 can no longer be connected to a Peer-to-Peer VTS since this support was withdrawn from marketing in April 2004. The IBM 3494-AX0, the predecessor of the VTC, and the IBM 3494 Auxiliary Tape Frame Model CX0 have also been withdrawn from marketing.

1.2 Components and functions

The function of the new models and features of the IBM TotalStorage Peer-to-Peer Virtual Tape Server is described in this section. Figure 1-2 on page 3 introduces some new terms used to describe the Peer-to-Peer configuration, which are explained later on.

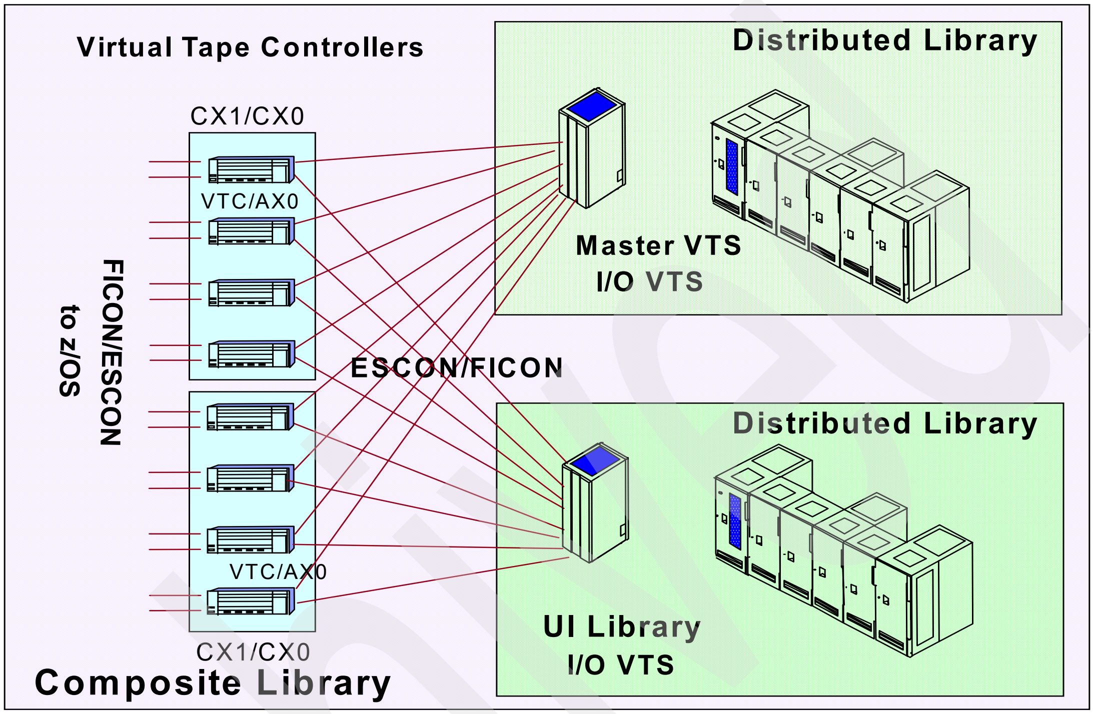

Figure 1-2 VTS Peer-to-Peer components

The Virtual Tape Controller (VTC)

The VTCs were introduced to provide the interface to the hosts and to connect two Virtual Tape Servers together. The model AX0 supported ESCON connections only. With the FICON support of the Peer-to-Peer VTS new VTC models and frames replaced the AX0 VTCs and CX0 frame. Each VTC is an independently operating, distributed node within the PtP VTS and continues to operate during scheduled or unscheduled service of another VTC.

The VTCs provide interconnection between two 3494 Model B10 or B20 Virtual Tape Servers with the Peer-to-Peer Copy features. It also provides two ESCON host attachments for sixteen drives of the Peer-to-Peer Virtual Tape Server. There must be four or eight VTCs in the configuration, with each attaching to both VTSs. The maximum number of virtual drives that can be host addressed is 256.

Each VTC provides synchronization of the copy of logical volumes, creates logical volume copies using large block transfers of compressed logical volumes, balances workload between the two Virtual Tape Servers, directs specific volume mounts to the Virtual Tape Server with a cached copy of the requested virtual volume, and displays the status and current configuration of the Peer-to-Peer VTS through a Web-based browser interface accessed through an Ethernet local area network connection.

Note: Effective December 19, 2003, the IBM TotalStorage Virtual Tape Controller (Model AX0 and the VTC frame CX0) were withdrawn from marketing. The AX0 VTC is replaced by Virtual Tape Controller features of the IBM TotalStorage Virtual Tape Frame (Model CX1).

The VTC Auxiliary Frame 3494 Model CX1

The Model CX1 provides the housing and power for two or four 3494 virtual tape controllers. Each Model CX1 can be configured with either two or four Model VTCs or two Model VTCs and two Model AX0s. There are two power control compartments, each with its own power cord to allow connection to two power sources. With the FICON support of the Peer-to-Peer VTS and the FICON enabled VTCs, the 3494 model CX0 frame was replaced by the CX1 frame.

Note: The AX0 VTC may still be used for ESCON PtP VTS or half ESCON half FICON installations

Peer-to-Peer Copy features

Special features installed on 3494 Model B10 and B20s in a Peer-to-Peer configuration provide automatic copies of virtual volumes. These features can be installed on existing VTS systems to upgrade them to a Peer-to-Peer VTS.

1.2.1 VTS Advanced Functions

As with a stand-alone VTS, the Peer-to-Peer VTS has the option to install some additional features and enhancements to existing features. These new features are:

- Outboard Policy Management
- Logical volume affinity (volume pooling)
- Large logical volume sizes up to 4000 MB
- Secure data erase
- Tape volume dual copy
- Peer-to-Peer Copy Control
- Peer-to-Peer Selective Copy Control
- Tape Volume Cache Management

For more detail on VTS Advanced Functions, refer to Section 2.5, “Advanced Policy Management (APM)” on page 40.

1.3 A new way to improve tape operations

The Virtual Tape Server or the Peer-to-Peer Virtual Tape Server resides between a z/OS host and a 3494 Tape Library or a TS3500 Tape Library. To the host processor's operating system, either device appears to be a series of automated 3490E tape drives. The use of virtual drives significantly reduces the number of real drives needed to support the processing environment. Unlike real tape drives, however, the Virtual Tape Server uses RAID 5 disk arrays to cache data, then efficiently fills tape cartridges and files them in the tape library. All data transfers with the host processor are performed to or from the Tape Volume Cache at disk access speeds.

In addition, utilizing the full capacity of the 3590 or 3592 cartridges can significantly reduce the requirement for additional tape drives, automation devices, media, and floor and shelf space.

1.3.1 Independence and investment protection

In the instance when one of the Virtual Tape Servers in a Peer-to-Peer Virtual Tape Server configuration is taken offline for maintenance, routine service, or even a system upgrade, all activity will be directed to the remaining Virtual Tape Server. When the Virtual Tape Server that was taken offline returns to service, each of the 3494 virtual tape controllers will resume copy operations to bring the logical volume copies in both Virtual Tape Servers back into synchronization.

Existing Model B10 and B20 Virtual Tape Servers can be upgraded to the new Peer-to-Peer Virtual Tape Server configuration by adding the incremental hardware and a second Model B10 or B20 with the Peer-to-Peer Copy features. Exiting Model B18 VTSs have to be upgraded to a model B20 VTS first. This protects the investment that many customers have already made in IBM TotalStorage technology.

1.3.2 Maximized data access

The Peer-to-Peer Virtual Tape Server is a major step in dealing with the ever-increasing need to maintain access to data 24 hours a day, 7 days a week, 365 days a year. Key components of the Peer-to-Peer Virtual Tape Server are duplicated to reduce single points of failure — a design that helps ensure that data remains accessible after a failure, during planned maintenance, or during system upgrades (Figure 1-3 shows a logical view of this concept).

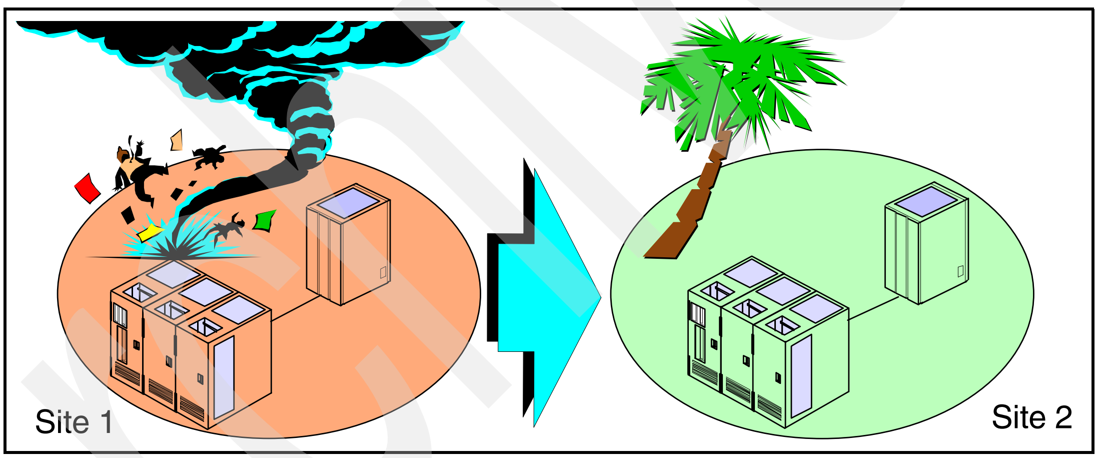

Figure 1-3 Peer-to-Peer design, logical view

Each Model B10 or B20 VTS in a Peer-to-Peer Virtual Tape Server configuration continues to operate during scheduled or unscheduled service of the other Model B10 or B20. After a repair action, the system initiates the creation of second copies to bring both Virtual Tape Servers back into synchronization.

1.3.3 Electronic vaulting and remote operation

Because one of the Virtual Tape Servers can reside in a geographically separate location from the other, operations can continue even if either one of the local or remote facilities is not operational. This design can virtually eliminate the need to physically transport data for disaster backup. The connections within a Peer-to-Peer Virtual Tape Server use ESCON or FICON channels, so Virtual Tape Servers can be physically separated by as much as the ESCON or FICON architecture allows.

1.3.4 Software support

The Peer-to-Peer Virtual Tape Server operates in the System z™ environment. Minimum level support is provided as a Small Programming Enhancement (SPE) for DFSMS/MVS 1.4 and 1.5 and for DFSMS 2.10. Toleration PTFs are provided for DFSMS/MVS 1.2 and higher. Additional SPEs and/or PTFs are required for some of the VTS Advanced Functions features. See Preventive Service Planning (PSP) buckets D/TVTS and D/TPTP for additional recommendations.

A Peer-to-Peer VTS may be attached to VM/ESA® or VSE/ESA™ guests of VM/ESA as long as there is also an z/OS host attached, with the full support installed. There is no support for open systems servers and TPF.

You monitor and control the Peer-to-Peer VTS using existing, standard host console commands and displays, the 3494 Library Manager panels, or the 3953-L05 Library Manager panels if the TS3500 is connected. Additional tools you can use include the IBM TotalStorage Enterprise Tape Library Specialist, which provides Web-browser accessed pages for monitoring the Peer-to-Peer configuration, and the IBM TotalStorage Enterprise Tape Library Expert, for detailed analysis of trends across many libraries.

2

Peer-to-Peer VTS Architecture

A single VTS presents itself to the host as one logical library containing 64, 128, or 256 virtual 3490E tape drives, if it is not integrated into a Peer-to-Peer configuration. In this book, we call a single Model B10 or B20 VTS that is not part of a Peer-to-Peer configuration a *stand-alone* VTS.

The Peer-to-Peer VTS also presents itself to the host as one logical library containing 64, 128, or 256 virtual 3490E tape drives, although the underlying hardware and the software implementation are different from the stand-alone VTS. It is essential that you understand the terminology, the general concepts, and the process flow to allow for detailed installation planning, easy implementation, and full integration of the Peer-to-Peer VTS into your disaster recovery plans:

- We describe the technical details and the internal logic of the Peer-to-Peer Virtual Tape Server.
- We explain the hardware components that make up a Peer-to-Peer VTS.
- We explain the attachment to the IBM TotalStorage 3494 and to the TS3500/3953 Tape Library, and the differences between 3494 and the TS3500/3953 Tape Library.
- We also give an overview of where the components of a Peer-to-Peer VTS can be located and which configurations are supported.
- We describe the terms and expressions that are introduced with the Peer-to-Peer VTS and used to describe the components and processes that allow for automated secondary copy of virtual volumes in the VTS.
- We discuss these processes as well as the control points that allow you to influence the Peer-to-Peer VTS internal processing.

For details about hardware features and software implementation of the Peer-to-Peer VTS, see Chapter 3, “Preinstallation planning” on page 63.

For a detailed description of the VTS feature codes related to a stand-alone VTS, and for more information concerning the tape library attachments for the VTS, please refer to:

- *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229

- *IBM TotalStorage 3584 Tape Library for zSeries Hosts Planning and Implementation,* SG24-6789
- *IBM TotalStorage 3494 Tape Library Introduction and Planning Guide*, GA32-0448

2.1 Peer-to-Peer VTS hardware components

Every VTS is always connected to a 3494 Tape Library or to a TS3500 Tape Library using the 3953 Library Manager, which provides all necessary elements to manage automated cartridge and physical tape drive handling. If the VTS is connected to a 3494 Tape Library, all the components needed for automation are integrated in the IBM 3494-Lxx Control Unit Frame. The configuration of a VTS consists of these hardware components:

- One IBM 3494-L22 frame, which contains the Library Manager and robotics.
- One IBM 3494-D12 frame installed inside an IBM 3494 Tape Library, including the IBM TotalStorage 3590 tape drives for exclusive use of the VTS. Up to six IBM TotalStorage 3590 Model B1As, Model E1As, or Model H1As can be installed in the D12 frame (two drive frames with 12 3590 drives are supported for B20), or
- One IBM 3494-D22 frame installed inside the 3494 Tape Library, including the 3592 tape drives for exclusive use of the VTS. Up to twelve 3592 Model J1A drives can be installed in the D22 frame, or
- Both frames exploiting the heterogeneous support, a D22 with up to twelve 3592 drives, and a D12 frame with up to six 3590 drives if connected to a B20.
- One IBM 3494-B10 or B20 frame installed up to 14 m, if there is a D12 frame connected and 25m if there has to be a D22 frame to be connected away from the IBM 3494 Tape Library.

For additional information on supported tape drive configurations, refer to 3.1.10, “IBM TotalStorage 3592 tape drives Model J1A and E05 in a 3494 library” on page 72.

If the VTS is connected to a TS3500 Tape Library, all the components needed for accessor control, all physical motions, inventory processing and error recovery are integrated in the Library Controller of the 3584-L22 frame. There may be up to twelve VTS dedicated 3592 drives installed in this frame. The Library Manager itself, 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 the 3953 frame.

The configuration of a VTS consists of these hardware components:

- One IBM 3584-L22 frame, which contains the library controller, the robotics, and which may house 3592 tape drives
- One or more 3584-D22 frames, if you want to spread the VTS dedicated 3592 drives over more frames, which is supported by the Automatic Library Management System (ALMS)
- One IBM 3953-F05 frame, which contains the IBM 3953-L05 Library Manager
- One IBM 3494-B10 or B20 frame installed up to 25 m away from the TS3500 Tape Library.

One VTS may have up to 500,000 logical volumes. Both the 3494 tape Library Manager and the 3953 Library Manager attached to a TS3500 Tape Library support up to 1,000,000 logical volumes.

In a Peer-to-Peer VTS environment, two VTSs residing in two different tape libraries must be interconnected to build the Peer-to-Peer configuration. All combinations of 3494 and TS3500 are supported. You may have one VTS connected to a 3494 Tape Library and the other VTS in a TS3500 Tape Library, or both VTSs are connected to a Tape Library of the same type. Additional hardware components are required to allow for automated secondary copy of virtual volumes, which are written to the Peer-to-Peer VTS.

Figure 2-1 shows all the hardware components required for a Peer-to-Peer VTS installation, in this case with four VTCs and one VTS connected to a TS3500 Tape Library and the second VTS connected to a 3494 Tape Library.

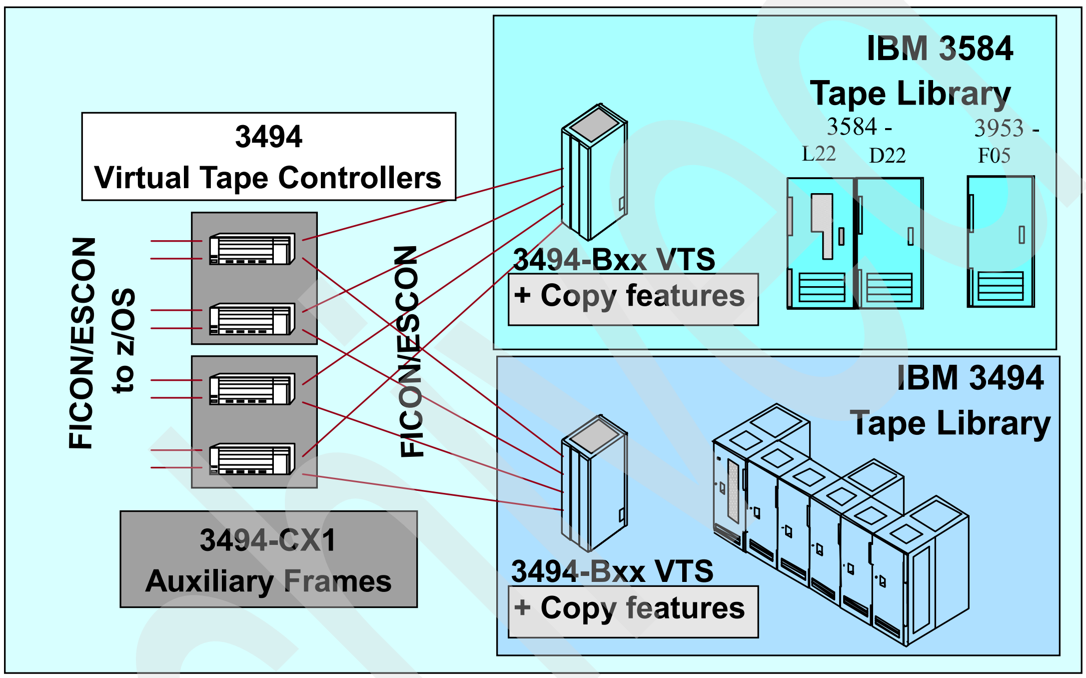

Figure 2-1 Peer-to-Peer VTS hardware components

The two boxes on the right-hand side of Figure 2-1 show the elements that are also required for the stand-alone VTS. The *copy features*, plus the components shown on the left-hand side of Figure 2-1, are required to integrate two stand-alone VTSs into a Peer-to-Peer VTS. A Model B10 or B20 VTS with the Peer-to-Peer Copy features installed is an independently operating distributed server node in a Peer-to-Peer VTS.

The minimum configuration of a Peer-to-Peer VTS consists of:

- Two IBM tape libraries, each of them containing one B10 or B20 VTS, which has the Peer-to-Peer Copy features installed, and one or more drive frames dedicated to each of the Model B10 or B20 VTS
- One to four Copy Features, FC4100 - FC4103 installed on each VTS
- Four IBM 3494 Virtual Tape Controllers (VTCs)
- One or two IBM 3494-CX0/1 Virtual Tape Frames (VTFs) housing the VTCs

2.1.1 The IBM TS3500/3953 Tape Library VTS attachment

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 available for mid-range and high-end open systems. Since its introduction, the library has been enhanced to accommodate different drive types and operating platforms, more recently including the attachment of System z (mainframe) hosts and tape drive controllers.

Previously, mainframe attached tape drives were housed in the IBM 3494 exclusively, which is an enterprise library designed for IBM 3590 (Magstar®) and IBM 3592 drive technology in both mainframe and open systems environments. Now the IBM TS3500 Tape Library is also available to connect drives to host systems with FICON or ESCON attachments (using the System z attachment capability covered in this book), as well as any combination of Fibre Channel and Ultra2/Wide Low Voltage Differential (LVD) SCSI.

The IBM TS3500 was designed and built on the strong foundation of the IBM 3494 robotics and microcode. Physically the designs look very similar (see Figure 2-2), but only approximately 3% of the hardware components are truly common to both machines. 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 3592 or LTO Ultrium high density cartridge technology, the IBM TS3500 provides a powerful robust tape storage solution for the whole enterprise, yet it is contained in a minimal footprint.

Figure 2-2 IBM 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.

System z attachment to the TS3500 Tape Library

Attaching System z to the IBM System Storage TS3500 Tape Library requires a number of specific hardware elements. In this section, we highlight basic conceptual differences in the requirements for mainframes and open systems hosts, and outline how the TS3500 Tape Library has evolved to include support for mainframes.

The tape drive type currently supported for System z attachment within the TS3500 Tape Library is the IBM TotalStorage 3592 Tape Drive or the IBM System Storage TS1120 Tape Drive, which can also be attached to open systems hosts. However, the System z attaches to the library differently from 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) within 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.

> **Note:** Any drives connected to a System z controller are dedicated for System z use and addressed only through the controller. Even though the IBM 3592 has additional native FCP adapters, they cannot be connected to open systems hosts in addition to the mainframe controllers.

The IBM Virtual Tape Server (VTS) can also be viewed in this context as a special type of controller. The VTS is attached directly to the mainframe FICON or ESCON host channels (perhaps through a director) and emulates a set of physical controllers and attached tape drives. The VTS itself has real physical tape drives attached to it, with which the host does not communicate directly (Figure 2-3). In this example the eight VTS tape drives are installed in the IBM 3584 L22 and D22 frames. The eight native 3592 drives are connected to the J70 controller installed in the 3953-F05 frame.

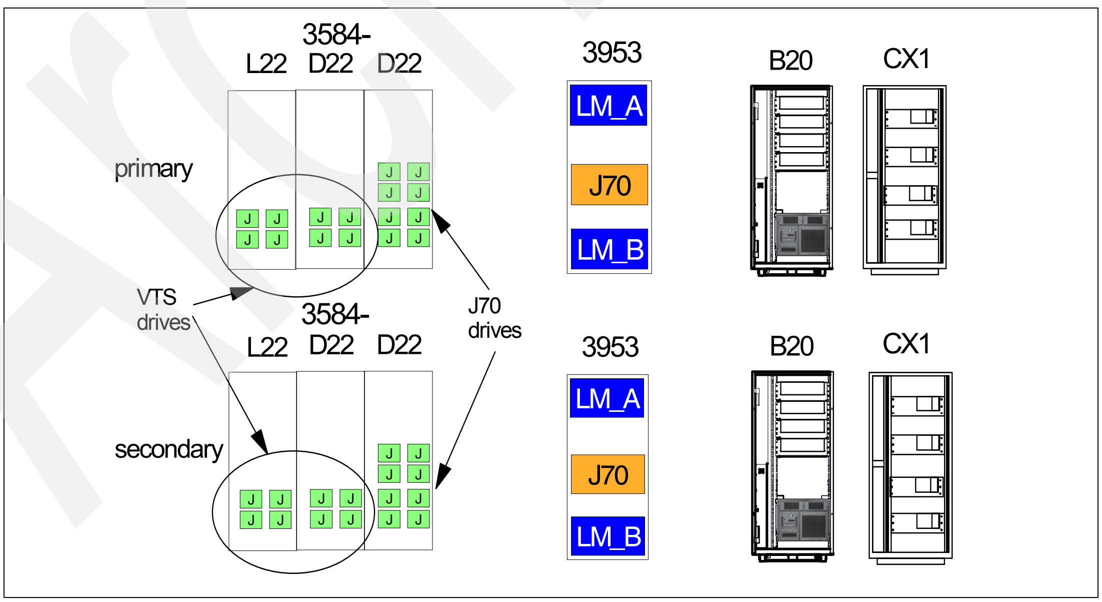

Figure 2-3 VTS attached to a TS3500 Tape Library

IBM TS3500 logical libraries and partitioning

The patented IBM Multi-Path Architecture enables simple SCSI medium changer libraries to be shared by like or unlike host platforms, without middle ware 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. Unlike 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 the System z controllers, (VTS and IBM 3592-J70), from the open systems drives. The System z host is allocated its own logical library partition within the TS3500 Tape Library; there may be:

- A single mainframe logical library defined
- A single open systems logical library
- Multiple mainframe logical libraries
- Multiple logical libraries, some for open hosts and some for System z hosts

The minimum requirement for a logical library definition is that it contain at least one drive and one cartridge slot. Since 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 libraries.

Each System z TS3500 Tape Library partition is managed on behalf of the host by a IBM 3953 Library Manager.

> **Note:** A Library Manager can be connected to only a single IBM TS3500 logical library.

Library control mechanisms

The IBM 3494 is controlled entirely by an integrated Library Manager Intel® processor. Open systems hosts connect directly to the Library Manager using an Ethernet LAN (or a direct serial connection) with an IBM-supplied device driver. The IBM 3494 Library Manager is responsible for maintaining the cartridge inventory, controlling the robotics, and sharing or partitioning the cartridges and drives between attached hosts.

The IBM TS3500 is controlled entirely by the attached host which uses its own application code and the IBM supplied device driver to send SCSI Medium Changer commands directly to the library robotics. The host also maintains its own cartridge inventory, and does any partitioning at the hardware level, allocating drives and slots to particular partitions in advance as defined by the user.

System z hosts control the IBM libraries by means of special commands, for example, the Perform Library Function (PLF) channel command word (CCW) in z/OS. In the IBM 3494 the host sends these commands using the data path to a controller which routes them internally to the integrated Library Manager. The IBM TS3500 has no such integrated server; so for System z attachment, a separate Library Manager, IBM 3953 Model L05, is provided which is physically external to the library. Similar (but not identical) in function to the Library Manager in the IBM 3494 library, the IBM 3953-L05 provides the appropriate System z interface to the IBM TS3500.

IBM TotalStorage 3953 Library Manager and frame

The IBM TS3500 Tape Library interface is managed on behalf of the System z hosts by an external Library Manager, IBM 3953 Model L05. This unit, along with the System z control units, and various switches, are housed in a special frame type, IBM 3953 Model F05, which is sited outside the library enclosure. The IBM 3953 Library Manager has common microcode with the Library Manager used in the IBM 3494 Tape Library (Figure 2-4).

Figure 2-4 IBM 3853-F05 frame with the IBM 3953-L05 Library Manager

The IBM 3953 Tape System provides additional flexibility and enhancements to enterprise customers. It expands the IBM TS3500 Tape Library with IBM 3592 Tape Drives from Open Systems Fibre Channel environments to System z Enterprise Systems ESCON and FICON environments.

You order the Library Manager for System z attachment to the IBM TS3500 as a separate device type and model number (IBM 3953 Library Manager Model L05). The Library Manager translates requests from the System z hosts into SCSI move medium changer commands to communicate with the TS3500 Tape Library. It shares common microcode with the Library Manager used in the IBM 3494 library. It is the same PC used in the IBM 3494; however the card compliment in the PC is different when it is used in an IBM 3494 compared with an IBM TS3500.

The 3953 Library Manager is supported with the following configurations:

- IBM 3494 VTS Model B10 or B20, and PtP VTS
- IBM 3592 tape controller Model J70
- Up to 16 subsystems per IBM 3953 Tape System (of which two can be VTSs)
- Up to four IBM 3953 Tape Systems per TS3500 Tape Library (a maximum of 8 VTSs)

Unlike the IBM 3494, this Library Manager does not (and does not need to) control any open systems drives.

The Library Manager is installed in an independent frame, the IBM TotalStorage 3953 Tape Frame Model F05. This frame sits outside the IBM TS3500 Tape Library enclosure, and can be physically located away from it. The IBM 3953 frame comes in a base user interface configuration and an expansion connectivity configuration. The base frame contains the IBM 3953 Library Manager or Managers, and up to one IBM 3592 tape controller. The expansion frame can house up to three IBM 3592 tape controllers. In one logical library subsystem there can be only one base frame and up to five expansion frames. Figure 2-5 shows a schematic drawing of the IBM 3584 frames and the IBM 3953 Tape Frame, in which the IBM 3953 Library Manager is installed.

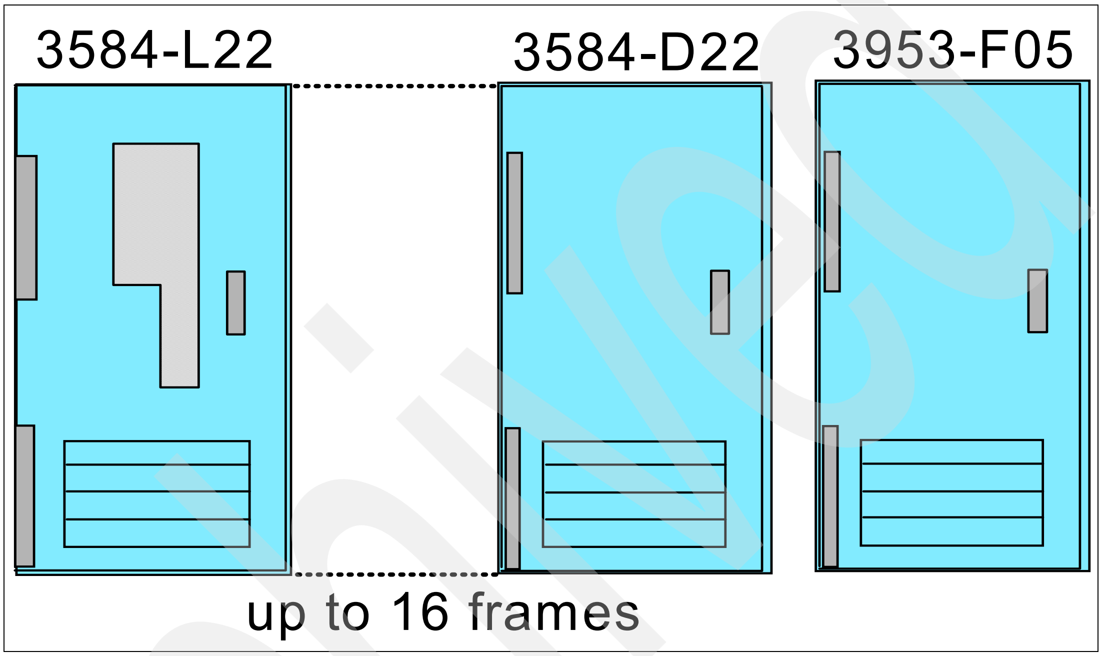

Figure 2-5 TS3500 Tape Library and IBM 3953 Tape System frame

Advanced Library Management System

The Advanced Library Management System (ALMS) for the TS3500 Tape Library is the next generation of the Multi-Path Architecture. It provides a user-friendly Web interface to facilitate the definition and management of multiple logical libraries.

> **Note:** ALMS is an optional feature for the IBM TS3500 in general; however it is a mandatory feature for the library when you attach System z, or when you install the IBM TS3500 Model HA1.

Figure 2-6 shows a simple diagrammatic illustration of the overall System z attachment to the TS3500 Tape Library.

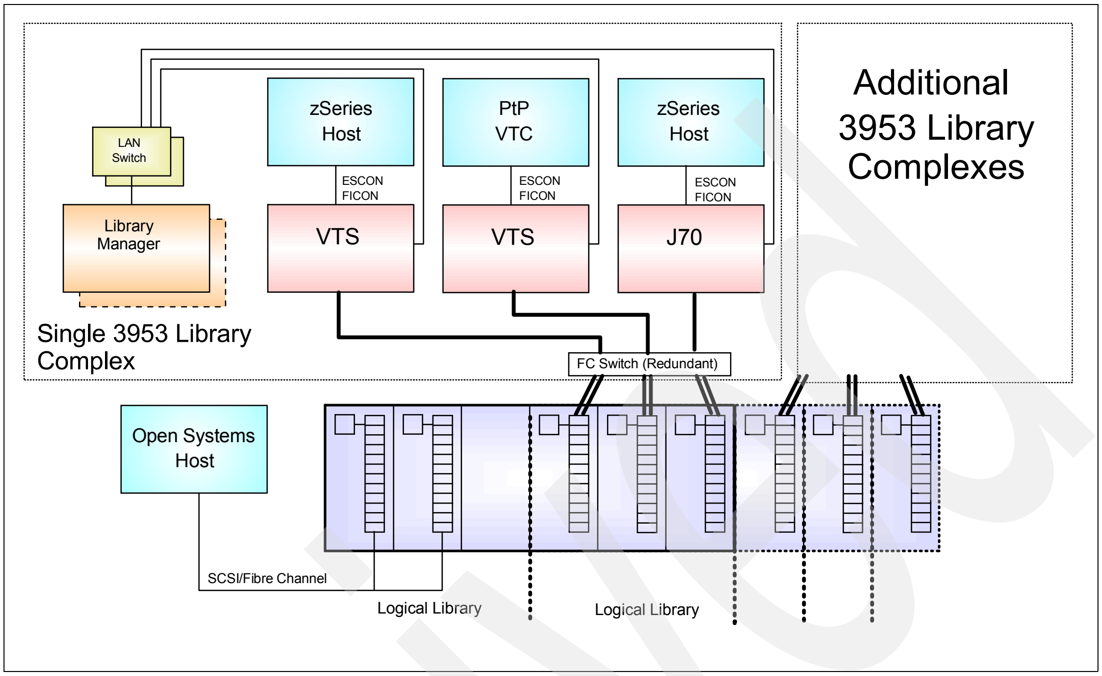

Figure 2-6 The 3953 connections to the TS3500 library

2.1.2 IBM 3494 versus IBM TS3500 with 3953 comparison

Here is a short description of the differences between the IBM 3494 and IBM TS3500 Tape Libraries and the System z attachment though IBM 3953 Library Manager and controllers for the System z connection to IBM TS3500 Tape Library.

You have multiple configuration options with the 3494 Tape Library and/or TS3500/3853 Tape Library. A single IBM 3494 Tape Library supports up to two VTS subsystems. The two scenarios shown in Figure 2-7 are *single Peer-to-Peer configurations* and *dual Peer-to-Peer configurations*, with 3494 Tape Library and TS3500 Tape Library.

A single 3953 library manager attached to a TS3500 library supports three logical library partitions:

- VTS_1
- VTS_2
- One J70 controller for native attached 3592 drives

As a single TS3500 Tape Library supports up to four 3953 Library Managers, up to eight VTSs may be attached to one physical TS3500 library.

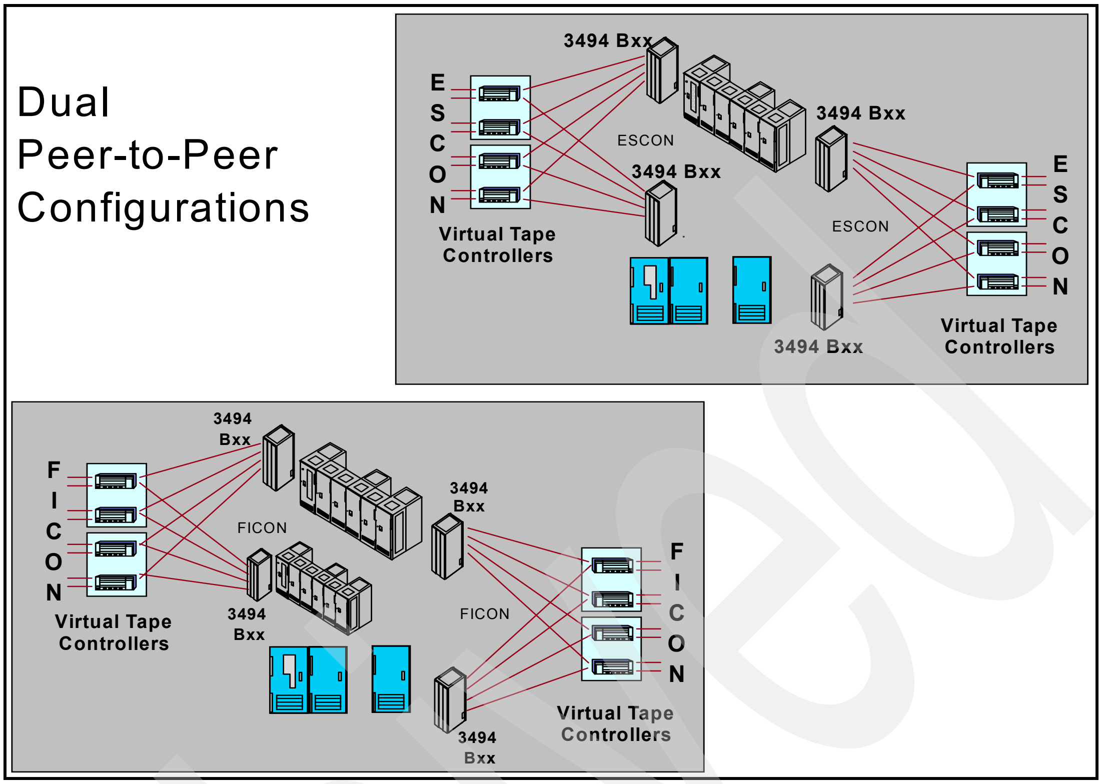

Figure 2-7 Dual Peer-to-Peer configurations

The installation option to choose depends on a variety of criteria such as, for example, the organization of your data center, the number of different locations, the distance between data centers, and data availability and disaster recovery requirements.

For details about integration of the Peer-to-Peer VTS configuration into existing disaster recovery and high-availability implementations, please refer to Chapter 9, “Disaster recovery scenarios” on page 321.

IBM 3494 Tape Library

The IBM TotalStorage 3494 Tape System z Library houses the Library Manager (LM). The first LM is LMA and is located inside the first Lxx frame. For daily operation the LM is located with its keyboard and monitor on the rear of the frame and there is a small operator panel on the front where you can set library in pause or auto mode. All IBM 3592-J70 tape controllers are also located inside the frame of the IBM 3494 Tape Library except the TS1120 Model C06 Tape Controller, the IBM 3494 VTS Models B10/B18/B20, and the IBM 3494 VTCs for communication between the VTS in a Peer-to-Peer environment; they are also in separate frames that are located outside of the IBM 3494 Tape Library frames.

Normally the VTS and VTC frames are standing side by side. If the feature for high availability (HA) is installed in the library, you will have an extra LM; this is called LMB and is located in the last frame, this last frame also houses the extra accessor B. Depending on the feature code in your tape library, the extra accessor is able to move tape cartridges inside the IBM 3494 frames, otherwise it is a backup accessor for accessor A. Although it is more than 10 years old, it is still robust, because the library, drive, controllers, and other features have been updated throughout time, and many new features have also been added:

- IBM 3590 Tape Drive Model B1A through H1A
- IBM controllers A00 through J70

- IBM 3592 Tape Drive
- IBM 3494 VTS B16 through B20
- TotalStorage Master Console
- High Availability (HA1) Dual Accessor and Dual Library Manager
- Library Manager is updated many times
- Peer-to-Peer
- ESCON attachment
- FICON attachment

IBM TS3500 Tape Library with IBM 3953 Library Manager

This library is only a few years old, but in the same way as the IBM 3494 Tape Library, this TS3500 Tape Library is also been redesigned over the years. Originally this library was for the open system environment and only had a connection for the IBM Ultrium LTO drive. Today the total footprint is smaller; the frame size is reduced 20% because the 3592 drive and LTO drive are smaller than the 3590 drive. The following features are included:

- High Availability (HA1) Dual Accessor
- IBM Ultrium LTO Drive from Generation 1 through 3
- IBM 3592 Tape Drive
- TotalStorage Master Console
- 20% smaller footprint
- System z attachment with IBM 3953 Library Manager and Frame
- IBM 3494 VTS B10 and B20
- Dual Library Manager
- Peer-to-Peer
- ESCON attachment
- FICON attachment

The IBM System Storage TS3500 Tape Library has handled tape for open tape system environments from the beginning. Now you can connect the TS3500 Tape Library to a System z host system as well. On front of the tape library there is an operator touch panel where all functions for the open systems hosts can be controlled and is controlled and tape library related operations for the System z. These functions can also be controlled through the TS3500 Tape Library Specialist; the rest of the System z operation is on the Library Manager in the IBM 3953 Base Frame. For this concept, all communication from the System z host will go though the IBM TotalStorage 3953 Library Manager Model L05 (LMA). The Library Manager is the same industrial PC as the PC inside the IBM 3494 Tape Library, and this Library Manager is mandatory inside the IBM 3953 Base Frame. From this LM, you run all daily operations, or you can use the Web inteface. The IBM 3953 frame houses all components for the connection to System z attachment. In addition to LMA, the base frame can house all of the following components:

- IBM 3953 Library Manager B Model L05 (LMB)
- TotalStorage Master Console (MC)
- KVM Switch for LM and MC
- House up to six fiber switches
- Two ethernet switches
- TotalStorage 3592 Model J70 Tape Controller

However, as with the IBM 3494 Tape Library, the IBM TS3500 Tape Library needs a Library Manager for communication between controllers and the attached System z environment.

The first IBM 3592 model J70 Controller for System z attachment is installed inside the IBM 3953 Base frame. If you need more than one IBM 3592 model J70 controller, they have to be installed inside an Expansion frame (no controllers inside the IBM TS3500 Tape Library). The IBM 3494 VTS Model B10 and B20, which are the supported VTS for System z attachments for the IBM TS3500 Tape Library, are located away from the IBM TS3500 frames as in the IBM 3494 Tape Library environment.

If the feature for high availability (HA) is installed in the IBM TS3500 Tape Library, you have an extra accessor, and this is located in the last frame and called accessor B. When installed with dual accessors, the IBM TS3500 Tape Library can perform a non-disruptive fail-over to the redundant accessor when any component of an accessor fails. Different from the IBM 3494, these accessors can operate in a System z environment at the same time and with only one Library Manager; this is called LMA and is located inside the IBM 3953 Base Frame.

For redundancies, you can install an extra LM inside the IBM 3953 Base Frame; this is called LMB and is the upper LM. The two Library Managers, second LMB as an option, controllers, and switches are placed away from the IBM TS3500 Tape Library just as the VTS is for the IBM 3494. If you wish to run the IBM TS3500 Tape Library with only one accessor and two Library Managers, this is also possible. Another possibility in the IBM 3953 Base Frame is the dual ethernet, also for redundancy. In the IBM TS3500 Tape Library it is only possible to run with an IBM 3494 VTS, an IBM 3592 J70 Controller, and an IBM 3592 Tape Drive to attach to a System z environment.

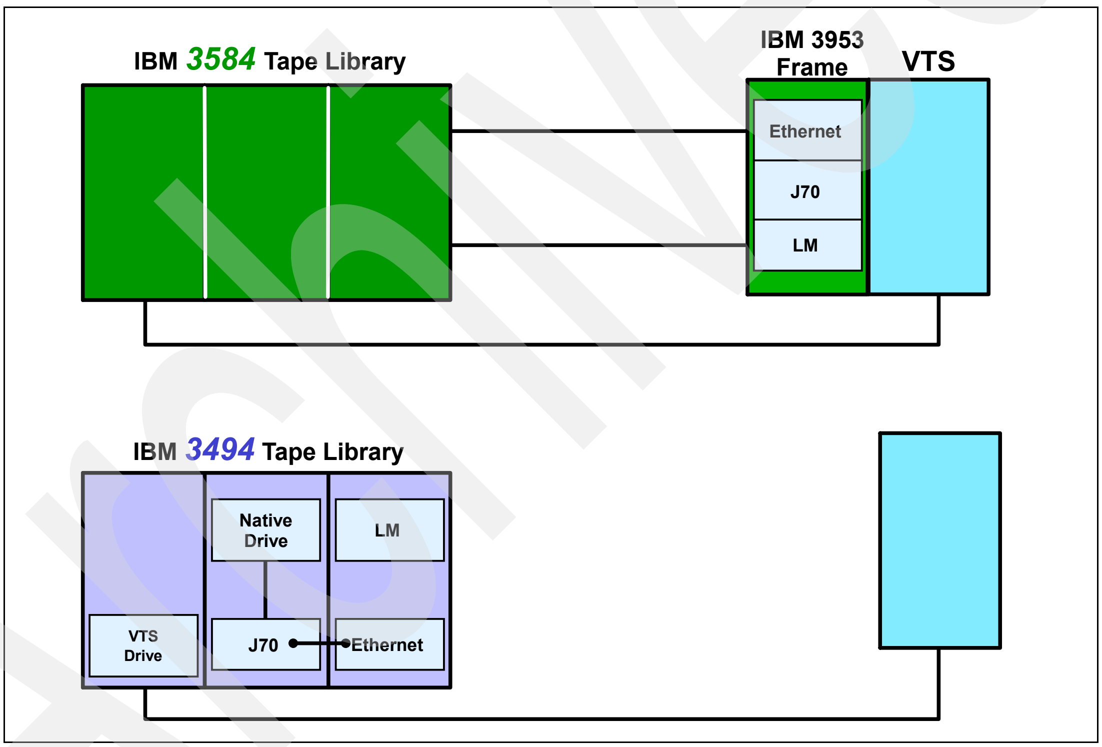

Figure 2-8 IBM TS3500 and IBM 3494 Tape Library connections

2.1.3 The IBM 3494 Model B10 and B20 VTS

The Peer-to-Peer Copy features enable the Model B10 or B20 VTS to operate in a Peer-to-Peer VTS configuration. For installation of the Copy features, each Model B10 or B20 VTS in the Peer-to-Peer configuration must fulfill these requirements:

- Each Model Bxx VTS must be installed in a separate IBM 3494 library or TS3500/3953 library.
- If you have four VTCs, four ESCON channels are required. If you have eight VTCs, you need additional ESCON channel attachment features, giving up to a total of sixteen channels. FICON channels are supported through either four or eight VTCs.

- With 36-GB disks installed, the TVC capacity can be 216 GB, 432 GB, 864 GB, or 1,728 GB. For more details on the available TVC configurations, please refer to “Tape Volume Cache features” on page 67.
- At least four IBM TotalStorage 3590 tape drives must be installed in the associated Model D12 frames. Four 3590s is the minimum for B10, six 3590s is the minimum for B20.
- If you are going to attach 3592 tape drives which are installed in the model D22 frame to the VTSs, at least four drives are required for the model B10 or B20. For supported configurations and an intermix of IBM 3590 and 3592 drives attached to the same VTS, refer to 3.1.10, “IBM TotalStorage 3592 tape drives Model J1A and E05 in a 3494 library” on page 72.
- SCSI attachment on any PtP candidate VTS is not supported and, if installed, needs to be removed for migration to a Peer-to-Peer VTS.
- The Import/Export function provided as part of the Advanced Function feature (feature 4000) on the Model Bxx VTS is not supported for a Peer-to-Peer VTS and, if installed, is disabled during migration to the Peer-to-Peer VTS.
- Advanced Policy Management (FC4001- FC4004) is an optional feature that provides a range of Library Manager policy configuration options, which have to be defined as identical in both attached Library Managers, and which enable you to:
  – Control PtP copy mode at the logical volume level via Management Class definitions
  – Do a selective copy of those logical volumes that request a copy
  – Create a VTS backup of a logical volume via Management Class definitions
  – Use Volume Affinity via Storage Group definitions
  – Control cache residency via Storage Class definitions
  – Control the logical volume size, from 400 MB up to 4000 MB
  – Use the secure date erase function for selective logical volumes in specific groups

In addition, we strongly recommend that the two VTSs have the same cache size, same number of physical drives, and same model of physical drives.

2.1.4 The Peer-to-Peer VTS Copy features

The Peer-to-Peer Copy features enhance the functionality of the Model Bxx VTS by providing the capabilities to combine two standalone VTS systems into one Peer-to-Peer VTS. Depending on the TVC size installed in the VTSs, one Copy feature (FC4010) to four Copy features (FC4010, FC4011, FC4012, and FC4013) are required. The Copy features provide the following functionality:

- Both Model Bxx VTSs in a Peer-to-Peer configuration store and manage the data (known as *tokens*) required to keep track of the status of a virtual volume.
- Additional internal control commands enable the Model Bxx VTS to process these tokens and to enable message broadcasting between the VTCs.
- Additional data transfer commands allow compressed large block transfers for the copy operation.
- Additional virtual devices are defined internally to be used for message broadcasting and copy operations.
- Tape volume cache usage is enhanced by providing expedited copying to the stacked volumes of the secondary copy, thus freeing up space in the tape volume cache.

One of the VTSs in a Peer-to-Peer configuration functions as a focal point for commands relating to volumes and drives in the Peer-to-Peer VTS. For more details, see 2.2, “Terminology and definitions” on page 24.

2.1.5 IBM TotalStorage Virtual Tape Frame (Model CX0 or CX1)

The IBM TotalStorage Virtual Tape Frame (Model CX1) provides the housing and power for two or four Virtual Tape Controllers. There are two power control compartments, each with its own power cord to allow connection to two power sources. The frame dimensions are the same as the Model CX0 frames.

2.1.6 AX0 virtual tape controller versus VTC

The IBM 3494 Model AX0 virtual tape controller is the original Virtual Tape Controller (VTC) component introduced with Peer-to-Peer VTS. Since FICON support was introduced, the virtual tape controllers became features of the associated CX1 frame. The *original virtual tape controllers* (3494-AX0) may still be used in ESCON environments only, and the CX1 features, ESCON or FICON attached, are called VTC. Figure 2-9 highlights the fact that the VTC is not installed in the B10 or B20 frame, but rather in a separate frame known as the CX1 frame.

A VTC in the Virtual Tape Frame 3494 Model CX1 provides interconnection between two VTSs with the Peer-to-Peer Copy features, and provides two host attachments for the PtP VTS. There must be four VTCs in a Model B10 and four or eight VTCs in a Model B20 PtP VTS configuration.

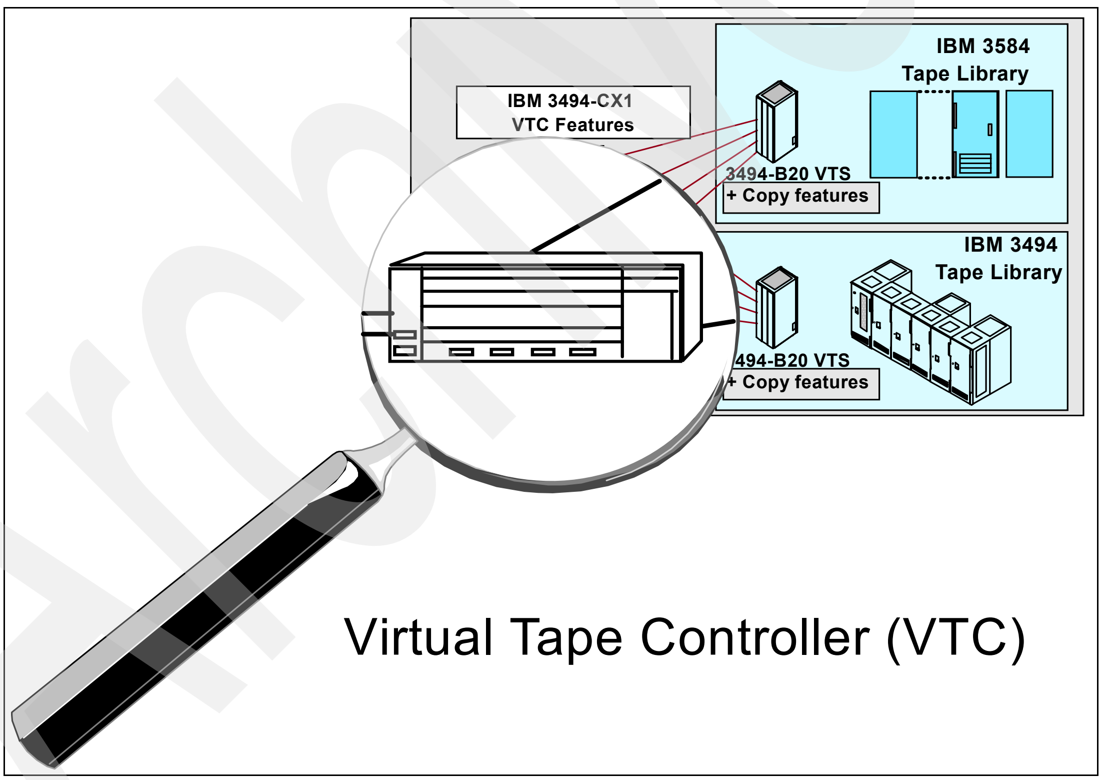

Figure 2-9 IBM 3494-CX1 Virtual Tape Controller

Each VTC provides a total of four ESCON or FICON channel attachments:

- Two ESCON/FICON channels attach to the host systems. The host systems no longer attach directly to the VTS, but to the VTCs instead.
- Two ESCON/FICON channels attach to VTSs in the Peer-to-Peer configuration. Each VTS attaches to each of the VTCs through one ESCON/FICON channel.
- Any mixture of ESCON and FICON within one single VTC is not supported.

2.1.7 IBM 3494 VTC virtual tape controller

The IBM 3494 VTC is a component introduced for the Peer-to-Peer VTS. Figure 2-9 highlights the fact that the VTC is not installed in the B10, or B20 frame, but rather in a separate frame known as the CX1 frame. The CX1 frames introduced feature codes for ESCON and FICON VTCs:

- Model conversion of CX0 to CX1 is allowed.
- Allows CX1 to contain two old AX0's and either two ESCON or two FICON VTCs.
- ESCON and FICON VTC combinations are allowed in CX1 frame B10 attached only.

A single VTC presents itself to the host as one or two 3490E tape control units each with16 3490E tape drives. In order to allow a single host to allocate 64,128 or 256 virtual device addresses of a Peer-to-Peer VTS, the host must be attached to all four or eight VTCs in the Peer-to-Peer configuration. The attachment options are:

- Direct attachment
- Attachment through ESCON directors
- Attachment using EMIF channels
- Attachment using the IBM 2029 Fiber Saver or similar DWDMs
- CNT or Inrange channel extenders (ESCON only)
- Attachment through FICON directors (as depicted in Figure 2-10)

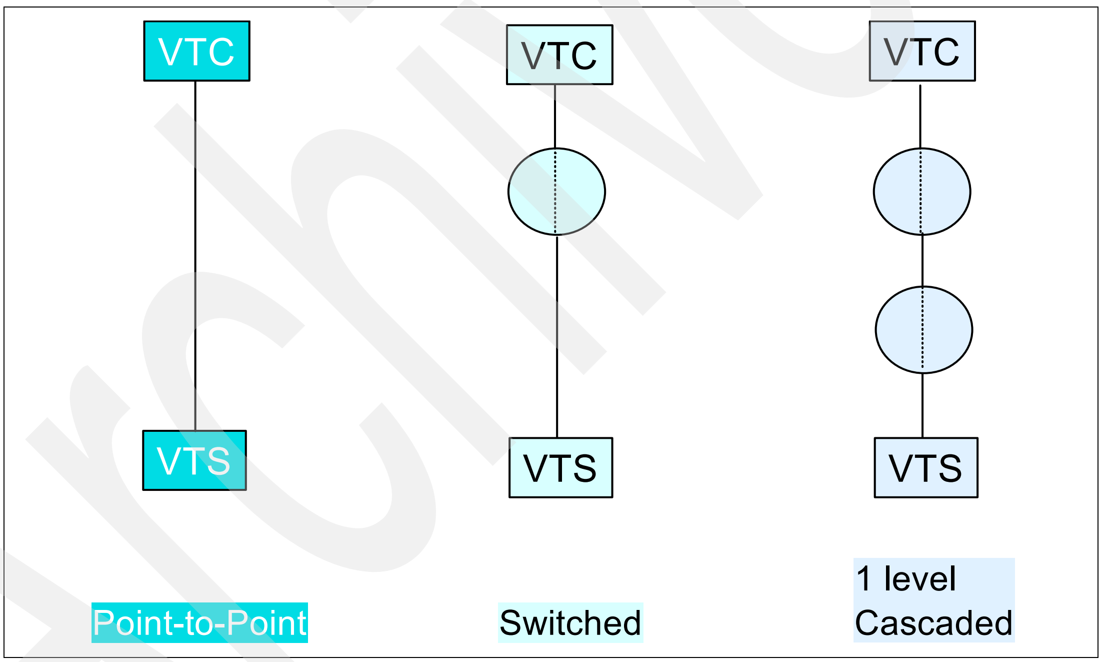

Figure 2-10 PtP VTS FICON attachment

> **Note:** Dense Wave Division Multiplexer (DWDM) is supported for the PtP VTS. Time Division Multiplexer (TDM) and Coarse Wave Divisions Multiplexer (CWDM) are not supported.

All communication between the System z host systems and the Peer-to-Peer VTS (including the distributed libraries) is through the VTC channel attachments (Figure 2-11). The VTCs combine the status of the underlying physical libraries to provide a logical view of the Peer-to-Peer VTS to the host.

There is a SMIT option available to set for high integrity fabric with a single switch or cascaded switches.

Cascading support is currently provided by:

- McData with McData
- Brocade with Brocade
- CISCO with CISCO

No intermix is supported at all. Neither different switch models nor different vendors' switches can be intermixed.

The VTCs use a broadcast mechanism to communicate with the other VTC in the Peer-to-Peer VTS configuration. The communication is through the VTSs. Each controller periodically broadcasts heartbeat information to the other controllers.

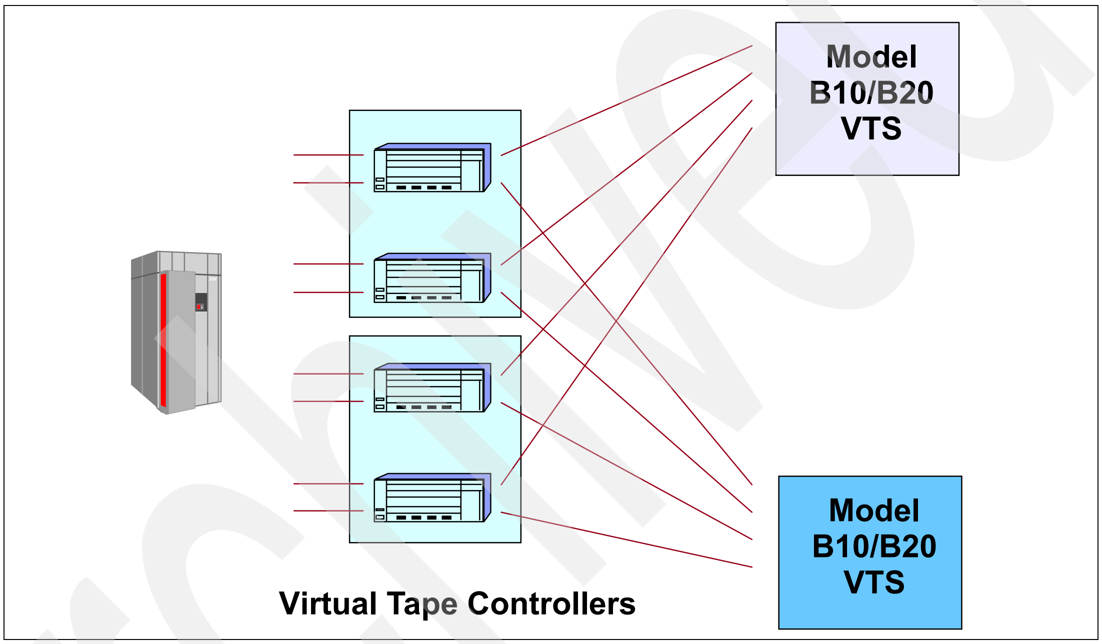

Figure 2-11 VTC channel attachment

Each VTC controller in a Peer-to-Peer VTS:

- Maintains synchronization of the secondary copies of logical volumes and maintains volume integrity between the two VTSs
- Creates logical volume copies using chained 64 KB block transfers of compressed logical volumes
- Selects, for each mount request, which VTS will process that mount request from the host
- Operates in non-preferred or in preferred mode when selecting the VTS on which to mount a virtual volume in response to a host mount request. The VTC:
  - Balances workload across both VTSs by directing host mounts to the VTS that is least busy, when operating in non-preferred mode
  - Directs the mount to the specified VTS when operating in preferred mode
- Directs all activity to the remaining VTS, if one of the VTSs is offline
- Resumes secondary copy operation to synchronize both VTSs after one of the VTSs has been brought online.
- Provides status displays on a customer-selected Web site via the customer's intranet

In addition to the tasks that each single VTC performs, there are some tasks that require more than one VTC. All VTCs in a Peer-to-Peer VTS, as a group:

- Elect one of the VTSs to be the preferred Master VTS of the Peer-to-Peer VTS
- Control the switch over to the other VTS, if the Master VTS fails
- Work in the same mode of operation for creation of secondary copies:
  - In *immediate copy mode*, where the secondary copy of a virtual volume is created, before completion of rewind/unload of the primary copy of the virtual volume
  - In *deferred copy mode*, where the secondary copy of a virtual volume is scheduled after completion of rewind/unload of the primary copy of the virtual volume

If your Peer-to-Peer VTS configuration has the Advanced Policy Management (APM) feature (FC4001-FC4004) installed, then the VTC performs additional tasks. These include:

- Ensuring that the construct names for logical volumes are equal on both LMs
- Ensuring that the logical volume properties (category and Advanced Policy Management constructs) are balanced across the two VTSs in the PtP

The advanced functions constructs are maintained by the VTC as a property of the logical volume, exactly as for the volume category. The token databases within the VTS are used to track the consistency of the constructs across the Peer-to-Peer configuration.

> **Note:** The VTCs have no visibility of the defined actions of the constructs in each LM.

Should you upgrade your current PtP VTS and enable APM, the VTCs will use specific channel commands to reconcile all the attributes of logical volumes only once all the components of the PtP complex indicate that they now support this advanced function. This is true even if the host does not have the PTFs for this support installed. Once the upgrade is complete, a non-APM capable VTC attempting to join the configuration will be rejected.

We explain the copy modes in detail in 2.3, “Creating the secondary copy of a virtual volume” on page 26.

Each Virtual Tape Controller (VTC) is an independently operating distributed node within the Peer-to-Peer VTS and continues to operate during the planned or unplanned outage of another VTC. For uninterrupted data access, at least one of the VTCs must be active. There is no way to access volumes stored in a PtP VTS without a VTC.

Standby Virtual Tape Controller

An additional 3494 Model CX1 with four VTCs can be installed as standby for four active VTCs in a PtP VTS configuration. The Standby VTCs at the remote site are intended for use where the VTSs are physically separated to provide disaster backup. If all the VTCs at the primary site become unavailable, then, with intervention by an IBM System SSR (SSR), the remote VTS and the Standby VTCs can be started in Read/Write Disconnected mode of operation to provide read and write access to volumes stored in the VTS. A host system must be configured and available to the remote site.

2.2 Terminology and definitions

With the introduction of the Peer-to-Peer VTS, changed processes and mechanisms are required to ensure virtual volume and data integrity as well as resynchronization of the Peer-to-Peer VTS after planned or unplanned outages of the different components. In this section, we explain the specific terms that are used with the Peer-to-Peer VTS.

Here we list and define the terms used with the Peer-to-Peer VTS as shown in Figure 2-12 on page 25.

2.2.1 Composite library

The composite library is the logical image of the Peer-to-Peer VTS, which is presented to the host. The host sees one logical tape library with four, eight, or sixteen 3490E tape control units, each of them with sixteen IBM 3490E tape drives and attached through two ESCON or FICON channel attachments. The composite library presents the same image to the MVS host as a single VTS does. It is defined to the host in a similar way to the definition of a stand-alone VTS. Note that the host does not have any knowledge of the VTCs and they do not need to be defined to the host.

2.2.2 Distributed library

A distributed library (or distributed VTS) is a physical IBM 3494 or TS3500/3953 library in the Peer-to-Peer VTS, which includes a Model B10 or B20 VTS with the Copy features installed. Two distributed libraries are required in a Peer-to-Peer configuration. 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 ISMF panels and has no tape devices defined; the tape devices were defined for the composite library.

2.2.3 Master VTS

The Master VTS is one of the VTSs in a Peer-to-Peer configuration. The designation of a Master VTS is necessary to serialize access to logical volumes within the Peer-to-Peer VTS. Certain commands must be sent to the Master VTS to ensure data integrity and volume consistency within the subsystem.

- These commands are referred to as *master required commands*:
  - Commands that cause a volume to be queued for a mount
  - Commands that change the Library Manager volume category assignment of a volume
  - Commands that change the device category assignment
  - Requests for volume or device data

The Master VTS is the focal point for synchronizing the mounting of volumes to virtual drives, and for the selection of a volume serial number (volser) from a scratch category in response to a scratch mount request. Most display commands from the host are also directed to the Master VTS, except for commands that are processed by the UI distributed library (insert category data). The Master VTS is determined at initialization of the Peer-to-Peer VTS by the VTCs. If one of the commands above cannot be processed through the Master VTS, a successful switchover must be completed and the command issued to the *new* master before the command can be completed.

With Advanced Policy Management (FC4001-4004), the copy mode for the logical volume is also controlled by the Master VTS/LM setting. This is true even if the other VTS is the I/O VTS. A diagram of this is shown later in this chapter, in Figure 2-24 on page 48.

2.2.4 I/O VTS

The I/O VTS is the VTS that processes the host I/O commands (such as read and write commands) for a certain virtual volume. The I/O VTS is determined by the VTC during mount processing (load-balancing function). Both VTSs may operate as I/O VTSs for different sets of volumes, determined at mount time for a logical volume. Hence, one VTS may be the I/O VTS for one mount of a logical volume and the other VTS could be chosen as the I/O VTS for a later mount, assuming that the volume is no longer in cache.

2.2.5 UI distributed library

The user interface (UI) distributed library is the distributed library that allows you to insert logical volumes into a Peer-to-Peer VTS. The time and date of all VTCs is synchronized with the date and time of the Library Manager of the UI distributed library. The UI distributed library is defined at installation time and can be either of the distributed libraries; typically, it would be the local library.

Figure 2-12 shows an example of how these terms apply to the hardware components of a Peer-to-Peer VTS with four VTCs.

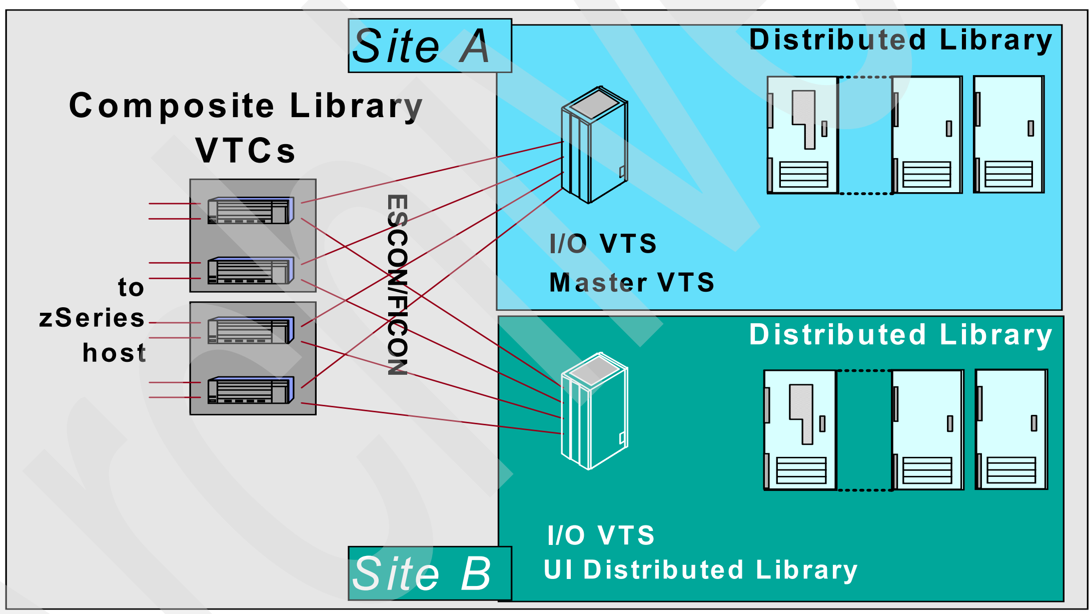

Figure 2-12 Peer-to-Peer VTS specific terms

The example in Figure 2-12 shows that the *distributed libraries*, located at site A and site B, are connected in a Peer-to-Peer VTS configuration. The VTCs present the logical view to the host as one *composite library*. The distributed library located at site B has been defined to function as the *UI distributed library*. The *distributed VTS* at site A has been determined to be the *Master VTS*. Both distributed VTSs are working as *I/O VTSs*.

In addition to the terms shown in Figure 2-12, the definitions in the following sections are used when describing the processes and components of a Peer-to-Peer VTS configuration.

2.2.6 Broadcast device

In each distributed VTS, four internal device addresses are defined and used by the VTCs to send broadcast messages via the distributed VTS. All VTCs together use one shared pool of four or eight broadcast devices per distributed VTS.

2.2.7 Copy device

In each distributed VTS, twelve internal device addresses are defined and used to perform the volume copies. Each VTC has three out of a total of twelve or twenty-four available copy devices assigned. Copy devices are not shared between the VTCs.

2.2.8 Token

A token is the data that is stored in the databases of both distributed VTSs, and that is used to keep track of the level and status of each volume within the Peer-to-Peer VTS.

Token reconciliation

Token reconciliation is the process of comparing tokens stored in each VTS to determine required copy updates and token synchronization. Token reconciliation is performed at power up and at periodic intervals.

Hot token list

The hot token list is held by each VTC. A hot token is created when a change has occurred to a virtual volume and the token could not be updated on one of the distributed VTSs. The hot token list is maintained to ensure that the latest version of a virtual volume is used even if one of the distributed VTSs is not operational. Refer to 6.3.7, “Logical volume status” on page 234, for a sample display of this information.

Token protected

This is a set of data transfer and tape control commands that cannot be performed until the tokens on both distributed VTSs have been updated. This prevents the Peer-to-Peer VTS from using a down level version of a logical volume if a failure or switchover has occurred.

2.2.9 Disconnected volume

A disconnected volume is a virtual volume that is in the process of having data on it changed, which causes a version of it on one distributed VTS to be down level until a secondary copy has been created.

Primary copy

This term is used to describe the primary copy of the logical volume. It is created according to host definitions.

Secondary copy

This term is used to describe the second copy of the logical volume.

2.3 Creating the secondary copy of a virtual volume

When a scratch logical volume is written by a host application in a PTP VTS configuration, it is first written to one of the VTSs in the configuration and then copied to the other one. That first VTS is typically referred to as the I/O VTS. The I/O VTS is determined by the Virtual Tape Controller (VTC) based on such factors as availability of both VTSs, workload on each, and customer specified defaults set during the installation of the system.

The mode by which the copy is made, immediate or deferred, is determined by the customer specified default, immediately during close processing (immediate copy mode), or asynchronously at a later time (deferred copy mode). Without Advanced Policy Management, immediate or deferred copy modes of operation are configured by the IBM System Service representative (SSR) during setup of the VTC and all new or modified data becomes copied to the second VTS.

The default copy mode, set when the VTCs are shipped, is immediate copy mode. The environmental copy mode of operation may be changed at any time; however, this must be done by a SSR using a service interface to a VTC. Any LM definitions for Management Class will override the VTC copy mode.

With VTS R7.2, the PtP selective dual copy function was introduced. It enables the individual selection of logical volumes that do not require a copy. The default mode can be overridden by an action specified for the Management Class construct assigned to the volume.

With the PTP Selective Dual Copy function, additional actions can be specified during the definition of a Management Class constructs to allow the Management Class to specify:

- Which VTS is to be used in writing data for a scratch mount
- That no copy is to be performed

These are in addition to the existing actions that specify that a second copy of the data is to be performed in the same VTS and the dynamic copy mode selection described above.

> **Important:** Running your PtP VTS under GDPS control, duplicates will be created regardless of any APM definitions.

All VTCs in one Peer-to-Peer VTS must use the same mode of operation. A mode change requires that all of the virtual devices of each VTC must be varied offline to allow the SSR to bring down a VTC, change the mode, and restore the VTC to normal operation.

> **Note:** The customer should ensure that the PtP copy mode control definitions and the selective dual copy definitions are the same for both distributed VTSs.

We describe planning considerations of the Advanced Policy Management (APM) in more detail in 3.2, “Planning for Advanced Policy Management exploitation” on page 82.

2.3.1 Immediate copy mode

In immediate copy mode, the copy of the new or modified logical volume to the other VTS (Peer-to-Peer VTS copy) is made prior to the completion of the rewind/unload of the original logical volume. When the volume completes close processing, the Peer-to-Peer VTS has completed performing the copy of the logical volume. No further Peer-to-Peer copy processing is required.

Immediate copy mode of operation (Figure 2-13) ensures that, under normal conditions, two identical copies of a virtual volume exist, as soon as the rewind/unload command has completed. You should consider using immediate copy mode if it is mandatory for your operation to have the highest level of protection for your tape data. Immediate copy mode is the default mode.

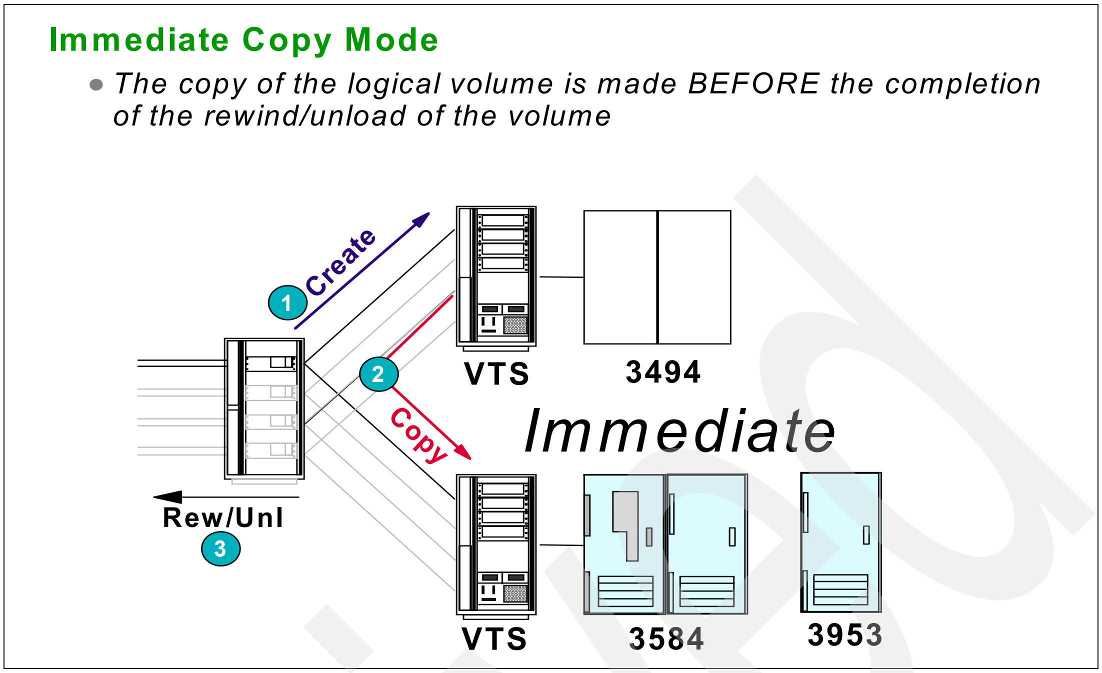

Figure 2-13 Immediate copy mode

In the immediate copy mode of operation in an unconstrained environment, the time required to copy a virtual volume will delay the completion of the rewind/unload command in a ESCON environment by around 80 seconds for an 800 MB volume that is full. This time will increase in a constrained environment and where larger volume sizes are used. With VTS R7.4 you have the choice of additional logical volume sizes by the use of the Data Class constructs. 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 by the use of a Data Class storage construct.

With VTS R7, we introduced some performance improvements for the immediate copy mode for ESCON and FICON installations:

- **Distributed Immediate Copies:**
  - Helps to balance copy workload across all available VTCs
    - Uses available copy devices on other VTCs
    - Reduces Immediate Copy times
  - Takes effect when the Rewind Unload (RUN) command is received on a device:
    - The copy will be put on the local VTC copy queue first
    - If copy devices are not available after 1 second, copy job will be distributed to other VTCs
    - Copy gets done by the first available VTC
    - The VTC performing the copy notifies the other VTCs that copy is complete
    - The originating VTC completes the RUN command
- **Immediate Mode Copy Throttling:**
  - FICON data rates allow host jobs to overpower copy jobs
    - There are up to 32 host devices per VTC but three copy devices only
    - Copies can exceed the 40 minute time-out
  - Throttling gives copy jobs priority over host jobs
  - VTCs collect and send copy workload related statistics to the VTS via Special Token
  - VTS analyzes data received from all VTCs and determines when it is appropriate to throttle host I/O

2.3.2 Deferred copy mode

In deferred copy mode, described in Figure 2-14, the secondary copy is created at some time after the host finished processing the primary copy. When a deferred copy operation is processed, the amount of channel bandwidth it is given between the VTCs and the distributed VTSs is such that it has minimal interference with host channel bandwidth requirements. Host bandwidth requirements are given priority over copy bandwidth needs. This allows the host to use maximum bandwidth during peak periods and the Peer-to-Peer VTS to catch up on the copies during periods of low host activity.

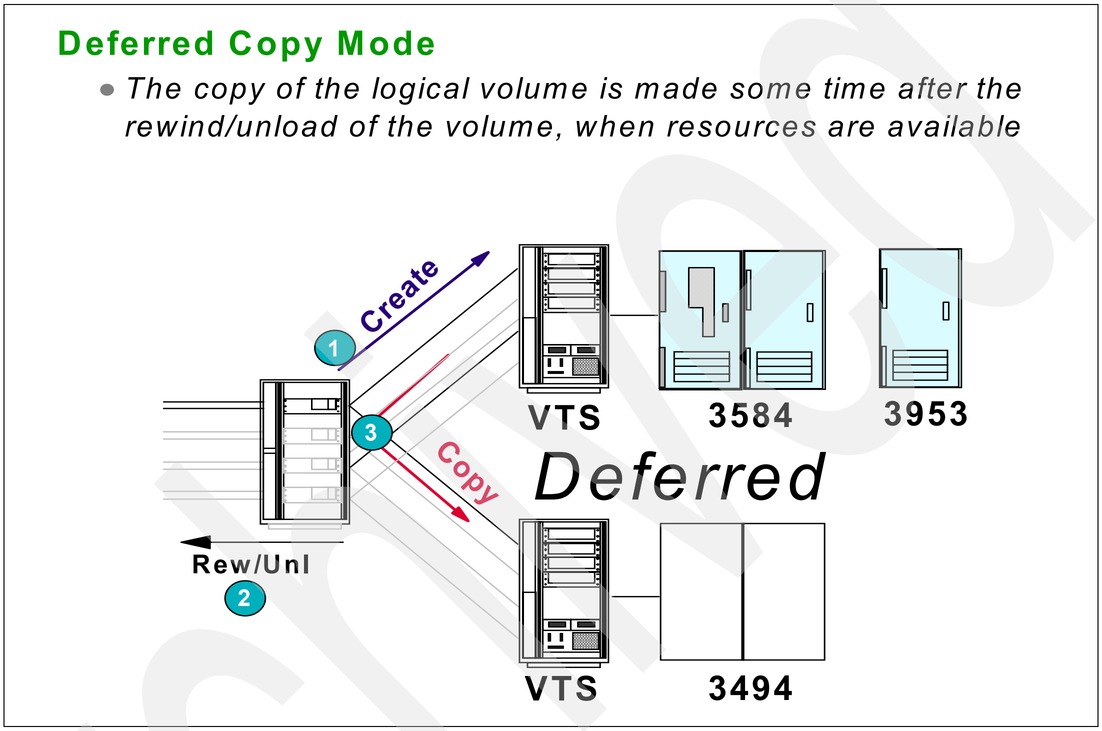

Figure 2-14 Deferred copy mode

The VTC, which has the task of making a copy of a particular deferred volume, checks the time that has elapsed since completion of the host rewind/unload operation for the volume. If this elapsed time is greater than the deferred copy priority threshold, the copy operation will occur with a priority that allows equal contention for processing cycles with current host I/O activity. The resolution of the elapsed time comparison is one minute. If the elapsed time is less than the deferred copy priority threshold, the copy will proceed with the normal deferred copy priority. During periods of high-host activity, deferred copy operations are performed slowly, and the number of queued copy operations grows. When the host workload decreases, the performance of the deferred copies increases, and the number of queued copy operations will decrease.

The low priority for copy operations continues until the space occupied by virtual volumes in the Tape Volume Cache of a VTS reaches a certain threshold. To avoid virtual volumes being removed from the Tape Volume Cache before they have been copied to the other VTS, the deferred copy operations will be given higher I/O priority.

Figure 2-15 illustrates two reasons for increasing the copy priority when operating in deferred copy mode of operation:

- If the Tape Volume Cache reaches a certain threshold for the space occupied by virtual volumes that have not had a secondary copy made in the other VTS, the completion of the rewind/unload commands following host write operations is delayed. This effectively throttles host writes and allows more secondary copies to be completed. This is indicated on the left-hand side of Figure 2-15.

  While the bandwidth available for host I/O operations is decreased during this time, copy operations continue at a higher priority until a lower threshold is hit, which removes the throttling delay described above.

- The copy priority for a single virtual volume is increased after the time defined for the *deferred copy priority threshold time* has passed. This is indicated on the right-hand side of Figure 2-15.

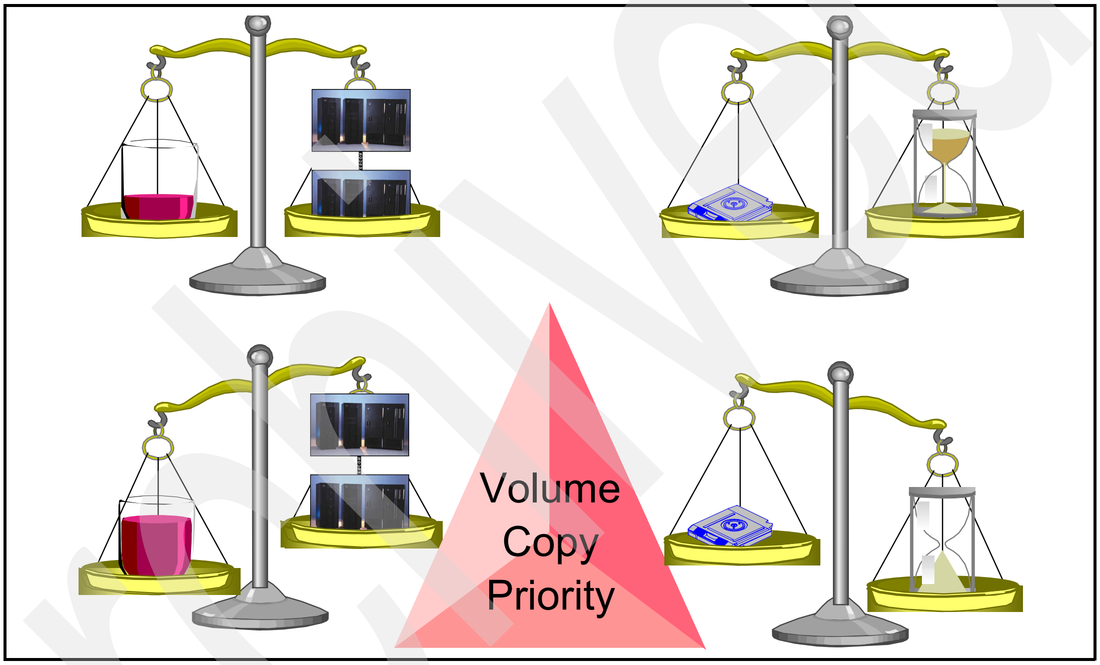

Figure 2-15 Increasing volume copy priority

The default for the deferred copy priority threshold time, indicated by the hourglass in Figure 2-15, is eight hours. It is set by the SSR during setup of the VTCs and cannot be changed dynamically.

You can select a value of zero to 24 hours for the deferred copy priority threshold. This value is set by the SSR.

In order to avoid recalls for virtual volumes that have not yet been copied, the size of the Tape Volume Cache should be configured to be large enough to allow all virtual volumes written or modified during the peak period to stay in the Tape Volume Cache. If the TVC size is too small, this may decrease the overall host throughput to the VTS.

Note: Regardless of the operational mode, the Peer-to-Peer VTS will perform as much of the background work (Peer-to-Peer copies and copies to physical tape) as possible with the resources available at the time. In Deferred Copy Mode, although the Peer-to-Peer copy is initially placed on a lower priority queue than host work, the Peer-to-Peer VTS will perform the copy as soon as resources are available.

2.4 Workload balancing and VTS preferencing

Workload balancing in a Peer-to-Peer VTS is performed on two levels:

1. On the host system, when allocating a tape drive. This allocation determines which VTC is used to perform the mount. We explain the allocation algorithm in 2.4.1, "Host-based workload balancing" on page 31.
2. On the VTC, when selecting the I/O VTS to process this mount. This selection determines where the primary copy of a virtual volume is written to. We explain the selection algorithm, and how you can influence it, in 2.4.2, "Peer-to-Peer VTS workload balancing" on page 33.

The overall objective of workload balancing is to maximize the overall throughput of the PtP tape subsystem from the using host's perspective. The workload to be considered within the Peer-to-Peer VTS is the workload from and to the host as well as the copy operations between both distributed VTSs.

2.4.1 Host-based workload balancing

The host allocation algorithm is important for a Peer-to-Peer VTS. With a stand-alone VTS, it is not necessary to distribute the mounts evenly across all defined virtual tape devices and tape control units to achieve the maximum throughput because all tape drives are available though all defined channel paths. Any control unit path can access the total bandwidth available from the stand-alone VTS. Figure 2-16 shows the logical view for the stand-alone VTS compared to a PtP VTS.

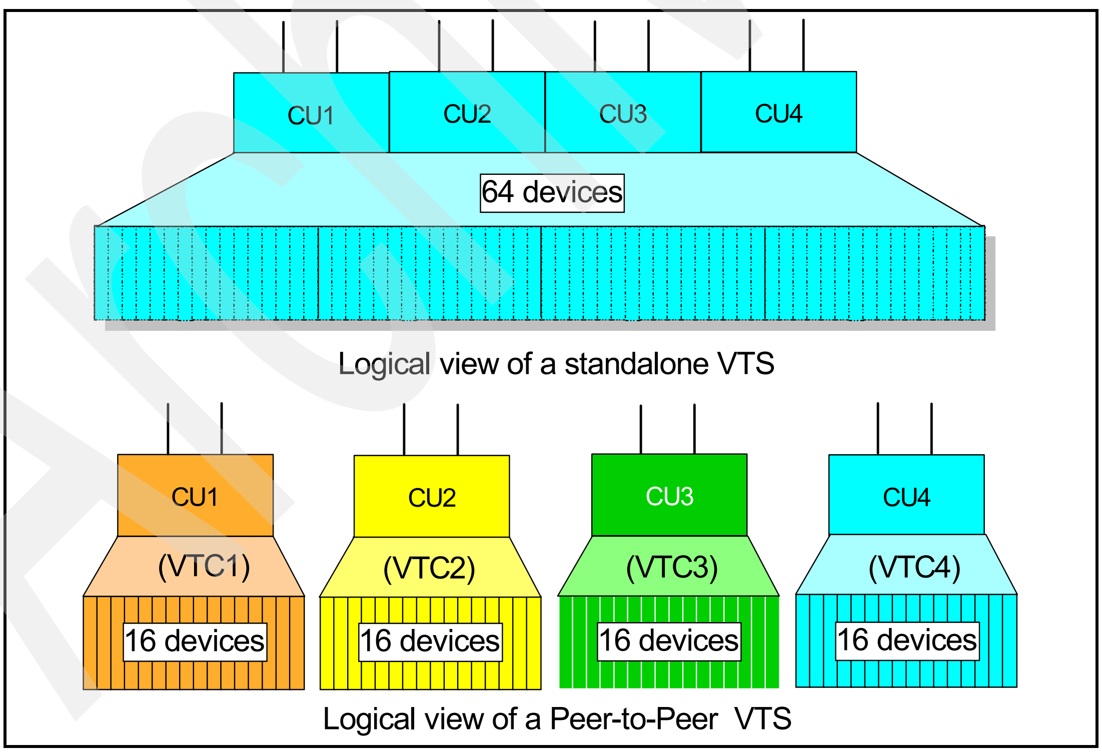

Figure 2-16 Stand-alone and Peer-to-Peer VTS virtual CUs and devices

Each VTC represents a single 3490 tape control unit with two channel paths and 16 tape drives, and only has access to a quarter of the total VTS bandwidth. In order to allow the peak performance of the Peer-to-Peer VTS, the workload must be distributed evenly across the four VTCs. The Peer-to-Peer VTS, in this respect, is similar to four physical tape control units, each of them having their own tape devices and channel paths. Therefore, the host allocation algorithm is important for the overall performance of the Peer-to-Peer VTS. We will look at the implications for JES2 and JES3 environments.

Tape allocation in JES2 environments

Under DFSMS/MVS with JES2 or JES3 with MVS-managed (non-JES3 managed) devices, MVS allocation routines are used for tape drive allocation.

For scratch mount requests, MVS allocation attempts to randomize the use of the drives in a library. For the Peer-to-Peer VTS, this refers to all 64, 128, or 256 virtual device addresses.

Specific mount requests are processed according to the same randomization techniques as a non-specific or scratch mount request. The four-VTC example shown in Figure 2-17 summarizes the process to direct a specific mount request to the VTC with the least drives currently used.

The first VTC (CU1) currently has five virtual drives allocated. It is the last one in the list because it has the most drives allocated.

The second and the fourth VTCs (CU2 and CU4) each have four drives allocated. They are listed in sequence because they have the same number of virtual drives currently allocated.

The third VTC (CU3) has three drives allocated. It is the first in the list because it has the fewest virtual tape drives allocated at this point in time.

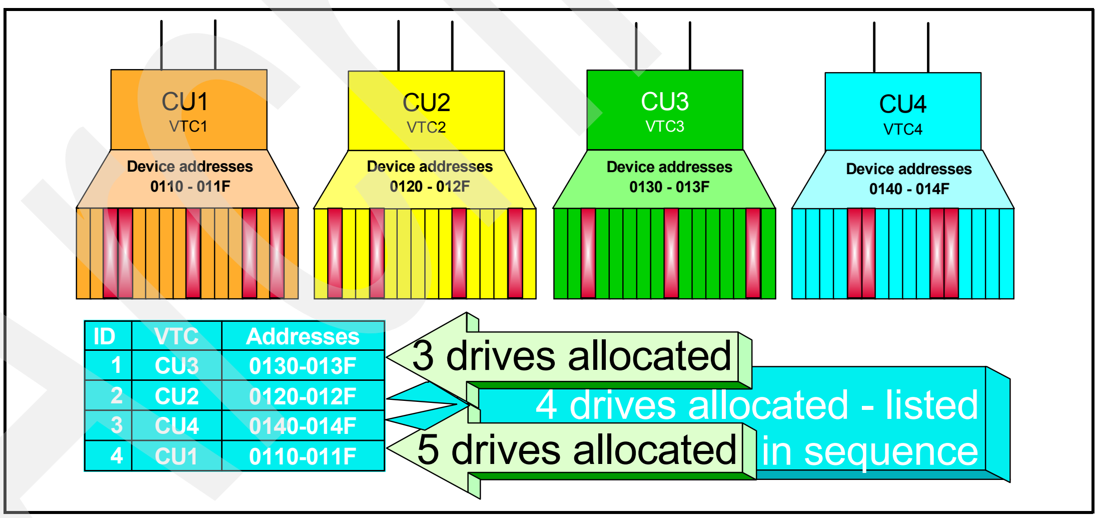

Figure 2-17 Subsystem affinity usage for a JES2 environment

Prior to issuing the mount request for a specific volume, the host obtains a subsystem affinity list of the composite library. The VTC processing the request builds the subsystem affinity list based on the current workload of each VTC.

The VTC with the least number of mounted devices is the first subsystem in the list. MVS will then try and allocate a device from the first subsystem in the list. If a device from that subsystem is not currently available, MVS allocation will try to allocate from the next subsystem in the list. Device addresses are not reported to the host, only the subsystem IDs representing the VTCs. Device addresses are shown for clarification only.

Tape allocation in JES3 environments

In a JES3 environment, where MVS allocation is not used, JES3 attempts to spread scratch mount requests across all library devices.

For specific mounts, subsystem affinity cannot be used. In a JES3 environment, where JES3-managed devices get assigned at pre-execution time before MVS allocation is invoked, a different approach can be used. JES3 can be set up, through the definition of `ADDRSORT=NO` in its INISH deck, to use the devices in the defined order, evenly distributing workload across the VTCs. Figure 2-18 shows the order in which JES3 DEVICE statements must be defined in a four-VTC Peer-to-Peer VTS with 64 virtual device addresses.

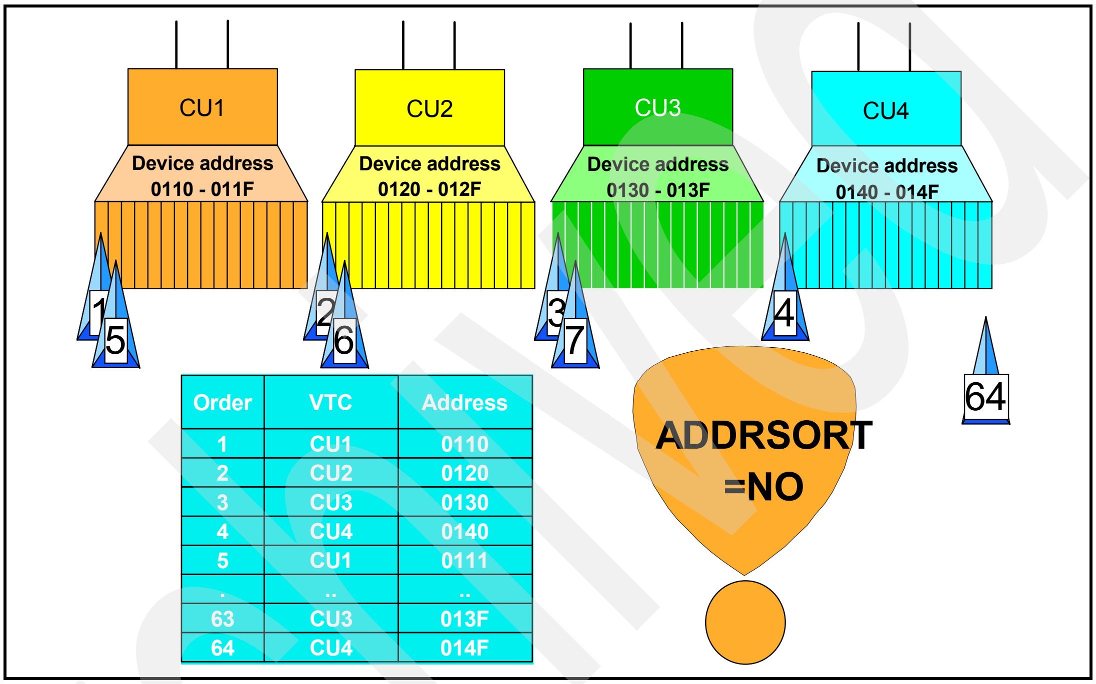

Figure 2-18 JES3 device order for specific mounts

As shown in Figure 2-18, the first device address in the list is 0110 from the first VTC, the second device address in the list is 0120 from the second controller, and so on. The last device address in the list is 014F for the highest device address of the fourth VTC. Based on this setup, specific mounts are spread across all VTCs as the device allocation starts from the top of the list until it finds one that is available.

For detailed information on defining the JES3 INISH deck, please refer to "JES3 initialization deck definition" on page 186.

2.4.2 Peer-to-Peer VTS workload balancing

During processing of a mount request, the VTC selects one of the distributed VTSs as the I/O VTS, which results in the primary copy of a virtual volume being directed to the selected distributed VTS. The selection of the I/O VTS is based on the following information in the listed order:

1. Component availability:

   If a distributed VTS or a path to a distributed VTS is not available, then the VTC will select the remaining VTS if it has valid data on it. This selection criteria will apply, if, for example, one of the distributed VTSs is in service preparation mode.

2. Volume tokens:

   The token data is read from both distributed VTSs as soon as the volser to be mounted is known. If one token indicates that it is the most current one, or that it is the only valid one, then the distributed VTS with this token is selected. If both distributed VTSs have the same level of the virtual volume, then the distributed VTS that has the virtual volume residing in the Tape Volume Cache will be favored.

3. Peer-to-Peer VTS preferencing:

   Please refer to 2.4.3, “I/O VTS preferencing” on page 34, where we explain in detail how you can influence the selection of the I/O VTS by preferring one distributed VTS over the other one.

4. VTS loading factor:

   The VTSs provide an integer value that indicates their current loading (VTS loading Factor). The VTC obtains the information from the VTSs when it obtains the token data. The value is calculated by each VTS approximately every minute. A VTS with a smaller value is performing less work than a VTS with a larger value. The loading factors for both distributed VTSs are compared to see which one is less busy and should attempt more of the work.

5. Count of mounts requiring recall:

   Each VTC keeps track of the number of mounts requiring recalls that are in progress on each of the distributed VTSs. If a mount request is for a specific volser, and the volser is not in the cache on either VTS, then the mount will be considered a recall mount, and the distributed VTS having the least recall mounts in progress will be favored.

6. Count of total mounts:

   Each VTC keeps track of the total number of mounts it has in progress or completed on each of the distributed VTSs. Based on the information from its own count, the VTC favors the VTS with the least number of mounts in progress as the I/O VTS.

2.4.3 I/O VTS preferencing

The VTCs in a Peer-to-Peer configuration can be located up to 100 km away (75 km with ESCON) from the hosts. The VTSs can be located up to 100 km away (50 km with ESCON) from the Model CX1 frames with the installed VTCs. Based on the distance specifications, two basic installation options are available:

Local installation

Both VTS subsystems and all VTCs are installed locally. All VTC are installed in at least one CX1 Auxiliary Frame. All hosts systems attach to all VTCs. This scenario is shown in the top part of Figure 2-20 on page 36.

Remote installation

The Peer-to-Peer configuration is installed in two separate sites. Each site has one VTS subsystem, at least one CX1 auxiliary frame, and four VTCs installed for a total of eight across the two sites. Host systems attach to the VTCs installed at their site. This scenario is shown in the bottom part of Figure 2-19.

Alternatively, you can install four or eight VTCs at the local site and have four additional Standby VTCs at the disaster site. This is described in 9.3, “Disaster Recovery testing with Standby VTCs” on page 341.

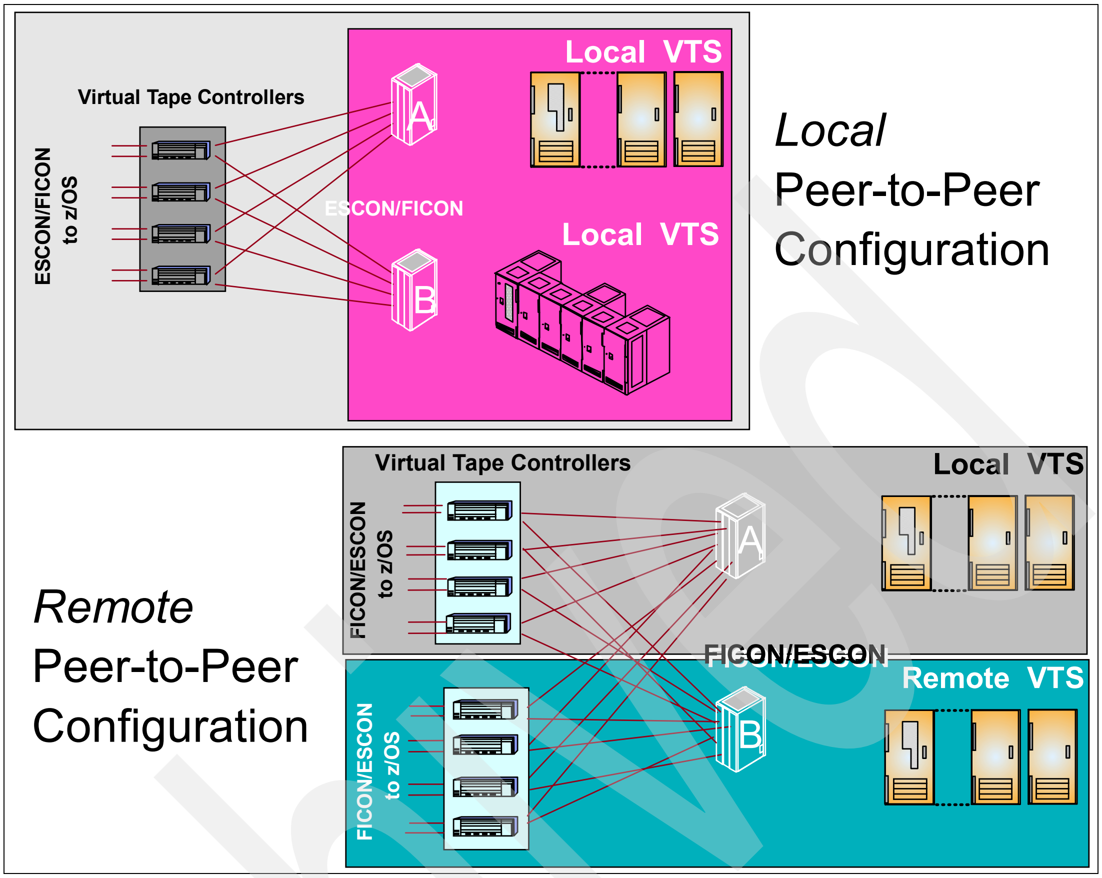

Figure 2-19 Local and remote Peer-to-Peer configuration

Note: For availability reasons, every host system using a Peer-to-Peer VTS should always attach to more than one VTC.

In remote Peer-to-Peer VTS configurations, you may wish to direct the primary copies of virtual volumes from a host at one site to the distributed VTS at either the same site, or to one at the other site. Although all VTCs of one Peer-to-Peer VTS must be operating in the same copy mode of operations, they can be set up differently for VTS I/O preferencing. Each VTC can be set up by the SSR to operate with no-preference or with preference of either VTS0 or VTS1 as the I/O VTS.

If all components of the Peer-to-Peer VTS are available, and if both distributed VTSs have a valid copy of the virtual volume, and if none of them or if both of them have a copy in the Tape Volume Cache, the I/O VTS will be selected based on the VTS preferencing setup you choose for the VTC.

No-preference selection

In this case, no specific distributed VTS is chosen. The next selection criterion applies, which allows the selection of the distributed VTS, which is less busy.

No-preference mode allows the spreading of the workload across both distributed VTSs and thus allows the maximum performance for the Peer-to-Peer VTS. The load gets balanced and therefore it is often called balanced mode. This mode should be chosen for Peer-to-Peer VTS configurations where both distributed VTSs are located at the same distance from the VTCs and the hosts. This applies to local and remote configurations and is shown in Figure 2-20 for a four-VTC Peer-to-Peer configuration.

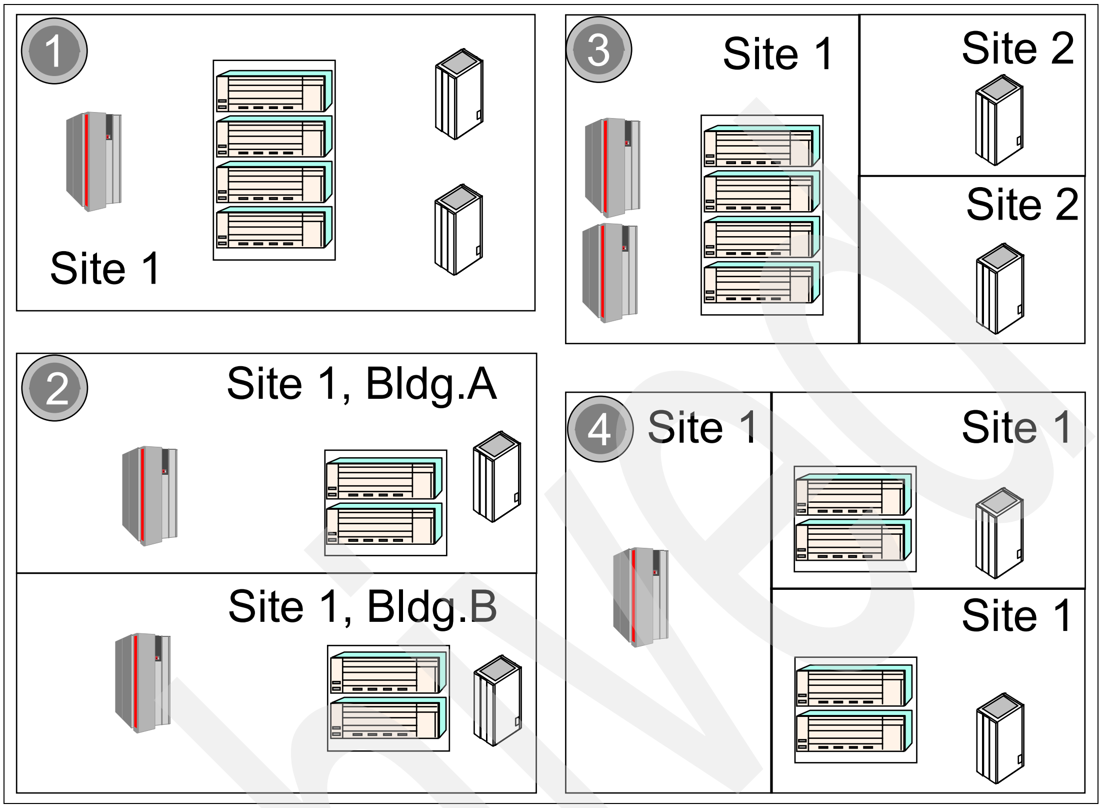

Figure 2-20 Scenarios for non-preferred VTS I/O selection

The scenarios shown in Figure 2-20 are theoretical scenarios which we describe to illustrate in which situations we would recommend using the no-preference I/O VTS selection:

1. In the first scenario, all components are located on the same site, basically in the same building such that no performance degradation would be expected.
2. In the second scenario, components are split into two locations on the same site, still within a distance that does not cause ESCON or FICON distance-related performance degradation.
3. In the third scenario, the host and all VTCs are installed at one site, whereas both distributed VTSs are installed at another site in different buildings the same distance from the host.
4. In the fourth scenario, the host is located at the same distance from both distributed libraries. The components are located in different rooms or buildings on the same site.

Where there are significant differences in distance between the components, using no-preference I/O VTS selection may not be the best choice.

Preferred VTS selection

Preferring one VTS over the other may be suitable in remote Peer-to-Peer configurations, where the distances from a VTC to the two distributed libraries are significantly different. Taking into account that ESCON performance degradation is expected to be close to two percent per kilometer (after the first nine km) of ESCON distance, and that the primary copy of a virtual volume is created sending uncompressed data to the distributed VTS, the overhead may be significant if the remote distributed VTS was selected as the I/O VTS from the local VTCs. FICON performance degradation is considered to be minimal over distances up to 100 km, and it is expected that overhead will be less of an impact if the remote distributed VTS was selected as the I/O VTS.

Figure 2-21 shows two scenarios where preferred I/O VTS selection may be suitable.

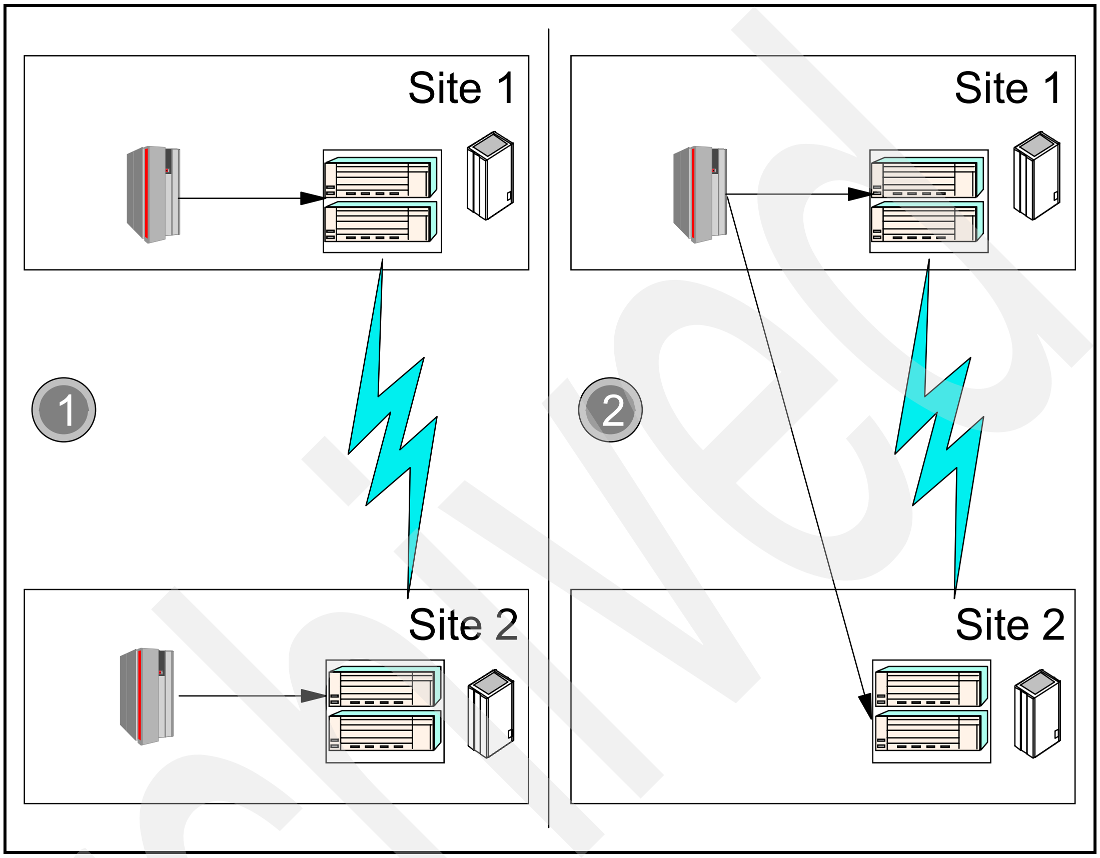

Figure 2-21 Scenarios for preferred I/O VTS selection

1. The first scenario shows a split configuration, basically a mirrored data center, where the host at each of the sites is only connected to the VTCs at the same location.

   The deferred copy mode of operation is recommended. If the secondary copy is created at a remote location, the distance may increase the time until rewind/unload completes, and it may prolong job run times significantly.

2. The second scenario shows one host connected to all VTCs at both the local site and the disaster recovery site. The same considerations as in the first scenario apply. You may consider using preferred I/O VTS selection (preferring the local distributed VTS) and use one of the business continuance enhancements introduced with code level 2.26.xx.x.

2.4.4 Primary mode within a GDPS environment only

The *primary VTS* is running in the VTS primary operation mode. This mode introduces some differences as compared to the "preferred" operation mode. For example, a primary VTS receives all host tape mounts and handles all read and writes while a VTS PtP in preferred mode allows host-VTS read I/Os to the secondary VTS. Let us assume a job with a modification request (`DISP=MOD`) onto a logical volume. The logical volume is stored in the cache of the secondary VTS but not in the primary. In a GDPS environment with a VTS PtP running in primary mode, the I/O accesses the primary VTS, which requests to copy the logical volume from the cache of the secondary VTS to the primary VTS before it is provided to the host for modification. In contrast, a PtP VTS in preferred mode would allow the host to access the secondary I/O VTS and provide the logical volume directly to the host for modification.

2.4.5 Business continuance enhancements

For customers who have installed or are considering the installation of a PTP VTS subsystem for business continuance, the standard management algorithms of the subsystem may not meet all of their business or recovery needs. To better meet those needs, the following aspects of the PTP VTS management algorithms can be tailored by the IBM SSR. Some of the controls are specific to each VTS and some are specific to each VTC.

Forcing scratch mounts to preferred I/O VTS

There are certain conditions which will result in a VTC selecting the non-preferred I/O VTS for the I/O VTS during scratch mount processing. The most likely condition is when the Delete Expired Volume Data function is incorrectly set up, resulting in the local volume being deleted prior to the volume at the remote VTS. The VTC is then forced to select the remote VTS for I/O operations because it has the only valid version of the volume. This places additional demands on the bandwidth of the communication links between the local and remote sites, which may increase the amount of time it takes to perform the workload. Also, in the event of a communication failure after the volume was created on the remote VTS, and before it is copied back to the local, it cannot be accessed by the local host until communications are restored. An option is needed to force the I/O VTS selection to the preferred I/O VTS independent of the status of the volume on either VTS.

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

- The default is to continue to select an I/O VTS based on VTS availability and logical volume version validity.
- If a VTS is specified as the Preferred VTS for I/O, an additional option forces the system to scratch all scratch mounts in order to use the VTS for I/O if that VTS is available.

If the preferred VTS is unavailable, then all I/O will go to the other VTS.

If the volume to be used to satisfy the scratch request is not valid or does not exist in the preferred I/O VTS, the VTC will create the volume as a re-initialized volume. This may cause problems with tape management systems that are not compatible with having a known volume being re-initialized outside of their control.

Accounting for slower remote link in I/O selection algorithm

For specific volume mounts, cache residency is highly weighted when determining which VTS to select as the I/O VTS as it avoids a recall. If the needed volume is resident only in the cache at the remote site, that VTS is selected for I/O even if the local VTS was specified as the preferred I/O VTS. However, if the communication links between sites is limited, a shorter overall job elapsed time may result in recalling the data from the local VTS instead of processing it from the remote VTS.

Based on customer requirements, the IBM SSR can set or modify this control through the service menu on each VTC. See 3.4.3, “Critical factors affecting the PtP performance” on page 93.

Copy Files Preferenced to Reside in Cache

Normally, the caches in both VTSs in a PTP are managed as one to increase the likelihood that a needed volume will be in cache. By default, the volume on the VTS selected for I/O operations is preferred to stay in the cache on that VTS, where the copy made on the other VTS 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 a PTP VTS running in non-preferred mode, both VTS are being selected as the I/O VTS and will have the original volumes (newly created or modified) preferred in cache, while the copies to the other VTS will be preferred to be removed from cache. The result is that each VTS cache is filled with unique newly created or modified volumes, thereby roughly doubling the amount of cache the host sees.

For a PTP VTS used for remote business continuation, particularly when all I/O is preferenced to the local VTS, this default management method may not be desired. In the case where the remote half of the VTS 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 for the remote VTS.

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 as preference group 1 volumes (LRU).

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 VTS is managed.

Recalls Preferenced for Cache Removal

Normally, a volume recalled into cache is managed as if it were newly created or modified, because it resides in the VTS selected for I/O operations on the volume. A recalled volume will displace other volumes in cache. In the case where the remote half of a PTP VTS 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 amount 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. Once a restore has been completed from a recalled volume, that volume is no longer needed. What is needed is to remove the recalled volumes from the cache after they have been 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 VTS:

- 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 VTS is managed.

2.5 Advanced Policy Management (APM)

August 2002 sees the General Availability of the Advanced Policy Management (APM) enhancements for the base VTS and the PtP VTS. Subsequent Microcode releases enhanced and added more functionality. Currently, the following policy management is available:

- Volume pooling with different reclamation policies and secure data erase
- Selective Dual Copy
- Peer to Peer Copy control
- Peer to Peer Copy mode control
- Tape Volume Cache Management
- Extended logical volume sizes
- Export/Import (for the base VTS *only*)

The early versions of the Virtual Tape Server (VTS) do not allow the customer to control how logical and physical volumes are managed within the subsystem. As it is a totally self-managing system with outboard management, there is no need for any external control. Furthermore, some customers want to better utilize the resources of a VTS or have a need to control the management of logical volumes. The Advanced function is available with FC4001 - FC4004. This includes all the features mentioned above. Existing users of FC4000 (Import/Export and Cache management) can upgrade to this feature. The Import/Export function has to be deactivated by the CE before the PtP copy features are installed.

2.5.1 APM implementation

The implementation is based on the expansion of the existing System Managed Storage (SMS) constructs to control these functions. The Automatic Class Selection routines have to be used to select specific constructs for individual mount requests. In addition, the construct names are defined at the Library manager level (outboard), where they are determine the actions to be taken for a logical volume.

Outboard policy management

The enhancement introduces the concept of Outboard Policy Management. This is defined as the functions and components on the IBM TotalStorage 3494 Library manager (LM). It is seen Outboard from the host system and its policies and controls. It is this that enables Advanced Policy Management and the VTS status panel on LM and Display Library command on the host will show if the VTS is capable of Outboard Policy management.

For OPM, a set of panels is supplied both at the LM and the 3494 Specialist. These panels enable the Storage Administrator to perform the following functions:

- Manage Data Classes:
  - Host - ACS routine used to pass media type and recording technology preferences.
  - Outboard - additional logical volume sizes
- Manage Storage Classes:
  - Host - ACS routines used to request SMS management of tape.
  - Outboard - VTS cache management based on IART values.
- Manage Management Classes:
  - Host - No effect. Pass construct name.
  - Outboard - Control of selective dual copy within a Peer-to-Peer VTS and Peer-to-Peer copy modes control for logical volumes.

- Manage Storage Groups:
  - Host - Use ACS routines to enable allocation of drives defined for the VTS.
  - Outboard - Associate logical volumes with a specific pool of physical volumes.

Table 2-1 shows the assignments of the DFSMS constructs with or without Advanced Policy Management installed.

Table 2-1 Pre/post APM SMS class comparison

SSM Construct	Without APM	With APM
Data Class	Assign media type and recording technology preferences	No change. Additional logical volume sizes are available
Storage Class	Request SMS managed tape. IART control of cache management	No change. However the control of cache management is now moved outboard
Management Class	Not used	Control of selective dual copy and Peer-to-Peer copy modes for logical volumes
Storage Group	Enable allocation to the 3494/VTS	No change. However, physical volume pooling enabled

Note: It is highly recommended that the definitions are the same for both VTSs in a Peer-to-Peer configuration.

Control flow

With APM enabled in the z/OS environment for SMS-managed tape, DFSMS construct names are passed to the LM. This flow of data and constructs passed:

- Is invoked with allocation of the first file on a volume:
  - All files on the same volume/set of volumes managed alike
- Uses a pre-ACS routine (already existing) if required:
  - Provides a way for a TMS or other programs to influence ACS routines
- Uses four construct names that are selected via the ACS routines:
  - Data Class name
  - Management Class name
  - Storage Class name
  - Storage Group name
- Is sent to VTS/Library with the mount request:
  - Upon acceptance of the mount request, the four construct names are stored in the Library Manager's inventory data base
- Is executed to perform a range of data management functions:
  - Based on actions defined at the Library Manager for individual construct names
  - When executed, depends on construct and action definition

Figure 2-22 details the file allocation, the data flow, and how the DFSMS constructs are passed to the Library Manager.

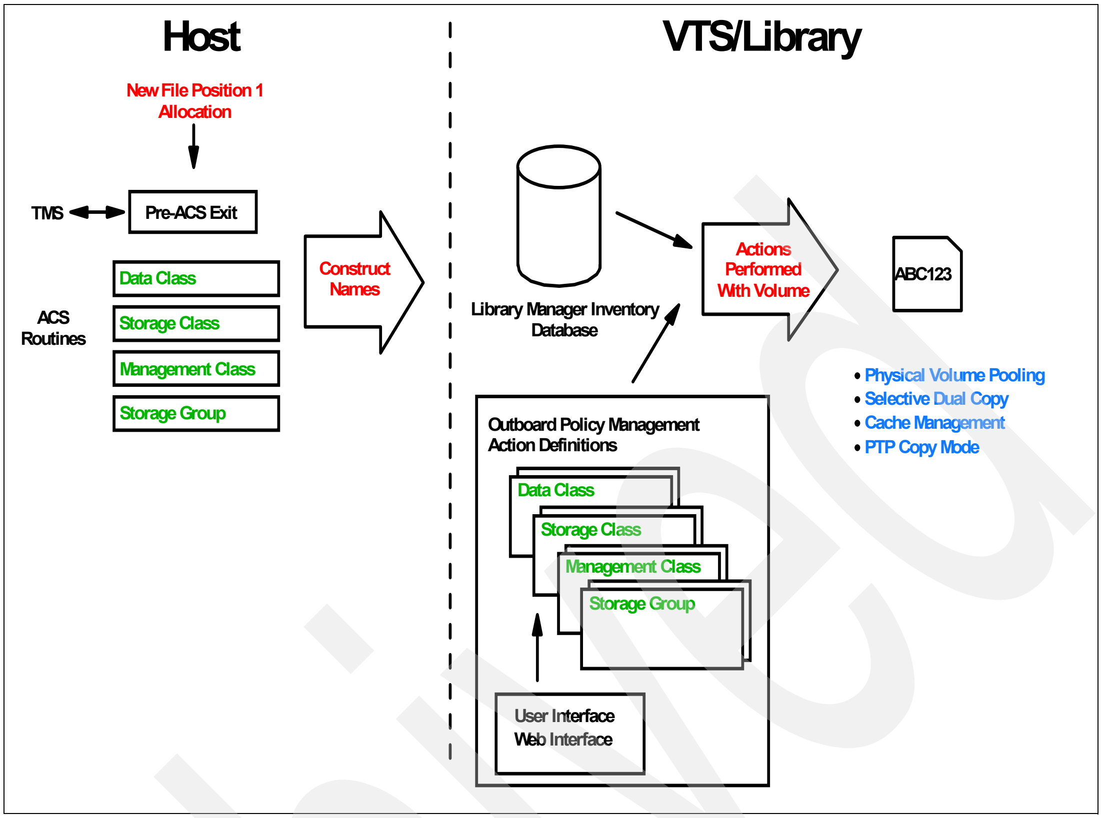

Figure 2-22 APM control flow

Enterprise Tape Library (ETL) Specialist

The IBM Enterprise Tape Library Specialist is the Web interface to the Library manager and the VTS. It is not a replacement for FC5226 (Remote Library manager Console) but its functionality makes the Remote Console obsolete.

The addition of Advanced policy management has also introduced a set of screens for the VTS component of the Enterprise Tape Library Specialist product. These screens will give administrators to the ability to create and change Volume pools and storage constructs, manage Physical and Logical volumes ranges and eject volumes. Security is controlled via a Web administration panel that can allow user Web access to only certain items

2.5.2 Physical volume pooling

With Outboard Policy Management enabled, we now have the facility where logical volumes can be assigned to selected Storage Groups. These Storage Groups point to primary storage pools. These 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 LM. Physical pooling of stacked volumes is identified by defining a pool number. This definition can be set through the LM or the ETL Specialist.

Volume pooling allows the administrator to define pools of stacked volumes within the VTS 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 (J, K, JJ or JA), mixed media or have a “first choice” and a “second choice”.

For inserting stacked volumes the volser range table has been changed to reflect the changes. 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 LM or ETL Specialist when defining stacked volume Pool Properties.

> **Note:** Category number FF03 is no longer used with the introduction of LIC level 2.26. Both scratch and private stacked volumes now have category FF04 assigned.

General purpose pools (Pool 1-32)

There are 32 General pools available per VTS. These pools can contain both physical scratch and private stacked volumes alike. All volumes, private and scratch, are now assigned category FF04. When initially creating these pools, it is important to ensure that the correct borrowing properties are defined to the pool.

> **Note:** Primary Pool 01 is the default private pool for VTS stacked volumes.

For those sites that are APM enabled, the following rules apply to the Dual Copy feature:

- A primary pool for one logical volume can be the secondary pool for another.
- Multiple primary pools can use the same secondary pool.

Without Outboard Policy Management enabled, all stacked volumes belong to one of the following pools:

- Scratch stacked volumes are assigned to pool 00 (non-active data).
- Private stacked volumes are assigned to pool 01 (active data).

Borrowing and returning

With the Advanced Policy Management, rules are introduced to manage out of scratch scenarios: the concept of *borrowing* and *returning*. With borrowing, stacked volumes can move from pool to pool and back again to the original pool. In this way the VTS can manage out of scratch and low scratch scenarios which occur within any VTS from time to time.

Stacked volume pool properties

Logical volume pooling has been enhanced to support drive type and cartridge type selection. This can be used to create pools of 3590 or 3592 cartridges, or pools of 60 GB 3592 cartridges and 3592 300 GB cartridges.

You may want to use the 60 GB JJ cartridge to provide fast access to applications such as HSM, Content Manager, and the 300 JA cartridges to address archival requirements such as full volume dumps.

Additional migration policies have been provided to facilitate the migration of data from 3590 cartridges to allow 3590 tape drives to be removed, and this may also improve the utilization of high capacity 3592 cartridges in future.

The parameters that can be used to control migration are based on characteristics such as:

- Percentage of active data on a cartridge
- Days without access
- Age of last data written
- Days without data inactivation

Reclaim policies are set to match the retention patterns for each individual pool. Reclaim targets can be set to select another media type as the target to support migration, and “secondary copies” in a PtP VTS, or “selective copies” in a stand-alone VTS can be created on a different media to the primary copy.

If Advanced Policy Management is not installed, you will need to order and install prior to upgrading a B20 VTS to attach to the 3592 tape drives. If Advanced Policy Management is installed, then the 3592 tape drives can be installed and new pools can be defined. At that point, new reclaim targets can be set and the migration process of data from 3590 to 3592 will commence. This will support a gradual migration of data to 3592 cartridges and allow for the eventual removal of the 3590 tape drives, but the customer may want to remove the 3590 drives, so we have provided two options to hasten this migration. You can find a detailed description of the migration process in a specific white paper for 3592 migration:

http://www-03.ibm.com/support/techdocs/

2.5.3 Secure Data Erase

With VTS R7.4 you can activate the *Secure data erase* function using the pool properties screen. This function adds physical volume erasure on a pool basis controlled by an additional reclamation policy. With the activation of this function, all reclaimed physical volumes in that pool are erased with a random pattern prior to being reused. A physical volume is not available as a scratch cartridge as long as it’s data is not erased.

If the value in the field *days before secure data erase* is non-zero, all reclaimed volumes in the pool are erased. A physical volume becomes eligible for reclaim when the number of days specified has elapsed since the *first logical volume* was invalidated on the volume after it became full. Zero means this policy is not to be applied; from 0 to 365 days can be specified.

The VTS internal ADSM code stores an “object” on tape, which is an instance of a logical volume. As a logical volume is written to again, a new instance is written to tape and the old instance is removed from the ADSM database, but still exists on a physical tape. With the delete expired volume function of the Library Manager, a logical volume’s instance is also removed from the ADSM database (once the expire time has elapsed and the Library Manager tells the VTS to delete the data, VTS deletes it from cache and reconcile then removes it from the ADSM database).

ADSM 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 what ADSM reports for percent of full is 1/10%, it rounds down, so even *one byte* of inactive data will drop percent to 99.9%. VTS keeps track of the time that the physical volume went from 100% full to less than 100% full:

- It checks on an hourly basis for volumes in a pool with a non-zero setting.
- It compares this time against the current time to determine if the volume is eligible for reclaim.
- When it is eligible for reclaim, it is subsequently erased.
- Read-only volumes will not be erased, even if they would have been eligible.
- Volumes tagged for erasure cannot be moved to another pool until erased.

- 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.
- Volumes tagged for erase can be ejected from the library, since such a volume is usually removed for recovery actions.

The hourly 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.

You should plan this function with care. It is explained in more detail within 3.2.4, “Secure Data Erase implications” on page 85.

2.5.4 Reclamation

The rules that govern reclamation have changed with the APM code level, to minimize the impact of volume pooling. Reclamation takes place for each pool and thresholds can be defined for each pool; this ensures that (unlike other implementations) the separation of logical volumes across physical pools is maintained.

However, the pool can be set up to move the active logical volumes to another pool after reclamation. The VTS follows this algorithm for concurrent reclaims under normal circumstances: The number of idle drives/2 minus 1, for example, 12 drives idle would result in a maximum of 5 reclaims. This means that up to 10 drives can be used for reclamation processes, leaving two drives for migration and recall.

The VTS has two states when the pool has a low number of scratch volumes, and this will increase the relative priority of reclamation:

Low scratch state: The pool has access to less than 50 but more than 2 scratch stacked volumes.

Panic scratch state: The pool has access to less than 2 scratch stacked volumes.

A key point here is that the pool has access to the common pool for scratch, and these volumes are taken into account when determining the scratch state. The common pool can be over committed; for example, the number of scratch volumes in the common pool is over 50 and there are 32 defined pools each ready to borrow from that pool.

2.5.5 Selective Dual Copy

With Advanced Policy Management, storage administrators have the facility to selectively create dual copies of logical volumes within a VTS. This function is also available in the Peer-to-Peer environment. At the site or location where the second Distributed Library is located, logical volumes can also be duplexed. This may end up with four identical copies of one logical volume. As we doubt that anyone has a need for this, we do not recommend this function for a PtP VTS.

2.5.6 Tape Volume Cache management

With the introduction of the Advanced Functions feature, FC4000, you had the ability to influence the time virtual volumes resided in the TVC. It was designed to use the Storage Class Initial Access Response Seconds (IART) attribute to determine whether a virtual volume (data set) should be preferentially deleted from cache once the VTS TVC threshold has been reached.

With the Advanced Policy Management (FC4001 - FC4004), you can either pass the Storage Class name to the Library Manager if you run in a DFSMS environment, or assign a Storage Class to a group of logical volumes if you run in a non-DFSMS environment; for example, an open system (the logical volumes are assigned a Storage Class at insert time).

The only difference between the two feature codes is that with FC4001 to 4004, the control has been moved outboard of the host. You can now override the host definition Storage Class attribute, should you choose to. While this can be done, it is not recommended for System z hosts.

Preference groups

DFSMS IART Storage Class policy management can be used to assign logical volumes to preference groups that further control when they are removed from cache on the I/O VTS. Other factors also control when they are removed from cache on the I/O VTS such as if they have been copied to the other VTS and to a physical tape (on the I/O VTS). There are currently two preference groups that can be assigned to a volume:

1. Preference Group 0:

   When space is needed in the cache for new or recalled logical volumes, volumes assigned to this preference group are selected first for removal from cache, ordered by largest volume first.

2. Preference Group 1:

   When space is needed in the cache for new or recalled logical volumes and there are no volumes in preference group 0 that are available for removal, volumes assigned to this preference group are selected to be removed from cache, ordered in least recently accessed order.

Improved recovery time

The objective here is to optimize the recovery process by keeping required volumes in cache and avoiding staging activity. It affects residency of volumes in remote VTS cache. For the logical volumes copied from an I/O VTS, their cache management is not influenced by DFSMS policy management. This function was announced in July 2002 via Technical Flash.

To increase the probability of cache hits in a recovery scenario, two options are available:

- Change “volume copy” preference group.

  Volume copies are placed in PG0 by default. This option will place copy volumes into PG1 (LRU).

- Preference “recalls” for cache removal.

  Recalled volumes are placed in PG1 (LRU) by default. This option will place recall volumes into PG0.

The default can be changed through a System Support Representative setting.

Figure 2-23 shows the CACHE settings available.

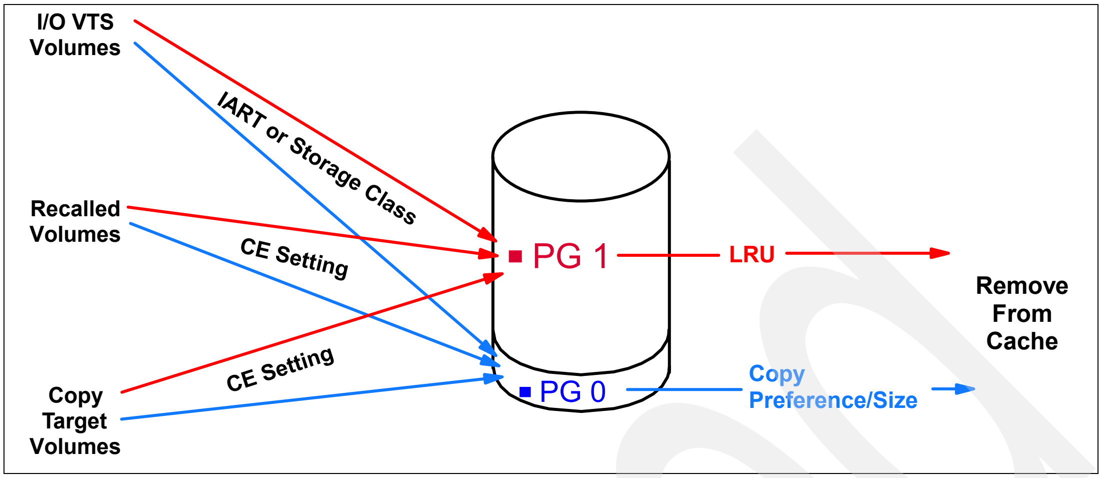

Figure 2-23 VTS CACHE settings enabled by the IBM SSR

2.5.7 Data Class usage considerations

With VTS R7.4, the default logical volume sizes of 400 and 800 MB are extended to 1000, 2000, and 4000 MB. The Data Class constructs passed by ACS routines are now honored by the Library Manager as far as they match with the definitions made there.

When a scratch mount is processed, the LM tells the VTS the logical volume size to use; the default is based on the media type or override in the Data Class. The VTS counts the compressed and formatted data that it is writing out to the file that represents the logical volume. The size is written to the logical volume header. When the count reaches the size the LM passed for the volume, the VTS indicates that logical end of volume (LEOV) has been reached and spannes the data to another logical volume.

These new large logical volume size still use the only two media types (MEDIA1 and MEDIA2), which are inserted with the “volume Insertion” during the installation and implementation process for the VTS. There is no change in the process a scratch request is handled: When a request for a scratch is issued to the VTS, the request specifies a mount from category. The Library Manager selects a virtual volser from the candidate list of scratch volumes in the category.

Note: Within a PtP VTS, when the Data Class is defined differently in both distributed libraries, the Data class definition of the I/O VTS is used. Copies on the other VTS do not use the Data class definitions — they are done to the same volume sizes as used in the I/O VTS.

2.5.8 PTP copy mode control

Advanced Policy management allows the users to control the copy behavior of individual logical volumes via the Management class. This means that critical volumes can be copied immediately and non-critical volumes can be deferred.

When the PtP VTSs have APM enabled, the copy mode will be set at *mount completion* time, as defined by the policies set at the Master VTS. Figure 2-24 shows how the process works.

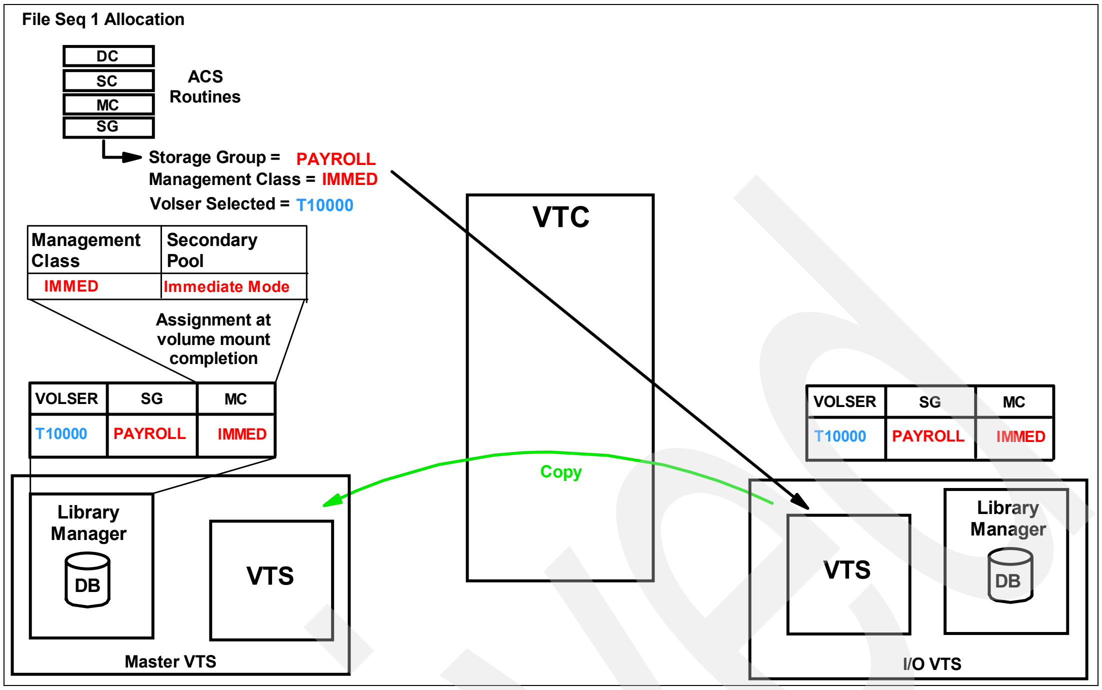

Figure 2-24 Copy mode assignment with Advanced Policy Management

2.5.9 Selective Peer-to-Peer Copy control

Prior to the addition of the PTP Selective Dual Copy function, when a scratch logical volume was written by a host application in a PTP VTS configuration, it was first written to one of the VTSs in the configuration and then copied to the other one. That first VTS is typically referred to as the I/O VTS. The I/O VTS is determined by the Virtual Tape Controller (VTC) based on such factors as availability of both VTSs, workload on each, and customer specified defaults set during the installation of the system.

There were a couple of primary uses for this function that were considered in its design. Applications are sometimes upgraded or new applications are created. These updated applications need to be tested and typically use a copy of production data. Once the testing has been completed, the data created is of no use. If the test data is written to a PTP VTS, it is also copied to the remote site, and since the data is not used at all in any kind of recovery, it is a waste of the inter-site bandwidth to make the copy. The data copied to the other VTS will also be migrated to physical tapes In addition, testing for disaster recovery at the remote site also generates data that would waste inter-site bandwidth and the storage resources at the local site. The PTP Selective Dual Copy function provides a way to eliminate the copy and save the inter-site bandwidth.

2.5.10 Performance and capacity

The Advanced Functions can have the ability to affect the performance and capacity of the VTS and 3494.

There are serious performance considerations that have to be taken into account before fully utilizing Advanced Functions. The library may need more resources cache, drives, cartridges, and library slots to return performance to the level previously enjoyed. Some of the potential problems that could occur are described below: Existing 3494 libraries were sized with a single pool of private/scratched tapes, the addition of up 32 pools can mean that the cartridge and slot requirement could increase due to more partially filled tapes across multiple pools;

insertion of extra scratch cartridges for separate pools. An existing VTS was sized with the assumption that all drives connected to a VTS were available for migrate or recall and during periods of heavy migration activity, the VTS would use more drives to cope with the extra workload. With additional pools defined, there is a real possibility of a drive not being available for a particular pool; this could ultimately affect the sustained performance of the VTS. More details are given in 3.2, “Planning for Advanced Policy Management exploitation” on page 82.

2.6 Peer-to-Peer VTS management

In this section, we summarize the mechanisms of the Peer-to-Peer VTS that allow for continuous data availability and operation. The internal management includes, for example, the decisions made during initialization, actions to be taken in case of a failure, synchronization between the components, and optimized usage of available resources.

2.6.1 Initializing the Peer-to-Peer VTS configuration

If all components are powered off, the initialization of the Peer-to-Peer VTS can only start after both distributed VTSs and distributed libraries and at least one of the VTCs are available. Establishing the Peer-to-Peer VTS is done through the VTCs.

The power on and initialization (POI) sequence ensures that the VTC is correctly configured within the Peer-to-Peer VTS, determines the overall status, and brings the devices controlled by this VTC to a known state before the tape control unit and the composite library are brought online to the attached host systems. The initialization steps are performed by each VTC in a defined sequence. The defined sequence guarantees the integrity of the overall Peer-to-Peer VTS.

As part of the POI sequence, the clock synchronization based on the UI distributed library and the selection of the Master VTS are performed.

Clock synchronization

The system clocks within the Peer-to-Peer VTS are synchronized to allow logs (which may be used for problem resolution) to show consistent time-stamps. The time setting does not affect the operation of the Peer-to-Peer VTS.

The time of the VTCs is synchronized with the time of the distributed VTS in the UI distributed library. If it is not available, the time of the distributed VTS in the other distributed library is used.

Preferred Master VTS

At any point in time one of the VTSs, and associated library, is the Master VTS. All library function requests (mount, category change, information query, etc.) go to the Master VTS first. A change in which VTS is the master occurs both in error or service situations. If a VTC cannot communicate with the Master VTS, it cannot continue to perform library functions (mounting volume being the key function).

In the case where a PTP VTS is split between a local and remote site, the customer generally wants the local VTS to be the master so that in case the communication links between the sites are interrupted, they can continue to perform work with the local VTS without requiring IBM SSR assistance. What is needed is to have the VTCs switch to the local VTS as the master whenever possible instead of basing the switch on an error condition or IBM SSR action.

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

- Default is no preference for a Master VTS
- VTS0 or 1 can be specified as the preferred Master VTS

If a preferred Master VTS has been specified, if it is not currently the master and conditions within the PTP VTS subsystem (primarily path and component availability) allow, the VTCs will initiate a switch to make it the Master VTS.

Important: This control is available at code level 2.27.xx or later.

Selecting the Master VTS

The initial selection of a Master VTS is performed during POI of the PTP VTS. For a distributed VTS to be selected as the Master, it must meet the following criteria:

1. All operational VTCs in the subsystem must be able to communicate with the VTS.
2. The Library Manager associated with the VTS must be in the online state.

If more than one distributed VTS library meet the criteria and a preferred Master VTS was specified for the configuration, it is selected as the Master. If a preferred Master VTS had not been specified, then the UI Distributed Library is selected as the Master.

Note: VTC preferencing should not use the "Preferred Master" setting in a system that is under GDPS control. Preferencing should be left at the "No Preference" setting for a PtP being used with GDPS.

2.6.2 Master VTS switchover

In the Peer-to-Peer VTS, the Master VTS needs to be available for normal processing. It is the focal point to which library commands that update a volume are initially directed. If the Master VTS fails, or if the connection between one or more VTCs and the Master VTS is interrupted, a switchover from the original Master VTS to the other VTS is initiated.

The switchover is initiated by the VTC that detects that it can no longer communicate with the Master VTS. A minimum number of VTCs is required to perform a Master VTS switchover, that is one-half plus one of the configured VTCs must be operational. In a remote installation where all the VTCs are at one site are no longer available, or if for any reason less than three out of four or five out of eight VTCs are operational, the switchover is not performed automatically. In this situation, the SSR can initiate a Master VTS switchover.

Once an VTC has determined that it cannot communicate with the Master VTS and a Master switchover attempt was not successful, it enters the master disabled state. It will remain in that state until:

- Communication with the current Master VTS has been re-established.
- A switchover to another Master VTS is subsequently successful.

2.6.3 Token processing

Data integrity in the Peer-to-Peer VTS is guaranteed by means of tokens that are stored in a database in each distributed VTS and synchronized upon every update. In addition, token data is stored with each logical volume when it is written onto physical tape.

Note: The existence of these tokens is the major reason why a distributed library cannot be used as a stand-alone VTS (for example, after a disaster at one site) and later be re-integrated into a Peer-to-Peer VTS: the tokens would be completely out of synchronization.

The token data for a virtual volume includes information:

- That indicates whether a virtual volume resides in the Tape Volume Cache of the distributed VTS or has already been copied to a stacked physical cartridge. The VTCs direct the mount request to the distributed VTS with the cached copy to improve recall efficiency.
- That indicates whether the distributed VTS contains a valid copy of a virtual volume. This prevents using an outdated version of a virtual volume.
- That prevents using the same I/O VTS again if a read operation has encountered an error. Subsequent mount requests for this virtual volume will be directed to the other distributed VTS.
- That indicates whether the copy of a virtual volume still needs to be created.
- That can be displayed by the ETL Specialist to provide up-to-date information about the logical volumes in a Peer-to-Peer VTS.

Tokens are initially created:

- During insert processing of a virtual volume
- When migrating an existing stand-alone VTS into a Peer-to-Peer VTS

During a recall of a virtual volume from physical tape into the TVC, the token stored with the virtual volume on tape and the token data in the token database are compared to ensure that the latest version of a virtual volume is used.

2.6.4 Dynamic Device Reconfiguration

Host Dynamic Device Reconfiguration (DDR) functions are used to allow jobs to continue execution after a failure where the host is reading data and, in some cases, even if only one copy of the data exists. Data availability from a PTP VTS is highly dependent upon:

- The copy of the virtual volume which is stored within the PTP VTS, and
- Re-run of host jobs which are abended due to a failure.

When jobs which were abended are re-run, the PTP VTS is able to bypass the failed element (VTC, VTS or connection between the controller and VTS) and provide access to the virtual volume. The host operator approval of DDR will result in a demount of the volume and allocation of another tape drive for a retry of the read operation.

When the read failure is a result of:

- Loss of the connection between the VTC and the I/O VTS, or
- Loss of the I/O VTS (B10 or B20)

The job may continue, depending upon the availability of copies of the virtual volume in both VTSs and the virtual drive selected via the allocation function of the host. Table 2-2 summarizes the cases when a mount of the virtual volume is successful and when unsuccessful, which results in a job abending, indicating Valid Volume Inaccessible.

Table 2-2 DDR scenarios

Original failure	Copy available		Copy not available	
	Same VTC	**Different VTC**	**Same VTC**	**Different VTC**
Connection between VTC and I/O VTS	Select VTS other than original I/O VTS	Select either VTS	Fail the mount	Select VTS with available copy
I/O VTS	Successful on VTS other than original I/O VTS		Fail the mount	

When the mount is successful after a DDR attempt, and the volume label is validated, the read command is retried after the DDR function positions to the block that previously failed.

2.6.5 Cache Management

Tape Volume Cache Management in the Peer-to-Peer VTS has three functions:

1. Keeping the host input operations and the copying of virtual volumes to stacked volumes in balance
2. Expediting the removal of the secondary copy of a virtual volume from the TVC (a process known as *fragmenting*)
3. Preventing the removal of a virtual volume from the TVC if it has not yet been copied to the other distributed VTS

Note that the APM Tape Volume Cache Management function described in 2.5.6, “Tape Volume Cache management” on page 45 also apply to a Peer-to-Peer VTS.

Balance host input and stacked volume copies

The Tape Volume Cache in the distributed VTSs is managed in the same manner regarding throttling as the TVC in a stand-alone VTS. If the distributed VTS has been running for some time at peak throughput rate, then the background operations (such as copying the virtual volumes from the TVC to the physical stacked volume) may require higher bandwidth to keep up with the data that is written from the host. To allow higher bandwidth for the background processes, host input is throttled. This is indicated in Figure 2-25.

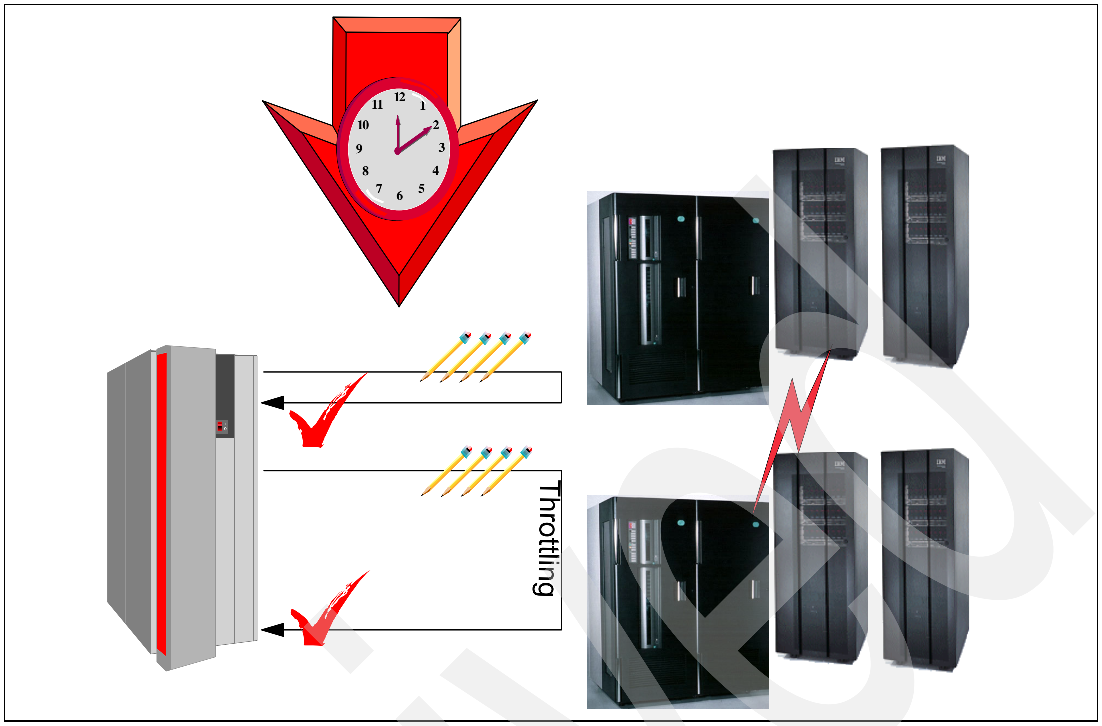

Figure 2-25 Host input throttling

Throttling is a fixed delay imposed on host write operations to allow the process of copying to stacked volumes to keep up with incoming data and hence prevent the TVC from filling. The first write operation indicated by the pencils in Figure 2-25 does not show throttling. The second write operation in the lower part of the diagram shows throttling as the time between starting the I/O and signaling completion has been increased by the throttling delay.

For a detailed discussion of VTS performance management and host throttling algorithms, please refer to *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

Expedite deletion of the secondary copy

Removal of the primary copy from the TVC (*fragmentation*) follows the same rules as in a stand-alone VTS, using the algorithms to keep the virtual volume in the TVC as long as possible.

The secondary copy is removed from the cache as soon as possible. This is achieved by adding the secondary copy to the top of the queue of volumes to be removed, thus allowing primary copies to stay longer in the TVC and improve recall efficiency.

Preventing removal from cache before copy

In deferred copy mode, virtual volumes that have not yet been copied to the other distributed VTS are preferentially held in the TVC to prevent them from being fragmented and later on recalled for copy. This is done to improve the overall performance of the Peer-to-Peer VTS.

If the TVC reaches a certain threshold of space occupied by uncopied virtual volumes, the copy priority is temporarily changed, as explained in 2.3.2, “Deferred copy mode” on page 29.

If virtual volumes have been removed from the cache before the secondary copy has been created, a special recall process allows multiple virtual volumes to be recalled from the same stacked volume in one operation. This improves the recall performance.

2.6.6 Distributed and composite library states

When all components of the Peer-to-Peer VTS are fully functional, the distributed libraries and the composite library are in *normal* state. If a failure occurs in one distributed library, it changes its state according to the failure. Multiple states can be active at the same time.

If *any* of the distributed libraries is in a state listed below, the composite library is also in this state:

- Automated mode operation
- Degraded operation
- Safety interlock open

If *all* distributed libraries are in one of the states listed below, the composite library is also in the same state:

- Paused mode operation
- Manual mode operation
- Vision system non-operational
- Library Manager Check 1 condition
- All storage cells full
- Out of cleaner volumes
- Secondary write disabled
- Out of mount resources
- Library Manager switchover in progress
- Out of Empty Stacked Volumes

Based on VTC status, the following states can exist for the composite library:

- Copy operation disabled:

  If all of the VTCs only have paths to one of the distributed libraries

- VTS operations degraded:

  If one of the components of the Peer-to-Peer VTS is not operational

- Immediate copy completions deferred:

  If one of the VTCs has outstanding secondary copies queued while operating in the immediate copy mode of operation

- Service preparation:

  If a distributed library is in the service preparation state. See 2.7, “Service considerations” on page 55, for more details

The state of the distributed library, the composite library, and all the VTCs can be displayed using the Peer-to-Peer VTS Specialist (see 6.3, “Monitoring the Peer-to-Peer VTS” on page 223) as well as by using the host console display commands (see “Host commands” on page 227).

2.6.7 Special VTC operational modes

With VTS R7.4, we introduce the customer control of special VTC modes in order to override the normal minimum configuration requirement in exceptional circumstances. For example, suppose that your remote site collapsed, a link failure occurred, and the master is not able to switch to the local site where your operation could continue. In this case you are allowed by this Web Specialist support to switch into the Read/Write Disconnected mode or the Read-Only mode. A special screen is displayed at the Web Specialist if VTC and at least one VTS is offline. Without this Specialist support, your IBM SSR has to be called to set this special mode.

- **Read/Write Disconnected mode:**

  Logical volumes may be read, modified, or rewritten unless the volumes are known to not be valid. This mode may only be set for one VTS; even if all links between two sites have been lost, you ***must not*** operate both sites in Read/Write Disconnected mode.

  When the second VTS is brought back into the configuration, the VTCs will then reconcile tokens, resume normal operations, and schedule logical volume copies.

- **Read-Only mode:**

  Logical volumes may be read unless they are known to not be valid. The host is prevented from writing to the volumes or from changing their category assignment (for example, private, scratch, or insert).

  After other components have been repaired or made accessible, the virtual drive addresses currently being used must be taken offline again before discontinuing Read-Only mode.

  Note: These modes have to be used with extreme care. They may be set, for example, in cases when, after a complete shutdown of the Peer-to-Peer VTS, only one VTS can be brought back online, or if all VTCs are in the master-disabled state. All the virtual drive addresses must be varied offline at the host systems before Read-Only or Read/Write Disconnected mode may be set.

- **Write protect mode**

  May be used in order to run a disaster recovery test for a specific application whilst the PtP VTS complex runs in normal production mode. You can set a selected VTC in write protect mode independent of the other VTCs via a Specialist screen. It has to be set while the VTC is ONLINE. While the specific VTC is in write protect mode, the following will fail:

  – Any scratch mount
  – A host attempt to modify the contents of a logical volume that was mounted while the mode was active
  – A host attempt to change the attributes for the volume (category or constructs)

PtP mode sense and statistics indicate the mode that is set for the VTCs. All three modes and their usage considerations are described in more detail in Chapter 9, “Disaster recovery scenarios” on page 321.

2.7 Service considerations

In this section, we cover the considerations related to Peer-to-Peer VTS service, including microcode upgrades of the various components. The key objective for the Peer-to-Peer VTS is to provide continued access to data stored within the configuration. In the event of service to a component of the Peer-to-Peer VTS, the component will be unavailable to perform any work; however, other components of the configuration will continue to provide full access to data.

2.7.1 Service preparation

When a VTS component of the Peer-to-Peer VTS needs to be serviced, the SSR must initiate the *Service preparation* state. The mode selected determines the impact on host operations. Service preparation does not apply to the VTCs, although any one VTC may be used for the initiation of the action.

VTC

Preparing the VTC for service is the same as preparing any tape control unit for service.

The following steps are performed by the SSR and the customer operator to prepare the VTC for service:

1. At the hosts attached to the VTC to be serviced, vary offline the tape drive addresses associated with the VTC.
2. Allow currently running jobs to complete or have the operator cancel them.
3. For long running jobs, swap the virtual tape drives being used to drives on another VTC using the SWAP command.
4. Verify that all of the virtual drives are unloaded and not ready.
5. Change the state of the VTC to offline.
6. The code can now be upgraded.

In order to allow permanent access to data, a single host should *always* be attached to multiple VTCs. If one VTC needs to be taken offline, the data can be accessed through the devices of another VTC.

VTS

When an operational VTS has to be taken offline for service, the Peer-to-Peer VTS must first be prepared for the loss of the resources involved. The main operational changes in the Peer-to-Peer VTS as it is prepared for service of a distributed VTS or library are:

- Make the VTS not being prepared for service the Master VTS.
- Direct workload to the remaining distributed VTS.
- Expedite data copies where the only valid copy is on the distributed VTS to be serviced.

Service preparation for the VTS is initiated by the SSR through one of the VTCs. Based on the type of service that is required, and based on the urgency of having it performed, the SSR selects the service mode. Only one VTS of a Peer-to-Peer VTS can be prepared for service at any time. Four service modes are provided. During the service preparation process, the mode may be changed to a more restrictive mode.

- Normal mode:

  This mode is used when service or a microcode update is to be performed on the distributed VTS or the distributed library with the least impact on performance of the Peer-to-Peer VTS. The following actions are taken:

  – If in immediate copy mode, copy operations in progress that target the distributed VTS to be serviced, complete normally.

  – Subsequent rewind/unload commands on the distributed VTS, which is not to be serviced, are treated as if in deferred copy mode.

  – Subsequent rewind/unload commands for modified volumes on the VTS to be serviced are treated as in immediate mode, and rewind/unload completes after the copy has been created. Copy operations compete with host I/O operations equally for bandwidth priority.

  – The VTS that is not to be serviced is selected as the I/O VTS for all workload except for when the only valid copy of a virtual volume is in the VTS to be serviced.

  – Copy operations from the distributed VTS to be serviced to the other VTS are selected prior to those that have the distributed VTS to be serviced as the target.

  – Copy operations to the VTS to be serviced are suspended.

- Expedite mode:

  This mode is used when the time to do service preparation has to be minimized, but continued access to data is required. The difference from normal mode is that copy operations from the distributed VTS to be serviced to the other VTS, are processed with a bandwidth priority *equal* to host I/O operations. Some impact to the Peer-to-Peer performance may be experienced.

- Immediate mode:

  This mode is used when there is an urgent need to service or perform an offline operation with the distributed VTS or its associated distributed library. The actions described under normal mode are changed so that:

  – All virtual volumes currently open on the distributed VTS to be serviced are allowed to process to completion.

  – All workload is routed to the VTS that is not to be serviced. If the only valid copy of a virtual volume that is subsequently required is in the VTS to be serviced, the job terminates abnormally with `valid volume inaccessible`.

  – All copy operations are suspended.

- Forced mode:

  This mode is used only after the other modes have failed to complete or if the jobs in progress could not be terminated to free up the virtual devices.

  Paths to the distributed VTS are terminated immediately regardless of copy and device status. Devices may be left in use and volume copies may fail. This mode is for emergency hardware service only and is not expected to be used for microcode upgrades.

Once a service mode has been selected, all attached hosts are notified of the state change. After service preparation has completed, each of the VTCs notifies the attached hosts. The SSR can then perform the normal steps for service of the distributed VTS or the distributed library.

During the time a VTS is being prepared for service, and for the duration of the service, host operations to the Peer-to-Peer VTS are degraded as listed for the different service modes. In addition, some host commands are restricted, such as Library Volume Audit, Library Volume Eject, and logical volume insert processing. These commands will not be performed. The VTC receiving the command will indicate `Function Incompatible` to the host.

The SSR can request information about the service preparation process in progress from any of the VTCs. The information received includes, for example:

- The time required to complete the copies
- The percentage complete of the secondary copy work
- The virtual drives in use by a host

The SSR may request that the operator issue a SWAP command to the same tape drive address. This will terminate the operation on the VTS being prepared for service and continue on the other VTS. In this manner, long running jobs can be moved to the VTS that is not being prepared for service.

Distributed library

If a distributed library requires service, the associated VTS must first be prepared for service. Once the VTS has completed service preparation, the SSR can continue the normal procedures for servicing of the tape library. When other tape drives or a VTS that is not part of the Peer-to-Peer VTS are installed in a tape library, other activities may need to be performed before servicing the tape library.

2.7.2 Microcode installation and activation

Customers who prefer and are able to halt all usage of the Peer-to-Peer VTS for microcode updates may find that the total process is simpler and less demanding of operations and SSR communications and interaction. It will avoid:

- Continued buildup of volumes to be copied
- Reduced performance during upgrade and after components are returned for use

Microcode upgrades should be considered a form of service, and a service preparation process as described in 2.7.1, "Service preparation" on page 55, will be performed by the SSR. Access to all logical volumes will be provided while microcode upgrades are performed.

If necessary, the Peer-to-Peer VTS may be returned to the customer at any time for data access without completing the microcode upgrades of all components.

2.7.3 Resuming operation after service

After the service activity has been completed, the component that was being serviced is brought back online again. If a VTC has been serviced, the virtual drives associated with this controller can be varied online again.

After service is completed for a VTS in a Peer-to-Peer VTS configuration, either because service was required on the distributed VTS itself or on the associated library, the Peer-to-Peer VTS will proceed to synchronize and perform copy operations that may have been suspended during the service period. While recovering to normal operational states, access to all virtual volumes is granted. Some impacts to Peer-to-Peer performance may be experienced due to accumulated workload of volumes to be copied. The amount of pending secondary copies is reported in field S94CLLVC of the SMF record 94 and can be displayed using the StorWatch Peer-to-Peer Specialist.

The VTS recovery processing includes these steps:

1. Verify that the Library Manager is online.
2. Resolve the hot token list. Tokens that have been created during the time of service are synchronized on both distributed VTSs.
3. Change the distributed VTS availability state to *available*.
4. Notify the host that the Peer-to-Peer VTS operation is no longer degraded.
5. Synchronize the mounts with the other VTS.
6. Check whether there are volumes in the insert category of the UI distributed library and generate an attention message if required.

2.7.4 Simple network management protocol traps

The 3494 can be attached to a variety of host systems over various connections. When errors or events occur that the host operations personnel should know about, the 3494 can notify them. This notification occurs using Transmission Control Protocol/Internet Protocol (TCP/IP) over a LAN. Connectivity to the LAN can be provided by Token Ring (FC5219) or Ethernet (FC5220).

SNMP is the message protocol used to transmit this information. The messages that the 3494 sends are called *traps*. They are sent to one or more (up to five) SNMP monitoring stations. From there, the monitor programs can take a specific action depending on the contents of the trap.

For additional information on setting up SNMP and on which traps are available, refer to the appropriate operator’s guide of your tape library:

IBM 3494 *IBM TotalStorage 3494 Tape Library Operator’s Guide*, GA32-0280

IBM 3584/3953 *IBM TotalStorage 3584 Tape Library Operator Guide,* GA32-0468

IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide, GA32-0473

2.8 Host-related processes

In this section, we describe how host- or operator-initiated processes are performed in the Peer-to-Peer VTS, for example, scratch mount requests from the host or the definition of virtual volumes at the Library Manager console.

2.8.1 Inserting logical volumes

In a Peer-to-Peer VTS, logical volume ranges can only be defined from the Library Manager console of the UI distributed library. For details on other definitions on the Library Manager console, such as defining physical volume serial ranges for the distributed VTS, please refer to Chapter 4, “Implementation” on page 127.

After virtual volumes are added in the UI distributed library during host insert processing, the virtual volume is first created in the other distributed library. Subsequently the volume category is changed in the Master VTS and in the other VTS to `PRIVATE` or `SCRATCH`, depending on the host request and the tokens for the virtual volume are created.

If either distributed library is not operational, no virtual volumes can be inserted. After both distributed libraries are operational again, the host is automatically informed about volumes in the insert category and the normal insert processing is performed.

2.8.2 Mount processing

The processing of mount requests in the Peer-to-Peer VTS depends on whether it is a specific or a scratch mount request. Writing data to a VTS, either stand-alone or Peer-to-Peer, you should use scratch mount processing to benefit from fast virtual mount times. DISP=MOD and DISP=OLD processing may require a recall of the volume into the TVC, unless the virtual volume still resides in the TVC.

Scratch mount and specific mount processing is summarized in the sections below.

Scratch mount processing

A scratch mount request is first sent to the Master VTS, which selects the next volume from the requested scratch category (MEDIA1 or MEDIA2).

After the mount is complete on the Master VTS and the validity of the volume in both distributed VTSs has been confirmed, the VTC selects the I/O VTS based on the criteria described in 2.4.2, “Peer-to-Peer VTS workload balancing” on page 33. A specific mount request is then sent to the non-Master VTS, using the Fast-Ready attribute to prevent recalling the virtual volume into the Tape Volume Cache, and the host is notified via the sending device end.

In the Peer-to-Peer VTS, it is not required that you define the Fast-Ready attribute for the scratch categories. However, we strongly recommend that you define the Fast-Ready attribute for scratch categories; otherwise, eject processing of virtual volumes will fail. A scratch mount request will automatically include the Fast-Ready indication and prevent the recall of volumes into the TVC even if they are newly written.

Specific mount processing

When a specific mount request from the host is sent to the Peer-to-Peer VTS, the VTC processing this request selects the I/O VTS based on the selection criteria described in 2.4.2, "Peer-to-Peer VTS workload balancing" on page 33.

2.8.3 Ejecting logical volumes

Eject processing results in deleting the logical volume from both distributed libraries. Therefore, eject processing requires both distributed libraries and VTSs to be operational in order to process the eject request.

Ejecting virtual volumes from the Peer-to-Peer VTS requires the virtual volume to be in the scratch category. The Fast-Ready attribute must be defined for the scratch category; otherwise, the eject will fail.A maximum of 1000 eject requests can be issued at a time. Additional eject requests will fail.

2.8.4 Creating Peer-to-Peer VTS statistics

As known from the stand-alone VTS, statistics for the Peer-to-Peer VTS are sent to the attached hosts once an hour. Because the Peer-to-Peer VTS has a logical image (composite library) as well as underlying physical components (distributed libraries), statistics are created for the composite library and for the distributed libraries. This results in three SMF94 records being created on the host once an hour. For details about the enhancements of the SMF 94 records for Peer-to-Peer VTS, please refer to Appendix D, "SMF type 94 record layout" on page 415.

The way in which the data is created and presented to the attached hosts is different from the way in which statistics are sent to the host for non-Peer-to-Peer libraries. The actions that are taken to create the three SMF type 94 records for the Peer-to-Peer VTS on a host can be summarized in six steps:

- Step 1: The Library Manager requests statistics; VTS returns statistics.

  On the hour, both Library Managers of the distributed libraries prepare a statistical record for the past hour. In preparation of the record, the Library Managers request data from the associated distributed VTS to be included in the record.

- Step 2: The Library Manager signals that the statistics are available.

  Once each Library Manager has created the data for its distributed library, it sends a message indicating that statistics are available to its distributed VTS. The VTS passes this request along to all VTCs.

- Step 3: VTC controllers exchange statistics.

  In response to this message, each of the VTCs prepares a data record containing the consolidated statistics of all VTCs and thus for the composite library.

- Step 4: The host receives notification and requests statistics.

  Once all VTCs have created their records of the statistics, each of them signals this to all attached hosts. In response, DFSMS/MVS requests statistical data for both distributed libraries and for the composite library.

- Step 5: VTC requests statistics from VTSs and routes them to the host.

  For host requests for statistics for the distributed libraries, the VTC requests statistics of the distributed libraries through the VTSs and routes the data to the requesting host. On the host, the device services component of DFSMS/MVS initiates creation of the SMF type 94 records.

- Step 6: VTC controller sends consolidated statistics to the host.

  In response to the request for statistical data for the composite library, the VTC returns the consolidated statistics for the composite library to the requesting host. On the host, the device services component of DFSMS/MVS initiates the creation of the SMF type 94 record. The steps listed above are summarized in Figure 2-26.

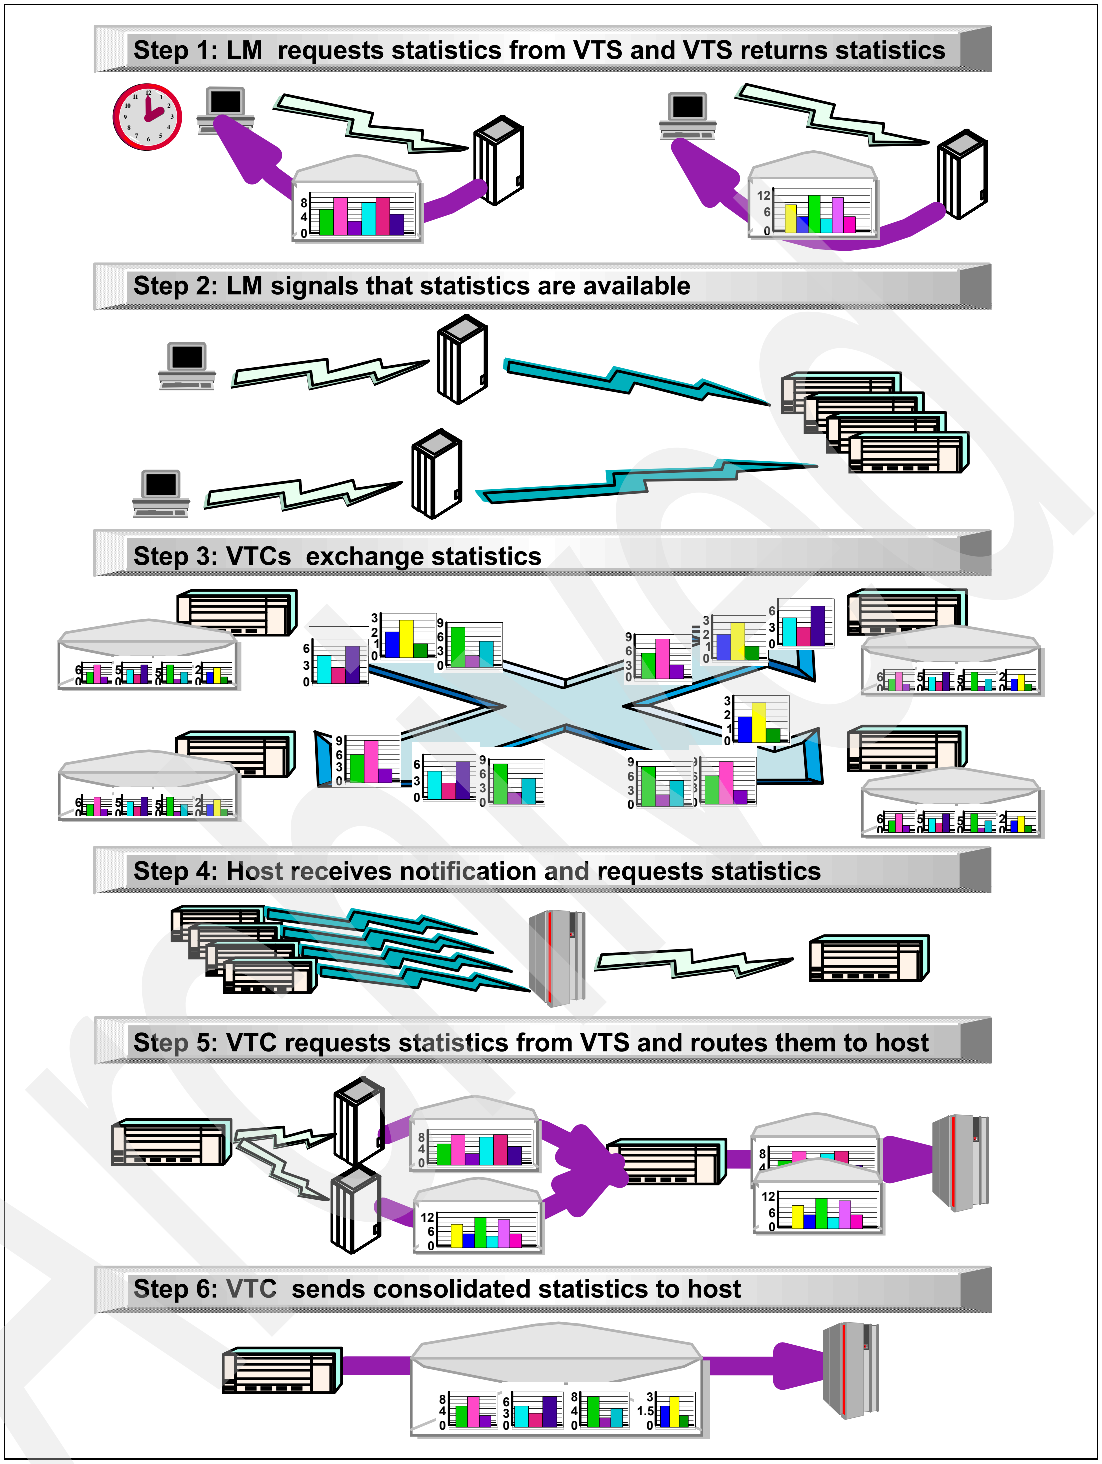

Figure 2-26 Creating Peer-to-Peer statistics on the host

3

Preinstallation planning

In this chapter, we describe the preinstallation tasks for customers who have no experience with a PtP VTS or who have experience with other virtualization solutions only. We cover both the basic implementation tasks for a VTS and the additional considerations for the Peer-to-Peer VTS. We also explain the hardware configurations, as follows:

- We describe the APM considerations.
- We discuss general performance considerations and present performance charts of different configurations:
  - PtP VTS in immediate versus deferred mode
  - Operating preferred and balanced
  - Benefits of the FPAF with 3592 tape drives
- We describe extended distance solutions with cascaded switches.
- We provide a comprehensive planning checklist, which contains planning tasks such as these:
  - Determine the workload for the VTS.
  - Size the VTS (TVC, storage slots, drives, and more).
  - Select the copy mode and operational mode
  - Plan Advanced Policy Management functions to be used.
  - Complete physical installation planning.

If you are upgrading an existing VTS subsystem or merging two existing VTS subsystems, refer to Chapter 5, “PtP VTS migration and relocation” on page 195. The software tasks are described in detail in Chapter 4, “Implementation” on page 127.

You can find information about the tasks that apply to a stand-alone VTS in *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

3.1 Hardware configurations

All of the stand-alone Model B10 or B20 VTSs and the Peer-to-Peer VTS configuration present themselves to the host as one single logical library containing up to 256 virtual 3490E tape drives (configuration dependent), although the underlying hardware requirements are different. To make it easy to identify the hardware components and configuration requirements for a Peer-to-Peer configuration, we have summarized these below. For a detailed description of the VTS feature codes, please refer to *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229, to the *IBM TotalStorage 3584 and 3953 for zSeries Hosts Planning, Implementation and Monitoring Guide, SG24-6789*, and to the *IBM 3494 Tape Library Introduction and Planning G*uide, GA32-0448.

3.1.1 Peer-to-Peer VTS configuration

The Peer-to-Peer VTS consists of two Model B10 or B20 VTSs and four or eight Volume Tape Controllers (VTC). These two VTSs only have either ESCON, FICON or a mix of both connections to the VTCs. The VTCs are installed in Model CX0 or Model CX1 frames, each of which can hold up to four VTCs. Each VTC has two host FICON or ESCON connections and a FICON or ESCON connection to each distributed VTS. You cannot have a FICON connection to the host and a ESCON connection to the VTS within the same VTC.

Both VTSs must be installed in separate tape libraries.

Table 3-1summarizes the PtP VTS configuration options.

Table 3-1 Possible PtP VTS configurations

VTS model	VTCs	FICON hosts	ESCON hosts	CU images	virtual drives
B10	4 4 4	0 4 8	8 4 0	4 4 4	64 64 64
B20	4 4 4 8 8 8 8 8 8	0 4 8 0 0 8 8 16 16	8 4 0 16 16 8 8 0 0	8 8 8 8 16 8 16 8 16	128 128 128 128 256 128 256 128 256

Each of the VTSs is associated with four to twelve 3592 tape drives that are installed in one D22 frame or four to twelve 3590 drives in D12 frames, if this VTS is attached to a 3494 Tape Library. Using a TS3500 Tape Library, only 3592 tape drives are supported. They do not have to be installed in one frame only. Table 3-2 summarizes the tape drive configuration options for each of the two physical VTSs which make up the PtP VTS. Although these tape drive configurations can be different for the two VTSs, we strongly recommend to keep them the same.

Table 3-2 VTS tape drive configuration

VTS model	With IBM 3494 Tape Libraries	With IBM TS3500 Tape Libraries
B10	4 to 6 IBM 3590 or 4 to 12 IBM 3592	4 to 12 IBM 3592
B20	4 to 12 IBM 3590 or 4 to 12 IBM 3592 or 4 to 6 IBM 3590 plus 6 to 12 IBM 3592	4 to 12 IBM 3592

We describe each component of a Peer-to-Peer VTS in the following sections.

3.1.2 VTS configuration requirements

In this section, we describe the specific features of the VTS in a Peer-to-Peer configuration. For detailed descriptions of the other VTS features, please refer to *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

Model B10 VTSs in a Peer-to-Peer configuration must meet the following requirements:

- At least four 3590 Model B1A, E1A or H1A tape drives, or
- At least four 3592 Model J1A tape drives
- Tape Volume Cache of at least 216 GB capacity
- VTSs attachment to separate tape libraries
- Peer-to-Peer Copy base feature (FC4010) and Peer-to-Peer Copy increments (FC4011, 4012, 4013) if required (refer to Table 3-3)

Model B20 VTSs in a Peer-to-Peer configuration must meet the following requirements:

- At least six 3590 Model B1A, E1A or H1A tape drives, or
- At least four 3592 Model J1A tape drives, or
- At least four IBM 3590 and four 3592 tape drives when using both device types. Section 3.1.10, “IBM TotalStorage 3592 tape drives Model J1A and E05 in a 3494 library” on page 72 describes the tape drive configuration options in more detail.
- The Tape Volume Cache must have at least 864 GB capacity.
- Both VTSs must be attached to separate tape libraries.
- Peer-to-Peer Copy base feature (FC4010) and Peer-to-Peer Copy increments (FC4011, 4012, 4013) if required (refer to Table 3-3) have to be installed.

Section 3.1.10, “IBM TotalStorage 3592 tape drives Model J1A and E05 in a 3494 library” on page 72 describes the tape drive configuration options in more detail.

3.1.3 Peer-to-Peer Copy features

In this section we describe the feature codes related to the Peer-to-Peer VTS configuration (Table 3-4).

Table 3-3 Peer-to-Peer Copy features

Feature	Maximum quantity	Prerequisite	Description
FC4010	1	—	Peer-to-Peer Copy Base
FC4011	1	FC4010	Peer-to-Peer Copy Increment 1
FC4012	1	FC4010 FC4011	Peer-to-Peer Copy Increment 2
FC4013	1	FC4010 FC4011 FC4012	Peer-to-Peer Copy Increment 3

Details on the feature codes listed in Table 3-3 are described below.

- **Peer-to-Peer Copy Base (FC4010):**

  This feature enables a B10 or B20 VTS to function in a PtP VTS configuration. When FC4000 (Advanced Function) is installed, a service representative must use the Library Manager service menus of both Peer-to-Peer VTS distributed libraries to disable the export and import of logical volumes. FC3422 (SCSI Host Attachment) may not be ordered with FC4010. If FC3422 is installed, it must be removed before FC4010 is installed.

  A B10 or B20 VTS with FC4010 can be used only in a PtP VTS configuration. For the B10 VTS, FC4010 requires two FC3412 (Extended Performance ESCON Channels), or one of FC3412 (Extended Performance ESCON Channels) and two FICON features (FC3415 or FC3416), or four FICON features (FC3415 or FC3416). The PtP VTS configuration must contain four VTCs.

  For the B20 VTS, FC4010 requires four of feature 3412 (Extended Performance ESCON Channels), or two of feature 3412 (Extended Performance ESCON Channels) and four FICON features (FC3415 or FC3416), or eight FICON features (FC3415 or FC3416). When FC4010 is installed, FC3418 (Activate Additional ESCON Channels) will be disabled. The PtP VTS configuration must contain four or eight VTCs.

- **Peer-to-Peer Copy Increment 1 (FC4011):**

  This feature enables the Peer-to-Peer VTS Copy function when more than 250 GB of usable disk capacity is installed. FC4010 (Peer-to-Peer VTS Copy Base) is a prerequisite. For the B10 VTS in a PtP VTS, this feature is required when FC3705 (288/432 GB Disk Storage) is installed. For the B20 VTS in a PtP VTS, this feature is always required.

- **Peer-to-Peer Copy Increment 2 (FC4012):**

  This feature enables the PtP function when more than 500 GB of usable disk capacity is installed. FC4010 (Peer-to-Peer VTS Copy Base) and FC4011 (Peer-to-Peer VTS Copy Increment 1) are prerequisites. For the B20 VTS in a PtP VTS, this feature is always required.

- **Peer-to-Peer Copy Increment 3 (FC4013):**

  This feature extends the VTS to function in a VTS Peer-to-Peer VTS configuration when greater than 1 TB of usable disk storage capacity is installed. This feature is required on a Model B20 in a VTS Peer-to-Peer VTS configuration when four of feature FC3705 are installed.

Table 3-4 shows the supported combinations of the Peer-to-Peer Copy features and the Disk storage features.

Table 3-4 Supported feature combinations

Disk storage feature	TVC native capacity (GB)	Peer-to-Peer Copy feature
1 x FC3704	216	FC4010
1 x FC3705	432	FC4010 FC4011
2 x FC3705	864	FC4010 FC4011 FC4012
4 x FC3705	1728	FC4010 FC4011 FC4012 FC4013

3.1.4 Tape Volume Cache features

The usable storage capacity of the TVC depends on the number of the VTS cache features installed. Table 3-5 shows the supported TVC disk storage capacity features for the Model B10 and B20 in a Peer-to-Peer VTS configuration.

Important: We strongly recommend that you install the same amount of Tape Volume Cache in both VTSs in a Peer-to-Peer VTS configuration.

Table 3-5 TVC disk storage capacity features

Features	VTS model	TVC native capacity (GB)[a]	Peer-to-Peer, total usable TVC capacity (GB)[b]
1 x FC3704	B10	216	1296
1 x FC3705	B10	432	2592
2 x FC3705	B20	864	5184
4 x FC3705	B20	1728	10,368

a. The column shows the native, uncompressed cache capacity per VTS.
b. The column shows the total capacity for both VTS subsystems in a Peer-to-Peer configuration, assuming 3:1 compression.

3.1.5 ESCON channel attachment

For the ESCON channel attachments, the following rules apply:

- If the Peer-to-Peer configuration has four VTCs, each of the Model B10 or B20 VTSs must have four ESCON channel attachments.
- If the Peer-to-Peer VTS has eight VTCs, each of the VTSs must have eight channel attachments.
- If your Peer-to-Peer configurations use Model B10s, then you will require two FC3412, and the configuration must contain four VTCs.There is a maximum of four VTCs within a B10 PtP configuration.

Before VTS R7.3, if your Peer-to-Peer configuration uses Model B20s, then you will require four FC3412s, and the configuration must contain eight VTCs. When FC4010 is installed, FC3418 will be disabled.

With R7.3 and R7.4, enhanced attachment options were introduced to be able to build a PtP VTS complex with B20 VTSs and four VTCs only. For additional details, see Chapter 8, “Upgrade scenarios” on page 289.

We summarize the ESCON channel attachment features in Table 3-6.

Table 3-6 Required channel attachment features

Feature combination	VTS model	Number of ESCON channels
2 x FC3412	B10	4
4 x FC3412	B20	8
4 x FC3412 4 x FC3418	B20	16

Each of these connections can use ESCON directors to extend the distance up to 26 km between the VTC and the VTS.

3.1.6 FICON channel attachment

Previously, the channels between the VTCs and host, as well as the channels between the VTCs and VTSs, were solely ESCON channels. Now connecting channels can be either ESCON or FICON. The FICON channels provide significantly higher data bandwidth than do the ESCON channels. The VTC processors supporting FICON are IBM System p servers (7028-6C1) that are designed to provide increased data bandwidth, complimenting the increased channel bandwidth provided by FICON attachment.

The Model B10 and Model B20 Peer-to-Peer VTS are able to exploit the FICON and/or ESCON capability of the Volume Tape Controller. FICON Peer-to-Peer VTS can provide aggregate peak rates at more than twice that of the current ESCON Peer-to-Peer VTS, the Peer-to-Peer VTS FICON enablement also supports greater distances than ESCON.

The VTS FICON attachments support 2-Gbps link speed when attached to a System z server, or an appropriate FICON/FC switch, with 2-Gbps FICON features and the appropriate levels of software to utilize the 2-Gbps capability. Extended distance can be provided with channel extenders certified for use by providers.

Migration from an ESCON-based to a FICON-based environment is a complex multi-staged process that involves partial shutdowns of the VTCs (see Figure 3-1). For an example of the migration to FICON, see Chapter 8, “Upgrade scenarios” on page 289.

From Feature	To Feature	Description
3412	3415	ESCON to FICON Long Wavelength Attachment
3412	3416	ESCON to FICON Short Wavelength Attachment
3418	3415	ESCON to FICON Long Wavelength Attachment
3418	3416	ESCON to FICON Short Wavelength Attachment
3415	3416	FICON Long wavelength to FICON Short Wavelength
3416	3415	FICON Short wavelength to FICON Long Wavelength

Figure 3-1 ESCON to FICON conversion

A FICON attachment can be made to either longwave (9um) or shortwave (50/62.5 um) optical fibers. When planning to upgrade or implement FICON for the Peer-to-Peer VTS, there are some basic things to remember:

- Match LW or SW laser at the other end.
- Match the cabling that you have between the VTC and the next active interface (Host/Switch/DWDM/VTS). 9/50/62.5 um.
- Match the physical size of the connector SC or LC on the end of the cable connected to the next physical plug location (patch panel/Host/Switch/DWDM/VTS).

Figure 3-2 shows the FICON feature codes for the longwave and shortwave cable options. Pick type of host connections and determine VTS connections; this determines FC for the VTC. All FICON VTC in this CX1 must have the same FC.

FICON feature codes for VTCs

Feature code	VTC to VTS S1	VTC to VTS S2	Host to VTC S3	Host to VTC S4
1001	E (2 ports)		E (2 ports)	
1011	L	L	L	L
1012	S	S	L	L
1013	L	S	L	L
1020	S	L	L	L
1014	L	L	S	S
1015	S	S	S	S
1016	L	S	S	S
1021	S	L	S	S
1017	L	L	L	S
1018	S	S	L	S
1019	L	S	L	S
1022	S	L	L	S

Figure 3-2 FICON feature codes for VTCs

Notes: Autosense matches the highest speed allowed on the other end (1 Gbps or 2 Gbps). It can be forced through `smit` to 1 Gbps (if the distance is longer than allowable on multimode for 2 Gbps). Also, LW or SW can change on the other side of a switch or repeater.

We support different distances for LW and SW attachments. See Figure 3-3 on page 70.

Supported native distance

VTC FICON Adapter	SingleMode Fibre (9 Micron) Supported Distance	Multimode Fibre (50 Micron) Supported Distance		Multimode Fibre (62.5 Micron) Supported Distance	
		1 Gbit	2 Gbit	1 Gbit	2 Gbit
Longwave	10 km	550 m		550 m	
Shortwave		500 m	300 m	500 m	300 m

Figure 3-3 LW/SW supported distances

3.1.7 3494 and TS3500 Tape Library attachment

Each of the VTSs must be attached to a separate 3494 Tape Library or to a TS3500 Tape Library (see Figure 3-4). Any combination is supported: 3494 with 3494 library, TS3500 with TS3500 library, and 3494 with TS3500 library. Each of the 3494 tape libraries or TS3500 logical libraries can have any of the supported configurations for VTS and can also be part of a second Peer-to-Peer VTS or may be attached to a stand-alone VTS.

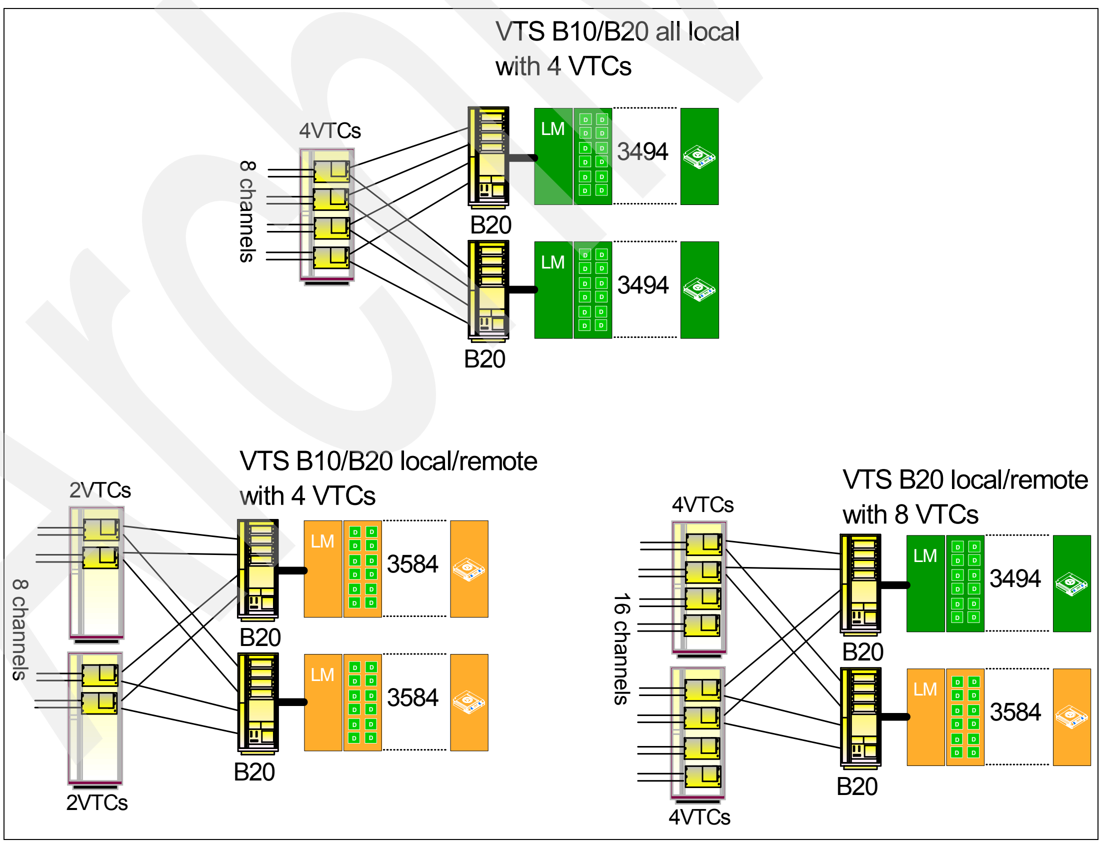

Figure 3-4 PtP VTS configurations

For detailed information about 3494 Tape Library attachment, logical libraries in a TS3500 Tape Library, sharing, and partitioning, refer to *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229, and *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning and Implementation*, SG24-6789.

3.1.8 Sharing and partitioning with multiple hosts

A Peer-to-Peer VTS can be shared by multiple z/OS or OS/390® (MVS) systems in the same way that a stand-alone VTS and a physical tape library can be shared. For details on sharing and partitioning a library, refer to *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229, and *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409.

In a stand-alone VTS, all logical devices and logical control units are associated with one VTS controller. Therefore, the bandwidth available and ESCON channel activities of the VTS do not depend on which logical devices and logical control units are in use.

In a Peer-to-Peer VTS, however, the logical devices and logical control units are associated with four or eight VTCs. Each VTC provides a fixed proportion of the Peer-to-Peer VTS addresses, that is, one or two logical control units with 16 or 32 logical devices, and two ESCON or FICON channels. The total bandwidth of the Peer-to-Peer VTS is divided among the four or eight VTCs.

Therefore, if you want to use the full bandwidth of the Peer-to-Peer VTS when sharing and partitioning a Peer-to-Peer VTS between multiple hosts, each host should attach to all VTCs. In any case, each host should attach to at least two VTCs to provide additional paths in the event of failure. For the purposes of balancing the workload of each VTC, each host should be assigned logical drives equally from each of the VTCs. Figure 3-5 shows a sample drive assignment when partitioning a four-VTC Peer-to-Peer VTS between four host systems.

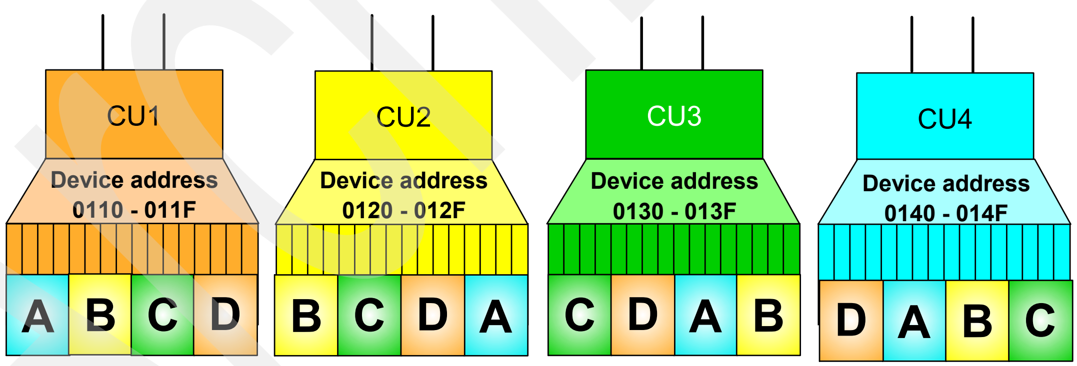

	Host A	Host B	Host C	Host D
CU1	0110-0113	0114-0117	0118-011B	011C-011F
CU2	012C-012F	0120-0123	0124-0124	0127-012B
CU3	0138-013B	013C-013F	0130-0133	0134-0137
CU4	0144-0147	0148-014B	014C-014F	0140-0143

Figure 3-5 Sample drive assignment in a partitioned Peer-to-Peer VTS

3.1.9 IBM TotalStorage 3590 tape drives models B1A, E1A, and H1A

Each VTS in a Peer-to-Peer configuration must have a minimum of four 3590 tape drives installed in the attached Model D12 frames. The Model B10 supports up to six drives and the B20 up to twelve drives. All tape drives in a frame or pair of frames attached to a VTS must be of the same model: either all 3590 Model B1As, Model 3590 E1As, or all 3590 Model H1As.

Using the maximum of twelve drives per VTS is advisable for maximum throughput.

It is not required that both VTS subsystems of the Peer-to-Peer VTS configuration have the same number and type of tape drives installed, especially when you have defined the PtP VTS to work with preference to one site. All writes will go to this site and because of that, more drives will be required there. However, we strongly recommend that you install the same number of tape drives in both VTSs in a Peer-to-Peer VTS configuration to avoid disparities in throughput. It is also helpful to have the same model of drive if possible.

Either length of 3590 cartridge (‘J’ - HPCT or ‘K’ - EHPCT) is supported in a VTS.

The 3590 drives are not supported within a TS3500 Tape Library.

3.1.10 IBM TotalStorage 3592 tape drives Model J1A and E05 in a 3494 library

With VTS Release 7.3 we introduced homogeneous and heterogeneous tape drives supporting the VTS. Homogeneous support is available for VTS Models B10 and B20 and includes only one drive type per VTS, either all 3590 or all 3592 drives. For the Model B20 VTS, we also have heterogeneous tape drive support, which means that you can have both 3590 and 3592 tape drives installed in the same VTS at the same time.

IBM 3590 drives reside in a 3494 Model D12 frame, and 3592 drives reside in a Model D22 frame. You cannot intermix both drive types within the same frame, nor can you intermix different 3590 models in the same D12 frame. All 3590 Tape Drives attached to a VTS need to be the same model.

The Fibre Channel 3592 tape drives must be attached to two SAN switches resident in the D22 frame, and these in turn are attached to two HBAs in the B10 or B20 VTS. The dual Ethernet LAN attachment features FC5245 are required, and in a HA1 environment the optional 3494 D22 dual PCC can be used to attach drives and SAN switches to different power sources to improve availability. Feature code 5238 provides the HBAs and the cables to attach the B10 or B20 VTS to the SAN switches installed in the 3494 D22 frame. Feature code 3486 provides the mounting hardware for the SN switches in the 3494 D22 frame, and two 3487 feature codes need to be ordered to add the two SAN switches. Drives are then attached by ordering one 3060 feature code per drive.The different maximum configurations are shown in Figure 3-6.

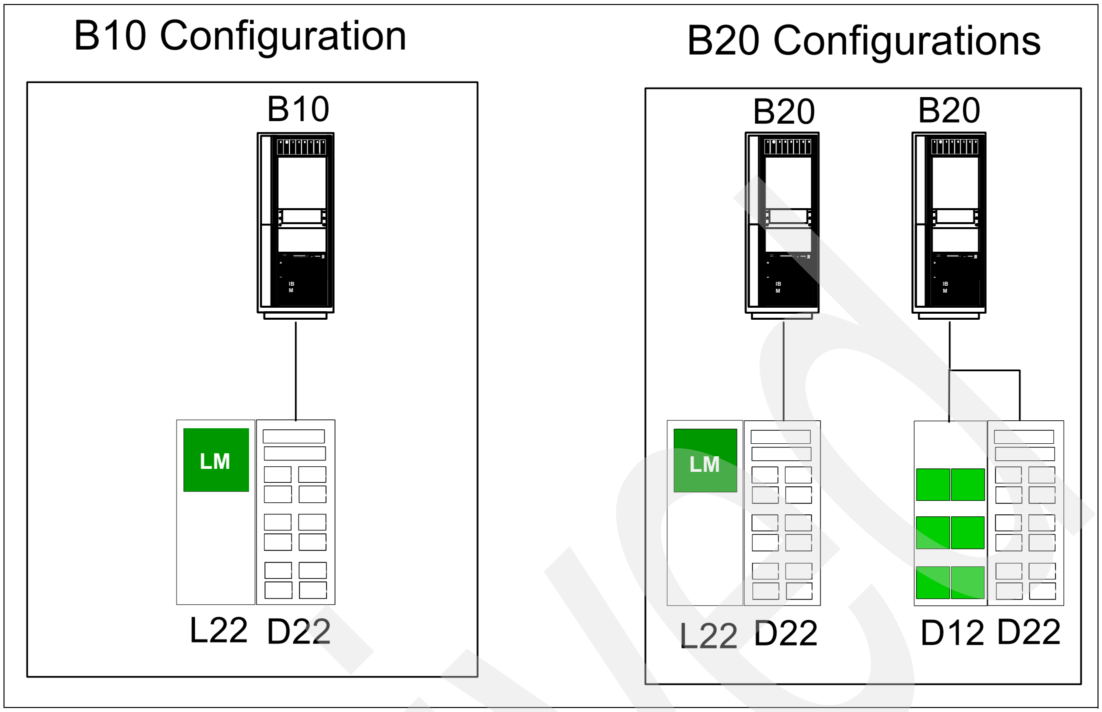

Figure 3-6 VTS Configurations with 3592

When only a single drive type and media type are installed and used, the physical volume pooling functions of the Advanced Policy Management feature are optional. When both media types (60 GB and 300 GB) are to be used, physical volume pooling can be used to separately manage each media type.

When both 3590 and 3592 drive types are installed, physical volume pooling of the Advanced Policy Management feature is required for their management. One or more physical volume pools are defined for each drive type and an individual pool can only support a single drive type. The definitions are made during the pool set up. The most important rules are that you have to have a minimum of four drives per device type and you cannot install more than 12 drives per device type, but the number of supported drive frames remains the same. This is one frame for the B10 VTS and two frames for the B20 VTS if both drive types are connected. You need to install all 3592 drives of a single VTS in one Model D22 frame. Currently, if a B20 VTS has more than six tape drives installed, the two D12 frames need to be installed adjacent to each other and at a maximum distance of 14 meters from the VTS. In a heterogeneous environment the D12 frame and the D22 frame do not need to be installed adjacent to each other. The D12 frame still cannot be installed more than 14 meters away from the VTS, but for the D22 frame a maximum distance of 25 meters is supported.

The Model B10 VTS can have up to six 3590 Tape Drives installed or up to 12 3592 Tape Drives. This doubles the number of tape drives supported on a B10 VTS. For new installs in environments with higher recall rates but without higher throughput requirements, the B10 VTS with 12 3592 drives might be an option. Upgrades from 3590 to 3592 are supported through the upgrade of the B10 to B20 VTS upgrade. For the B20 VTS, a maximum of 18 tape drives will now be supported as the maximum configuration in a heterogeneous environment. Installation options for B20 VTS subsystems, the 3590 drives only remain the same as before. When only 3592 drives are to be used, the B20 VTS can attach up to 12 drives. Existing B10 VTS installations cannot be upgraded to 3592 in the field, as mentioned before already. For the Model B20 VTS, there are different ways to attach additional drives and frames to the installation. See Chapter 8, “Upgrade scenarios” on page 289.

If a B20 VTS has more than six 3590 drives attached already you can either convert an existing D12 frame, VTS attached to a D22 frame, or you can add a D22 frame. The second D12 frame in the 3590 drives, which will no longer be attached to the VTS after the upgrade, can be de-installed or reused as required.

You can also add a D12 frame and up to six 3590 drives to a B20 VTS, which has only 3592 drives installed. This might be required when merging workloads using import/export, for example.

Here is a summary of the configuration rules:

- A minimum of four drives of any drive type is supported.
- A maximum of twelve drives of any one type is supported:
  - 3590 or 3592 for B20
  - 3592 only for B10 — with 3590, a maximum of six drives is supported
  - For B20, only a mixture of twelve 3592 drives and six 3590 drives
- A maximum of twelve 3592 drives is supported:
  - All in one frame, if the VTS is attached to a 3494 library.
  - Distributed across D22 frames, if the VTS is attached to a TS3500/3953 library.
- A maximum of six 3590 drives can be installed per frame.
- IBM 3590 drives are all SCSI attached.
- A maximum of two drive frames for the B20 is supported in a 3494 library:
  - D12 frames must be adjacent.
  - Mixed D12 and D22 frames do not need to be adjacent.
- A maximum of one drive frame for the B10 is supported inside a 3494 library.
- Both IBM 3592 models are supported and can be intermixed behind one single VTS. The 3592-E05 attached to a VTS always operates in J1A emulation mode, providing similar performance and capacity to J1A drives installed.
- Existing Model J1A drives cannot be replaced by 3592-E05 drives. However, you can add 3592-E05 drives to a VTS which already has 3592-J1A drives attached.
- A maximum of twelve 3592 drives can be installed independent of the drive model.

In the VTS, both the 3590 J and K cartridges are supported. With 3592 drives installed, the 3592 JA cartridge with 300 GB capacity is supported as well as the JJ cartridge, which provides a capacity of 60 gigabytes uncompressed. In the VTS, neither WORM cartridges nor performance scaling or cartridge segmentation are supported. The full capacity of 3592 JA and JJ cartridges is used in the VTS.

3.1.11 IBM 3592 in an IBM TS3500/3953 Tape Library

Support for the IBM 3953 Tape Frame requires a functional microcode update to an installed IBM 3592 Tape Controller Model J70 and IBM 3494 VTS Models B10/B20. Any System z IBM 3592 Tape Controller Model J70 or VTS that needs to attach to IBM 3592 Tape Drives Model J1A in the IBM TS3500 Tape Library requires a microcode firmware upgrade. This microcode firmware level will be shipped on all new IBM 3592 Tape Drives Model J1A, IBM 3592 Controllers Model J70 and VTS. Already installed IBM 3592 Tape Drives Model J1A or IBM 3592 Controllers Model J70 or VTS can have the new firmware update installed by ordering the following features:

- For installation of System z attached IBM 3592 Tape Drives Model J1A in an IBM TS3500 Tape Library, the Library Manager Functional Enhancement Field feature (FC0500) should be ordered for each IBM 3592 Tape Drive Model J1A.
- For installation of an IBM 3592 Controller Model J70 in an IBM 3953 F05 Frame, the Functional Enhancement Field feature (FC0520) should be ordered for each IBM 3592 J70. This feature will also update the drive microcode firmware on any attached IBM 3592 Tape Drive Model J1A.
- For attachment of an existing IBM 3494 VTSs to IBM 3592 Tape Drives in a IBM TS3500 Tape Library, the VTS Functional Enhancement Field feature (FC0521) should be ordered for each VTS.

For each System z tape controller or VTS:

- You should set up a minimum of four drives, where there are four or more attached drives.
- You should set up all the drives, where there are one through three attached drives.
- You may not set up more than eight drives.

Note: The optimum number of control path drives per tape controller or VTS is *four*.

For additional information on the installation, refer to the redbook, *IBM TotalStorage TS3500 Tape Library for zSeries Hosts: Planning and Implementation*, SG24-6789.

3.1.12 Virtual Tape Controller

The Virtual Tape Controller (VTC) is one of the unique components of the Peer-to-Peer VTS configuration. A VTC is an IBM System p server that is used to connect the two VTSs in the Peer-to-Peer VTS configuration with the host system. Table 3-7 shows the connectivity options available within a CX1 frame.

Table 3-7 ESCON & FICON Options for CX1

VTS model	ESCON	FICON	VTCs	CX1 frames
B20	8		8	2 - 4
B20		8	8	2 - 4
B20	4	4	8	2 - 4
B10	4		4	1 - 2
B10		4	4	1 - 2
B10	2	2	4	1 - 2

If you require eight VTCs, you need to have eight channel attachments on each VTS (and two CXx frames) as described in “ESCON channel attachment” on page 67.

All FICON VTCs in a CX1 must have the same configuration. B10/B10 mixed ESCON/FICON must have ESCON VTCs or AX0s in the top two positions in CX1. B20/B20 mixed ESCON/FICON ***must*** have all ESCON VTCs in one CX1 frame and all FICON VTCs in another CX1 frame.

VTC LAN connection for the PtP Specialist

Each VTC provides an Ethernet connection for the customer LAN. When this connection is used, each of the VTCs provides the ETL Specialist Web interface, which is a Web page for monitoring the status of the Peer-to-Peer VTS. This is the only interface to monitor the Peer-to-Peer VTS except by using the SMF type 94 record. We highly recommend that you use the Peer-to-Peer VTS Specialist to monitor the state of the Peer-to-Peer VTS.

> **Note:** The VTC does not support a token ring LAN connection for providing the ETL Specialist Web interface.

Standby VTCs

If you configure a Peer-to-Peer VTS with one VTS remote from the other, you can choose between two types of configuration for the VTCs:

- CXn frames at each site, with VTCs distributed between the two sites
- CXn frames at the local site with all the VTCs (either four or eight)

If you choose the second option and you have host systems at the remote site, you can optionally place an additional CXn frame at the remote site with four Standby VTCs in it. The Standby VTCs are not normally attached to any host or to the distributed VTS, but they have been set up with TCP/IP information to allow ETL Specialist access. Before they can be used, they must be cabled. If you also have one or more host systems at the remote site and you have copies of the tape inventory data sets, you can make use of Standby VTCs in one of these ways:

1. In the event of a disaster at the local site, the remote distributed VTS can be accessed using the Standby VTCs in Read/Write Disconnected mode. Read/Write Disconnected mode is described in 2.6.7, “Special VTC operational modes” on page 54, and can only be used if the local site is not operational.
2. If you want to test your disaster recovery capabilities at the remote site, you can do so without interrupting the production work at the local site. The Standby VTCs must be started in Read-Only mode only. Read-Only mode is described in 2.6.7, “Special VTC operational modes” on page 54.
3. Another way of preventing writing to the local VTS is to use the stand-by VTCs in Write-Protect mode. Write-Protect mode is described in 2.6.7, “Special VTC operational modes” on page 54.

The procedures for exploiting Standby VTCs are described in Chapter 9, “Disaster recovery scenarios” on page 321.

Model CX0 and CX1 frames

The Model CXn frame is another of the unique components in a Peer-to-Peer VTS configuration. It provides the housing and power for two or four VTCs. Two VTCs in a Model CXn frame must be within the same Peer-to-Peer VTS. An additional two VTCs in the Model CXn frame can be used for another Peer-to-Peer VTS, or all four VTCs in a Model CXn frame can be for the same Peer-to-Peer VTS.

The Model CXn frame contains two power control compartments (PCC), each with its own power cord, which allows connection to two power sources. Figure 3-7 shows the Model CXn frame.

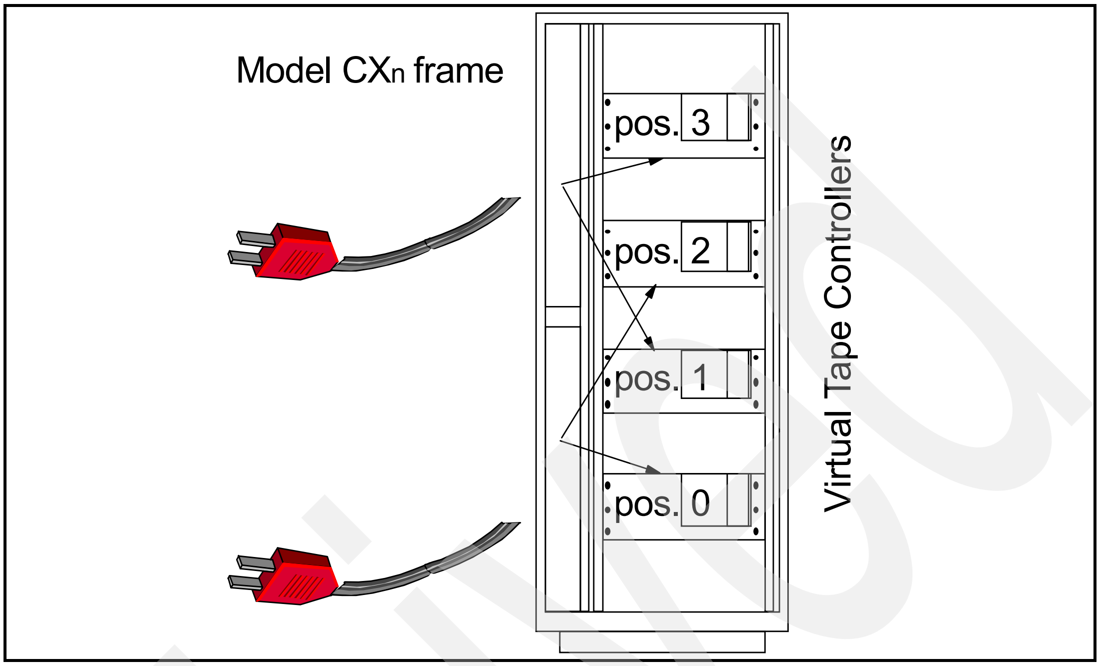

Figure 3-7 Model CXn frame and power scheme

The Model CX0 and Model CX1 frames do have different specifications. The CX1 was introduced with the availability of FICON for the VTS. There are certain rules to be considered with regard to the CX0 frames when upgrading to FICON capable VTCs. The first of these is that the frame must be upgraded to a CX1 frame.

Model Cx0 frame

The position numbers for the AX0s are 0, 1, 2, and 3 (see Figure 3-7). Position 0 is the lower-most position and position 3 is the upper-most position. When only two Model AX0s are in the Model CX0, they are in positions 0 and 1. The lower PCC powers AX0s in positions 0 and 2. The upper PCC powers Model AX0s in positions 1 and 3.

Model CX1 frame

This frame can be used to house both the FICON capable VTCs and the ESCON only VTCs (formerly known as AX0s). If the configuration allows a mixture in the CX1 frame then the AX0 devices must occupy the top two positions of the frame. The mixture of devices is only supported in the Model B10 VTS. All VTC models must be identical in the CX1 frames of a Model B20 VTS.

3.1.13 FICON Performance Accelerator Feature (FPAF)

The FICON Performance Accelerator Feature - FPAF (FC5250) was introduced with the availability of FICON for the B20 VTS. It adds two more processors to the VTS and increases the total throughput capability.

The B20 with the FICON Performance Accelerator Feature showed the greatest potential performance benefit with 3592 drives. Up to the maximum of twelve total 3592 drives provided benefit for host data-intensive workloads. Further, in these laboratory measurements, for a given number of active drives, the FPAF B20 usually provided for higher data rates than did either the B10 or base B20. When considering the question of whether fewer 3592 drives can be configured to give performance equivalent to 3590 drives, remember to factor in all types of VTS activity that utilize physical drives, including reclaim and import/export. It is also wise to provide a safety margin to allow for occasional drive outages. The use of VTS pooling may also dictate having more drives than aggregate performance needs would indicate.

3.1.14 TotalStorage Specialist configuration

The IBM TotalStorage Enterprise Tape Library Specialist and the IBM TotalStorage Peer-to-Peer Virtual Tape Server Specialist are members of the TotalStorage product family and provide comprehensive status, capacity and performance information. Enterprise Specialists, an IBM Enterprise Storage Resource Management (ERSM) solution, are members of a growing software family whose goal is to enable storage administration to efficiently monitor storage resources from any location within the enterprise. Widely dispersed, disparate storage resources will ultimately be viewed and managed through a single, cohesive control point. The TotalStorage Specialist family includes a number of storage “Specialists,” interfaces that enable the customer to monitor and manage specific storage devices. The following Specialists are available for the library and VTS:

- IBM TotalStorage Enterprise Tape Library Specialist enables you to monitor the 3494 and the TS3500/3953 Tape Library. It is a system independent, Web based user interface to the Library Manager, which enables the customer to monitor the Library Manager and attached Virtual Tape Servers from a remote location.
- IBM TotalStorage PtP VTS Specialist is the Web user interface used to view the current status and configuration of the PtP VTS. This Specialist requires a customer supplied Ethernet network which is connected to the Virtual Tape Controllers (VTCs). Commonly used browsers such as Microsoft® Internet Explorer or Netscape Navigator can be used to view the information. Information applying to all attached 3494 VTS systems is available on a read-only basis and is updated periodically by the subsystems.

There are appropriate worksheet(s) available that should be completed and supplied to service personnel prior to installation of the 3494 or TS3500/3953.

http://w3-03.ibm.com/support/assure/assur30i.nsf/WebIndex/SA185

The Specialists have been tested with Microsoft Internet Explorer 5 and Netscape Navigator 4.7, both with Java™ and JavaScript™ enabled.

IBM 3494 and 3953 ETL Specialists

The ETL Specialist functions for Advanced Policy Management are described in detail in *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229. The ETL Specialist has:

- A home page
- A set of Library Manager pages
- A set of VTS pages (if a VTS is installed)
- Links to the 3494 Peer-to-Peer VTS Specialist (if a Peer-to-Peer VTS is installed)

The Library Manager must be connected to your LAN using either a Token Ring Adapter or an Ethernet Adapter. If the 3494 Tape Library is part of a Model HA1 high availability configuration, the stand-by Library Manager also must be connected to the LAN. We recommend the use of the Ethernet Adapter to avoid the installation of any bridges between token ring and Ethernet because you have no choice between Token Ring and Ethernet for the VTCs.

During the installation process, the service representative 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 must provide the following information before the installation starts (items marked * are optional):

- TCP/IP host name for the Library Manager
- TCP/IP address for the Library Manager
- Subnet mask (or network mask)
- Router address (or gateway address)*
- Domain name*
- Name server address*

Peer-to-Peer VTS Specialist

The VTC provides an optional Web server for the display of status on your Web browser by way of your intranet. In a Peer-to-Peer VTS, each of the VTCs can be configured to provide the Web server function. The Peer-to-Peer VTS Specialist only provides monitoring functions. There is no capability to make changes or issue commands. Refer to 7.1.1, “Monitoring using the IBM Specialists” on page 248.

> **Note:** The Peer-to-Peer VTS Specialist does not allow any access to user data contained in the logical volumes.

You must provide the following hardware and information to the service representative before the installation:

- An Ethernet/twisted pair cable (category 5, 4-pair stranded, PVC, RJ45 connector) for connection to a 100/10 Base T connector on each VTC
- Standard Ethernet Network Interface (en0)
- Network names and addresses are required for each VTC (items marked * are optional):
  - Host name for each VTC
  - Internet address of each VTC
  - Network mask
  - Name server Internet address *
  - Name server domain name *
  - Default gateway address*

Figure 3-8 shows the 3953 connections to the LAN and the TSMC.

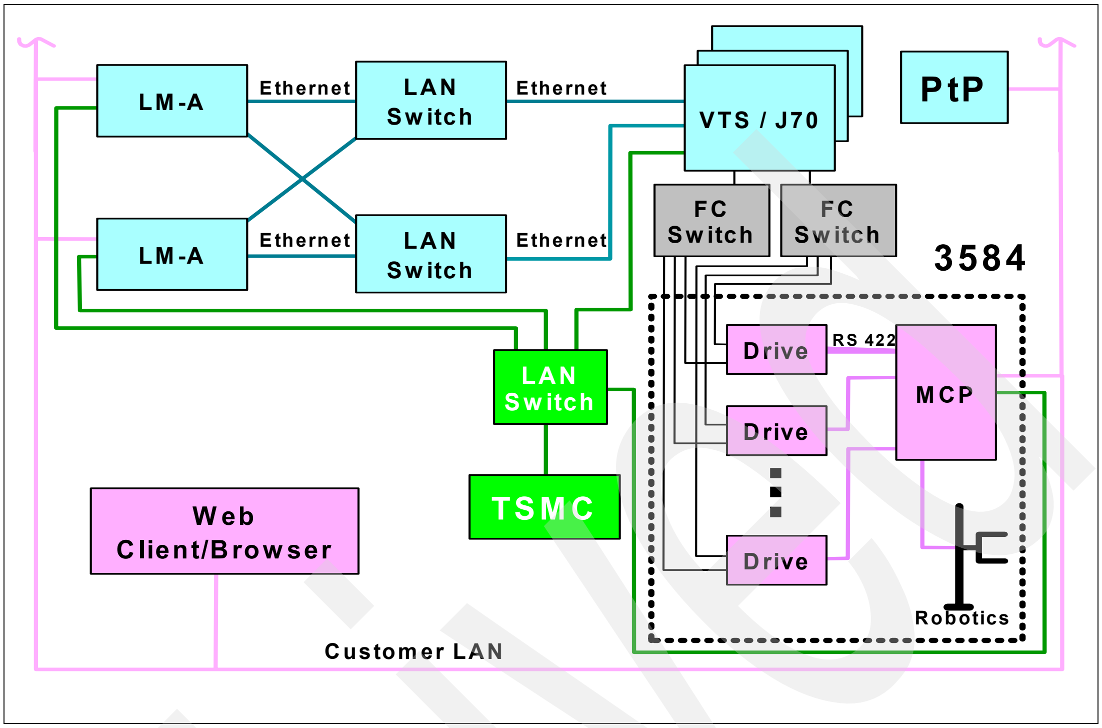

Figure 3-8 3953 interfaces to LAN and TSMC

3.1.15 IBM System Storage TS3000 System Console

The IBM System Storage TS3000 System Console (TSSC) replaces the previously used TotalStorage Master Console (TSMC). IBM VTS models have a function to call the IBM support center with summary information when they detect unusual conditions. Remote Support Feature (FC2710, FC2711, or FC2712) has been introduced from the first generation of the VTS and it is a required feature for the reporting of errors and for remote diagnostic support.

VTSs can be configured to automatically execute a *health check* and send the results to the TSSC for transmission to IBM for review for preventive action. The TSSC integrates service monitoring of up to forty-three VTSs and/or Virtual Tape Controllers. It centralizes microcode maintenance, service terminals and enhances remote support capability.

The TS3000 System Console (FC2713, FC2714 or FC2715 + FC2716) is designed to protect customers' investment in Virtual technology by centralizing the IBM support and monitoring aspects of single or multiple Virtual subsystems. This feature is designed solely for the use of the IBM service representative. Figure 3-9 describes TSSC connections.

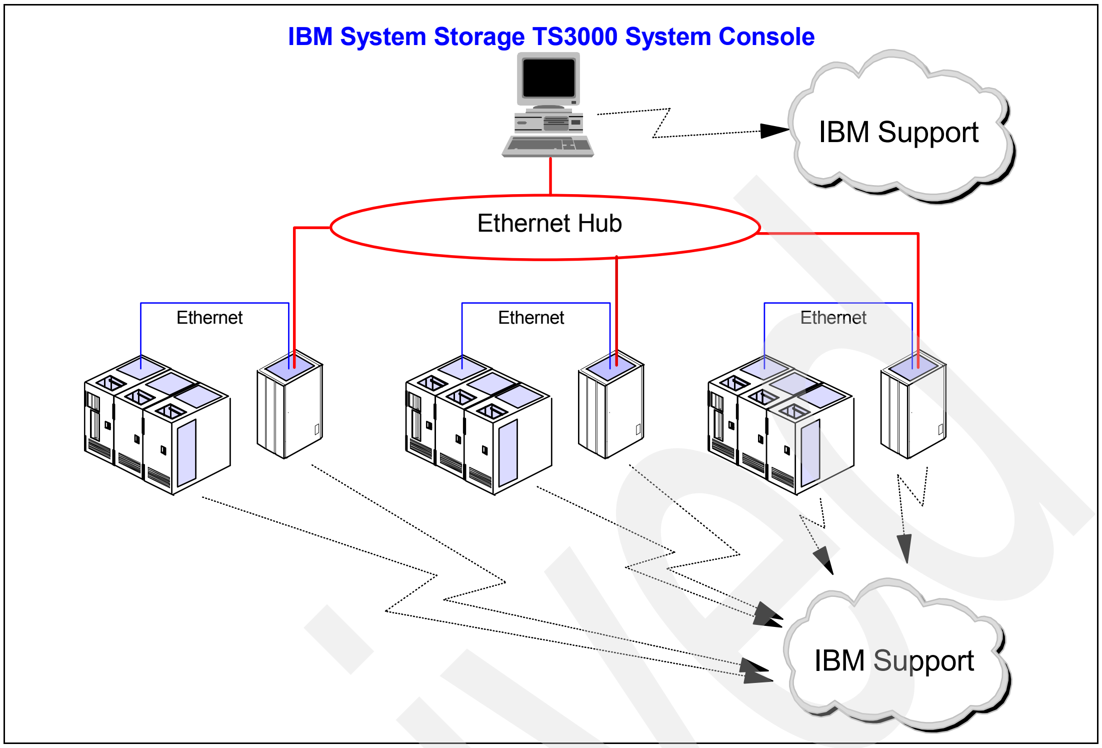

Figure 3-9 IBM System Storage TS3000 System Console connectivity

Attachment to the TSSC is an installation requirement for any VTS or 3592 Models J70 and C06 attached to an IBM TS3500 Tape Library. The TSSC can provide Remote Support for as many as 43 attached tape systems using feature-provided switches and ethernet cabling.

The IBM 3953 Library Manager is enhanced to support the TSSC and can join the dedicated local area network. The TSSC attaches to each installed IBM 3953 Model L05 Library Manager, IBM 3592 Model J70 or C06 Controller, and IBM 3494 B10/B20 VTSs. The TSSC also attaches to the IBM TS3500 Tape Library. The TSSC monitors these components for early detection of unusual conditions. Error information is sent automatic to IBM support called Remote Technical Assistance Information Network (RETAIN®).

Remote Support capabilities

The TSSC support includes:

- Call Home problem reporting capability with staged, error specific Data Gathering for support (via modem)
- Call-In capability with authenticated access (via modem) including file transfer and multiple connections with attached systems
- Simultaneous call home and call in capability using dual modems (stand-alone Master Console only)
- Automatic health checking for attached systems
- Ability to defer disposition code 2 call home activity to specific business hours
- Automatic archival of log files for most attached tape systems for subsequent support reference
- Automatic download and storage of new tape drive code images via modem connection with IBM Remote Technical Assistance Information Network (RETAIN)

Additionally, the TSSC provides a convenient focal point for local service activities within the data center. The TSSC provides the following local service tool applications:

- Ability to telnet to multiple tape systems and simultaneously perform multiple service tasks from the Master Console
- Graphic user interface for tape system and tape drive service diagnostic utilities
- Ability to broadcast control unit and tape drive code images to tape systems for subsequent activation from the Master Console
- Diagnostic tools for verifying communications with tape systems and IBM RETAIN
- Graphic user interface for configuring, backing up, and restoring Master Console settings
- One-time authentication for logging in to multiple attached subsystems

Feature FC2718 for the IBM TotalStorage 3953 Tape Frame Model F05 provides the rack-mountable TS3000 System Console, an Ethernet switch for the TSSC, and a 1.0 m cable and connector for connection from the TSSC to the Ethernet switch in the IBM 3953 Model F05 frame. Feature FC5505 on the IBM 3953 Model F05 frame, as well as feature FC5511/FC5510, are required.

Attached subsystems

The same TSSC may be shared between the following units, up to a maximum of 43 units:

- IBM 3494 B10, B18, B20, HA1, L10, L12, L14, or L22 Frame
- VTC in a IBM 3494 CX1 Frame
- IBM 3590 A60 Controller
- IBM 3592 J70 Controller
- IBM 3953 Library Manager
- IBM TS3500 Tape Library

Existing TSSC or TotalStorage Master Console

If you already have a TSSC or TotalStorage Master Console attached to your tape system, it is possible to connect it to your IBM 3953 Tape System, but you might not be able to use the keyboard, video, and KVM switch inside the IBM 3953 Base Frame because the existing Master Console is located outside the IBM 3953 frame. Perhaps it is located on a table near your IBM Tape system, and this Master Console is a normal PC running Linux®, so you have to use the existing Master Consoles keyboard, video, and mouse. If so, consider this solution.

FC2714 provides for attachment of the IBM 3953 Model L05 Library Manager to an existing TSSC. This feature provides a 50-foot attachment cable, rack-mountable Ethernet switch, and associated mounting hardware to attach to an existing external TSSC.

3.2 Planning for Advanced Policy Management exploitation

Preparing for the exploitation of the multiple functions offered by APM requires careful planning. Investigate first, which of the functions can help you manage and improve your tape production.

The feature codes required for APM are dependent on the amount of cache feature codes you have installed (see Table 3-8)

Table 3-8 Advanced Policy Management disk capacity feature requirements

Disk storage feature	TVC native capacity (GB)	Advanced Policy Management features
1 x FC3704 or 3 x FC3705	Up to 250 GB	FC4001
1 x FC3705 or 4 x FC3702	Up to 500 GB	FC4001 FC4002
2 x FC3705	Up to 1000 GB	FC4001 FC4002 FC4003
4 x FC3705	Up to 2000 GB	FC4001 FC4002 FC4003 FC4004

Note: APM (FC4001 - 4004) must be installed on both VTSs in a PtP Subsystem.

3.2.1 Logical volume affinity considerations

Pooling of logical volumes may have an effect on the throughput of your VTS. If you are introducing pooling at your site, consider the following possibilities:

1. The possible increase of concurrent 3590 and/or 3592 drive usage within the VTS. Depending on the number of pools and the amount of logical volume data being created per pool, you need to plan for the following requirements:
   - Ensure that sufficient drives are available to handle:
     - Premigration 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

     See Figure 3-10 for an explanation of the changed process.
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 that you maintain a sufficient amount of physical VTS 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. VTS processing/cache capacity.

3.2.2 Selective Dual Copy within a VTS

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

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

Consider that dual copy within a PtP VTS complex means that you get four copies of each logical volume. It may make sense if you want to have dual copy within that VTS, which is defined to store the volumes that have been excluded from being copied automatically by the PtP selective copy function. If you really need more than two copies, you have to plan for additional throughput and capacity.

3.2.3 Larger logical volumes support

With the exploitation of larger logical volume sizes, fewer logical volumes are needed for large data sets that exceed the 800 MB barrier of today. If you need five logical volumes, each with 800 MB today for a specific large data set, you can store this onto one logical volume with 4000 MB now. In case of a recall, only one recall action is necessary, and one physical tape mount instead of five recalls (and in the worst case, five physical mounts), which interrupts the processing of the job if there is no spanning of this large volume.

Changes with larger volume sizes: The probability of more spanning on the physical volume is increased as you select larger volume sizes. As an example, with 3590 model E drives and J media, the uncompressed capacity is 20 GB. So with 2000 MB logical volume sizes, we would get 10 volumes per physical volume and we would expect one out of 10 accesses to require a span compared to one out of 25 with the 800 MB logical volume size.

There is a higher probability that a recall will require two physical mounts, which is still less than today for large data sets that span five or more logical volumes.

Furthermore, the cache premigration thresholds needed to be changed because more space is needed to open large virtual volumes residing in cache:

- For logical volumes with 2000 MB maximum size, the threshold is now the installed cache size minus 400 GB
- For 4000 MB volumes, it is set to 700 GB.

The result is that less data can be written to the cache before pre-migration throttling may occur.

Attention: The change of the premigration thresholds will take place with the activation of the larger volume size support by the CE. Even without any use of this function, you may experience some performance degradation.

We recommend not to use these large volume sizes where it makes no sense at all. You may plan to use only one Data Class with 4000 MB volume size, as no physical space on the cartridge itself is wasted but the overall performance of your system will be affected.

The 4000 MB size is available only for B20 configured with maximum cache and 3592 drives because if all 256 virtual drives are mounted with 4000 MB volumes, that would exceed the amount of space needed with only 864 GB of cache.

The maximum logical volume size is set by the CE during microcode installation. The maximum size of 4000 MB can only be set if the configuration requirements are met: this will be checked by the microcode.

The definition of a Data Class is restricted to the logical volume sizes which are supported in your installation. The size of a logical volume can only be changed when it is mounted as a scratch volume. If you mount a logical volume with DISP=OLD (specific mount), and you open a data set on this volume for “Write,” then the size will remain as it was.

3.2.4 Secure Data Erase implications

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. Fewer physical drives may 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.

As the first logical volume that expires triggers the physical volume to be erased, almost a full physical cartridge will be first reclaimed and secondly 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 subsequently erasure operations take place. Grouping by expiration date may help to reduce unnecessary reclamation and erasure activity.

3.2.5 500,000 logical volumes support

With VTS R7.4 you can expand the number of logical volumes to 500,000 per VTS, or up to 1,000,000 per Library Manager (3494 or 3953).

To get the first portion of additional 50,000 logical volumes, you need the logical volume expansion feature (FC4036). Each feature increases the logical volume capacity by 50,000. Up to five features can be factory or field installed. The first feature is disruptive because the file and metadata systems must be reconfigured to account for the increase in the number of files managed to 1,000,000 (500,000 are needed for the primary logical volumes file and 500,000 for the secondary). The first feature configures to the maximum so subsequent expansion features are concurrently installed.

The model B20 VTS looks for added features 4036 and calculates the maximum number of logical volumes supported and provides a supported count field to the Library Manager. The Library Manager allows volumes to be inserted up to the count provided by the VTS. Insertion of more than the allowed number of logical volumes is prohibited.

Feature code 4036 requires larger databases and larger file systems. If you are planning for a data migration, the following considerations apply:

- If the source VTS has FC4036 installed, then the target VTS *must* also have it installed.
- If the target VTS has FC4036 installed, the source VTS may or may not have FC4036. The migration diskettes will expand file systems in this case.

This support is reduced to model B20 with the dual control path feature in a base or Peer-to-Peer configuration. It requires 3590 mode E or H drives or 3592 drives. For database backup, an additional 250 MB per cartridge is used.

3.2.6 Cache management advantages

There are no negative effects for the overall VTS performance expected. Nevertheless, plan with care which data you want to be kept in cache and which data you want to be candidates for immediate “deletion”.

Identifying suitable data for cache management

It would be wise to move data gradually to Storage Classes that enable cache management. The objective is to identify data that need not stay in cache. We can break data into three categories:

- **Expected to be suitable for early removal from cache:** We would expect to include database logs, image copies, archival data, and long-term backups.
- **Not certain whether suitable for early removal or not:** This is data that does not fall into the other two categories. It is likely that the effort of identifying candidates in this category would outweigh the benefits.
- **Not suitable for early removal from cache:** This data would have a high chance of reuse. It could include DFSMShsm migration data or application data written directly to the VTS.

Tools

The appropriate analysis tool, VOLREUSE, is available via anonymous FTP from:

`ftp://ftp.software.ibm.com/storage/tapetool`

It will take SMF type 14, 15, 21, and 30 records as input and will show what percentage of VTS write activity (at both an hourly and daily level) is never read again and therefore is an ideal candidate for early removal from cache. The TGROUP control statement allows you to keep data together for address ranges that you specify. This will allow you to produce separate reports for different VTSs.

A *dsname* filter list can be generated to identify the data sets to be marked as immediately eligible for flushing from the TVC. At a later stage, this list can be provided to BatchMagic to allow a simulation showing the improvements in read hits by using cache management.

3.2.7 PtP VTS copy mode control

In a Peer-to-Peer VTS, there are two types of copy modes; Immediate where a copy is scheduled as soon as the original is created and the application waits until the second copy is created before ending; Deferred where the second copy is queued but not created and the application is allowed to end. In the previous releases of PTP VTS, this copy mode is defined globally with the virtual tape controllers (VTCs) and could only be changed by CE.

Advanced Policy management allows the users to control the copy behavior of individual logical volumes via the Management class. This means that critical volumes can be copied immediately and non-critical volumes can be deferred. In an environment that uses Immediate copy, this could effectively increase the available bandwidth for critical volumes, improving the performance of the PTP VTS. Settings for the PTP copy are Immediate, Deferred, No Copy, and VTC defined.

The option to create no copy in the other VTS is called PtP VTS Selective Dual Copy. With the PTP Selective Dual Copy function, additional actions can be specified during the definition of a Management Class constructs to allow you to specify:

- Which VTS is to be used in writing data for a scratch mount
- That no copy is to be performed

There are several steps to be taken before the PTP Selective Dual Copy function can be used:

1. Review the PTP configuration(s) that are installed. The two VTSs that are part of a PTP configuration are labeled distributed VTS 0 and distributed VTS 1. You will need to understand this relationship when you define the Management Class constructs later on. If the library has two VTSs associated with it, the two VTSs are labeled as VTS 1 and VTS 2, however, that labeling has no relationship to the distributed VTS labeling.
1. Decide on the Management Class construct names to be used to control the function. We suggest that the names relate to the desired function. For example, if the desired action is to have the local VTS be the selected I/O VTS for scratch mounts and the data is not to be copied, a name like "MCLOCLNC" would be a good choice. "MC" for Management Class, "LOCL" for local VTS and "NC" for no copy. Likewise, if the desired action is to have the remote VTS be the I/O VTS for scratch mounts and the data is not to be copied, a name like "MCRMOTNC" would be a good choice. It is also suggested that a Management Class construct name be selected for the default action of letting the VTC determine the I/O VTS and to make the copy. An example could be "MCVTCCPY".

Keep the following considerations in mind:

- The Management Class construct names and their associated actions for using PTP selective dual copy must be identical on both libraries; otherwise inconsistent results may occur.
- If there are more than one VTS installed in the library, the construct definitions apply to both VTSs.
- Specifying a secondary pool has no influence on whether or not a copy is made between VTSs in a PTP configuration. If a secondary pool is specified, a second copy of the logical volume is made within the VTS associated with the Library Manager used for the definition of the Management Class.
- A PTP Copy Control selection other than VTC defined will override any default VTC setting or a setting established as part of GDPS support.
- A PTP I/O VTS selection other than VTC defined will override any default VTC setting or a setting established as part of GDPS support.

Define the Management Class construct names and their associated actions through the Library Manager console or Specialist (WEB) interface. This is done through the *Manage Management Classes* panel selected from the Systems Management pull down on the Commands menu. Refer to Chapter 4, "Implementation" on page 127.

You should take this into consideration when you do your sizing using the BatchMagic tool. BatchMagic honors the different copy modes.

PtP Selective Dual Copy was introduced with VTS code level 2.30.720.40 or later and Library Manager code level 529.01or later which were made available in May 2004. Although existing host software levels that support APM can exploit the function, it is recommended that either the SPE for 3592 additional media type support (the WORM capable media, as well as the economy cartridge) be installed, or the following APAR (depending on OS level) if the additional media type support is not needed or desired:

- For OS 390 V2 R10 or z/OS V1 R1 or V1 R2 — OA06497
- For z/OS V1 R3 or V1 R4 — OA06492
- For z/OS V1 R5 and R6 — OA07105

3.3 Host software requirements

For z/OS, the Peer-to-Peer VTS is supported only in a DFSMS/MVS system-managed tape environment. The software levels required to support the Peer-to-Peer VTS are z/OS V1R4 and above. Support was initially provided at earlier release levels, however, those release levels are no longer in service. Additional PTFs are required for some of the VTS Advanced Functions features. Refer to PSP buckets D/TVTS and D/TPTP for additional recommendations.

A Peer-to-Peer VTS may be attached to hosts running either VM/ESA or VSE/ESA as a VM/ESA guest as long as the Peer-to-Peer VTS is also attached to an z/OS host. The levels of VM/ESA and VSE/ESA supported are the same as for a stand-alone VTS.

Support for IBM TS3500/3953

Support for System z attachment of the IBM TS3500 through the IBM 3953 Library Manager is provided for the following operating systems:

- **z/OS:** Recommended support for the TS3500 Tape Library with the 3953 Library Manager is provided at z/OS V1R4 and above. For more information on the support provided, refer to APARs OA09751, OA10566, OA09753, and OA10566.
- **VSE/ESA and z/VSE™:** Support for the TS3500 Tape Library with the 3953 Library Manager is provided at VSE/ESA 2.6 and above.
- **z/VSE 3.1:** Support for all 3592 J1A configurations without APARs in all automation offerings, including 3494 and TS3500, is provided.
- **z/VM:** Support for the TS3500 Tape Library with the 3953 Library Manager is provided for z/VM® 3.1.0, 4.4.0, and 5.1.0 with DFSMS/VM FL221 (function level 221).

Additional software support may be required if Advanced Policy Management is installed on the VTS. Support (both Exploitation and Toleration) for this advanced function is covered in *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

Note: Support for the Peer-to-Peer VTS is not available for BTLS and TPF.

3.3.1 VTS and Library Manager microcode levels

We summarize the microcode levels for VTS and Library Manager and their major improvements in Table 3-9.

Table 3-9 VTS and LM code levels

VTS Code-Level or Title	VTS code	LM-code	Description
Denali	2.26.24.xx	526.xx	Advanced Policy Management
R6	2.27.600.xx	527.19	3590 model H support, additional business continuance definitions
R7	2.28.700.xx	528.xx	FICON for P2P
R7.1	2.30.720.xx	528.xx	Additional FICON upgrade options
R7.2	2.30.720.40	529.01	Selective P2P copy, BVIR
R7.3	2.31.730.xx	530.xx	3592 support, B20 with 4 VTCs
R7.4	2.32.740.xx	531.xx	TS3500/3953 support, APM enhancements, VTS upgrade options

3.4 Peer-to-Peer VTS performance considerations

The Peer-to-Peer VTS is also creating, copying, and managing a duplicate copy of every logical volume; therefore, it must perform a greater amount of work to process the same amount of host data as a stand-alone VTS. For this reason, and because of the basic architecture of the Peer-to-Peer VTS, it has unique performance considerations.

The Peer-to-Peer VTS shares many of the performance characteristics of a stand-alone VTS. The performance of the individual VTSs in a Peer-to-Peer VTS configuration is affected by the same factors that govern the performance of a stand-alone VTS:

- Peer-to-Peer Configuration:
  - The number of Virtual Tape Controllers
  - The distance between the hosts, VTC features and VTSs
- The size of the Tape Volume Cache
- Number of attached Tape Drives
- Copy or Synchronization Mode:
  - Deferred Copy
  - Immediate Copy
- Workload Balancing:
  - No Preference
  - Preferred
- Customer Workload Attributes:
  - Blocksize and compressibility
  - Processor and channel configuration
- Business Continuance Settings

The most significant Peer-to-Peer VTS performance consideration is the mode of operation, either immediate copy mode or deferred copy mode. The choice of mode has, in addition to significant availability ramifications, a very considerable impact on the throughput capability of the Peer-to-Peer VTS. Much of the following discussion is devoted to a detailed examination of the effect of deferred copy mode on Peer-to-Peer VTS performance.

The performance information discussed in this chapter is based on modeling and on laboratory performance measurements. These are explained in more detail within the Performance White Paper, which is available at:

http://www.ibm.com/support/techdocs

Specific customer environments may experience different results. Please also note that performance characteristics can be changed by microcode changes.

3.4.1 VTS throughput management processes

In the body of this chapter, reference is made to various TVC capacity thresholds and the action taken at those points in managing host write input. Here is a general description of the principal processes that are invoked:

1. 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 VTS from peak mode to sustained mode.

2. 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.

3. VTS Copy Management:

   This applies only to PtP subsystems and comes into effect when the amount of un-copied data in TVC reaches a predefined threshold. It in particular applies 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 copied to physical tape prior to being copied to the other VTS, which could result in a recall operation prior to the Peer-to-Peer Copy.

4. VTC Copy Time Management:

   This also applies only to PtP subsystems, and in particular to immediate copy mode. It comes into effect only if specifically invoked via 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, thus causing the job to terminate prior to the copy finishing.

3.4.2 Peer-to-Peer operational modes

A Peer-to-Peer VTS has two modes of copy operation for creating dual logical volume copies, either immediate copy mode or deferred copy mode. The default mode is immediate copy mode. We already gave you detailed information on the copy mode operations, please refer to 2.3, “Creating the secondary copy of a virtual volume” on page 26.

The mode of the operation that is specified affects the peak performance of the Peer-to-Peer VTS.

In immediate copy mode, while both VTSs are active, a VTC will initiate the creation of a dual copy when a rewind/unload is received and will signal completion when the dual copy operation is complete. This mode provides the highest level of data protection.

In deferred copy mode, a VTC will schedule creation of the dual copy operation when a rewind/unload is received and will signal completion immediately. The time needed to complete the dual copy operation will depend on the activity of the VTSs.

In deferred copy mode, you also need to specify a deferred copy priority threshold. This is a time parameter that determines when the deferred copy will be processed with a priority for processing bandwidth equal to host I/O activity. The time has a range of 0–24 hours in one hour increments. The default deferred copy priority threshold is eight hours.

In deferred copy mode, when a longer deferred copy priority threshold is specified, the Peer-to-Peer VTS can maintain a high level of performance for a longer time without the dual copy operation overhead. However, each VTS in the Peer-to-Peer VTS must have enough TVC capacity to keep in cache the logical volumes that do not yet have a dual copy or have not been copied to a stacked volume. If their TVCs do not have enough capacity, the VTS subsystems may reach a state where the logical volumes created that have not been copied to the other VTS must be copied to a stacked volume and removed from the TVC to make room for newly created or recalled logical volumes by the host. Then, when their dual copy operation is to be performed, the logical volumes would have to be recalled back into the TVC.

In a remote Peer-to-Peer VTS configuration, based on the distance of the connections, the data transfer rate between the VTSs may be slower than would be the case for a local configuration. It may take more time to complete the dual copy operation. In immediate copy mode, every write may take a longer time to complete the dual copy operation. In deferred copy mode, it is not necessary to wait for the completion of the copy operation.

This choice is a trade-off between the level of the data protection and the performance in a Peer-to-Peer VTS configuration.

Without Advanced Policy Management, the mode of operation is specified in each of the VTCs. All the VTCs in the Peer-to-Peer VTS configuration must use the same mode of operation. With Advanced Policy Management the copy mode has to be set as well, but now you have the ability to define a Management Class to control the copy mode for any logical volume created. Your choice of copy modes is:

- VTC defined
- Immediate
- Deferred
- No Copy

Details of how to configure your Management Class with Advanced Policy Management can be found in 4.3.3, “Creating Management Classes” on page 172.

Immediate copy mode

The copy of the new or modified logical volume to the other VTS (the Peer-to-Peer Copy) is made prior to the completion of the rewind/unload of the original logical volume. When the volume completes close processing, the Peer-to-Peer VTS has completed performing the copy of the logical volume. No further Peer-to-Peer Copy processing is required. The impacts on throughput are:

- The elongation of the job time
- The reduction in potential host bandwidth due to the resources consumed by the copy work being performed

Deferred copy mode

The Peer-to-Peer Copy of the new or modified logical volume is queued, and occurs some time after the volume completes close processing. When the Peer-to-Peer Copy is performed depends on several factors, including:

- The current VTS workload
- The length of time since the original volume was created or modified
- The amount of uncopied data in the Tape Volume Cache of the VTS that contains the original volume

The impact on throughput is reflected in a reduction of host bandwidth when the copy process is ultimately performed.

It is important to note that, regardless of the operational mode, the Peer-to-Peer VTS will perform as much of the background work (Peer-to-Peer copies and copies to physical stacked volumes) as possible with the resources available at the time.

This should be kept in mind when considering some of the information in the following sections, especially the information concerning deferred copy mode. In deferred copy mode, although the Peer-to-Peer Copy is initially placed on a lower priority queue than host work, the Peer-to-Peer VTS will perform the copy as soon as resources are available.

Peer-to-Peer host write throughput in immediate copy mode

The behavior of Peer-to-Peer VTS write throughput in immediate copy mode is similar to that of a stand-alone VTS. In practical terms, the level of host write performance is dependent upon a single threshold, the amount of data in the Tape Volume Cache that has not been copied to stacked volumes.

Assuming a continuous maximum host write workload, and beginning with a cache whose volumes have all been previously copied to physical tape, there are two distinct performance data rates: a peak rate and a sustained rate.

During the initial period, the *peak rate* is achieved since only few resources are being used for copying logical volumes to stacked volumes. Tape volume cache space is created for new volumes by fragmenting (removing from cache) those volumes that were already copied to stacked volumes. The duration of this peak rate period will depend on the host write activity, the data compression achieved, and the size of the cache.

The *sustained rate* occurs when the uncopied data threshold is exceeded. This will cause older logical volumes in the cache to be copied to physical tape. The sustained rate can continue indefinitely, with the host write activity that is allowed, balancing the copying of logical volumes to stacked volumes.

Note that the Peer-to-Peer Copy process is being performed through both the peak and sustained rate periods.

Peer-to-Peer host write throughput in deferred copy mode

The behavior of Peer-to-Peer VTS write throughput in deferred copy mode is more complex than for immediate copy mode because the Peer-to-Peer VTS must also manage the queue of outstanding deferred copies.

Instead of a single threshold, the Peer-to-Peer host write performance level is dependent on three separate thresholds or triggers:

- The deferred copy priority threshold
- The amount of data in the Tape Volume Cache not yet copied to the other VTS
- The amount of data in the Tape Volume Cache not yet copied (pre-migrated) to stacked volumes

Deferred copy priority threshold

The deferred copy mode peak rate will begin at a higher level than the immediate copy mode peak rate since the Peer-to-Peer volume copy activity is being deferred. This peak rate period will continue until the amount of data that has not yet been copied to the other VTS exceeds the Peer-to-Peer Copy threshold.

When the Peer-to-Peer Copy threshold is exceeded, the Peer-to-Peer VTS will begin to copy new and modified logical volumes from each VTS to the other VTS in a manner similar to immediate copy mode. At the same time, the priority of the Peer-to-Peer Copy process will be raised. At this point, the host write rate will be reduced since additional bandwidth is being devoted to the Peer-to-Peer Copy process.

When the amount of data in a Tape Volume Cache in a Peer-to-Peer VTS configuration that has not been copied to stacked volumes exceeds the uncopied data threshold, then that VTS, just like a stand-alone VTS, will cause logical volumes in the cache to be copied to physical tape. This will also cause a reduction in the host write data rate that can be achieved. This is referred to as the *deferred copy mode sustained rate*, since it can continue indefinitely, with the host write activity that is allowed balancing the copying of logical volumes to stacked volumes. Thus, the Peer-to-Peer VTS will alternate between the peak and reduced performance level until the host write throughput demand drops.

As can be seen, the level of throughput of the Peer-to-Peer VTS in deferred copy mode is dependent on a complex set of conditions. If you choose to run the Peer-to-Peer VTS in deferred copy mode, you need to consider several factors.

3.4.3 Critical factors affecting the PtP performance

This section deals with the basic operational and design decisions that, for the most part, will determine the overall performance “envelope” of the PtP configuration.

Immediate versus deferred copy mode

The copy mode decision *must* be driven by your *availability* requirements, *not* performance. If you have the requirement that the PtP must guarantee that the logical volume copy exists in the second VTS before end-of-job, then immediate mode is the only alternative. The fact that the immediate mode decision may require a larger PtP configuration to meet performance requirements, or that Immediate mode may result in longer job run times versus deferred mode must not be allowed to override the availability requirement.

If you are willing to accept the fact that in deferred mode, the possibility exists that a failure of a VTS before the logical volume is copied will result in a loss of access to that logical volume until repairs are performed, then deferred mode will likely provide for a smaller configuration, or greater overall performance.

Deferred copy priority threshold hours setting

If the availability decision is that deferred will be the proposed copy mode, then the next decision is what deferred copy priority threshold hours setting should be recommended. There has been some confusion about the `deferred hours = 0` setting, and exactly what availability and performance it provides versus the immediate mode, or other possible deferred hours settings.

The short answer is that it does *not* provide equivalent availability to immediate mode, and will result in less overall throughput than a higher deferred hours setting.

The deferred copy priority threshold tells the PtP, for each logical volume, at what point the copy of that logical volume should be raised to the same priority as host work. In the case of `deferred hours = 0`, the copy of a logical volume will, from the time it is written, be sharing the bandwidth of the VTCs along with any host work. That means that host work bandwidth will be reduced by copy work. This is probably not the most effective use of the PtP bandwidth with well-defined peaks and valleys in host workload. To maximize host work throughput, the deferred copy priority threshold should be set to *slightly longer* than the duration of the peak host workload, assuming this peak period is short enough to allow the backlog of copies to be completed prior to the next peak period. If you are in doubt as to whether the copy queue backlog can be completed, you should contact IBM support for assistance.

Deferred Hours = 0 and multi-volume considerations

It is important to realize that in deferred mode, the copy of a logical volume can start as soon as the logical volume is closed — not at job end. So in deferred mode, when deferred hours = 0, the copy of the first logical volume will likely occur during the writing of the second logical volume. The copy of the second logical volume will compete for bandwidth with the writing of the third, and so on. The effect of bandwidth contention in this case may result in a considerable increase in elapsed job time.

The immediate copy and deferred copy modes are the only user selectable copy modes available on the PtP VTS. In addition to this mode selection, the observed throughput performance can depend on the initial state of the TVC, the write content and compressibility of the workload, and how long the operation has been sustained.

For each of the copy modes, we define a peak throughput, observed at the beginning with a TVC all of whose new or updated volumes have been copied to the other VTS and to physical tape. We define as a sustained throughput one that is observed after sufficient operation with a high workload, after which it can be verified that the content of the TVC is in dynamic equilibrium, with the rate of copying to tape equal to the rate at which data are being written to the VTS. As shown in performance tests, the sustained throughput is approximately the same for immediate copy and deferred copy operation. There can also be periods of other, intermediate, throughput in the transition from peak to sustained throughput.

Regardless of the copy mode, the internal algorithms are designed to do as much of the background work (Peer-to-Peer copies and copies to physical tape) as possible with any excess bandwidth that is available. Thus, unless there is a strict requirement to keep the VTSs synchronized, the best performance can typically be observed with the PtP VTS in deferred copy mode, within the constraints detailed in the section on Peak Write / Recovery Time and TVC Capacity Planning. When the maximum write input is occurring, much of the asynchronous background copy work can be suspended in order to handle read/write traffic with the host. The other extreme is the immediate copy mode, which is designed to provide a copy on each of the peer VTSs before rewind/unload complete is presented to the host.

Peak Write / Recovery Time and TVC Capacity Planning

The peak period duration for both deferred and immediate copy modes is determined by the point at which premigration throttling sets in. During a peak period host data are written into the TVC, but not immediately backed up, or premigrated, to physical tape as well. When the amount of non-premigrated data reaches a certain threshold, the VTS is designed to begin to limit, or throttle, host input data to an extent necessary to reach a balance between what comes into the cache from the host processor and gets written out of the cache to physical tape. As stated previously, this reduced balanced data rate is the sustained data rate. The larger the TVC, the higher the premigration threshold, and thus the longer the peak period. Because deferred copy mode has a higher potential peak rate than immediate copy mode, it also has generally shorter peak periods than immediate copy mode, since it can reach the premigration threshold sooner. In addition to the premigration threshold, the peak rate and duration for deferred copy mode is affected by two other factors.

- **Deferred copy space threshold:** This is a threshold for data in the cache that have not yet been copied to the peer VTS. Since for immediate copy mode, copies are completed prior to job end, this threshold can usually be reached only for deferred copy mode. Laboratory tests have shown that the premigration threshold is typically reached at about the same time as or earlier than the copy threshold. In these cases, reaching the copy threshold has had the effect of just lowering the already existing post-peak data rate.
- **Deferred copy time threshold:** This relates to the maximum age in the TVC of data un-copied to the peer VTS (in integral 0 to 24 hours). When this hours age has been reached by a tape volume, its priority for being copied is increased. Although reaching this threshold may not actually terminate the peak period, since host write activity is not being throttled, tests have shown the peak rate to diminish and come closer to the peak rate for immediate copy mode.

If a management goal for the PtP VTS is to have maximum throughput available on demand in deferred copy mode, the hours parameter should be kept large enough so that most Peer-to-Peer copies occur naturally before reaching the deferred copy priority threshold. Thus the hours parameter should be greater than the expected daily peak period duration. Approximate peak periods can be calculated by applying peak write rates to the appropriate premigration thresholds. These thresholds are approximately 30 GB less than the total cache capacity, except for the largest, four-drawer cache, where the premigration threshold is set at 1,000 GB (1 TB) and except for those installations which exploit larger logical volumes that implies lower thresholds.

Preferencing

Each VTC can be specified to send I/O activity to a specific VTS if it is available. The selections that you are allowed are VTS0, VTS1, and no-preference. VTS0 and VTS1 indicate one of the two VTSs. There is no default selection; you must set the desired mode at installation time.

If *no-preference* is selected, a VTC will try to balance the workload between the two VTSs. Therefore, no-preference is recommended for a local Peer-to-Peer configuration. Figure 3-10 shows the workload balanced by all four VTCs.

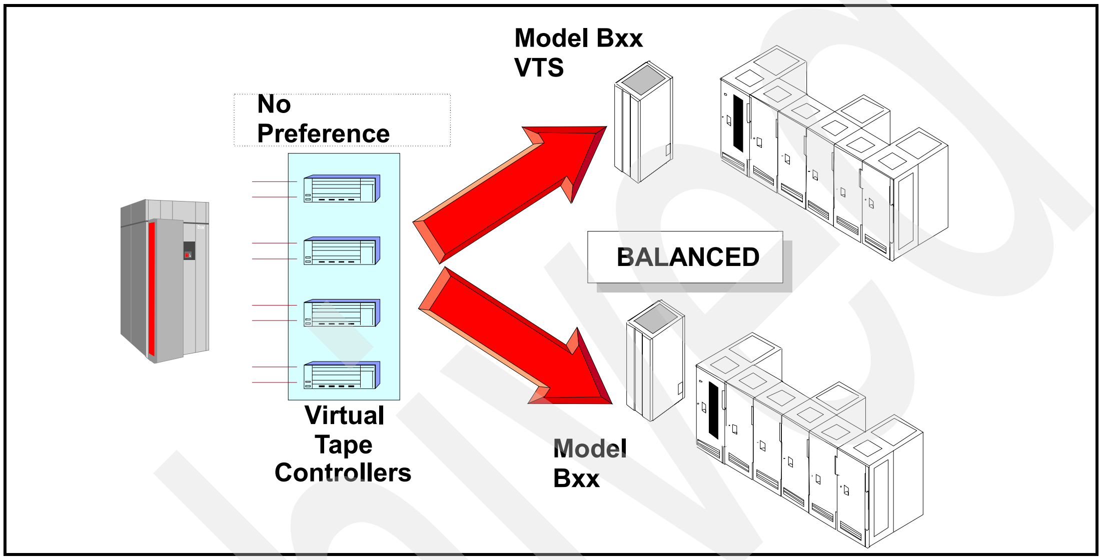

Figure 3-10 Balanced workload

In a remote Peer-to-Peer configuration, you can choose to have the VTCs local to each VTS prefer VTS as the I/O VTS. However, current measurements show that this provides worse performance than the no-preference case because of the imbalance between VTSs.

Four versus eight VTCs

With VTS R7.4 you can reduce the number of VTCs for a B20 VTS from four to eight. You should examine the modeling results of BatchMagic with care before you order this upgrade:

- The ESCON peak data rate capability will drop to approximately 50% of that of a similarly configured PtP with eight VTCs.

  One trade-off for reduced cost offered by the four VTC B20 Peer-to-Peer is in peak performance. For the ESCON Peer-to-Peer the host ESCON channel adapters are the bottleneck for peak performance. Each ESCON VTC has two host port adapters. So, comparing eight VTCs to four VTCs reduces the total ESCON ports installed from 16 to eight, or 50 percent. Any significant skew in the workload across the VTCs has the potential to degrade performance further.

- The FICON peak data rate capability will be between 50% and 90% of that of a similarly configured PtP with eight VTCs, depending on features and numbers of FICON channels.

  For FICON, it is not so simple. An installation where an LPAR or processor has pathing to only one FICON channel per VTC is where one might expect 50 percent performance compared to eight FICON VTCs. If the same processor is able to access two channels per VTC, one might expect 90 percent performance. With FICON VTCs each VTC can handle about 25 percent of the total throughput of the Peer-to-Peer. Therefore, with eight VTCs there is a great deal of built-in tolerance of workload skew across the VTCs. Since the eight VTCs can handle 200 percent, with just four VTCs the VTCs can handle about 100 percent of the Peer-to-Peer throughput. So, with any significant workload skew, the peak performance will degrade accordingly.

VTC data compression

Since the initial availability of PtP VTS B10 and B20, a version of Licensed Internal Code (LIC) that enables data compression at the host interface of the VTC. This LIC was introduced with LM 526 and VTS 2.23.29.0 or higher and is intended to improve the efficiency with which the data is then transferred to the peer VTSs. These features help lower the overhead of the VTC and VTS processors in handling data and thereby increase the bandwidth of the ESCON/FICON channels connecting the PtP VTSs when dealing with compressible data.

Link performance characteristics

Links between each VTC and the each VTS are assigned a rating. There are different ratings for direct ESCON / FICON attachment and ratings for supported Channel Extenders and DWDM products.

- The Default is no weighting of communication link performance.
- For each link between the VTC and the two VTSs, a relative link value can be specified. The following values are defined based on the technology of the link:

Value	Link technology
1	ESCON < 5 km, FICON < 30km
2	ESCON < 10 km, FICON < 80 km
3	McData, CISCO < 1000 km, DWDM < 10 km, FICON < 100 km
4	McData, CISCO > 1000 km, DWDM < 15 km
5	DWDM extenders > 15 km

If a preferred I/O VTS is specified and each link has a different relative link value, cache residency is ignored and the preferred VTS is selected if the non-preferred VTS link has a value greater then two more when compared to the link value for the preferred I/O VTS.

You can use the link performance characteristics to force specific mounts to a certain VTS, for example, in disaster testing together with using Write Protect mode. See 9.4.2, “Test steps” on page 348 and 9.5.4, “GDPS DR testing consideration” on page 361 for further discussions.

3.5 Peer-to-Peer VTS performance measurements

The performance measurement charts presented on the following pages show what performance you can expect, given certain hardware configurations like ESCON versus FICON, the copy modes, and the preference used.

3.5.1 ESCON versus FICON performance comparisons

Performance tests demonstrate that FICON connectivity can provide a substantial boost in PtP performance over ESCON. Figure 3-11 shows basic write and read maximum data rates for (1) PtP VTS ESCON, (2) stand-alone VTS FICON (for reference purposes), and (3) PtP VTS FICON. Data compression for these comparisons is 2.66:1. For the PtP VTSs, non-preferred performance is shown for both deferred and immediate copy modes. The B20 VTSs in each case are base B20s with standard processors. Also available are B20s with the FICON Performance Accelerator Feature (FPAF) (Feature Code 5250), which contain enhanced processors. Comparisons between base and FPAF B20s appear later in the paper.

A further characteristic of FICON PtP as seen in laboratory tests is that an individual FICON VTC can handle approximately twice the relative write workload of an ESCON VTC (approximately 25% versus 12.5% of the total VTS throughput potential), meaning that a FICON PtP may be better able to handle periodic host workload imbalances, or skew, than can an ESCON PtP.

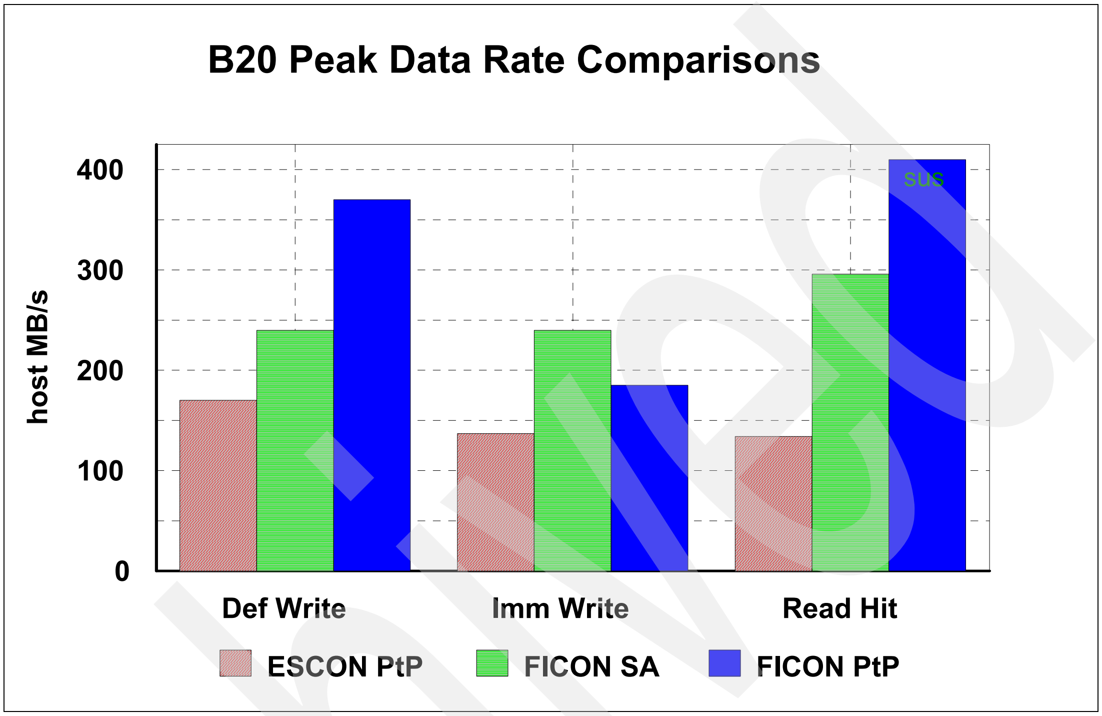

Figure 3-11 Maximum data rate at 100 percent write - No preference

3.5.2 FICON performance

In this section we present data for the case when all PtP VTS components, VTCs and VTSs, are local and the VTCs are operating in non-preferred mode.

Balanced mode

Figure 3-12 shows the modelled performance of the PtP VTS B20 and B10 under a 100% write workload for VTS models with the maximum FICON channel configurations.

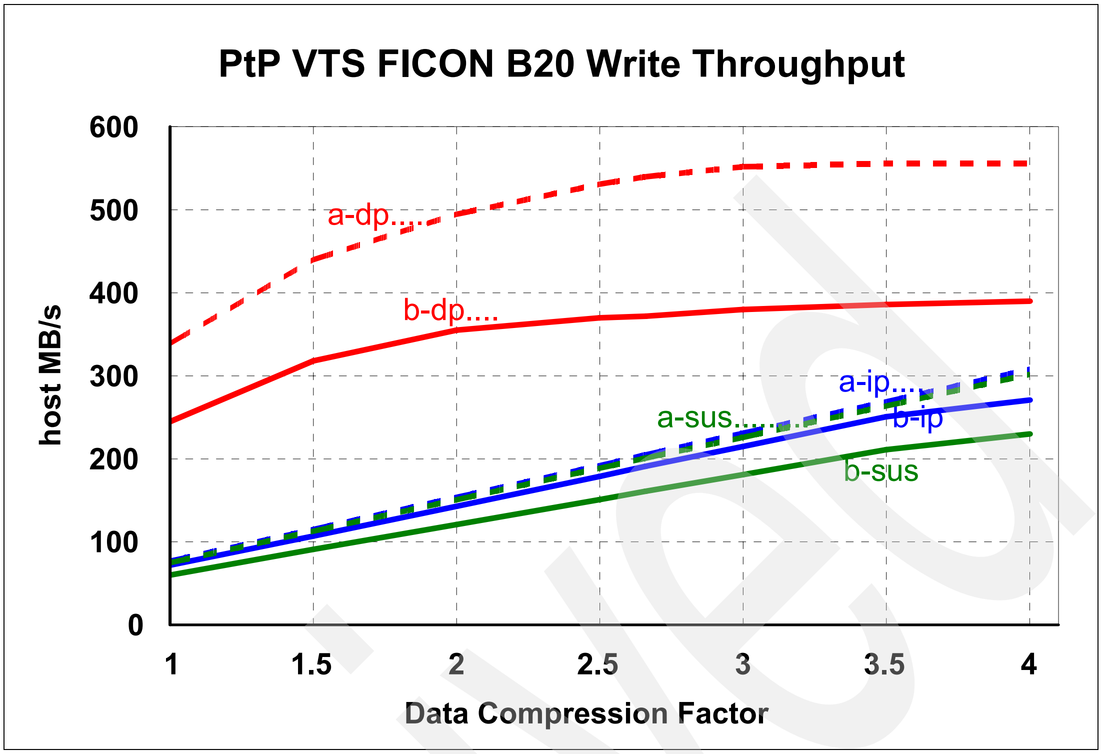

Figure 3-12 Maximum data rate at 100 percent write - non-preference

The maximum throughput, marked *dp* for *deferred-copy/peak* is the performance with copies from one VTS to the other being deferred in favor of maximum host write throughput. In this mode, Peer-to-Peer copies make up the balance of the PtP VTS workload if the host input is not at peak bandwidth. The next level of throughput performance is obtained in the *immediate copy* mode, with the copying of data to tape not necessarily keeping up with the rate at which new data are being written to the VTS. This is the *immediate-copy/peak* mode designated as *ip* in the figures. If the *dp* mode throughput is higher than that of the *sus* mode, data can build up in the TVC that need to be copied to the peer VTS, as well as needing to be copied to tape in the first VTS and then copied to tape in the peer VTS. In the *ip* mode the Peer-to-Peer copies are essentially immediate, but there can still be buildup of data in both the TVCs that need to be copied to tape.

When the amount of uncopied data reaches a fixed threshold in the TVC, the VTS begins to enforce a policy of not accepting new data until a corresponding amount of space has been released after copies have been made. Some detail of these actions is described in 3.4.3, "Critical factors affecting the PtP performance" on page 93 in the subsection, *Peak Write / Recovery Time and TVC Capacity Planning.* This state of the VTS is called the *sustained* (or *sus*) throughput state, in that it could be maintained for an extended, theoretically infinite, period.

Important: The compression factor influences the total throughput of a PtP VTS significantly (see the following charts). This has to be considered doing the tape analysis in preparing the PtP VTS implementation.

Preferred mode

In the default mode, a VTC will refer host I/Os to either VTS in a manner that will keep the load on the two VTSs approximately equal. A VTC can be set, however, to refer host I/Os to one of the two VTSs exclusively. The latter is termed the *preferred* VTS mode of operation.

The principal motivation for implementation of preferred VTS operation is reduced and consistent response time performance in remote PtP VTS configurations. The choice of *no preference* or *preferred* VTS mode is made at the VTC. It is a static choice requiring a power-down of the VTC to alter.

A version of *preferred* mode, *primary*, is required for Geographically Dispersed Parallel Sysplex™ (GDPS) implementation; it forces all host I/Os to go to and from the primary VTS without exception under non-failure conditions. GDPS also requires operation in *immediate copy* mode.

Preferred mode operation reduces the total throughput of the PtP VTS because it forces all the data to pass through the preferred VTS while the secondary VTS acts mostly as a receiver of copies. On host read hits, the data can still come from either VTS; in *primary* mode, however, data from the secondary VTS is first copied to the primary VTS. On reads that require a recall from tape, the recall is handled at the preferred VTS.

LIC 2.23.29.0 can significantly improve the throughput in preferred and primary mode operation. In these modes there is significantly more data traffic on the channels between the VTCs and the preferred/primary VTS than to the secondary VTS. Since with this LIC this traffic is all in compressed data, performance is improved. The modeled write throughput in preferred mode operation is shown in Figure 3-13 for the PtP VTS B10 and B20, with FICON channels and FC5250 (FPAF).

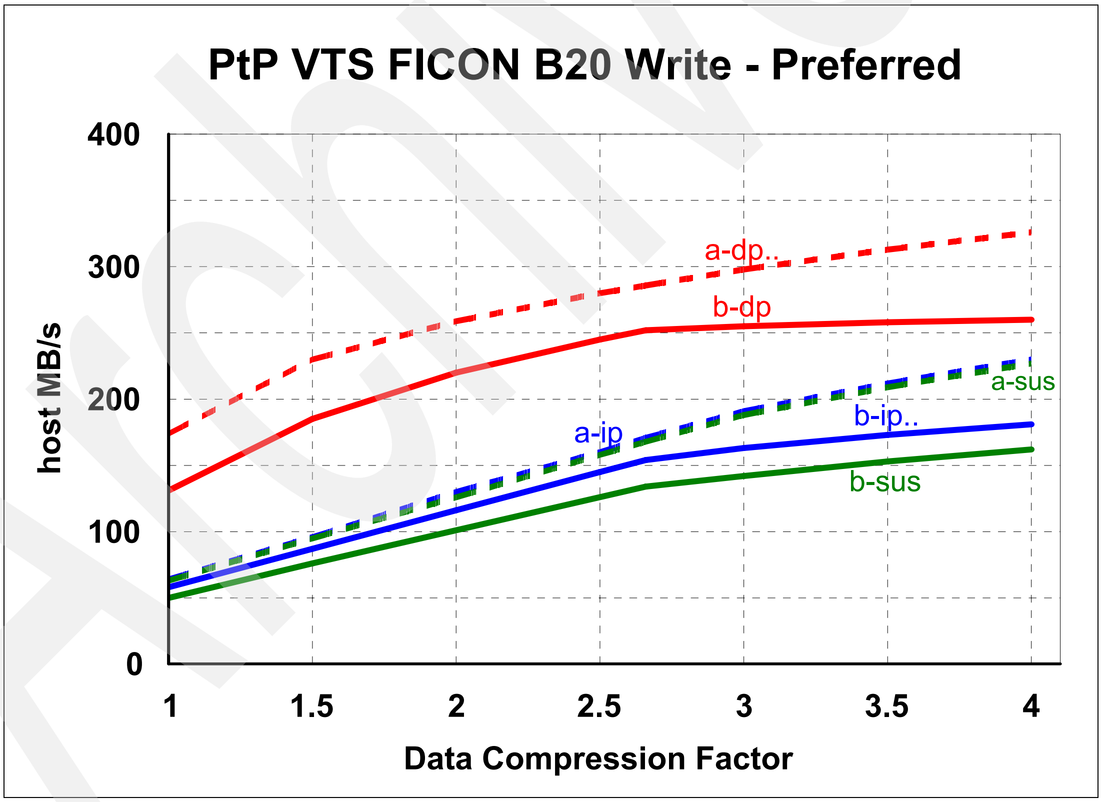

Figure 3-13 Maximum data rate at 100 percent write - preferred

Figure 3-13 shows model *projected* data. The labels are:

dp For deferred copy mode peak
ip For immediate copy mode peak
us For sustained operation

The prefix b- is for the base B20, and prefix a- is for the B20 with FICON Performance Accelerator feature. Note that the a-ip and a-sus plots are nearly *coincidental*. Sustained data rates for these plots assume the maximum number of premigration drives, which is eleven for the B20s.

Preferred versus no-preference selection: Split host sites

In Peer-to-Peer VTS configurations where host(s) and one VTS are located at one site, and other host(s) and the other VTS are located at a remote site, the decision of whether the VTCs will be configured for preference of a particular VTS or for no-preference is more complex.

Modeling has shown that, even in a split site configuration, no-preference will theoretically provide the highest host write data rate. However, there will be a response time advantage when a VTS is configured to be preferred. By having all tape I/O served locally, the deferred copy mode writes and reads will be more likely to have a shorter open time. In this environment, the requirement for throughput should be weighted against the need for faster mount time, and the preference set accordingly.

The effect of distance

The effect of distance on data rate is generally much less for FICON than for ESCON. With FICON connections, distances of up to 100 km are supported via Dense Wave Division Multiplexer (DWDM) channel extenders. Figure 3-14 reflects measurements taken with DWDMs at 100 km. It shows individual VTC data rates versus number of concurrent jobs. As can be seen from the charts, distance performance improved as the number of concurrent jobs increased. Also shown is the VTS-to-VTS copy performance at various distances. Local-to-remote copy throughput is an important metric when running in preferred mode.

For distances of less than 100 km, interpolate between the 100 km and 0 km data rates. Data rates shown are with no other PtP activity. For these measurements, DWDMs were connected to host and VTC, or to VTC and VTS, by FICON switches. The VTC was connected to the host via two 2-Gbit (FICON Express) channels.

This chart shows rates at which data was copied by a single VTC from one VTS to the other VTS. The copy operation was either local-to-remote (L—>R), remote-to-local (R—>L), or local-to-local (L—>L). Since each VTC had a maximum of three copy tasks, the copy rate did not increase after three concurrent jobs. These copy rates were with no other VTC or VTS activity. Local-to-remote copy rates measured the same as local-to-local copy rates. Even without using channel extenders, FICON distances of up to 10 km may be achieved using long wave, single-mode fibre with direct attach, or up to 30 km with appropriate cascaded (dual) switches.

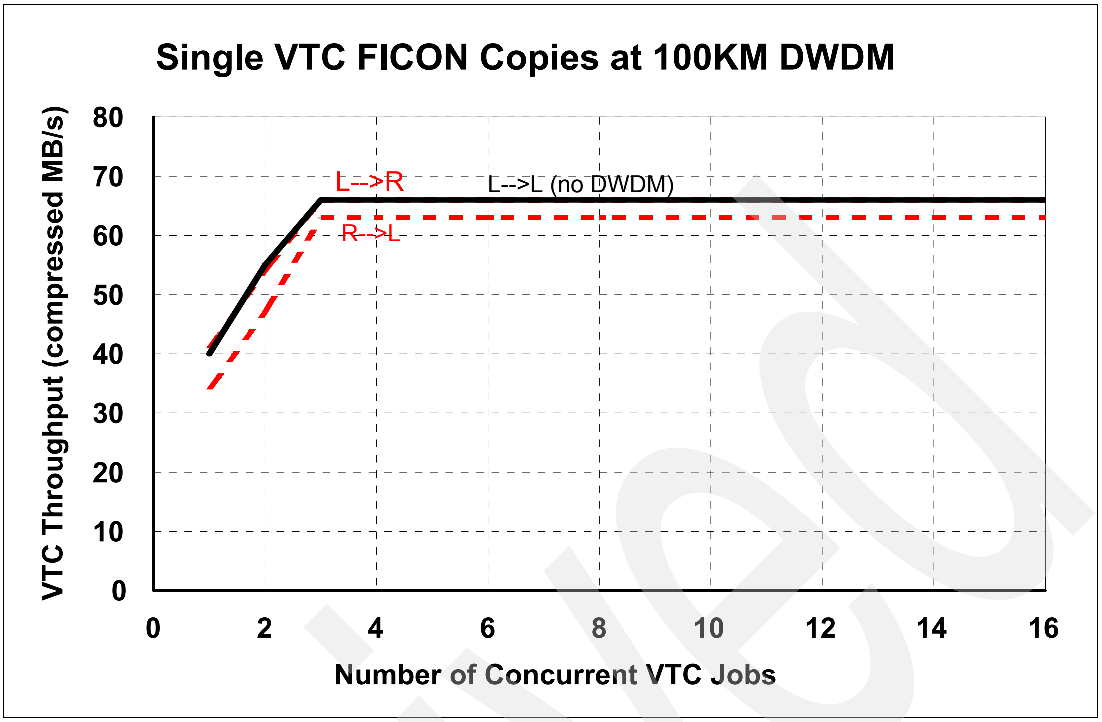

Figure 3-14 Extended distance effect on VTS-VTS copy operations

3.5.3 Performance considerations for 3592 drives

The 3592 drives can provide potential performance benefits in five metrics of VTS performance:

- Sustained data rate
- Premigration (cache recovery) data rate
- Recall data rate
- Reclaim rate
- Import/Export rate (which does not apply for PtP VTS)

The existence and degree of any performance improvement in these five metrics is dependent on several factors, including VTS model and workload characteristics. One metric showing no observable performance difference between 3590 and 3592 drives is peak data rate. During the peak data rate period, the main VTS activity is writing into the Tape Volume Cache (TVC), with premigration to physical tape occurring on only a limited basis as determined by the incoming host data rate. The “sustained data rate” is that rate that occurs when the incoming host data rate is in equilibrium with the outgoing data rate onto physical tape. In laboratory measurements, use of 3592 J1A drives resulted in increases in maximum sustained data rate in several cases compared to 3590 E1A (3590E) drives, depending on VTS model and number of back-end drives.

There are two maximum back-end drive configurations for the VTS B10 and three maximum back-end drive configurations for the VTS B20:

- VTS B10:
  – Six 3590 drives (homogeneous)
  – Twelve 3592 drives (homogeneous)
- VTS B20:
  – Twelve 3590 drives (homogeneous)
  – Twelve 3592 drives (homogeneous)
  – Six 3590 and twelve 3592 drives (heterogeneous)

The following examples show only homogeneous drive configurations (that is, they do not consider configurations with combinations of 3590 and 3592 drives).

Sustained throughput B20 (without FPAF — feature code 5250)

Figure 3-15 shows graphically the sustained data rate comparisons for 3590 E1A and 3592 J1A drives on a VTS B20 without the FICON Performance Accelerator Feature (FPAF), available via feature code 5250. The 3592 drives provide a higher data rate than 3590 drives, although the advantage is small beyond three active drives. Note also that for both 3590 and 3592 drives, the sustained data rate does not increase beyond five drives devoted to premigration activity concurrent with host data input.

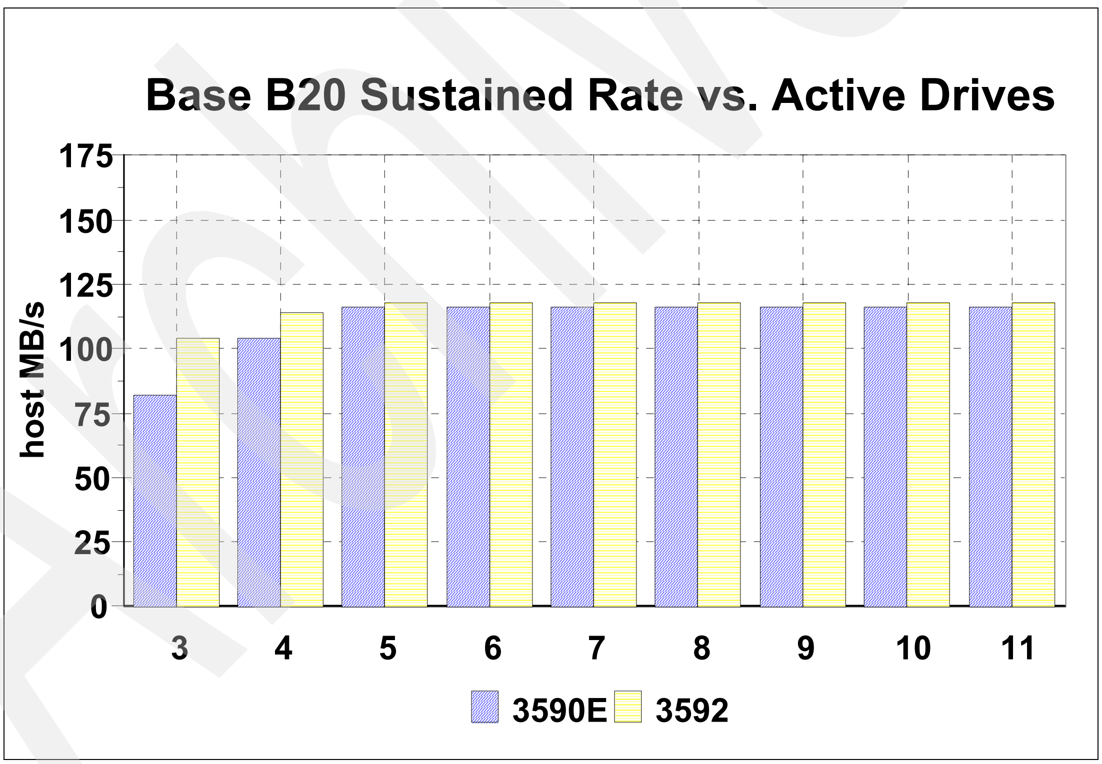

Figure 3-15 Base B20 Sustained Data Rate

This chart compares maximum sustained data rate for different numbers of 3590 and 3592 drives involved in premigration activity concurrent with equal host data input. The number of drives involved in premigration activity is limited to *one less than the total number of online drives.*

Sustained throughput B20 (with FPAF — feature code 5250)

Figure 3-16 shows graphically the sustained data rate comparisons for 3590 E1A and 3592 J1A drives on a VTS B20 with the FICON Performance Accelerator Feature (FPAF), available via feature code 5250. The 3592 drives show an improvement in data rate over 3590 drives for up to eight premigration drives. Note that no increases in data rate are seen for either 3590 or 3592 drives beyond nine drives performing premigration during host input.

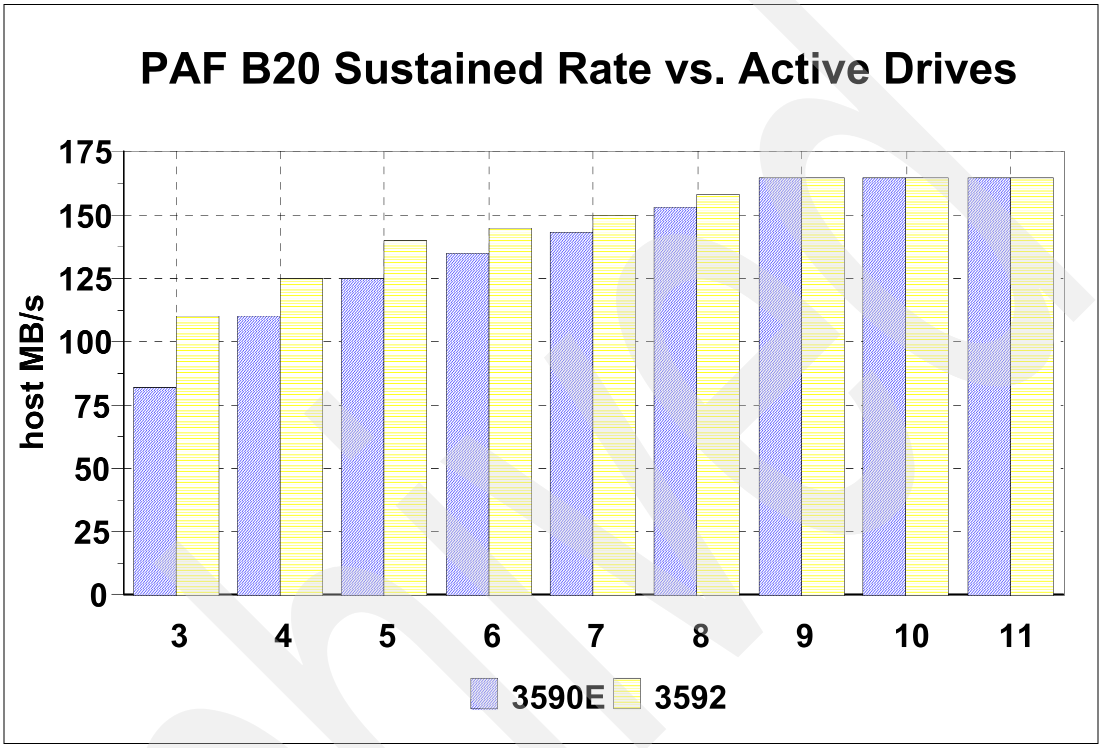

Figure 3-16 PAF B20 Sustained Data Rate

This chart (Figure 3-16) compares maximum sustained data rates for different numbers of 3590 and 3592 drives involved in premigration activity concurrent with equal host data input. The number of drives involved in premigration activity is limited to *one less than the total number of online drives*.

Premigration data rate for B20 with PAF

Premigration data rate is similar to sustained data rate, but with no concurrent host input activity. The only VTS activity occurring during these measurements was copying of volumes from the TVC onto back-end physical tape drives. This data rate can serve as an indicator of the minimum recovery period between daily batch windows. Actual recovery time would be determined by a data rate somewhere between that shown in the following premigration-only charts and that in the previous sustained rate charts. These measurements did show a slight data rate advantage for JA cartridges (300 GB) over JJ cartridges (60 GB). However, the difference was typically only around 3%. These premigration charts show the average data rate of the two cartridge types.

As the performance benefit for a base B20 is not as high as for a B20 with PAF, we only show these results in Figure 3-17. It shows graphically the premigration data rate comparisons for 3590 E1A and 3592 J1A drives on a PAF VTS B20 (with feature code 5250). Here moderate to substantial data rate improvement is seen for 3592 versus 3590 drives up to the maximum of eleven drives performing premigration. Note that for both 3590 and 3592 drives, there is no data rate increase for more than nine drives performing premigration activity.

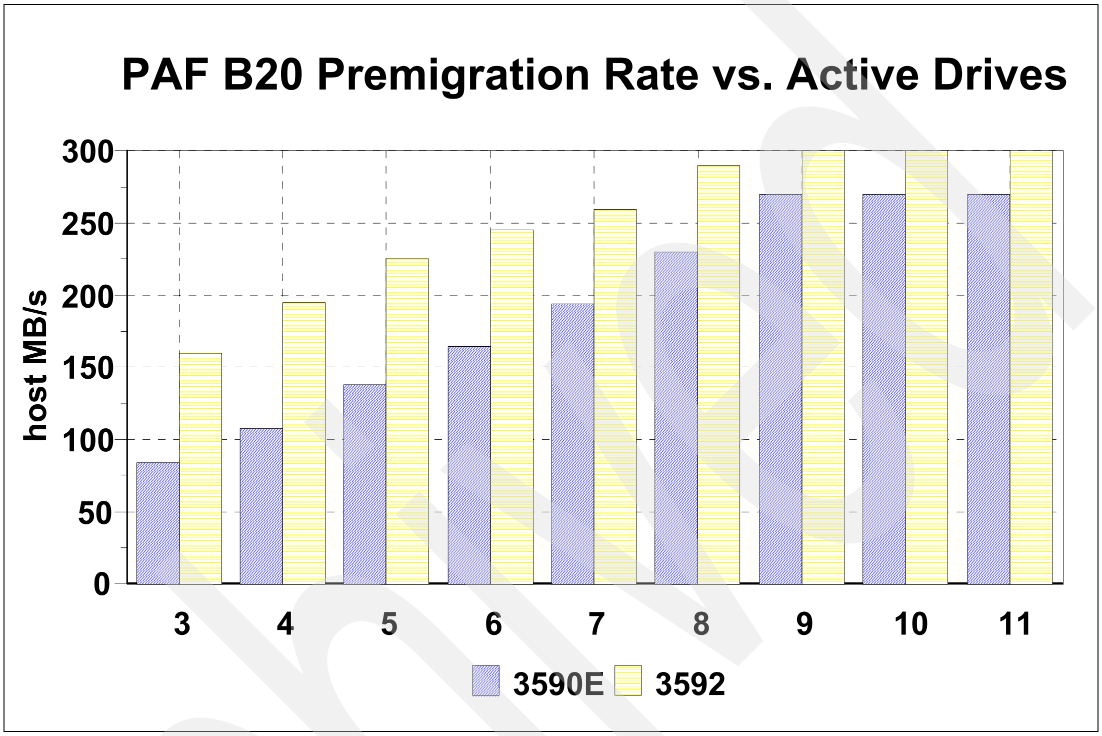

Figure 3-17 PAF B20 Premigration Data Rate

This chart compares maximum premigration data rate for different numbers of 3590 and 3592 drives involved in premigration with no other VTS activity. The number of drives involved in premigration activity is limited to *one less than the total number of online drives*. To facilitate comparison with other charts in this paper, data rates are shown as “host view”, which is prior to data compression upon entering the VTS. To arrive at actual Tape Volume Cache to physical tape data rates, divide the rates shown by 2.66, which is the data compression factor for the workload used for these measurements.

Recall data rate for B20

Recall data rate is the most difficult VTS data rate to characterize, since recall performance is highly dependent on the distribution of logical volumes over physical volumes. This procedure results in a high number of logical volumes recalled per physical tape mounted, and so recall performance is highly influenced by locate times.

Base B20

Figure 3-18 shows graphically the recall data rate comparisons for 3590 E1A and 3592 J1A drives on a base VTS B20. With recalls, 3592 data rate measurements in many cases show a higher data rate for the 60 GB JJ cartridges than for the 300 GB JA cartridges, as a result of shorter average locate and rewind times. Both types of 3592 cartridges give higher data rates than do 3590 cartridges, which were predominately the shorter J cartridges. Little or no recall data rate increases exist with greater than nine 3592 drives devoted to recall activity.

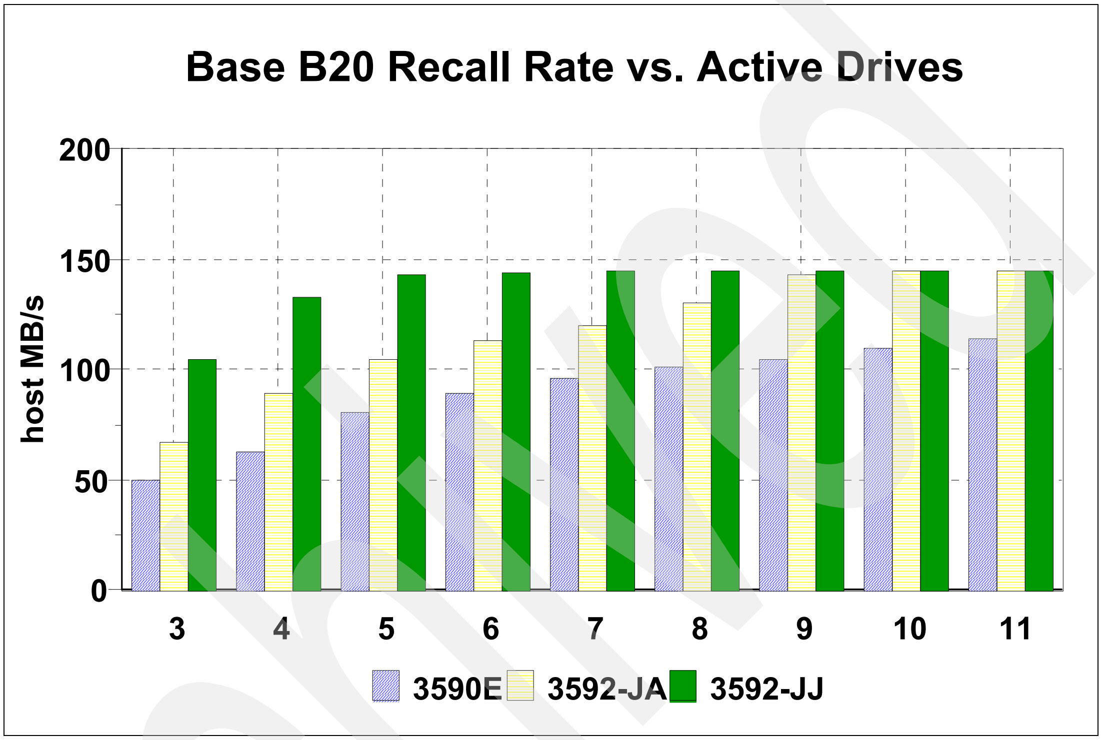

Figure 3-18 Base B20 Recall Data Rate

This chart compares recall data rate for different numbers of 3590 and 3592 drives involved in recall with no other VTS activity. Data rates are shown for two types of 3592 cartridges: JA (300 GB) and JJ (60 GB). The 3590 drives had a mixture of J (20/30 GB) and K (40/60 GB) cartridges, but predominately J cartridges. The number of drives involved in recall activity is limited to *one less than the total number of online drives*.

Attention: These recalls involved multiple logical volume reads per physical mount, and so while providing a basis for comparison within this paper, may not be representative of actual customer recall experience.

PAF B20

Figure 3-19 shows graphically the recall data rate comparisons for 3590 E1A and 3592 J1A drives on a PAF VTS B20 (feature code 5250). With recalls, 3592 data rate measurements in many cases show a higher data rate for the 60 GB JJ cartridges than for the 300 GB JA cartridges, as a result of shorter average locate and rewind times. Both types of 3592 cartridges give faster data rates than do 3590 cartridges, which were predominately the shorter J cartridges.

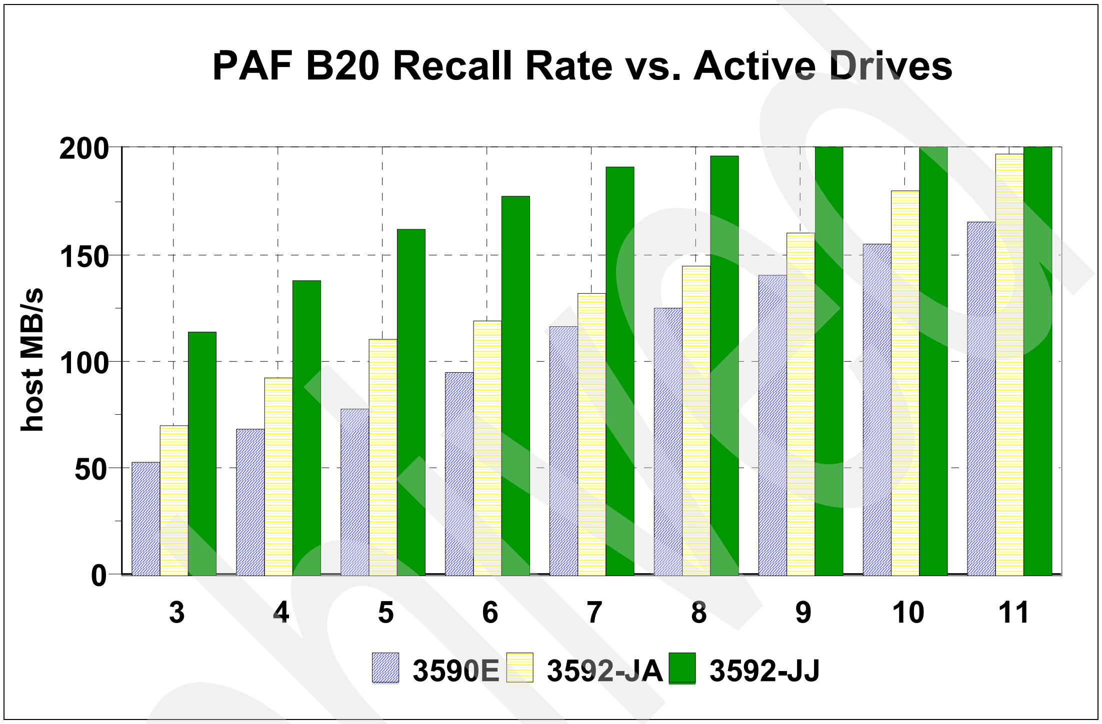

Figure 3-19 PAF B20 Recall Data Rate

This chart compares recall data rate for different numbers of 3590 and 3592 drives involved in recall with no other VTS activity. Data rates are shown for two types of 3592 cartridges: JA (300 GB) and JJ (60 GB). The 3590 drives had a mixture of J (20/30 GB) and K (40/60 GB) cartridges, but predominately J cartridges. The number of drives involved in recall activity is limited to *one less than the total number of online drives*.

Attention: These recalls involved multiple logical volume reads per physical mount, and so while providing a basis for comparison within this paper, may not be representative of actual customer recall experience.

VTS model comparisons for 3592 drives

This section repeats data presented in earlier charts, but on a VTS model basis and just for 3592 J1A drives. In general, measurements show that the B20 models provide higher levels of performance with 3592 drives than does the B10 model, and the PAF B20 model provides higher levels of performance than does the base B20 model.

3592 sustained data rate by VTS model

Figure 3-20 shows graphically the sustained data rate comparisons for 3592 J1A drives for VTS models B10, Base B20, and PAF B20. Note that the sustained data rate reaches a maximum at five drives for both a B10 and a base B20, although both B20 models provides for a consistently higher data rate than for the B10. Increases in sustained data rate occur for up to eight drives for the PAF B20, at which point the PAF B20 attains a substantially higher data rate than do either the B10 or base B20.

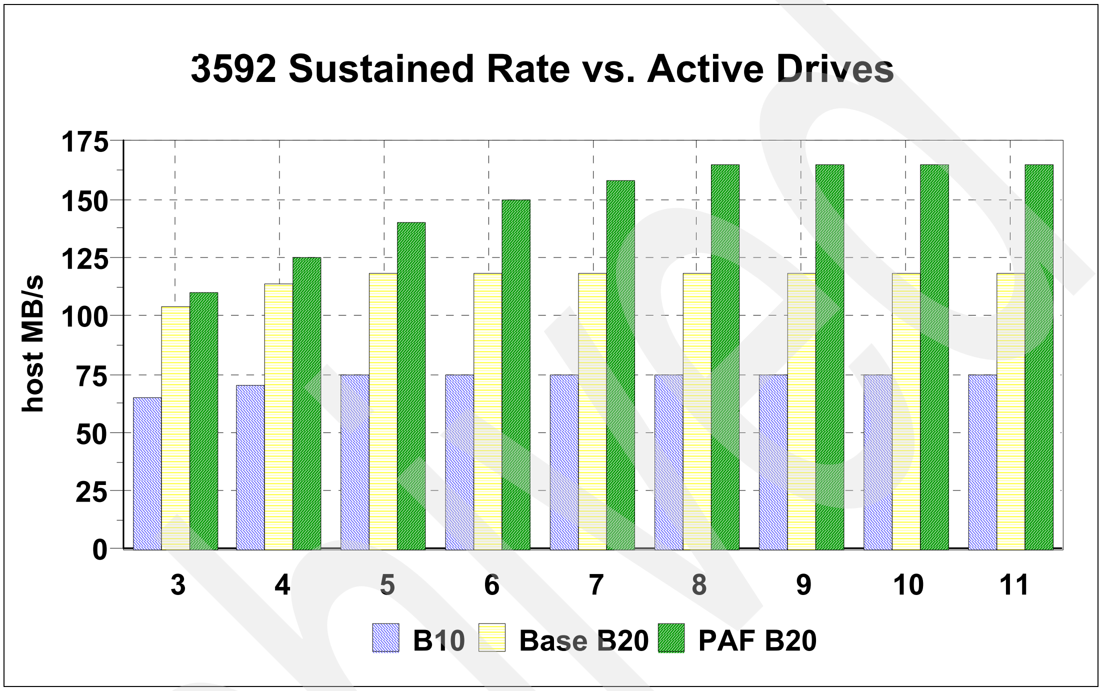

Figure 3-20 3592 Sustained Data Rate

This chart compares maximum sustained data rate by VTS model for different numbers of 3592 drives involved in premigration activity concurrent with equal host data input. The number of drives involved in premigration activity is limited *to one less than the total number of online drives*.

Summary of the performance considerations with 3592 drives

Within laboratory measurements, the 3592 drives provided generally increased VTS performance compared to 3590 drives. Here are some observations for the three VTS models.

B10 VTS

Little if any data rate increases were seen for more than five 3592 drives engaged in either sustained, pre-migration, or recall laboratory activity. Laboratory experience also shows that some maintenance type functions can take substantially longer when more than five drives are engaged in host data-intensive production activities. Therefore, at least for host data-intensive workloads, a total of six 3592 drives might be considered as an optimal number. Other, non-host-data-intensive workloads, such as single-volume recalls or import/export activity, might well justify up to the full complement of twelve 3592 drives on a B10.

Base VTS B20

In these measurements, up to eight total 3592 drives provided performance benefits for host data-intensive workloads. The maximum of twelve drives could also bring benefits for other workloads such as single-volume recalls, reclaim, and import/export.

VTS B20 with FPAF

The B20 with the FICON Performance Accelerator Feature showed the greatest potential performance benefit with 3592 drives. Up to the maximum of twelve total 3592 drives provided benefit for host data-intensive workloads. Further, in these laboratory measurements, for a given number of active drives, the FPAF B20 usually provided for higher data rates than did either the B10 or base B20. When considering the question of whether fewer 3592 drives can be configured to give performance equivalent to 3590 drives, remember to factor in all types of VTS activity that utilize physical drives, including reclaim and import/export. It is also wise to provide a safety margin to allow for occasional drive outages. The use of VTS pooling may also dictate having more drives than aggregate performance needs would indicate.

Finally, at least for these laboratory locate-intensive recalls, the shorter 3592 JJ cartridges in many cases showed higher recall data rates than did the longer JA cartridges. JJ cartridges might be considered for installations that could benefit from the JJ cartridge performance characteristics and do not need the data space concentration provided for by the JA cartridges. For further information and detailed explanations, look in the Performance White for 3592 which is available with the link:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100488

3.5.4 Other performance considerations

We have concentrated on the throughput of a Peer-to-Peer configuration. However, there are some further performance differences between Peer-to-Peer and stand-alone configurations that we discuss in this section.

Mixed Peer-to-Peer configurations

The standard PtP VTS configurations are symmetric; that is, they comprise two B10s or B20s. However, a number of mixed configurations are supported. The specifics are available from IBM tape storage Specialists and IBM business partners. The performance of such configurations can be evaluated on a case-by-case basis. For example, if a current installation has a B10 VTS and were to purchase a B20 VTS, could that be configured as a mixed PtP VTS B20/B10?

The answer is yes, but the problem is that the performance for much of the time would be equivalent to that of a PtP VTS B10, except for the deferred copy write period, and the recovery period would be extended in time. However, there are situations where such mixed configurations might make sense. For example, if the intention is to run the PtP VTS in preferred mode, then with the B20 local and the B10 remote the projected performance is estimated to be approximately equivalent to a symmetric PtP VTS B20 (assuming a compression factor on the order of three). That is, in the preferred VTS mode, the B10 can approximately keep up with the task of handling copies from the B20 in sustained operation.

Return to scratch time

The time taken to return volumes to scratch is greater for a Peer-to-Peer VTS than for a stand-alone VTS. Laboratory measurements showed that a stand-alone VTS returned volumes to scratch at about four times the rate of a Peer-to-Peer VTS.

Mount times

Measurements of a laboratory workload show that the Fast-Ready mount time for a virtual volume is about three seconds for a stand-alone VTS using 64 drives concurrently. In a Peer-to-Peer VTS, the additional processing needed elongates this mount time to about six seconds.

VTC configuration and channel throughput limitation

In a stand-alone VTS, all virtual drives are accessible through all channel adapters. The MVS I/O definition at the host, for example, is coded as if all the virtual control units were daisy-chained on the same ESCON channels. This means, in effect, that the throughput of the VTS is not dependent on which virtual control units or drives are in use — the total bandwidth of all the ESCON channels is available to all drives.

In the Peer-to-Peer VTS, however, the virtual drives are associated with individual VTCs, each having 16 or 32 virtual drives, and two physical channel attachments (see “Sharing and partitioning with multiple hosts” on page 71.) The effect of this architecture is that, for example, in an configuration with 64 virtual drives only one quarter of the total Peer-to-Peer VTS bandwidth is available to each VTC and its associated virtual drives. In addition, in order for a host to achieve maximum throughput to a VTC, both channel attachments must be utilized.

In cases where multiple hosts will be attached to the Peer-to-Peer VTS, careful consideration must be made to ensure that each host has channel attachments to a sufficient number of VTCs to provide for the peak throughput requirements of that host.

Delete expired volume data

In using the Delete Expired Volume Data function with a PtP VTS, there are additional operational considerations that must be included in determining the expire time settings for the function because of the manner in which the VTCs select which VTS will be used in processing the I/O for a scratch mount. The following I/O VTS selection criteria applies here:

- Unless the VTC option for forcing scratch mounts to the preferred I/O VTS is enabled, the VTC will always select the VTS that has a valid version of a logical volume, independent of whether it has been expired by the host (assigned to a scratch category).
- If both VTSs have a valid version, then the preferred VTS is selected, if that mode of operation is specified.
- If a logical volume has been deleted from a VTS, it is no longer valid and that VTS cannot be selected for the I/O VTS unless the mount is a scratch mount and the option for forcing scratch mounts to the preferred I/O VTS is enabled.
- But, if the volume has been deleted from both VTSs, then the preferred VTS is selected, if specified, and the VTC creates the volume as a newly initialized volume.

In setting up a PtP VTS ***and*** when using the delete expired volume data function, the following considerations apply, depending on the I/O selection mode of operation:

- For no preference:

  Ensure that the non-zero expire time on *both* Library Managers is set to the same value. If not, the VTS that expires the data first will not be used for scratch mounts.

- For preferred VTS:

  Enable the option for forcing scratch mounts to the preferred I/O VTS. As long as the preferred VTS is available, it will be selected as the I/O VTS independent of whether it has a valid version of the volume. If the volume had been deleted by the delete expired volume data function, an initialized volume is created for the mount.

Note: If the option for forcing scratch mounts to the preferred I/O VTS was *not* enabled, ensure that a *non-zero* expire time is set on *both* Library Managers *and* that the non-preferred VTS/Library is set to delete the data *before* the preferred VTS/Library.

Setting a non-zero expire time on only the preferred library or setting it to delete data before deletion on the non-preferred one can cause scratch mounts to use the non-preferred VTS for I/O. We recommend that the data on the non-preferred VTS/Library be set to be expired a minimum of 48 hours ahead of the data on the preferred VTS. An additional amount should be added to account for the amount of time the VTSs are to be disconnected for a disaster recovery test. For example, if the maximum disconnected time for a test is 24 hours, then the setting on the preferred VTS/Library should be 72 hours more than the setting on the non-preferred VTS/Library.

Single host — job performance

FICON can provide a significant reduction in job times compared to ESCON, including single job times. Elapsed job times, if running in immediate copy mode, must also account for the Peer-to-Peer copy time which follows the host to VTS transfer when it is complete, and so will always be longer than for a corresponding non-PtP VTS. Even when running in deferred copy mode, the VTC-VTS connection results in some additional job time, but FICON can reduce this additional time in comparison to ESCON.

Figure 3-21 shows some measured single write and read job times for ESCON PtP, stand-alone FICON, and FICON PtP. Write job times are shown for both deferred and immediate copy mode. The job wrote or read 800 MB of data using 32 KB blocks, 2.66:1 data compression, and QSAM BUFNO=20. Laboratory tests indicate that the job time differential between FICON PtP and stand-alone FICON would increase for block sizes less than 16 KB.

Job times are for single jobs with no other VTS activity. Job time is from job start to job end times as shown in MVS job output.

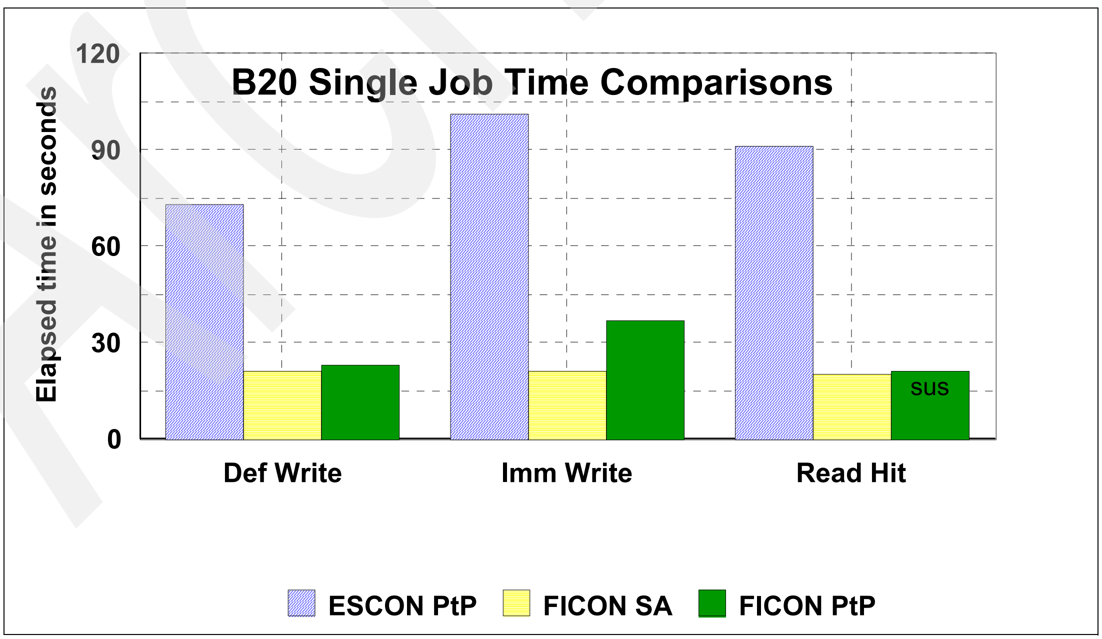

Figure 3-21 Single Job Time Comparisons

Critical job elapsed time requirements

No PtP modeling or sizing tools have the capability to estimate the job time elongation typically experienced when moving workloads to the PtP from native tape or stand-alone VTSs. Although the PtP may provide sufficient total throughput, individual jobs will probably run longer, and sometimes significantly longer on a PtP. This run-time elongation applies to both immediate and deferred mode, and does not include the time in immediate mode to copy the logical volume to the second VTS. The VTS microcode level 2.23.19.1 and higher has improved single job runtimes as opposed to previous code levels, but even with newer levels, runtimes can be elongated.

If you are anticipating runtime-critical production jobs on the PtP, you should involve IBM VTS specialists in the PtP configuration process.

Immediate mode application considerations

Applications like TSM and HSM, which enqueue on SYSZTIOT, or otherwise have several tasks waiting on the completion of a tape rewind/unload, will likely experience significant delays when running in immediate mode compared to a stand-alone VTS or the PtP in deferred mode. In immediate mode, any application that mounts, appends, and demounts the same logical volumes multiple times will also run significantly longer than compared to a stand-alone VTS or a PtP in deferred mode, because a copy will be done at each unload of the volume.

Using FICON instead of ESCON channels will help to reduce these delays.

Immediate mode and 256 virtual devices

Each VTC represents 16 or 32 virtual devices to the hosts and provides internally three copy devices and one broadcast device. The copy devices are used for all the copies which have to be created. If you are going to use 256 devices, each VTC represents 32 virtual devices to the hosts but still 3 copy devices only. This over-commitment may result in copy throttling, if most or even all of the 32 devices are allocated for write activity.

Block size and BUFNO effects on throughput performance

In laboratory measurements, data rates increased significantly with each doubling of block size, from 1 KB (1024 byte) up to 64 KB blocks, after which results were mixed. There was more than an order of magnitude difference in single job data rates between 1 KB and 32 KB blocks. Increasing BUFNO values beyond the QSAM default of five also has been observed to increase data rates for single jobs, although data rate increases were less pronounced for multiple concurrent jobs.

Small volume effects on throughput performance

The performance information presented in this paper has been based on measurements performed with either 800 host MB, or in the modeling for the Mixed Workloads, 250 host MB logical volumes. There are special performance considerations that need to be made if the average volume size in the workload is significantly smaller.

These “small volume” limits are:

- **Logical scratch mount per hour limit:** The rate at which logical scratch mounts can be performed can become the VTS throughput limit if it is asked to handle predominantly small volumes. A FICON PtP VTS is capable of up to approximately 2,500 logical scratch mounts per hour.
- **Logical volume processing limit:** There is a VTS overhead associated with processing each logical volume that can limit the maximum achievable data rate for small volumes.

- **Logical volume closing limit:** There is an overhead associated with closing a volume in the TVC even when the volume has previously been copied to tape. If these volumes are small, it may be necessary to close a number of them before a new host volume can be received.

Workload imbalance across VTCs

The modeled, measured, and published bandwidth numbers for the PtP all assume that the PtP workload is evenly distributed across the VTCs. In actual environments, this is rarely the case. Natural workload dynamics will inevitably result in short-term imbalance across the VTCs, and is to be expected. However, long-term, consistent imbalance is a symptom of an allocation problem. In JES3 systems, be sure that the INISH deck is ordered properly for the PtP. In JES2 systems, a problem has been identified that will prevent allocation from randomizing device allocations to the PtP, resulting in chronic imbalance across the VTCs. You should insure that the fixes for APAR OW52873 are applied.

3.6 Remote installations and switch support

The VTS can be used for any sequential data. You have to evaluate placing data that needs to be taken out of the library, such as off-site backups and interchange data on the VTS. Make sure that the time requirements meet your needs. A second VTS at the off-site location may also be used for off-site copies.

ESCON distances supported are shown in Table 3-10.

Table 3-10 VTS supported distances

	ESCON	ESCON with DWDM	FICON with DWDM
Host server to VTC	43 Km	75 Km	100 Km
VTC to VTS	26 Km	50 Km	100 Km

FICON supported distances are diagrammed in Figure 3-22.

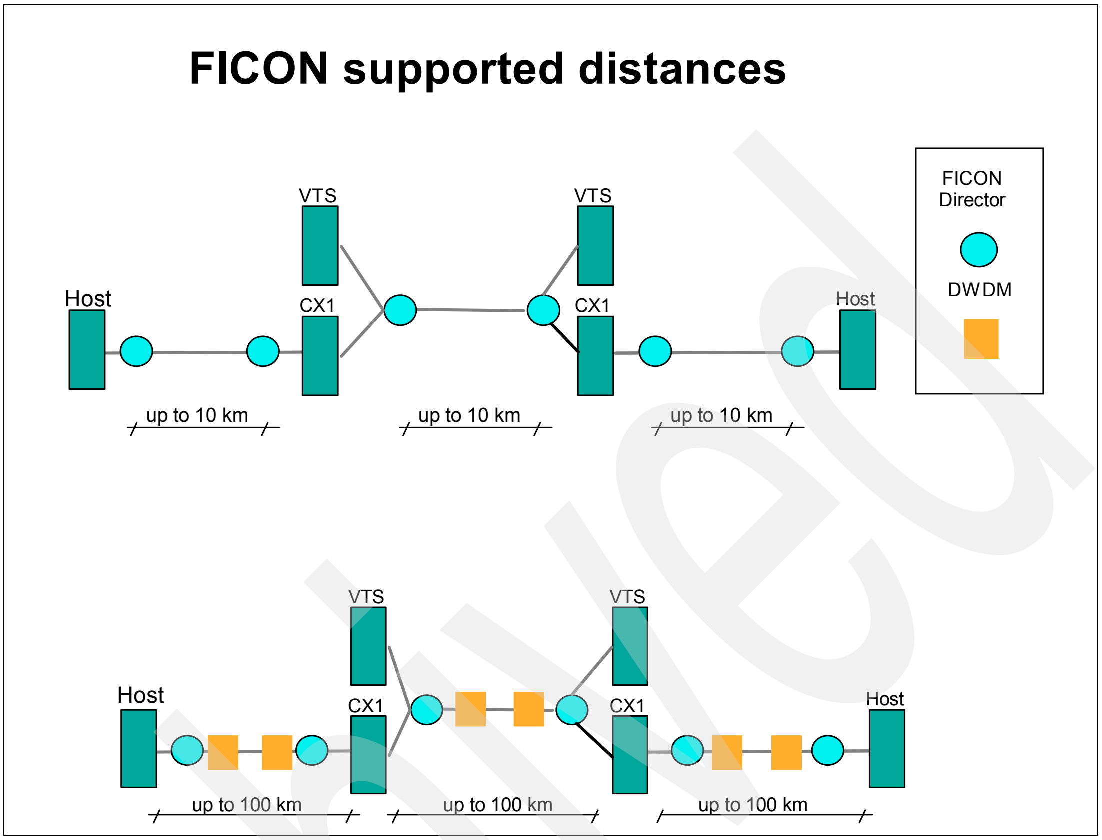

Figure 3-22 Supported FICON distances

3.6.1 Fabric support

There are three basic types of switch connections that can be used between the host and VTC and VTS without any DWDM or channel extenders:

- Direct connect
- Single switch
- Cascaded switch (two switches)

All switches are supported for both stand-alone and Peer-to-Peer VTS with 1Gb or 2Gb links. The components will auto negotiate to the highest speed allowed (1Gb/2Gb). You cannot mix different vendors (McData (CNT and Inrange; CISCO and Brocade) but you may mix models of one vendor. See the switch Web pages for specific intermix combinations supported.

3.6.2 FICON transfer modes

FICON channel extenders are available working in one of the 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 will be 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 are emulating 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.

Channel extender support for R7.4 is planned for stand-alone and Peer-to-Peer VTS for: USD-X and the Edge Router from CNT (now McData) and for CN2000 from CIENA.

3.6.3 FICON support of cascaded directors

FICON support of cascaded directors means that a Native FICON (FC) channel or a FICON CTC can connect a server to a device or other server via two (same vendor) directors in between. This support is for a two switch, single hop configuration only. This type of cascaded support is important for disaster recovery and business continuity solutions because it can provide high availability connectivity as well as the potential for fiber infrastructure cost savings by reducing the number of channels for interconnecting the 2 sites. This cascaded director function is generally available for McDATA (CNT, INRANGE) FC/9000, and IBM TotalStorage SAN Switch M12.

The primary value of cascaded directors would be its potential for cost savings. Two cascaded directors can allow for shared links leading to the reduction in number of channels between sites, and therefore improved utilization of inter-site connected resources and infrastructure (see Figure 3-23).

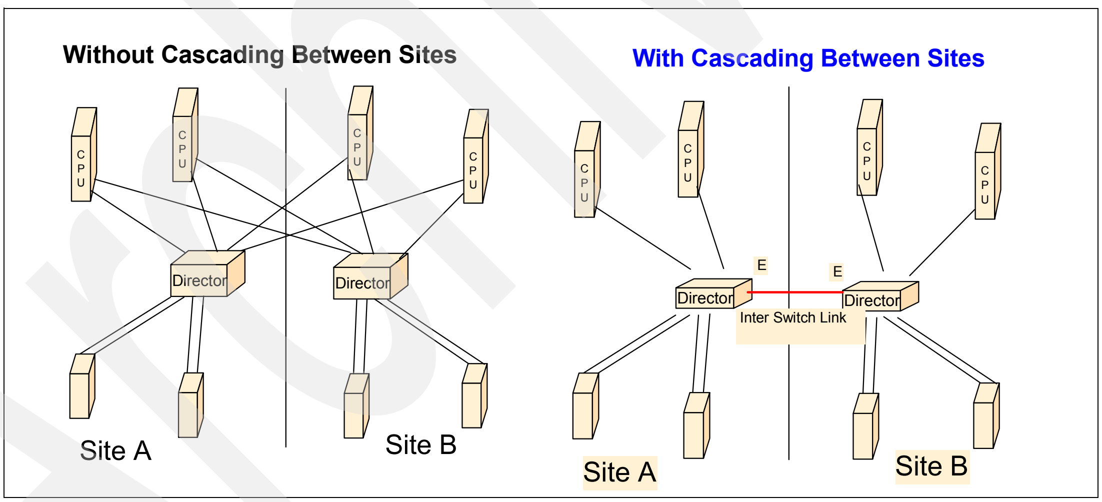

Figure 3-23 PtP VTS without and with cascading

Even more channel consolidation can result from the use of 2Gbps inter-site links between the cascaded directors. Solutions such as Geographically Dispersed Parallel Sysplex (GDPS) can benefit from the reduced inter-site configuration complexity that cascaded directors provide. While specific cost savings vary depending upon infrastructure, generally customers who have data centers separated between two sites may reduce the number of cross site connections by using Cascaded Directors. Further savings may be realized in the reduction of the number of channels and director ports. Care should be taken when consolidating FICON channels onto inter-switch links (ISLs), the FICON links between two switches/directors).

Another potential benefit of cascaded directors is that the maximum unrepeated distance between data centers could be extended, providing more choice for business continuity solutions. With an additional FICON director in the channel configuration the unrepeated distance between two sites could be extended to 30 km (10 km between each server and director) or up to 60 km with RPQ (20 km between each server and director) at 1 Gbps link speeds. At 2 Gbps link speeds, this same RPQ has a maximum of 36 km (12 km between each server and director).

Furthermore, FICON support of cascaded directors introduces integrity features within the FICON cascaded switch fabric that detect and report any mis-cabling actions and prevent data from being delivered to the wrong end point, thus preserving data integrity.

FICON cascading is available to z800, z900, z890_and z990 for all native FICON (FC) channels on both the previous-generation FICON cards (feature codes 2315, 2318) and the newer FICON Express cards (feature codes 2319, 2320). No additional System z hardware is required.

FICON Cascading requires z/OS V1R3 with PTFs or later (see the PSP bucket in Retain for applicable microcode support).

FICON support of cascaded directors requires the use of multiple directors, specifically the McDATA 6140/6064 or (CNT, INRANGE) 9000 (models 064,128,and 256). In addition, these directors must be equipped with their own high integrity features. The McDATA 6140/6064's require the latest 4.1 code level and the optionally priced SANtegrity software. The (CNT, INRANGE) 9000 requires the latest 4.0.1.H code level and the purchase of its FICON Cascading feature.

A general “rule of thumb” might be to have one 2 Gbps (1 Gbps) ISL for each four 2 Gbps (1 Gbps) FICON Express (or FICON) channels whose traffic is to be routed over the ISLs. In special cases where the channels may need to run at high data rates, this number may have to be reduced to avoid exceeding the capacity of the ISL (e.g., 140 MB/s for I/O traffic predominantly in one direction on a 2 Gbps ISL with large block traffic where distance is not a problem). Where the peak I/O workload is known to have low bandwidth requirements, it might be possible to allow more than four FICON Express channels per ISL. However, it is very important in all cases to keep in mind the possible limitations on routing of traffic across multiple ISLs discussed above.

See the Technote, *Performance Considerations for a Cascaded FICON Director Environment* at the following URL:

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

3.6.4 Cascading example

The following diagrams guide you through an example for the definition and implementation of cascading for a PtP VTS Model B10. Cascading applies to stand-alone and PtP VTS. For stand-alone VTS, switches can be cascaded between host and VTS. For the PtP VTS, switches can be cascaded between host and VTC and between VTC and VTS.

Examples shown here (Figure 3-24) are also covered in the redbook *FICON Native Implementation and Reference Guide*, SG24-6266.

We describe cascading for a Model B10 PtP VTS with four VTCs and also four switches for redundancy. To illustrate the HCD definitions, we use the IOCP notation. Figure 3-24 illustrates the installation with two sites with symmetric configurations.

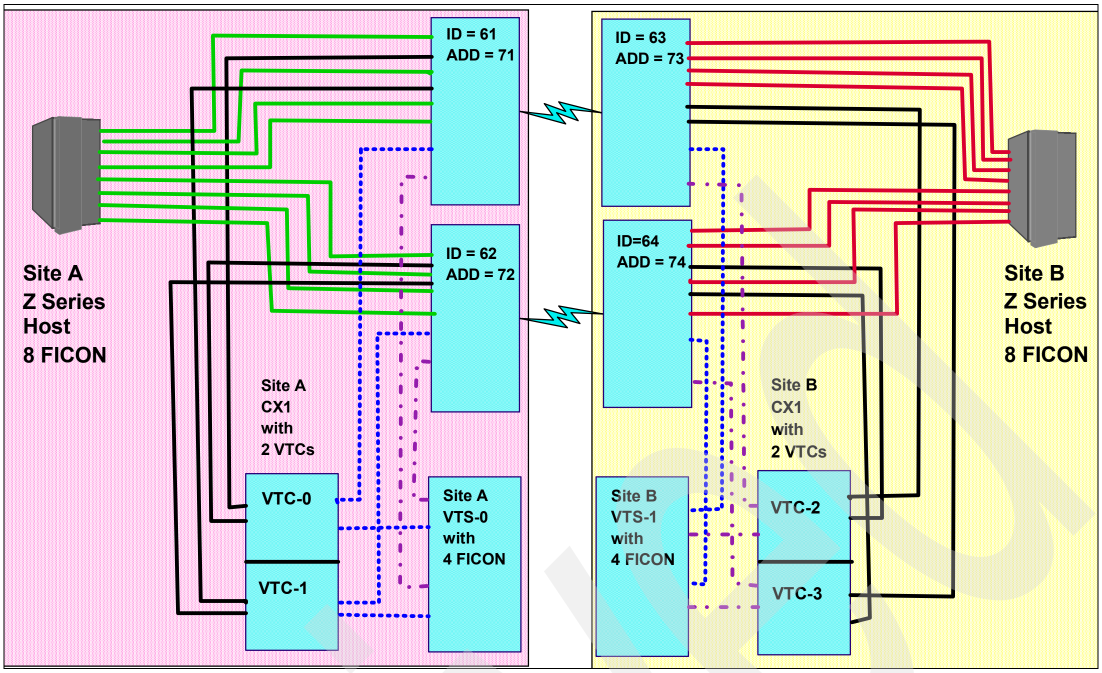

Figure 3-24 Cascading example for a B10 PtP VTS

The following four steps describe the different views of the whole cascading example shown in Figure 3-24.

Host-to-switch definitions for Site A

We start with the host to switch definitions on Site A as shown in Figure 3-25. We use two directors for redundancy and eight host paths for maximum peak throughput.

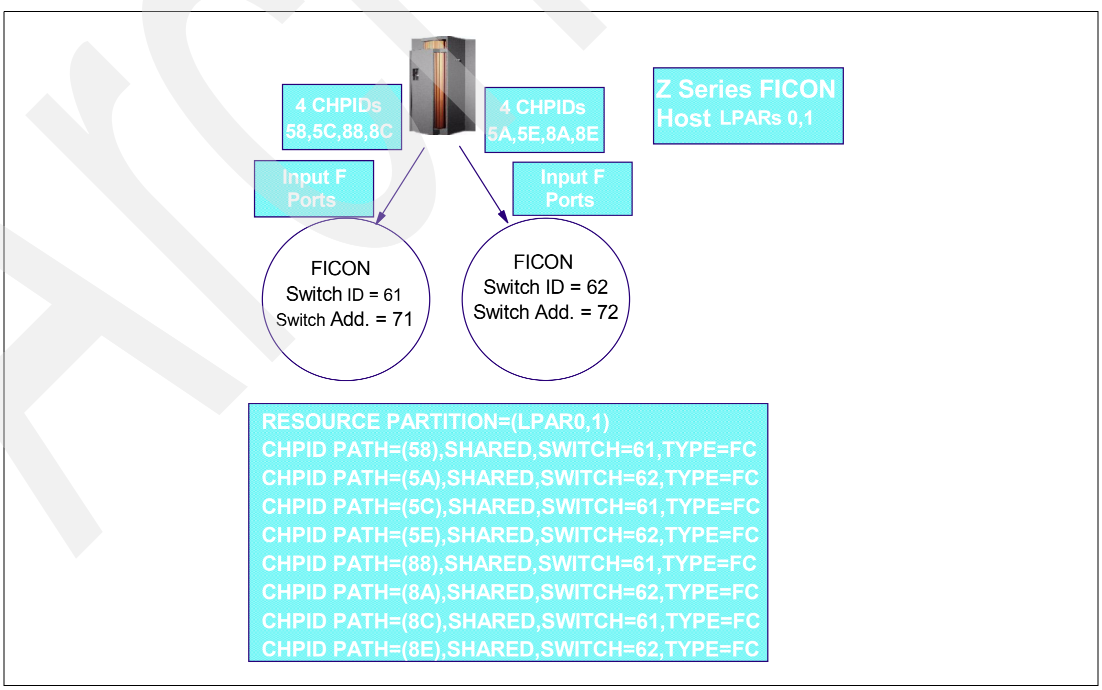

Figure 3-25 Host to switch definitions

Host-Side switch ports are identified in the HCD “Add Channel Path” dialog; these are *F* ports. Note that the director address may be different than the director ID. You must use the hexadecimal director ID in CHPID definition in HCD.

Switch to switch definitions

We continue with the switch to switch definitions from Site A to Site B (see Figure 3-26). The lines represent Inter-Switch links (ISLs); these are *E* ports.

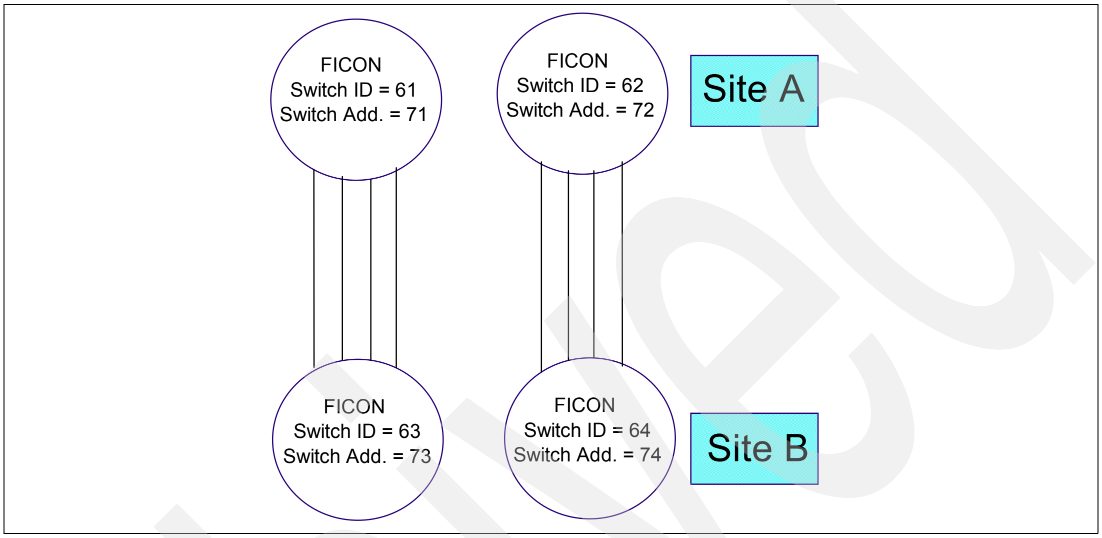

Figure 3-26 The switch-to-switch definitions

ISLs are bidirectional. They are defined using GUI switch setup, which we do not explain here, as the switch setup is manufacturer and model dependent. The number of E ports affects the total bandwidth, therefore you should make sure that you have enough ISLs not to create a performance bottleneck.

Switch to VTC definitions

We continue with the switch-to-VTC definitions on Site B (see Figure 3-27). One 16-device VTC is shown for simplicity, but the other ones are defined similarly.

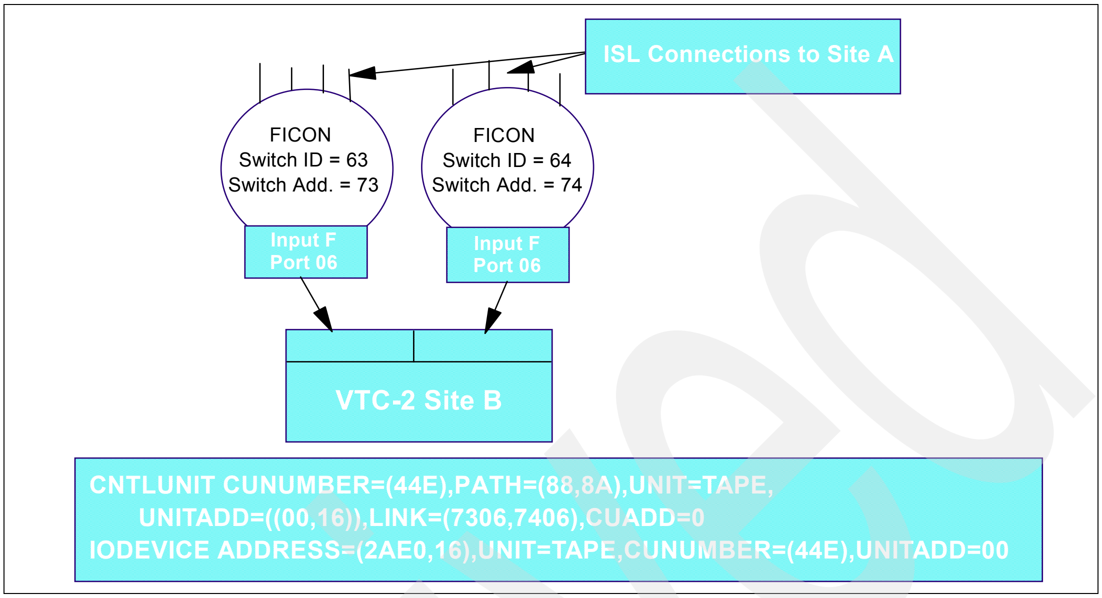

Figure 3-27 HCD definitions for the switch-to-VTC connection at Site B

The CU-side switch port is identified in the HCD link parameter; these are *F* ports. The Link parameter is a 2-byte field containing the switch address and the port number. LINK=(7306) defines port 06 of the switch with switch address 73.

Remember to use the hexadecimal switch address in the CNTLUNIT definition, not the Switch ID.

VTC to VTS definitions

The last step consists of the VTC to VTS definitions on Site B (see Figure 3-28). These definitions are not made at the host because the connections are not known at the System z host. These connections are Fibre Channel connections. The Fibre Channel fabric is in discovery mode, which means that any path enabled is used.

The VTS at Site B directly attaches to the VTC at Site B. This connections could go through the switch as well.

The connection from the VTCs at Site B to the VTS at Site A is back through the same cascaded switches. These connections are configured in the VTC configuration panels using *wwnn* for the VTS FICON adapter; they are not defined in HCD. Traffic back to Site A flows over the same ISL connections as traffic to Site B.

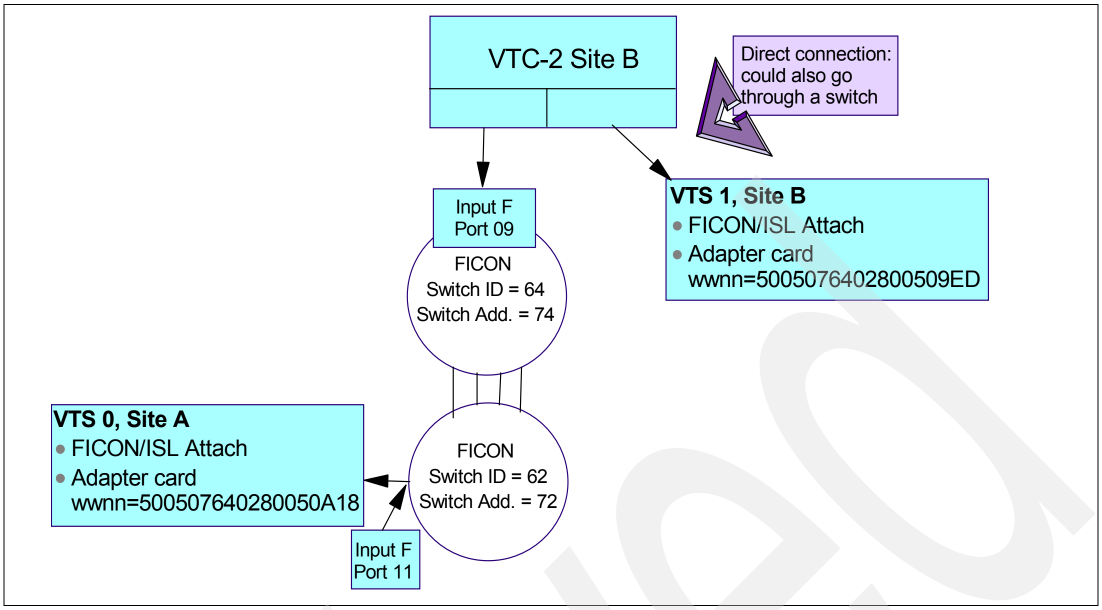

Figure 3-28 Connections from VTC to VTS at Site B

Common rules for cascading definitions

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

- **Director Switch ID:**

  This is defined in the director GUI setup.

  The inboard director switch ID is used in 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 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 in the LINK parameter on 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 may 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 in 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.

3.6.5 Buffer credits

Here we describe some factors that will affect distance.

Fibre Channel distances depend on many factors and 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 loss — connectors, cables, and even bends and loops in the cable can result in dB signal loss

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 will provide a signal boost that can potentially extend distances 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 on 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 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 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 will heavily influence the throughput that is attained within the Fibre Channel. Fibre Channel architecture is based on a flow control that ensures a constant stream 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 the redbook, *Introduction to SAN Distance Solutions*, SG24-6408.

3.7 Peer-to-Peer VTS planning checklist

After you are sure which functions of the PtP VTS you want to exploit, you should set up a checklist with all “to do” items. Such a list might include hardware, software, and operational changes, and could incorporate performance and operational considerations. For example:

- Determine the workload for the PtP VTS.
- Decide which of the Advanced Policy management functions should be used:
  - Logical volume affinity to group logical volumes
  - Dynamic setting of immediate or deferred copy mode by Management Class constructs
  - Cache management by the definition of preference groups and the associated Storage Class
  - Dual copies of logical volumes in each distributed library in addition to the PtP copies (do not do that)
  - Secure data erase, which will influence the usage of back-end drives
  - Larger logical volume sizes defined by Data Classes, which will reduce the peak write time
  - Selective PtP copy to prevent those volumes from being copied which do not request one
- Decide which operational mode has to be selected:
  - No preference or preference (preferred mode or balanced mode)
- Decide which default copy mode has to be selected:
  - Immediate or deferred
- Prepare a tape study:

  The analysis and sizing tasks needed for a Peer-to-Peer VTS are basically the same as those for a stand-alone VTS.

  The tape study based on SMF records of your current installation is necessary for a sound sizing of the PtP VTS. It should consider all data that has to be copied to a remote VTS, any data that should not be copied (selective PtP copy), and the copy mode (immediate or deferred) to project the required and the available bandwidth and storage capacities of the target configuration. Also, the operation mode (preference or no-preference) has to be defined for the study. Furthermore, you should consider the grouping of logical volumes.

  The analysis and sizing will be accomplished by your IBM representative or BP.

  As soon as the sizing is completed, the detailed planning tasks have to be finalized.
- Determine the number of logical volumes:

  Sizing the number of logical volumes and the number of stacked volumes uses the same methods as for a stand-alone VTS. Naturally, you will need to double the number of stacked volumes for a given set of logical volumes since you will be writing the data to two libraries.

- Determine the logical volume serial numbers:

  The logical volumes are defined for the composite library to be used by the host system. Each distributed library will contain the logical volumes defined for the composite library. However, as with a stand-alone VTS, system-managed tape requires that the logical volume serial numbers must not duplicate other volume serial numbers in your systems.

- Determine the number of stacked volumes:

  Similarly, the same considerations apply for your choice of volume serial number ranges for stacked volumes.

- Determine the stacked volume serial numbers:

  We recommend that the stacked volume serial numbers be unique to clearly distinguish each set of the two distributed library's stacked volumes.

- Library name:

  Each of the libraries, composite and both distributed libraries, in the PtP VTS complex must have a name that you specify. We recommend that you use names which let you recognize each library easily.

- Library ID:

  Each of the libraries has its own library ID. Usually, the convention is to use hardware serial numbers of the library frames or VTSs as library IDs. It has to be a 5-digit HEX string (00-FF). For the composite library in a Peer-to-Peer configuration, you could use the serial number of the local CX0 frame. But you may prefer to use a more descriptive library ID such as a number that you assign.

- Configuration options:

  – Cache sizes:

    You should carefully review the need for larger Tape Volume Cache sizes to avoid removal of volumes from the cache before completion of the dual copy during:

    • Extended peak periods with deferred copy mode:

      You should have enough cache to prevent removal of volumes from the cache during peak workloads.

    • Extended maintenance action:

      You should have enough cache to prevent removal of volumes from the cache during a planned outage so that dual copies can be made efficiently when all the components of the configuration are back in service.

      We strongly recommend that both VTSs have the same TVC capacity as well as the same number and, ideally, model of 3590 tape drives. If one of the VTSs has less TVC capacity or fewer 3590 tape drives, its sustained throughput will be less than that of the other VTS. The VTCs balance the workload between the VTSs so the less capable VTS will tend to do less work than the more capable VTS, thereby reducing the combined throughput.

  – Number and type of back-end drives

  – Usage of different media

  – The number of VTCs according to the number of virtual drives needed

  – ESCON or FICON channel attachment

  – Usage considerations of switches

  – Cascading option for infrastructure simplification and reduction for the remote VTS

- Configuration layout:

  This is a detailed configuration plan, which has to include all hardware components, channel connections, installed switches, port IDs, channel extenders, LAN attachments, IP-addresses and remote support lines to the IBM service center.

- Software maintenance:

  Although no additional software is required for the PtP VTS, you should check the basic levels and PTFs.

- HCD and JES3 INISH deck:

  Select the address ranges for the virtual addresses of the PtP VTS to define them by HCD. If JES3, consider the INISH deck changes.

- Determine VTC settings for additional cache management (see business continuance settings)
- Selection of the preferred master VTS
- Selection of the U/I VTS
- Selection of the user interface (UI) distributed library:

  We recommend that you use the distributed library that allows easiest access as the UI distributed library.

  > **Note:** This designation is static and can only be changed by taking the entire Peer-to-Peer VTS offline and performing CSR reconfiguration.

- Plan the VTS management policies:
  - Fast-ready category
  - Expired volume management
  - Reclaim thresholds
  - Free storage threshold
  - Inhibit reclamation schedule
- Plan all the software definitions for DFSMS system-managed tape:
  - OAM address space
  - SMS CDSs
  - TCDB
  - PARMLIB changes
  - SMS constructs and ACS routines
- Plan the education of data management and operating staff.
- Update the operational procedures and documentation.
- Plan monitoring procedures with the specialists.
- Plan monitoring procedures with VTSSTATS based on SMF94 records.

3.8 Education and training on the VTS

The amount of education and training your staff requires on the VTS depends on a number of factors, including these:

- Are you installing the PtP VTS in an existing library?
- Are both the VTS and the library new to your site?
- Are you using BTLS or SMS tape to manage your library?
- Is the library and the VTS shared among multiple systems?
- Do you have existing tape drives at your site?

3.8.1 Adding a VTS to build a PtP VTS

If you are adding a VTS to an existing 3494 containing a VTS in order to build a PtP complex, the training needed for your operators, systems programmers, and storage administrators is minimal. Although the Library Manager posts operator interventions that are specific to the VTS, 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 VTS-specific interventions and the actions necessary to resolve the conditions.

We recommend that training for your operations staff include the following items:

- IBM VTS as a closed store for IBM VTS without the Advanced Function feature
- Operator intervention conditions and their resolution
- Proper procedures for opening and closing frames
- VTS-specific Library Manager functions
- Advanced Function (Export/Import)

We recommend that your storage administrators and systems programmers receive the same training as the operations staff, plus these items:

- Software choices and how they affect the IBM VTS
- Disaster recovery considerations

3.8.2 New VTS, 3494 library, or TS3500 Tape Library with 3953 Tape System

If you are putting in a new 3494 or TS3500/3953 with a VTS, the education and training for your staff should include the VTS training items listed in 3.8.1, “Adding a VTS to build a PtP VTS” on page 124, 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 libraries
- Handling operator intervention requests
- Performing Manual Mode operations
- Hands-on training with the Web Specialist

We recommend that you review:

- *IBM TotalStorage 3494 Tape Library Introduction and Planning Guide*, GA32-0448
- *IBM TotalStorage 3494 Tape Library Operator’s Guide*, GA32-0449
- *IBM TotalStorage Tape Library Guide for Open Systems*, SG24-5946
- *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229 (Appendix C: VTS implementation step-by-step)
- *IBM TotalStorage Enterprise Automated Tape Library (3494) Systems Assurance Product Review (SAPR) Guide*, for additional information about education and training. (See“Related publications” on page 453 for bibliographic information for this publication.)

Education is a service offering of IBM Global Services (IGS). Also available are the services for library and VTS implementation described in the following sections.

IBM implementation services

A range of services is available to assist you with your VTS. 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:

  – Virtual Tape Server implementation provides technical consultation, software planning, and assistance and operation education to customers implementing an IBM VTS. 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 VTS 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, VTS, or Peer-to-Peer implementation.

- Migration:

  IBM Virtual Tape Server Migration Services provides VTS data migration for virtually all system platforms from any IBM VTS to another IBM VTS. This can be used with data center consolidations or when upgrading to newer VTS technology.

- 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.

4

Implementation

In this chapter, we describe the steps and tasks required to implement a new Peer-to-Peer VTS. Again, we assume that the reader is already familiar with basic VTS concepts, and thus we concentrate specifically on information for the PtP VTS. However, if you are new to the IBM TotalStorage 3494 Tape Library or the IBM System Storage TS3500 Tape Library with IBM 3953 Tape System, then we recommend that you first review the *IBM TotalStorage Enterprise Tape A Practical Guide*, SG24-4632 or the *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning and Implementation*, SG24-6789. If you are new to the IBM TotalStorage VTS, then we also suggest that you review the *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring*, SG24-2229 redbook to familiarize yourself with the VTS concepts before implementing a PtP VTS.

The PtP VTS Implementation can be logically separated into three sections:

1. **Hardware I/O configuration definitions:** This segment relates to the system generation. It consists of processes such as FICON/ESCON attachment to the host, HCD/IOCP definition, and Missing Interrupt Handler (MIH) settings.
2. **VTS definitions from the Library Manager:** This segment focuses on defining volumes, policies, and categories using Library Manager operations. With the addition of Advanced Policy Management (APM), there are also considerations for the definition of Data Class (DC), Management Class (MC), Storage Class (SC), and Storage Group (SG) construct names and policies.
3. **VTS software definitions:** This segment describes how to define the new PtP to the host operating systems:[1]:
   - **System z:** For DFSMS/MVS and SMS Tape this would include updating the DFSMS Automatic Class Selection (ACS) routines, Object Access Method (OAM), and your tape management system. With APM, you must also consider that the DC, MC, SC, and SG construct and policies that you define will be passed to the PtP VTSs.
   - **VM/guests:** VM only supports the PtP VTS in VTS emulation mode. It also requires that an MVS system also be attached to the same PtP VTS.

   The tasks we outline in these sections are meant to give you an overall concept of the tasks required for a PtP VTS installation. The task lists provided are not all-inclusive, and they are not intended to replace any current documentation for the PtP VTS.

[1] Open Systems, TPF, and native VSE operating systems do not support the IBM TotalStorage Peer-to-Peer VTS.

4.1 Hardware I/O configuration definition

From a host perspective, a Peer-to-Peer (PtP) VTS looks like four, eight, or sixteen IBM TotalStorage 3490E tape control units, each with sixteen devices attached through ESCON or FICON channels.

In a stand-alone VTS, all logical devices and logical control units (LCUs) are associated with one VTS controller. To define a stand-alone VTS with 64 virtual drives, it is necessary to define four tape control units of four strings, 16 drives each, using logical unit address (CUADD) equal to 0 to 3. To define a stand-alone VTS with 256 virtual drives, it is necessary to define sixteen tape control units using `CUADD = 0` to `15`.

In a Peer-to-Peer VTS, the logical devices and logical control units are associated with four or eight independent Virtual Tape Controllers (VTCs). Each VTC for a PtP VTS is an independent tape control unit with its own ESCON/FICON channel adapters. For a model B10 PtP VTS with 4 VTCs and 64 virtual drives, or a model B20 PtP VTS with 8 VTCs and 128 virtual drives, the CUADD parameters of the CNTLUNIT statements are coded as CUADD=0 for each VTC. Each VTC in these two configurations has only a single logical control unit (LCU).

However, in the case of a model B20, with 4 VTCs and 128 virtual drives or 8 VTCs and 256 virtual drives, each of the VTCs has <u>two</u> logical control units. Hence for these last two PtP VTS configurations the CUADD parameters of the CNTLUNIT statements are coded as CUADD=0 for the first LCU in each VTC, and CUADD=1 for the second LCU in each VTC.

Table 4-1 shows the differences regarding logical tape control unit definitions between a stand-alone VTS and a PtP VTS.

Table 4-1 Difference between Peer-to-Peer VTS and stand-alone VTS

VTSs	**Logical tape control unit definition**		
	# of physical ESCON paths	**# of physical FICON paths**	**CUADD**
Stand-alone VTS	4, 8, or 16	2, 4, or 8	0 - 3 0 - 7 0 - 15
Peer-to-Peer VTS	2 (per VTC)	2 (per VTC)	0 0 - 1

Before you can use the PtP VTS, you need to define it to the MVS system and processor through HCD. Virtual 3490E tape drives are defined just like a physical IBM 3490-A10 controller with 16 addresses through the HCD panels. You must define them through the HCD dialog, specifying `LIBRARY=YES`. We strongly recommend that you also specify LIBRARY-ID and LIBPORT-ID. You may need to update missing interrupt handler (MIH) values as well. In this section we discuss the steps required to configure a Peer-to-Peer VTS.

If ESCON or 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).

4.1.1 Logical path considerations

The VTS can be attached to System z hosts and logical partitions through two, four, or eight FICON channels or four, eight, or sixteen ESCON channels.

A PtP VTS with four VTCs attaches to the hosts through eight ESCON channels, eight FICON channels, or four ESCON and four FICON channels. A PtP VTS with eight VTCs provides twice the number of channels and attaches to the hosts through sixteen ESCON channels, sixteen FICON channels, or eight ESCON and eight FICON channels.

The following formula can be used to calculate the number of logical paths required in an installation:

(# of hosts) x (# of CUs) x (# of ESCON/FICON channels) = number of logical paths

Each *FICON* channel attachment provides 128 logical paths, which makes up a total of 256 logical paths in a two-channel configuration, a total of 512 logical paths in a four-channel configuration, and 1024 logical paths in an eight-channel configuration.

Each *ESCON* channel attachment provides 64 logical paths, which makes up a total of 128 logical paths in a two-channel configuration, a total of 256 logical paths in a four-channel configuration, 512 logical paths in an eight-channel configuration, and 1024 logical paths in a sixteen-channel configuration.

Define one control unit (CU) macro in the HCD dialog for every 16 virtual devices. Up to eight channel paths may be defined to each control unit. A logical path may be thought of as a three element entity: a host port, a VTS port, and a logical control unit in the PtP VTS.

If the number of logical paths needed exceeds the maximum number of paths supported, then the number of logical paths must be reduced. It may be assumed, that not all of the hosts have the same throughput requirements. Therefore the best solution is to adjust the attachment definitions to the throughput requirements of individual hosts. Define more channels and control units for hosts with high throughput requirements, and define fewer channels or control units for hosts that have less throughput requirements or require less virtual devices.

ESCON versus FICON HCD definitions

The examples in this chapter are for FICON attachments. ESCON attachment definitions will be identical, with two exceptions:

- For ESCON channels, the TYPE parameter on the CHPID definition must specify TYPE=CNC
- The maximum number of VTS channel attachments for ESCON is sixteen as shown in Table 4-1.

FICON cascading HCD definitions

FICON support of cascaded directors means that a Native FICON (FC) channel or a FICON CTC can connect a server to a device or other server via two (same vendor) directors in between. This support is for a two-switch, single-hop configuration only. This type of cascaded support is important for disaster recovery and business continuity solutions because it can provide high availability connectivity as well as the potential for fiber infrastructure cost savings by reducing the number of channels for interconnecting the two sites. This cascaded director function is generally available for McDATA (CNT, INRANGE) FC/9000, and IBM TotalStorage SAN Switch M12.

Note: The HCD examples in this chapter are for *regular* FICON. For FICON cascading HCD examples, refer instead to 3.6.4, “Cascading example” on page 115.

HCD support for library and port IDs

LIBRARY-ID and LIBPORT-ID are z/OS HCD parameters which 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.

> **Important:** We strongly recommend that you specify 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 may protect you from having to IPL when you make changes to your I/O configuration.

LIBRARY-ID

Each IBM 3953 partition (Native, VTS1, VTS2) is assigned a unique five hex-character "Sequence Number" by IBM service representative during installation. This sequence number is arbitrary, and can be selected by the customer. If a number is not selected by the customer, you may use the last five digits of the hardware serial number. For a native TS3500/3953 library, this would be the last five digits of the serial number of the L05 frame. For each VTS distributed library, it would be the last five digits of the serial number of the VTS.

Each logical library in the IBM 3494 (Native, VTS1, VTS2) is also assigned a unique five hex-character sequence number by the IBM service representative during installation. Again, the number is arbitrary, but you can choose to use the last five digits of the serial number. (Associated VTS for the two IBM PtP distributed VTS libraries, or Lxx frame for native)

Again, the composite library ID defined by your IBM service representative is arbitrary. We have no specific recommendation for what to use for this five hex-character sequence number. You may wish to consider using a name or number that makes it easily recognizable as being associated with a composite library.

> **Important:** These sequence numbers must match the LIBRARY-ID numbers used in the HCD library definitions and the library IDs listed in the ISMF Tape Library Define panels.

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 find out what to specify. All tape drives (logical or physical) connected to a given logical control unit (LCU) 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**

DS	device service
QT	query tape
devnum	device address
1	number of devices to be displayed
RDC	read device characteristics

Figure 4-1 details the output of a DS QT system command:

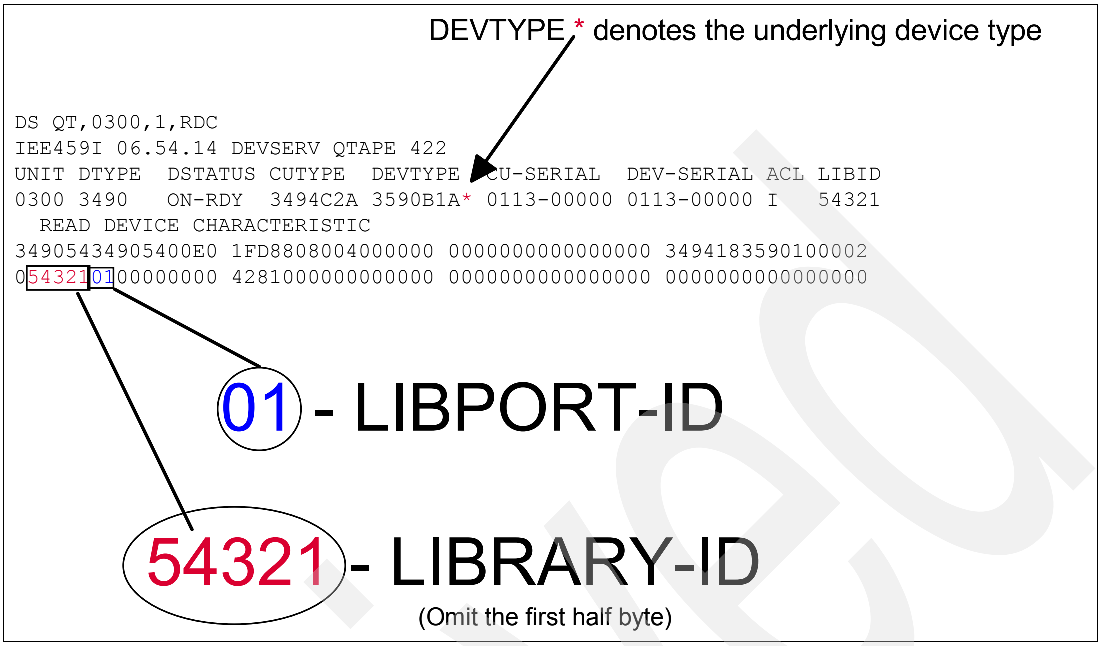

Figure 4-1 Sample DEVSERV QT command output

4.1.2 Defining devices through IOCP

Although you cannot use IOCP to define a Peer-to-Peer VTS (you must use HCD), you may find the IOCP statements helpful in understanding how the Peer-to-Peer VTS has to be defined in HCD. The four examples that follow are currently the only four IBM PTP VTS configurations possible. The first two configurations have a single LCU for each VTC. The last two configurations have two LCUs for each VTC.

Model B10 PtP VTS with 4 VTCs and 64 virtual drives

This IOCP example and Figure 4-2 shows a typical Peer-to-Peer VTS model B10 configuration that has four virtual tape controllers (VTCs), eight FICON channel paths to the host, and four logical tape control units (LCUs) with a total of 64 virtual tape drives.

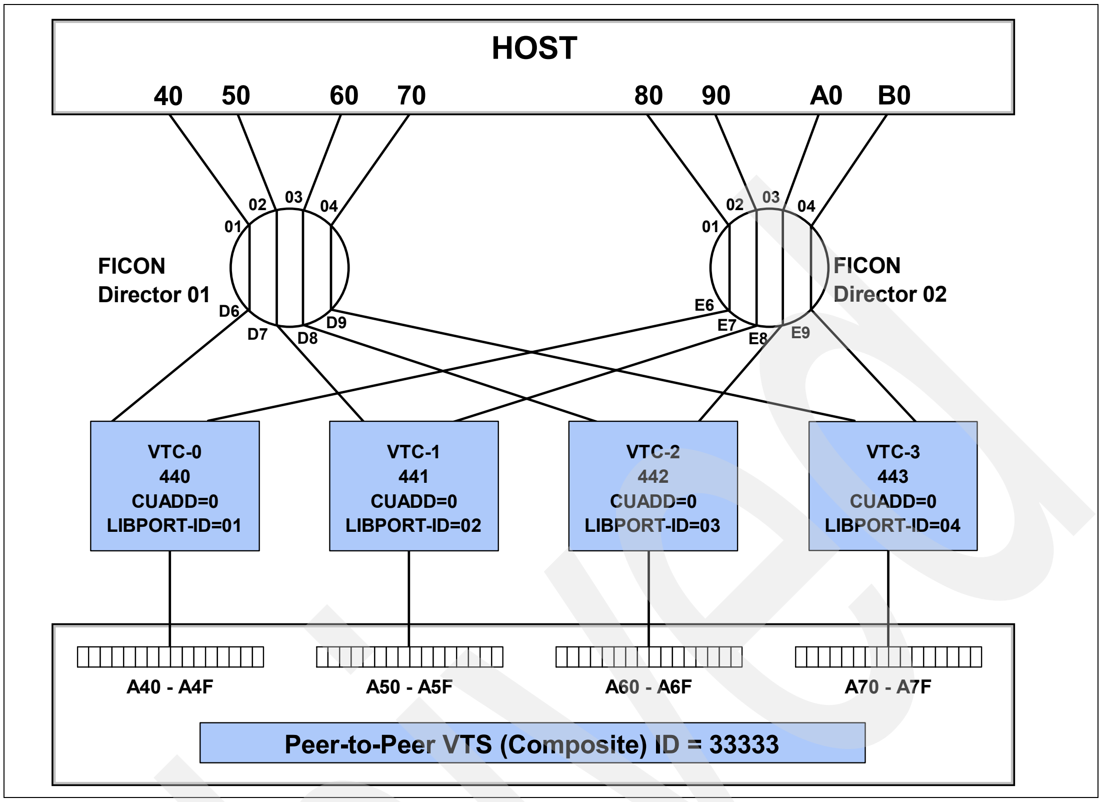

Figure 4-2 PtP VTS Model B10 with 4 VTCs and 64 virtual drives

The IOCP definitions for the configuration shown in Figure 4-2 are listed in Example 4-1.

Example 4-1 IOCP definitions for B10 PtP VTS

```
CHPID PATH = 40,TYPE=FC,SWITCH=01
CHPID PATH = 50,TYPE=FC,SWITCH=01
CHPID PATH = 60,TYPE=FC,SWITCH=01
CHPID PATH = 70,TYPE=FC,SWITCH=01
CHPID PATH = 80,TYPE=FC,SWITCH=02
CHPID PATH = 90,TYPE=FC,SWITCH=02
CHPID PATH = A0,TYPE=FC,SWITCH=02
CHPID PATH = B0,TYPE=FC,SWITCH=02
*
CNTLUNIT CUNUMBR=440,PATH=(40,80),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D6,E6),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=01
IODEVICE ADDRESS=(A40,16),UNIT=3490,      X
CUNUMBR=(440),UNITADD=00
*
CNTLUNIT CUNUMBR=441,PATH=(50,90),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D7,E7),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=02
IODEVICE ADDRESS=(A50,16),UNIT=3490,      X
CUNUMBR=(441),UNITADD=00
*
CNTLUNIT CUNUMBR=442,PATH=(60,A0),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D8,E8),CUADD=0,                     X
```

```
LIBRARY-ID=33333,LIBPORT-ID=03
IODEVICE ADDRESS=(A60,16),UNIT=3490,      X
CUNUMBR=(442),UNITADD=00
*
CNTLUNIT CUNUMBR=443,PATH=(70,B0),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D9,E9),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=04
IODEVICE ADDRESS=(A70,16),UNIT=3490,      X
CUNUMBR=(443),UNITADD=00
```

Model B20 PtP VTS with 8 VTCs and 128 virtual drives

This IOCP example and Figure 4-3 on page 133 shows a typical Peer-to-Peer VTS model B20 configuration that has eight virtual tape controllers (VTCs), sixteen FICON channel paths to the host, and eight logical tape control units (LCUs) with a total of 128 virtual tape drives.

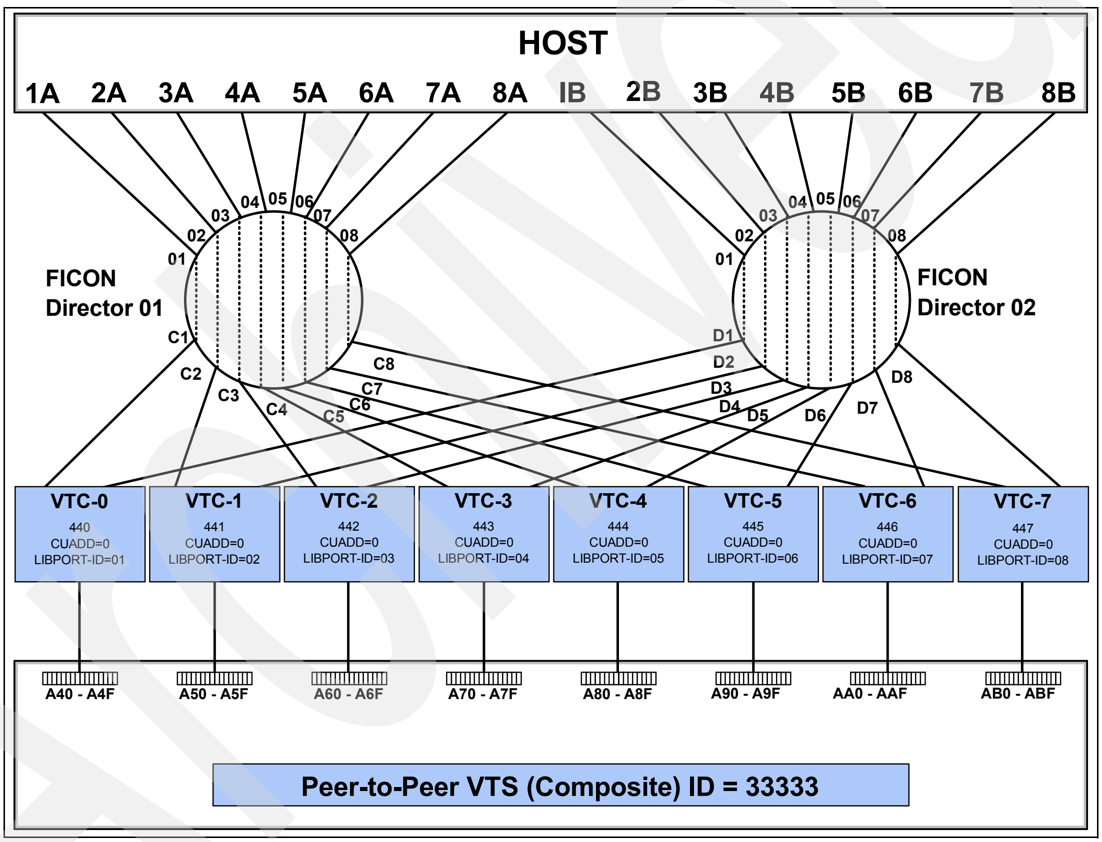

Figure 4-3 PtP VTS Model B20 with 8 VTCs and 128 virtual drives

The IOCP definitions for the configuration shown above in Figure 4-3 are listed below in Example 4-2.

Example 4-2 IOCP definitions for a B20 PtP VTS with 8 VTCs and 129 devices

```
CHPID PATH = 1A,TYPE=FC,SWITCH=01
CHPID PATH = 2A,TYPE=FC,SWITCH=01
CHPID PATH = 3A,TYPE=FC,SWITCH=01
CHPID PATH = 4A,TYPE=FC,SWITCH=01
CHPID PATH = 5A,TYPE=FC,SWITCH=01
CHPID PATH = 6A,TYPE=FC,SWITCH=01
CHPID PATH = 7A,TYPE=FC,SWITCH=01
```

```
CHPID PATH = 8A,TYPE=FC,SWITCH=01
CHPID PATH = 1B,TYPE=FC,SWITCH=02
CHPID PATH = 2B,TYPE=FC,SWITCH=02
CHPID PATH = 3B,TYPE=FC,SWITCH=02
CHPID PATH = 4B,TYPE=FC,SWITCH=02
CHPID PATH = 5B,TYPE=FC,SWITCH=02
CHPID PATH = 6B,TYPE=FC,SWITCH=02
CHPID PATH = 7B,TYPE=FC,SWITCH=02
CHPID PATH = 8B,TYPE=FC,SWITCH=02
*
CNTLUNIT CUNUMBR=440,PATH=(1A,1B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C1,D1),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=01
IODEVICE ADDRESS=(A40,16),UNIT=3490,      X
CUNUMBR=(440),UNITADD=00
CNTLUNIT CUNUMBR=441,PATH=(2A,2B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C2,D2),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=02
IODEVICE ADDRESS=(A50,16),UNIT=3490,      X
CUNUMBR=(441),UNITADD=00
*
CNTLUNIT CUNUMBR=442,PATH=(3A,3B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C3,D3),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=03
IODEVICE ADDRESS=(A60,16),UNIT=3490,      X
CUNUMBR=(442),UNITADD=00
CNTLUNIT CUNUMBR=443,PATH=(4A,4B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C4,D4),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=04
IODEVICE ADDRESS=(A70,16),UNIT=3490,      X
CUNUMBR=(443),UNITADD=00
*
CNTLUNIT CUNUMBR=444,PATH=(5A,5B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C5,D5),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=05
IODEVICE ADDRESS=(A80,16),UNIT=3490,      X
CUNUMBR=(444),UNITADD=00
CNTLUNIT CUNUMBR=445,PATH=(6A,6B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C6,D6),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=06
IODEVICE ADDRESS=(A90,16),UNIT=3490,      X
CUNUMBR=(445),UNITADD=00
*
CNTLUNIT CUNUMBR=446,PATH=(7A,7B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C7,D7),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=07
IODEVICE ADDRESS=(AA0,16),UNIT=3490,      X
CUNUMBR=(446),UNITADD=00
CNTLUNIT CUNUBR=447,PATH=(8A,8B),         X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C8,D8),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=08
```

```
IODEVICE ADDRESS=(AB0,16),UNIT=3490,      X
CUNUMBR=(447),UNITADD=00
```

Note: Even if you are installing a 128-address Peer-to-Peer configuration, we recommend that you consider the impact of increasing the number of addresses in the future.

Model B20 PtP VTS with 4 VTCs and 128 virtual drives

This IOCP example and Figure 4-4 on page 135 shows a typical Peer-to-Peer VTS model B20 configuration that has four virtual tape controllers (VTCs), eight FICON channel paths to the host, and eight logical tape control units (LCUs) with a total of 128 virtual tape drives.

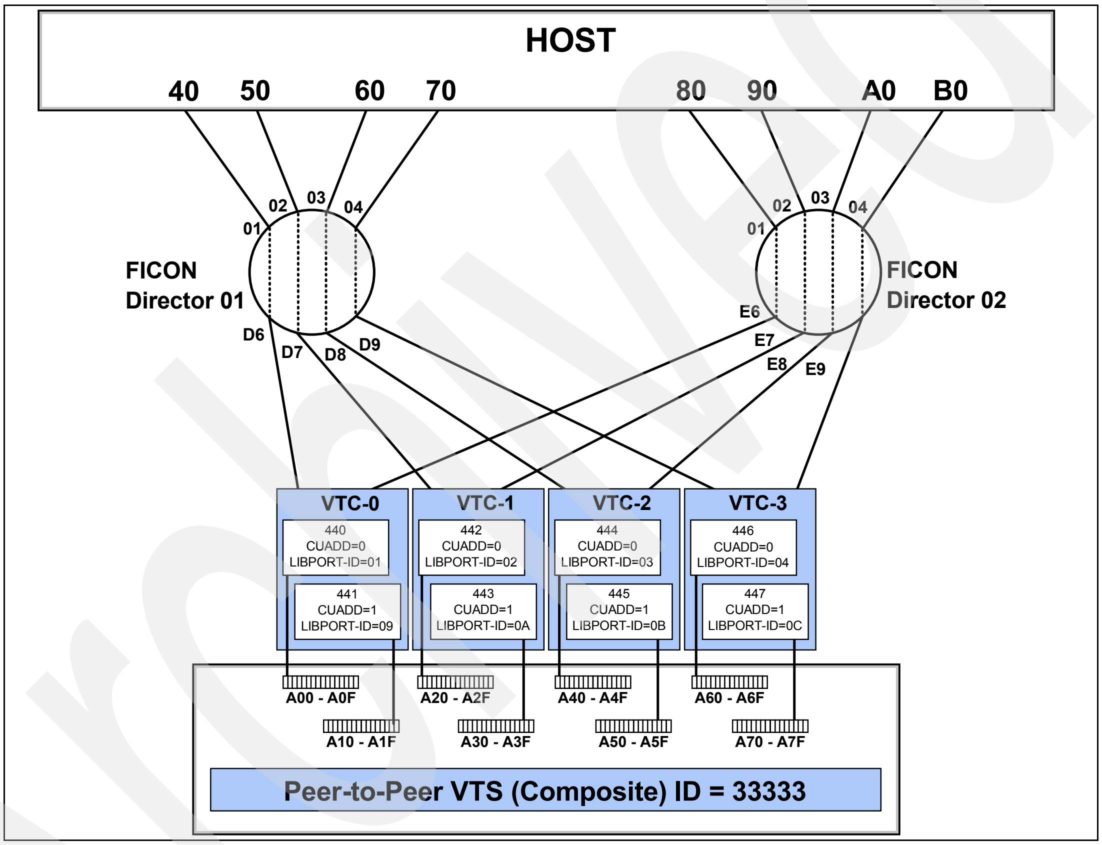

Figure 4-4 PtP VTS Model B20 with 4-VTCs and 128 virtual drives

Note that in this PtP VTS configuration, each VTC consists of two LCUs, the first coded as CUADD=0, the second as CUADD=1. Also note that the LIBPORT-IDs must be assigned as shown.

Important: The range of device addresses within each LCU must be contiguous, starting with xxx0 and continuing through xxxF.

Although not required, it is recommended that the 16-device address ranges for the two LCUs in each VTC be contiguous, and that the device address range of the first LCU in the first VTC be xx00-xx0F, as shown. This addressing scheme will map to the VTS internal device addressing, and can facilitate any required problem determination.

Example 4-3 shows the IOCP definitions for the configuration described in Figure 4-4.

Example 4-3 IOCP definitions for a B20 PtP VTS with four VTCs

```
CHPID PATH = 40,TYPE=FC,SWITCH=01
CHPID PATH = 50,TYPE=FC,SWITCH=01
CHPID PATH = 60,TYPE=FC,SWITCH=01
CHPID PATH = 70,TYPE=FC,SWITCH=01
CHPID PATH = 80,TYPE=FC,SWITCH=02
CHPID PATH = 90,TYPE=FC,SWITCH=02
CHPID PATH = A0,TYPE=FC,SWITCH=02
CHPID PATH = B0,TYPE=FC,SWITCH=02
*
CNTLUNIT CUNUMBR=440,PATH=(40,80),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D6,E6),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=01
IODEVICE ADDRESS=(A00,16),UNIT=3490,      X
CUNUMBR=(440),UNITADD=00
CNTLUNIT CUNUMBR=441,PATH=(40,80),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D6,E6),CUADD=1,                     X
LIBRARY-ID=33333,LIBPORT-ID=09
IODEVICE ADDRESS=(A10,16),UNIT=3490,      X
CUNUMBR=(441),UNITADD=00
*
CNTLUNIT CUNUMBR=442,PATH=(50,90),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D7,E7),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=02
IODEVICE ADDRESS=(A20,16),UNIT=3490,      X
CUNUMBR=(442),UNITADD=00
CNTLUNIT CUNUMBR=443,PATH=(50,90),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D7,E7),CUADD=1,                     X
LIBRARY-ID=33333,LIBPORT-ID=0A
IODEVICE ADDRESS=(A30,16),UNIT=3490,      X
CUNUMBR=(443),UNITADD=00
*
CNTLUNIT CUNUMBR=444,PATH=(6A,A0),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D8,E8),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=03
IODEVICE ADDRESS=(A40,16),UNIT=3490,      X
CUNUMBR=(444),UNITADD=00
CNTLUNIT CUNUMBR=445,PATH=(60,A0),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D8,E8),CUADD=1,                     X
LIBRARY-ID=33333,LIBPORT-ID=0B
IODEVICE ADDRESS=(A50,16),UNIT=3490,      X
CUNUMBR=(445),UNITADD=00
*
CNTLUNIT CUNUMBR=446,PATH=(70,B0),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D9,E9),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=04
IODEVICE ADDRESS=(A60,16),UNIT=3490,      X
CUNUMBR=(446),UNITADD=00
CNTLUNIT CUNUMBR=447,PATH=(70,B0),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(D9,E9),CUADD=1,                     X
```

```
LIBRARY-ID=33333,LIBPORT-ID=0C
IODEVICE ADDRESS=(A70,16),UNIT=3490,      X
CUNUMBR=(447),UNITADD=00
```

Note: Even if you are installing a 128-address Peer-to-Peer configuration, we recommend that you consider the impact of increasing the number of addresses in the future.

Model B20 PtP VTS with 8 VTCs and 256 virtual drives

This IOCP example and Figure 4-5 shows a typical Peer-to-Peer VTS model B20 configuration that has eight virtual tape controllers (VTCs), sixteen FICON channel paths to the host, and sixteen logical tape control units (LCUs) with a total of 256 virtual tape drives.

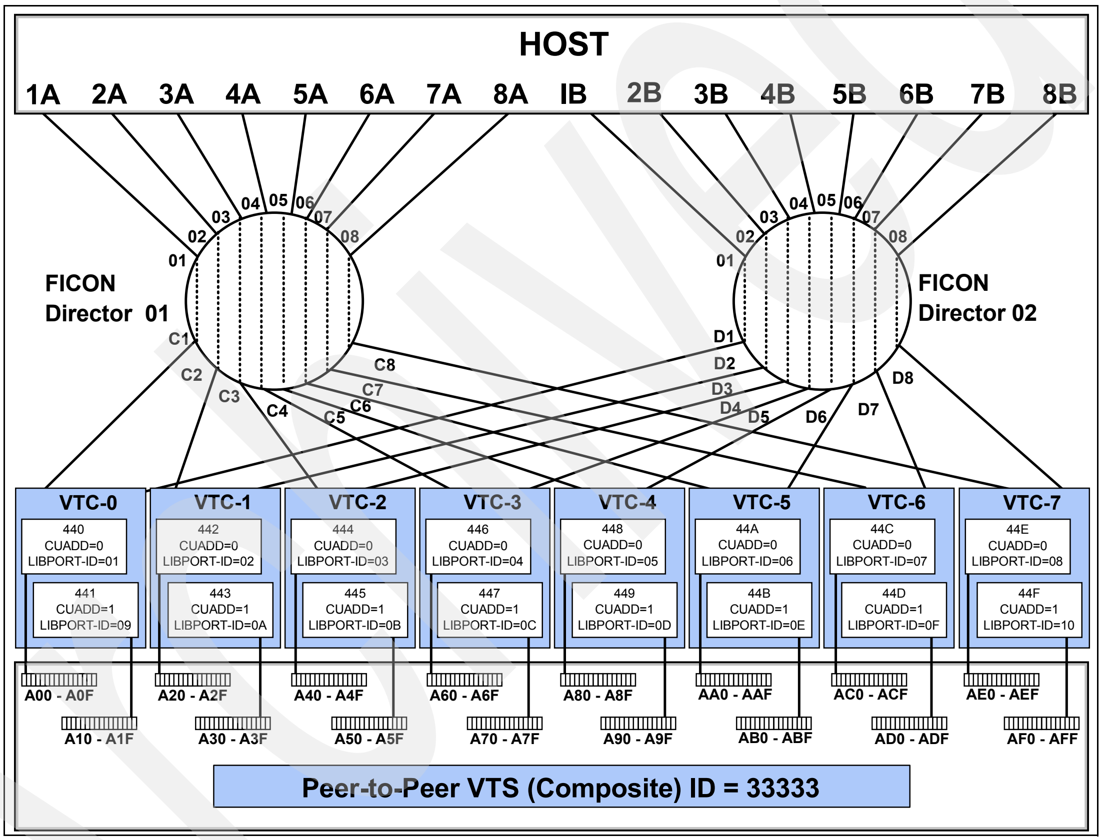

Figure 4-5 PtP VTS Model B20 with 8 VTCs and 256 virtual drives

Note that each VTC consists of two LCUs, the first coded as CUADD=0, the second as CUADD=1. Also note that the LIBPORT-IDs must be assigned as shown.

Important: The range of device addresses within each LCU must be contiguous, starting with xxx0 and continuing through xxxF

Although not required, it is recommended that the 16-device address ranges for the two LCUs in each VTC be contiguous, and that the device address range of the first LCU in the first VTC be xx00-xx0F, as shown. This addressing scheme will map to the VTS internal device addressing, and can facilitate any required problem determination. See Example 4-4.

Example 4-4 IOCP definitions for a B20 PtP VTS with 256 devices

```
CHPID PATH = 1A,TYPE=FC,SWITCH=01
CHPID PATH = 2A,TYPE=FC,SWITCH=01
CHPID PATH = 3A,TYPE=FC,SWITCH=01
CHPID PATH = 4A,TYPE=FC,SWITCH=01
CHPID PATH = 5A,TYPE=FC,SWITCH=01
CHPID PATH = 6A,TYPE=FC,SWITCH=01
CHPID PATH = 7A,TYPE=FC,SWITCH=01
CHPID PATH = 8A,TYPE=FC,SWITCH=01
CHPID PATH = 1B,TYPE=FC,SWITCH=02
CHPID PATH = 2B,TYPE=FC,SWITCH=02
CHPID PATH = 3B,TYPE=FC,SWITCH=02
CHPID PATH = 4B,TYPE=FC,SWITCH=02
CHPID PATH = 5B,TYPE=FC,SWITCH=02
CHPID PATH = 6B,TYPE=FC,SWITCH=02
CHPID PATH = 7B,TYPE=FC,SWITCH=02
CHPID PATH = 8B,TYPE=FC,SWITCH=02
*
CNTLUNIT CUNUMBR=440,PATH=(1A,1B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C1,D1),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=01
IODEVICE ADDRESS=(A00,16),UNIT=3490,      X
CUNUMBR=(440),UNITADD=00
CNTLUNIT CUNUMBR=441,PATH=(1A,1B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C1,D1),CUADD=1,                     X
LIBRARY-ID=33333,LIBPORT-ID=09
IODEVICE ADDRESS=(A10,16),UNIT=3490,      X
CUNUMBR=(441),UNITADD=00
*
CNTLUNIT CUNUMBR=442,PATH=(2A,2B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C2,D2),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=02
IODEVICE ADDRESS=(A20,16),UNIT=3490,      X
CUNUMBR=(442),UNITADD=00
CNTLUNIT CUNUMBR=443,PATH=(2A,2B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C2,D2),CUADD=1,                     X
LIBRARY-ID=33333,LIBPORT-ID=0A
IODEVICE ADDRESS=(A30,16),UNIT=3490,      X
CUNUMBR=(443),UNITADD=00
*
CNTLUNIT CUNUMBR=444,PATH=(3A,3B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C3,D3),CUADD=0,                     X
LIBRARY-ID=33333,LIBPORT-ID=03
IODEVICE ADDRESS=(A40,16),UNIT=3490,      X
CUNUMBR=(444),UNITADD=00
*
CNTLUNIT CUNUMBR=445,PATH=(3A,3B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C3,D3),CUADD=1,                     X
LIBRARY-ID=33333,LIBPORT-ID=0B
IODEVICE ADDRESS=(A50,16),UNIT=3490,      X
CUNUMBR=(445),UNITADD=00
CNTLUNIT CUNUMBR=446,PATH=(4A,4B),        X
UNIT=3490,UNITADD=((00,16)),              X
LINK=(C4,D4),CUADD=0,                     X
```

```
LIBRARY-ID=33333,LIBPORT-ID=04
IODEVICE ADDRESS=(A60,16),UNIT=3490,       X
CUNUMBR=(446),UNITADD=00

CNTLUNIT CUNUMBR=447,PATH=(4A,4B),         X
UNIT=3490,UNITADD=((00,16)),               X
LINK=(C4,D4),CUADD=1,                      X
LIBRARY-ID=33333,LIBPORT-ID=0C
IODEVICE ADDRESS=(A70,16),UNIT=3490,       X
CUNUMBR=(447),UNITADD=00
CNTLUNIT CUNUMBR=448,PATH=(5A,5B),         X
UNIT=3490,UNITADD=((00,16)),               X
LINK=(C5,D5),CUADD=0,                      X
LIBRARY-ID=33333,LIBPORT-ID=05
IODEVICE ADDRESS=(A80,16),UNIT=3490,       X
CUNUMBR=(448),UNITADD=00

CNTLUNIT CUNUMBR=449,PATH=(5A,5B),         X
UNIT=3490,UNITADD=((00,16)),               X
LINK=(C5,D5),CUADD=1,                      X
LIBRARY-ID=33333,LIBPORT-ID=0D
IODEVICE ADDRESS=(A90,16),UNIT=3490,       X
CUNUMBR=(449),UNITADD=00
CNTLUNIT CUNUMBR=44A,PATH=(6A,6B),         X
UNIT=3490,UNITADD=((00,16)),               X
LINK=(C6,D6),CUADD=0,                      X
LIBRARY-ID=33333,LIBPORT-ID=06
IODEVICE ADDRESS=(AA0,16),UNIT=3490,       X
CUNUMBR=(44A),UNITADD=00
*
CNTLUNIT CUNUMBR=44B,PATH=(6A,6B),         X
UNIT=3490,UNITADD=((00,16)),               X
LINK=(C6,D6),CUADD=1,                      X
LIBRARY-ID=33333,LIBPORT-ID=0E
IODEVICE ADDRESS=(AB0,16),UNIT=3490,       X
CUNUMBR=(44B),UNITADD=00
CNTLUNIT CUNUMBR=44C,PATH=(7A,7B),         X
UNIT=3490,UNITADD=((00,16)),               X
LINK=(C7,D7),CUADD=0,                      X
LIBRARY-ID=33333,LIBPORT-ID=07
IODEVICE ADDRESS=(AC0,16),UNIT=3490,       X
CUNUMBR=(44C),UNITADD=00
CNTLUNIT CUNUMBR=44D,PATH=(7A,7B),         X
UNIT=3490,UNITADD=((00,16)),               X
LINK=(C7,D7),CUADD=1,                      X
LIBRARY-ID=33333,LIBPORT-ID=0F
IODEVICE ADDRESS=(AD0,16),UNIT=3490,       X
CUNUMBR=(44D),UNITADD=00
*
CNTLUNIT CUNUMBR=44E,PATH=(8A,8B),         X
UNIT=3490,UNITADD=((00,16)),               X
LINK=(C8,D8),CUADD=0,                      X
LIBRARY-ID=33333,LIBPORT-ID=08
IODEVICE ADDRESS=(AE0,16),UNIT=3490,       X
CUNUMBR=(44E),UNITADD=00
CNTLUNIT CUNUMBR=44F,PATH=(8A,8B),         X
UNIT=3490,UNITADD=((00,16)),               X
LINK=(C8,D8),CUADD=1,                      X
LIBRARY-ID=33333,LIBPORT-ID=10
```

```
IODEVICE ADDRESS=(AF0,16),UNIT=3490,      X
CUNUMBR=(44F),UNITADD=00
```

4.1.3 Defining devices through HCD panels

In this section, we describe the process of defining the Peer-to-Peer VTS through HCD panels. These are the most important points to observe:

- HCD definitions are required for SMS tape.
- Four or eight 3490 tape control units have to be defined, with 16 3490E drives each.
- Keep the link address blank when no ESCON director is used.
- Specify `LIBRARY=YES` when using system-managed tape.

Remember, when defining a Peer-to-Peer VTS, it is necessary to define four, eight, or sixteen independent tape control units, and four, eight, or sixteen strings of 16 3490E drives each. In a stand-alone VTS, the definitions of the four or eight tape control units use a logical unit address of `CUADD = 0` to `3`, or `0` to `7`, or 0 to 15 respectively. In a Peer-to-Peer VTS, however, the CUADD parameters of the CNTLUNIT statement are coded differently.

For a model B10 PtP VTS with 4 VTCs and 64 virtual drives, or a model B20 PtP VTS with 8 VTCs and 128 virtual drives, the CUADD parameters of the CNTLUNIT statements are coded as CUADD=0 for each VTC. Each VTC in these two configurations has only a single logical control unit (LCU).

However, in the case of a model B20, with 4 VTCs and 128 virtual drives or 8 VTCs and 256 virtual drives, each of the VTCs has *two* logical control units. Hence, for these last two PtP VTS configurations the CUADD parameters of the CNTLUNIT statements are coded as *CUADD=0 for the first LCU* in each VTC, and *CUADD=1 for the second LCU* in each VTC.

Coding the HCD for configurations that have two LCUs on each VTC

In the case of the last two configurations, PtP VTS Model B20 4-VTC 128 drives and PtP VTS 256 drives, the values for the LIBPORT-IDs associated with the Logical Control Units (LCUs) ***must*** be coded as shown in Table 4-2 and Table 4-3. This numbering scheme is designed to facilitate upgrading from a 128-device configuration without modifying the existing LIBPORT-IDs, but must be used even if the 256-device configuration is initially installed.

Although not required, we recommend that the device address ranges for the two LCUs in each VTC be contiguous, and that the device address range of the first LCU in the first VTC be xx00-xx0F, as shown in Table 4-2 and Table 4-3. This addressing scheme will map to the VTS internal device addressing, and can facilitate any required problem determination.

Table 4-2 LIBPORT-ID Values for B20 PtP VTS with 4 VTCs and 128 Devices VTC

VTC	LIBPORT-ID of first LCU in this VTC	Recommended device address range	LIBPORT-ID of second LCU in this VTC	Recommended device address range
VTC-0	01	xx00-xx0F	09	xx10-xx1F
VTC-1	02	xx20-xx2F	0A	xx30-xx3F
VTC-2	03	xx40-xx4F	0B	xx50-xx5F
VTC-3	04	xx60-xx6F	0C	xx70-xx7F

Table 4-3 LIBPORT-ID values for model B20 PtP with 256 devices

VTC	LIBPORT-ID of first LCU in this VTC	Recommended device address range	LIBORT-ID of second LCU in this VTC	Recommended device address range
VTC-0	01	xx00-xx0F	09	xx10-xx1F
VTC-1	02	xx20-xx2F	0A	xx30-xx3F
VTC-2	03	xx40-xx4F	0B	xx50-xx5F
VTC-3	04	xx60-xx6F	0C	xx70-xx7F
VTC-4	05	xx80-xx8F	0D	xx90-xx9F
VTC-5	06	xxA0-xxAF	0E	xxB0-xxBF
VTC-6	07	xxC0-xxCF	0F	xxD0-xxDF
VTC-7	08	xxE0-xxEF	10	xxF0-xxFF

Sample HCD Panels

The screens in Figure 4-6 and Figure 4-7 show the important HCD control unit definition panels. For the purposes of this HCD example, we use the 4-VTCs 128 virtual drive configuration shown in "Model B20 PtP VTS with 4 VTCs and 128 virtual drives" on page 135. This configuration has 2 FICON switches.

We now use HCD to define the same first control unit as in Example 4-3 on page 136.

HCD Panels (Add 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 02 __ __ __ __ __ __ +
Ports . . . . . . . . . . . D6 E6 __ __ __ __ __ __ +
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 4-6 Add Control Unit screen (1)

After you have specified the control unit number (here, 0440) and the type (virtual so here, 3490) in the first panel, and any FICON switches (here, 2 FICON and LINK=(D6,E6)), then press Enter. The panel shown in Figure 4-7 will be displayed in order for you to choose the processor to which the control unit is to be connected.

```
 -------------------------- 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 80 __ __ __ __ __ __ +
Link address . . . . . . D6 E6 __ __ __ __ __ __ +

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 4-7 Add Control Unit screen (2)

As shown in Figure 4-7, the CUADD for control unit number 0440 is **0**. For control unit number 0441, in our example, CUADD would be **1** (indicated by the arrow).

Note that if the VTS is attached through ESCON directors or FICON switches, the link address fields specify the director's port number. (In this example LINK=(D6,E6) PATH=(40,80)) If the PtP VTS is not being attached through ESCON directors or FICON switches, then the link address fields are blank.

Remember: For configurations with *two* LCUs per VTC, you must code CUADD=0 for the first LCU in the VTC and CUADD=1 for the second LCU in the VTC! Also, you must follow the LIBPORT-ID chart shown in Table 4-3 on page 141 and Table 4-3 on page 141 when specifying your LIBPORT-ID values.

You should define the remaining VTS tape control units, specifying the logical unit address CUADD=0 or CUADD=1, in the same way as the first tape control unit definition in the CBDPCU12 panel. Refer to "Model B20 PtP VTS with 4 VTCs and 128 virtual drives" on page 135 for the IOCP statements generated from this PtP VTS 4-VTC 128 device HCD definition and use them to assist you with your PtP VTS HCD definitions.

HCD Panels (Add Device)

The number of drives to be specified for a Peer-to-Peer VTS control unit is 64, 128, or 256 and you must specify these in strings of 16 devices, with each string associated with one of the previously defined control units. To define the IBM 3490E 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 VTS virtual drives, press **PF11**. The screen is shown in Figure 4-8.

```
 ------------------------------ Add Device ---------------------------
CBDPDV10
Specify or revise the following values.
Device number . . . . . . . . 0A00 (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 4-8   Add Device screen

After entering the required information and specifying to which processors and operating systems the devices are connected, you can update the device parameters using the screen shown in Figure 4-9.

```
 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  33333          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 4-9   Define Device Parameters screen

Notes: When defining library parameters (see Figure 4-9 on page 143):

1. You must specify `LIBRARY=YES`.
2. We strongly recommend that you specify the composite library ID in the LIBRARY-ID field.
3. LIBPORT-ID is dependent on the device number that was assigned to the logical drive:
   – For the first two configurations, B10 with 64 drives and B20 with 8-VTCs:

     The first 16 devices that have the lowest device numbers must specify `LIBPORT-ID = 1`, the second 16 devices must specify `LIBPORT-ID = 2`, the third 16 devices must specify `LIBPORT-ID = 3`, and so on, until you have defined LIBPORT for all the logical drives.

   – For the last two configurations, B20 4-VTCs with 128 drives and B20 256 drives:

     Refer to the chart in table Table 4-3 on page 141 and Table 4-3 on page 141. While the first 16 devices will use LIBPORT-ID = 1, the second group of 16 will use LIBPORT=9, etc.
4. If more than one MVS system is sharing the virtual drives in the VTS, specify `SHARABLE=YES`. This will force OFFLINE to `YES`. It is up to the installation to ensure proper serialization from all attached hosts. We recommend that you review the IBM Redbook, *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409, before attaching a second host to the IBM PtP VTS.

To define the next 16 VTS 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 VTS virtual drives, press **PF11**. The CBDPDV10 screen is displayed again.

4.1.4 Set values for Missing Interrupt Handler (MIH)

An MIH time of 45 minutes is recommended for the all PtP VTS drives. The PtP VTS emulates 3490E devices and does not automatically communicate the MIH timeout values to the host operating system in the Read Configuration Data Channel Control Word (CCW). You must specify the MIH time-out values for IBM 3490E devices. The value applies only to the virtual 3490E drives and not to the real IBM 3590/3592 drives that the VTS manages. Remember that the host only knows about logical 3490E devices. Table 4-4 summarizes recommended minimum values that may need to be increased, depending on specific operational factors.

Table 4-4 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
3494 VTS 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 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 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 timeouts, but you cannot change the secondary timeout. Those are delivered by the self describing values from the device itself.

More information on MIH settings can be found in the *MVS Initialization and Tuning Reference*, SA22-7592-08

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 may present their own MIH timeout values, via the *primary/secondary* MIH timing enhancement contained in the self-describing data for the device. The *primary* MIH timeout value is used for most I/O commands, but the *secondary* MIH timeout value may 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 timeout 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 timeout value may adversely affect MIH recovery processing for the device or device class.

Please refer to the specific device's reference manuals to determine if the device supports self-describing MIH time values.

Note: The TS3500 support APAR OA09751 for OAM and OA09753 for Device Support says:

“This occurs before Missing Interrupt Handler gets control in the Recovery path. The algorithm for selecting a device to send an order to a tape library includes logic that selects a recovery device following an I/O error. Some errors are detected when a 4-minute timeout occurs (I/O has not completed after 4 minutes).

Following such errors, an alternate device is selected. Currently, an alternate device will continue to be selected until all devices have been attempted. The logic is changed by this APAR in two ways.

First, a 4-minute timeout is reported by the new message:

```
IEA439I TAPE LIBRARY(lllll),DEVICE(dddd) FOUR MINUTE I/O TIMEOUT
```

Second, after two such errors for the same I/O operation, the I/O is failed.

If using RMM as your tape management system, install OA10566 (RMM takes advantage of the new CBRXLCS I/O timeout detection reason code).”

4.1.5 Activating the I/O configuration

After the HCD definitions for the PtP VTS have been made, you must activate the IODF with the new definitions. To create the required control blocks, you can either IPL your host systems — which is not the preferred solution — or you can follow the steps described below, which are also described in APAR II09065.

The procedures to follow are different for the installation of a new PtP VTS and for modifications to an existing PtP VTS.

New installation

1. Activate an IODF that defines all of the devices.
2. MVS Vary online the devices in the library. This will create some control blocks, and you will see the message:

   ```
   IEA437I TAPE LIBRARY DEVICE(DDD), ACTIVATE IODF=XX, IS REQUIRED
   ```

3. Then do a final activate. This activate is required to build the Eligible Device Table (EDT) for MVS Allocation.

Modification of an existing configuration

When new control units are added to a library, that will cause the address of any of the string’s devices to change, you must delete the devices first before you can define them. The steps of this procedure are:

1. Activate an IODF deleting all devices from the library.
2. Activate an IODF that defines all of the devices.
3. MVS Vary online the devices in the library. This will create some control blocks, and you will see the message:

   ```
   IEA437I TAPE LIBRARY DEVICE(DDD), ACTIVATE IODF=XX, IS REQUIRED
   ```

4. Then do a final activate.

After that, the virtual devices are defined and ready to use.

An alternative procedure to define new devices or modify existing configurations makes use of the DS QL command (see APAR OA07505) to delete the control blocks:

DS QL,*nnnnn*,DELETE

Where *nnnnn* is the LIBID. Depending on whether you have defined LIBID and LIBPORT or not, use one of the following two procedures to define and activate the devices:

1. **With LIBID and LIBPORT defined:**
   a. Use QLIB LIST to display that the INACTIVE control blocks have been deleted.
   b. Use ACTIVATE IODF to redefine the devices.
   c. Use QLIB LIST to display that the ACTIVE control blocks are properly defined.
2. **Without LIBID and LBPORT coded:**
   a. MVS Vary online the devices in the library. This will create some control blocks and you will see the message:
   `IEA437I TAPE LIBRARY DEVICE(DDD), ACTIVATE IODF=XX, IS REQUIRED`
   b. Use ACTIVATE IODF to redefine the devices.
   c. Use QLIB LIST to display and verify that the ACTIVE control blocks are properly IPL the system

4.1.6 D SMS command

The *D SMS* command can be used to display and check the settings within your DEVSUPxx member. Specifically, it will show you the current hardware scratch categories assigned to the MVS host it was issued from. This enhancement to display the PtP VTS hardware scratch categories was provided through OW54054 (for APM VTS release 7.2).

> **Note:** DISPLAY SMS,LIBRARY(xxxxxxxx),DETAILS

In addition, several new library status lines were added, including:

```
"Library supports outboard policy management."
"Library supports import/export."
```

The following status line now includes the distributed library:

```
"Service preparation occurring in distributed library x."
```

Example 4-5 shows the new display output, but in this case, the OPM function is not supported by this library. The display shows that MEDIA2 uses category 0002.

Please note that the scratch count of MEDIA2 must not match with the number of scratch volumes of your tape management system when you use the *Expire Hold* function in the VTS. OAM displays the scratch count it gets by the Library Manager.

Example 4-5 Display SMS,LIB

```
DISPLAY SMS,LIB(AVTS),DETAIL:
CBR1110I OAM library status: 230
TAPE        LIB   DEVICE     TOT    ONL     AVL      TOTAL     EMPTY  SCRTCH  ON   OP
LIBRARY     TYP    TYPE      DRV    DRV     DRV       SLOTS    SLOTS  VOLS
AVTS        VL  3494-L10     64     64       60      3622       817     4505  Y     Y
----------------------------------------------------------------------
MEDIA     SCRATCH     SCRATCH        SCRATCH
TYPE        COUNT       THRESHOLD    CATEGORY
MEDIA2      4505            0            0002
----------------------------------------------------------------------
OPERATIONAL STATE:  AUTOMATED
ERROR CATEGORY SCRATCH COUNT:             1
SCRATCH STACKED VOLUME COUNT:            75
PRIVATE STACKED VOLUME COUNT:          1388
----------------------------------------------------------------------
Convenience I/O station installed.
Convenience I/O station in Output mode.
Convenience I/O station Empty.
Bulk input/output not configured.
```

4.1.7 DEVSERV QUERY LIBRARY command

The DEVSERV query library function should always be used to query your PtP VTS library configuration before and after an activate of your IODF to define the PtP VTS, or any time changes are made to the PtP VTS library environment.

DS QL,LIST

Figure 4-10 shows an example of how to list all libraries. Note that some of the libraries have an asterisk next to them. Those libraries are actually attached to the host.

```
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 library's that are actually attached to the host.
```

Figure 4-10 DEVSERV QLIB,LIST

DS QL,libid,DETAIL

Figure 4-11 shows a detailed List of one single Library. Use this display to check that no duplicate LIBPORT IDs are listed and that each Port has 16 devices, which is the correct number for a VTS or PTP VTS composite library.

```
DS QL,BA094,DETAIL
IEE459I 15.02.46 DEVSERV QLIB 774
LIBID   PORTID       DEVICES
BA094   05           298B* 298C* 298D* 298E* 298F* 2980* 2981* 2982*
                     2983* 2984* 2985* 2986* 2987* 2988* 2989* 298A*
        06           29AB* 29AC* 29AD* 29AE* 29AF* 29A0* 29A1* 29A2*
                     29A3* 29A4* 29A5* 29A6* 29A7* 29A8* 29A9* 29AA*
        07           29CB* 29CC* 29CD* 29CE* 29CF* 29C0* 29C1* 29C2*
                     29C3* 29C4* 29C5* 29C6* 29C7* 29C8* 29C9* 29CA*
        08           29EB* 29EC* 29ED* 29EE* 29EF* 29E0* 29E1* 29E2*
                     29E3* 29E4* 29E5* 29E6* 29E7* 29E8* 29E9* 29EA*
        0D           2990* 2991* 2992* 2993* 2994* 2995* 2996* 2997*
                     2998* 2999* 299A* 299B* 299C* 299D* 299E* 299F*
        0E           29B0* 29B1* 29B2* 29B3* 29B4* 29B5* 29B6* 29B7*
                     29B8* 29B9* 29BA* 29BB* 29BC* 29BD* 29BE* 29BF*
        0F           29D0* 29D1* 29D2* 29D3* 29D4* 29D5* 29D6* 29D7*
                     29D8* 29D9* 29DA* 29DB* 29DC* 29DD* 29DE* 29DF*
        10           29F0* 29F1* 29F2* 29F3* 29F4* 29F5* 29F6* 29F7*
                     29F8* 29F9* 29FA* 29FB* 29FC* 29FD* 29FE* 29FF*

NOTE: asterisks indicate devices that are currently attached.
```

Figure 4-11 DEVSERV QLIB,libid,DETAIL

Attention: Please see Appendix C, "DEVSERV QLIB command" on page 407 for the complete QLIB syntax and description.

4.1.8 Library LMPOLICY command

Use the LIBRARY LMPOLICY command[2] 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. See "Displaying Tape Volume Status" (see Figure 6-18 on page 246).

[2] This function was originally shipped with APAR OW54054 that also provided OPM Support and the tape management Eject Notification function.

The syntax of the LIBRARY LMPOLICY command to assign or change volume policy names is as listed in Example 4-6.

Example 4-6 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.

The following parameters are Optional Parameters:

- **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

4.2 VTS definitions from the Library Manager

After the hardware installation of a Peer-to-Peer VTS, you must define the environment for each of the two distributed libraries.

After your IBM service representative installs the IBM PtP VTS hardware, you must define the environment for each of the two distributed libraries. The VTS definitions for both distributed VTSs are performed from the respective Library Manager console except for the define (insert) logical volumes operation. The insert logical volumes operation can only be performed on the UI distributed library.

For further details on operating a 3494, refer to the *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-00449. For details on operating the TS3500 Tape Library, refer to *IBM System Storage TS3500 Tape Library Operator Guide, GA32-0560*, and for the IBM 3953 Tape System, refer to *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473.

Table 4-5 shows operations that must be performed at both the Library Managers.

Table 4-5 Library Manager operations

Operations	**Distributed libraries**	
	VTS 0 (UI library)	**VTS 1**
Define Advanced Policy Management outboard constructs[2.] Storage Groups for Stacked Media Pooling. Storage Classes for TVC management. Data Classes. Management Classes for dual copies and PtP Copy modes.	X	X
Define Stacked Volume ranges.	X	X
Define (insert) logical volumes [1.]	X	-
Define Fast-Ready categories.	X	X
Define VTS management policies. Inhibit reclaim schedule. Reclaim threshold percentage. Free storage threshold.	X	X
Make sure the Send Interventions to Host Console button is activated on both Library Manager operator intervention screens.	X	X

Notes:
1. The insert logical volumes operation is allowed only on the UI distributed library.
2. While the settings for each construct can be different, this is not advisable.

The procedure for defining the VTS from the Library Manager includes defining initial volser ranges and management policies for the VTS operation and the 3590/3592 cartridge insertion for VTS stacked volume preparation. The following tasks have to be performed at the Library Manager:

- Define stacked volume ranges. See 4.2.2, "Define stacked volume ranges" on page 153.
- Define (insert) logical volumes. See 4.2.3, "Define (insert) logical volumes" on page 154.
- Define Fast Ready categories. See 4.2.4, "Define Fast Ready categories" on page 156.
- Define VTS management policies. See "Reclamation and reconciliation" on page 161.
- Define Inhibit reclaim schedule. See "Inhibit Reclaim Schedule" on page 163.
- Define Reclaim threshold percentage. See "Reclaim Threshold Percentage" on page 165.
- Define Free storage threshold.See "Free Storage Threshold percentage" on page 166

- Define cleaning schedule. See 4.2.7, “Cleaning schedule pop-up window” on page 166.
- Define SNMP traps if required.

These tasks are most likely to be performed once the system IPL has been completed and/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 PtP VTS, the Library manager definitions and the System z software definitions are intertwined.

4.2.1 Navigating the Library Manager

Before you can insert physical 3590 and/or 3592 cartridges for VTS use into a PtP VTS library, you must define them through the Library Manager (LM) console. You must also define a beginning set of volser ranges for your Virtual Volumes.

Once your PtP VTS LM is at the Licensed Internal Code (LIC) level 527 or above, the LM panels will change. In order to maintain a level of currency in this book, all the panels listed will be post LIC level 527. As you read on, you may notice that significant changes have occurred. Refer to the *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring*, SG24-2229 redbook, Implementation section, for a discussion of the pre-527 versus 527 panel changes.

In addition, please note that we will refer to 3590 media type when discussing the panels. However, the PtP supports both 3590 and 3592 media for its physical volumes.

Use the **Commands - System management** pull-down menu (see Figure 4-12) to work with the tape library command invocation to define the VTSs in a tape library. When you define a VTS in a tape library soon after the physical installation, you must use the **System management** pop-up window for VTS definition work.

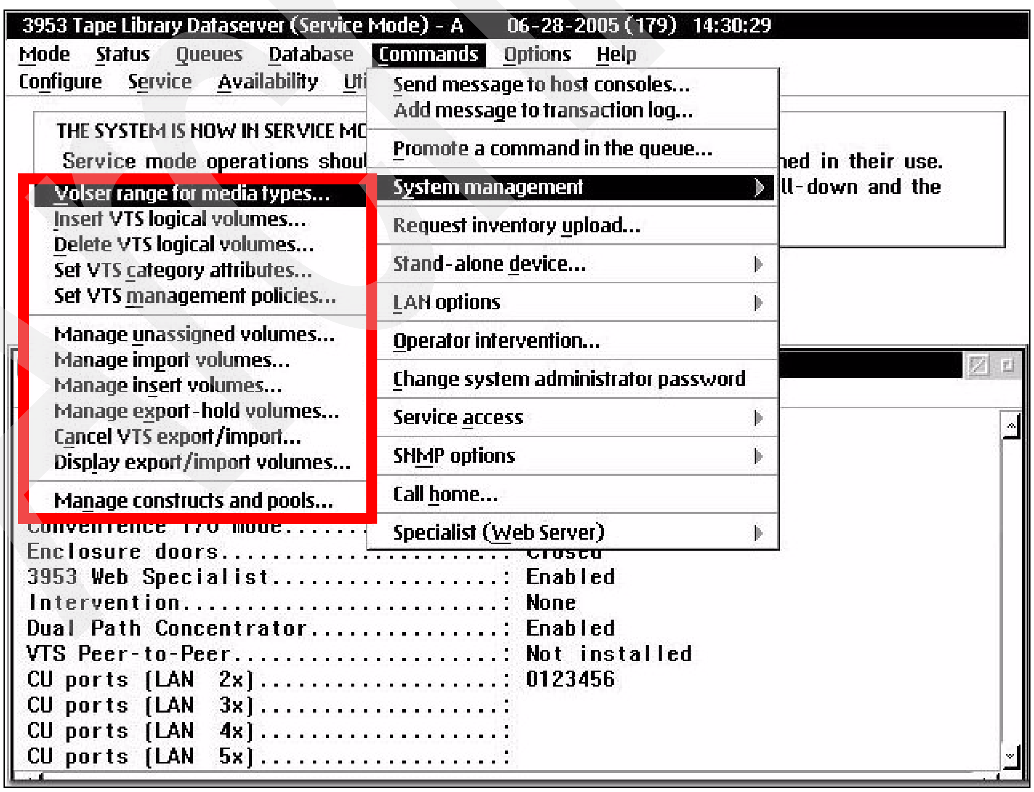

Figure 4-12 Commands pull-down menu

Click **System management**. The System management pop-up window (Figure 4-12) has the following selections:

- **Volser range for media types**: Use this panel to define the volser range of VTS 3590 cartridges to use for the VTS. These 3590 cartridges are used for stacked volumes in the VTS.
  - On this panel you can enter up to 50 or 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 to insert logical volumes into a VTS library.
- **Delete VTS logical volumes:** Use this 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).
- **Set VTS category attributes:** Use this panel to assign the Fast Ready attribute to categories.
- **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:** This allows you to reevaluate the physical volumes in the Insert category for 3590 native use. By redefining the volser ranges, you can move the volumes to the Insert categories for the VTSs. You can also eject the volumes from the 3494.
- **Manage constructs and pools:** This provides 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.

Note that all functions related to Import/Export are not available for the PtP VTS.

A detailed review of the System management pop-up window begins with 4.2.2, “Define stacked volume ranges” on page 153.

4.2.2 Define stacked volume ranges

You have to tell the Library Manager which 3590 or 3592 cartridges are to be used as VTS stacked volumes before you actually insert these cartridges in the 3494.

The volser range definition is required for stacked volumes and must be a unique range of volsers within the 3494. We highly recommended that the range be unique to all hosts within the customer environment.

With Advanced Policy Management (APM) you can also assign your stacked volumes (and / or ranges) to specific volume pools. Details on how this is done are provided within Figure 4-13.

Although you can define new volser ranges to add new stacked volumes at a later time, there is no penalty in defining a range larger than there are available cartridges for.

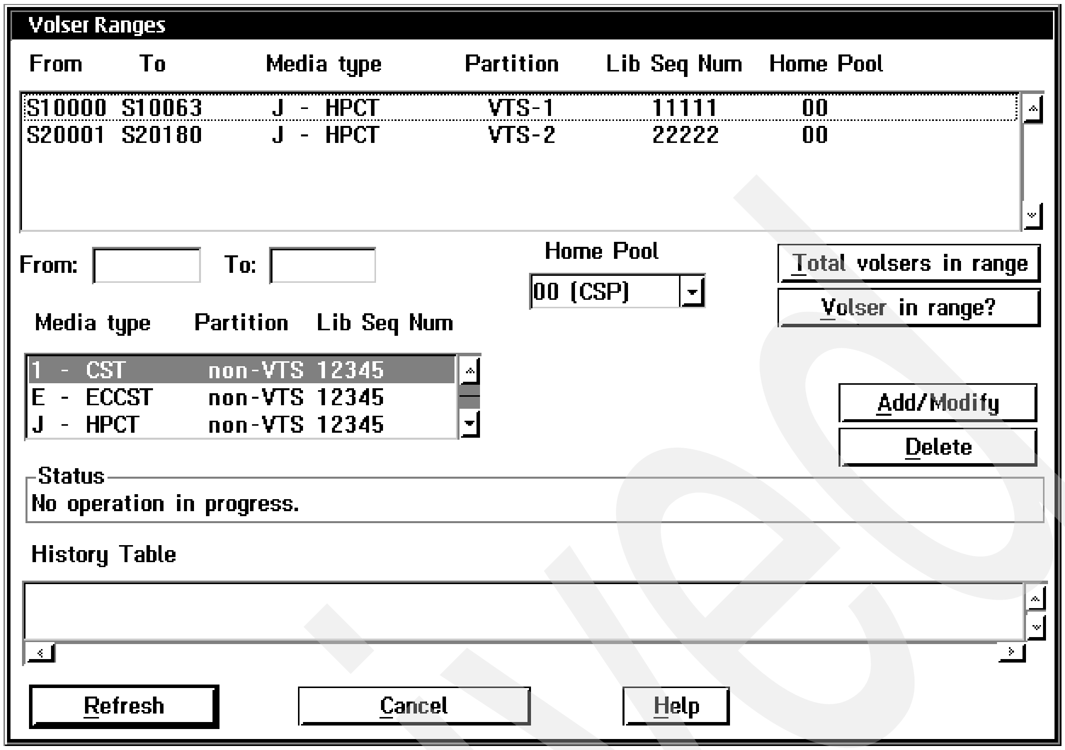

Figure 4-13 Define Volser Ranges

4.2.3 Define (insert) logical volumes

You insert logical volumes in the VTS through the Library Manager console by defining a range of volsers for the logical/virtual volumes and their media type. You can insert up to 500,000 logical volumes per VTS, but this same 500,000 limit exists for the total PtP VTS.

The insert process

This process is managed by Object Access Method (OAM). OAM provides the management facilities for the physical movement and tracking of the tape volumes used within the tape library. OAM provides many services for the support of tape libraries and in particular, cartridge entry, eject, audit, mount, and demount processing.

When you define one or more logical volumes to the VTS via the Library Manager, the following processes take place.

The Library Manager notifies all hosts that there are logical volumes to process in the insert category. For each logical volume, OAM invokes the cartridge entry installation exit (CBRUXENT) to:

- Approve or disapprove entry of a cartridge into a library.
- Determine TCDB volume record contents.
- Ensure volume serial uniqueness for each entered volume.

The logical volume serial has to be unique within all host systems that are connected to the same TCDB.

Note: CBRUXENT is supplied by DFSMSrmm. If your installation is not using DFSMSrmm, then your tape management vendor will supply the equivalent exit. This exit may require modification to specify the volsers acceptable in the host system.

If you have systems connected to the VTS which do not share a GRS ring, all hosts get notified and start to process the same list of entered volumes. You cannot influence which host will process the insert processing. The first host to move the volume from scratch to private status will get the volume added to its TCDB. If the volumes in the library are to be partitioned across multiple hosts, the RC=12 (ignore) return code from the cartridge entry installation exit must be used to make sure that the appropriate host and TCDB process the volume for entry. When this option is used, the following messages are issued:

```
CBR3602I Enter request rejected by the cartridge entry installation exit
CBR3621I Enter request ignored by the cartridge entry installation exit CBRUXENT
```

When systems share a GRS ring, the same processing occurs. However, each host within the GRS ring will wait until its turn to process the list of entered volumes. For those volumes that had already been processed and moved out of the insert category, the remaining hosts will have nothing to do. Therefore, for those volumes that had been previously ignored (RC=12 from the entry exit), the remaining hosts will have the opportunity to process these volumes and add to their respective TCDB.

Note: If all LPARs reject the volume, then the volume will remain in the insert category and OAM will retry the insert process with all attached LPARs every time it is reloaded, the CBRUXENT exit is refreshed or additional volumes are inserted into the same VTS.

Though entry processing will work without the usage of a GRS ring, it is recommended that a GRS ring be used for connecting the systems. This then enables OAM to send a systems level ENQUEUE around the ring, limiting the processing of the insert category to one host at a time.

Note: Within a SYSPLEX where all systems are part of a GRS ring or in GRS star mode, you may see that the insert process is very fast. This is mostly a JES2 complex. A JES3 complex is often a plex beyond the SYSPLEX boundary. The GRS ring does not communicate with the other hosts. The customer may use other ISV products for serialization. That is why you may see that entry processing could take much longer in JES3. We recommend that you disable the CBRUXENT for the other JES3 hosts when you define new logical volumes and enable it again when finished.

If you are defining large ranges of logical volumes at a time and are partitioning your volumes across multiple hosts with the usage of RC=12 (Ignore) from the entry exit, the messages mentioned previously may appear multiple times until an owning host processes the volumes for entry.

Note: When you are inserting many logical volumes into the VTS via the Library Manager console, we advise you to add volumes only to one LPAR at a time to prevent the operator console from being flooded with the ignore-type messages.

The following options can assist you in this effort:

- Suppress the messages via an automation package.

- Only have OAM up on the LPAR that you are adding logical volumes during the insert process. This should only be done if you are sure that no tape processing will take place when you are inserting logical volumes.
- Issue the command `Library Disable,CBRUXENT` on the attached LPARs where you do not want to process the logical volumes. This will turn off the entry process for the LPAR in question. Once you have finished your logical volume insert processing, then issue the command `Library Reset,CBRUXENT` for each of the LPARs. This will reset the Exit and make it enabled.

Attention: If you are partitioning a VTS and do not enforce a return code 12 from CBRUXENT on all systems which are not the primary owners of specific volumes, then these volumes may have their category updated by a subsequent process.

Insert tasks

Be aware that you can only delete or change the media type of logical volumes from a VTS as long as they are in an insert or scratch (Fast Ready) category. That means if a volume was in use once, it can only be deleted if it is expired and returned to the scratch category. You should therefore carefully plan what volsers and media types to insert.

The logical volumes inserted here should already have been defined to the local tape management system.

At this point, insert your initial set of logical volumes in the VTS. We recommend that you insert only the number of logical volumes that your application needs in the near future. You can insert new logical volumes whenever more scratch volumes are needed.

Important: Only define 10,000 volumes at a time and then the next 10,0000, when the first 10,000 volumes finish with the insert processing.

Should you have the additional feature Advanced Policy Management (APM) and are operating the VTS in a MVS environment, we also recommend that you insert your logical volumes ***without*** any assigned constructs. Any construct name will be assigned by your defined DFSMS environment. With this feature in an open systems environment, this is the point where you could assign the logical volumes to a specific stacked media pool via a defined Storage Group, thereby invoking volume affinity outboard of the host.

After the 3494 completes all initialization operations (including the teach operation by the IBM service representative and the inventory operation) and enters the online state for the first time, the host software requests an upload of the volume inventory. The information from the Library Manager database is uploaded to the attached hosts before host applications can access the tape cartridges inside the tape library.

Uploading this volume inventory information requires no operator action and, depending on the customer environment, may require some planning to avoid uploading to the wrong host or swamping the host consoles with “rejected” messages as has been described before.

4.2.4 Define Fast Ready categories

To take advantage of the scratch mount performance in the VTS 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 VTS 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 console. Figure 4-14 shows the Define Fast Ready Categories window.

The actual category hexadecimal number depends on the software environment and on the definitions done in the SYS1.PARMLIB member DEVSUPnn for library partitioning. Also, the DEVSUPnn member must be referenced in IEASYSnn to be activated.

Note: Any category up to and including to X'FEFF may be overridden in DEVSUPnn, but be aware that categories may only be overridden if the VTS will not be accessed by the owning operating systems described in Appendix B, "Library Manager volume categories" on page 397.

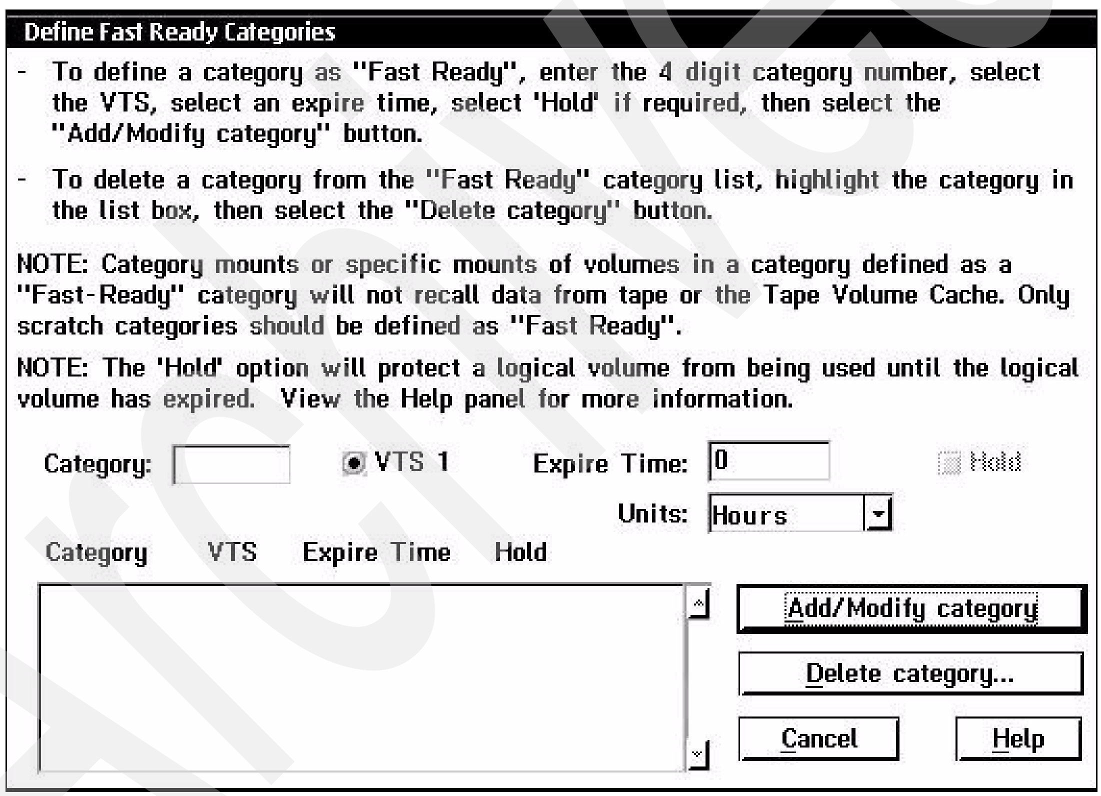

Figure 4-14 Set VTS category attributes pop-up window

In this panel you will add the category to the list of categories in the tape library that have the Fast Ready attribute set. To reset a category's Fast Ready attribute, on the Library Manager panel you highlight one of the categories in the list with the Fast Ready attribute set and then request that it be reset with the **Add/Modify category** button.

To define a category as Fast Ready, on the Define Fast Ready Categories window, enter the 4-digit hexadecimal category number, select the desired VTS, and select the **Add/Modify category** button. You should specify the scratch mount categories of MEDIA1 and MEDIA2 as VTS Fast Ready when you use DFSMS/MVS. These scratch mount categories are 0001 and 0002, respectively if you use the DFSMS/MVS default values. They are ‘0031’ and ‘0032’ in Figure 4-14 for VTS 1.

In addition to this, the expired volume management assignment can be set to a category with the Fast Ready attribute. This will ensure that volumes returned to scratch will have an extra “grace period” before they are removed from the VTS inventory. Refer to the redbook, *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring*, SG24-2229, for additional information on expired volume management.

> **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 397 for the scratch mount category for each software platform. In addition to the DFSMS/MVS default value for the scratch mount category, you can define your own scratch category to the 3494 Library. In this case, you should also add your own scratch mount category to the VTS Fast Ready category list.

To delete a category from the Fast Ready list, highlight the category in the list box and select the **Delete category** button. You are prompted to confirm the delete operation. Select **Yes** to continue the deletion operation. Select **No** to cancel the delete operation. You can click the **Help** button to display the Set VTS Category Attributes help panel. Once all updates have been completed, select **Cancel** to exit this option. Such an action would only serve to elongate the mount wait time for scratch volumes within the particular category. This should only be considered as a “clean-up” activity once a category is to be removed from service.

4.2.5 Physical cartridge insertion

At this point, insert your cartridges in the tape library. As this is the first time you are inserting stacked volumes for your VTS, you may have a large number of cartridges. The easiest way to insert larger numbers of cartridges is to open the frame doors and insert them in empty cells in the library. Once you have inserted all of the cartridges and closed the frame doors, you have to perform inventory update, if it is not done automatically. See 6.4.7, “Adding stacked volumes” on page 240 for information about inserting cartridges in the library.

If you are attaching the VTS to an IBM 3494 Tape Library, once you have inserted all of the cartridges and closed the frame doors, you have to perform inventory update, if it is not done automatically. See 6.4.7, “Adding stacked volumes” on page 240 for information about inserting cartridges in the library.

If you are attaching the VTS to an IBM TS3500 Library Manager partition, you must first make sure that the cartridges are assigned to the proper TS3500 logical library, as the cartridge insert process requires an additional step. Figure 4-15 sketches the different phases of insert processing of physical cartridges in an IBM TS3500 Tape Library.

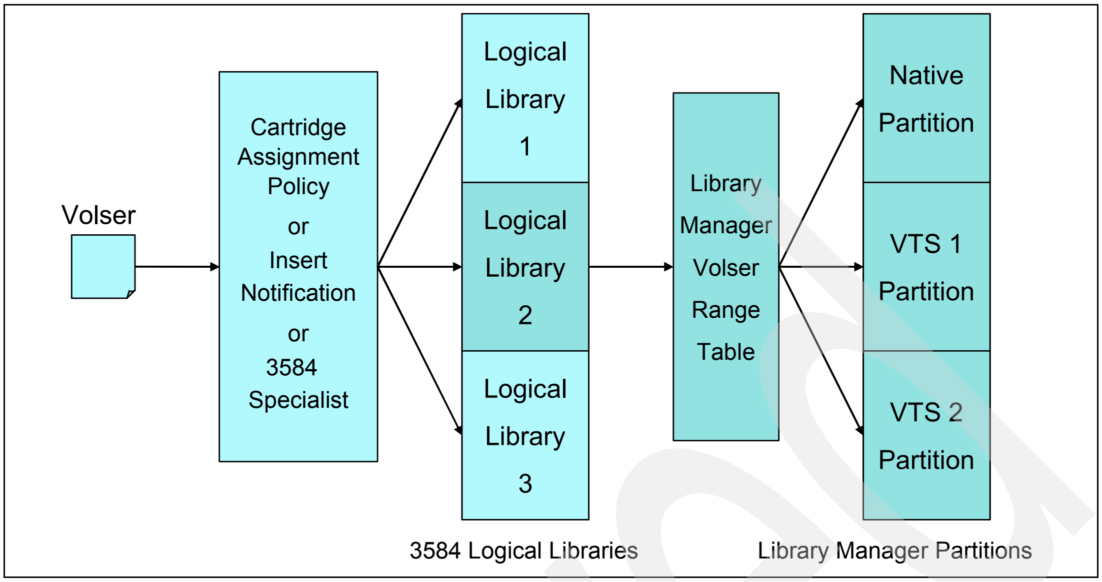

Figure 4-15 Insert Processing for TS3500/3953 Physical Volumes

There are two ways to insert physical cartridges:

- Insert volumes directly into storage cells by opening library doors:
  - The TS3500 Cartridge Assignment Policy defines which volumes are assigned to which logical library. See Figure 4-16.
  - Library Manager is notified after doors are closed and the TS3500 performs inventory. Library Manager updates its database accordingly.
  - The TS3500 Web Specialist can be used for unassigned volumes.
- Insert volumes via Convenience I/O Station (Insert notification is required to be enabled on the TS3500):
  - The TS3500 Cartridge Assignment Policy (CAP) defines which volumes are assigned to which logical library.
  - Operator identifies 3953 logical library as destination using Insert Notification. This covers volumes that are not defined by the CAP.
  - Library Manager sees that I/O Station has volumes assigned to it.
  - Library Manager moves the volumes to empty cells.
  - The TS3500 Specialist can be used to assign unassigned volumes.

Figure 4-16 TS3500 Cartridge Assignment Policy

For more information on the setup of the IBM TS3500 Tape Library, please refer to the *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning, Implementation, and Monitoring,* SG24-6789.

Initial cartridge insertion: summary

Physical cartridges can be loaded into the tape library after the hardware installation is complete. The loading of cartridges can take place before the IBM service representative teach operation and the inventory operation.

> **Note:** You must define the physical volser range to the VTS before initial physical cartridge insertion. See 4.2.2, "Define stacked volume ranges" on page 153 for proper procedures on performing this function.

The initial cartridge insertion phase of VTS should include the following actions:

- Define stacked volume range.
- Insert stacked volumes to be managed by VTS.
- Insert (define) logical volumes.
- Insert volumes to be used for native drives, if any.
- Insert cleaner volumes.

The following types of cartridges can be loaded:

Stacked volumes: If a VTS is installed, the 3590 HPCT or 3590 EHPCT that it will use to store and manage logical volumes cannot be loaded into the tape library without first:

- Having the IBM service representative perform an initial teaching of the tape library.
- Setting up one or more volser ranges that identify the physical volumes that the VTS will manage. See 4.2.2, "Define stacked volume ranges" on page 153 for a description of how to enter the volser ranges.

Logical volumes: If a VTS is installed, logical volumes are inserted into the tape library through the Insert Logical Volumes panel.

Customer volumes: Customer volumes are the initial set of data and scratch volumes to be automated. Any number of cartridges can be added to the tape library up to its maximum available cartridge storage cells.

Note: These customer volumes are not VTS-owned volumes.

Cleaner volumes: One or two cleaner volumes should be installed in the tape library for each type of tape subsystem (3490E/3590) that is resident in the tape library. Two cleaner cartridges are considered a minimum. The cleaner cartridges can be placed in any available cell.

Note: If the cleaner cartridge is ejected, the mount count of a tape cartridge is reset to zero.

The external volser must match the mask value that is provided; otherwise the inventory operation will treat the cleaner cartridge as a normal customer volume. A cleaner cartridge should have a unique volser.

Service volume: The IBM service representative installs one or two service volumes in the tape library, depending on its configuration.

4.2.6 Define VTS management policies

The **Set VTS** Management Policies window is where you can define the inhibit reclaim schedule, the reclaim threshold percentage, and the free storage threshold policies.

Reclamation and reconciliation

To minimize the effect of VTS 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.

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 VTS.

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 which does not use the TVC.

All volumes with active data less than the Reclaim Threshold Percentage are eligible to go through the space reclamation process. The default Reclaim Threshold Percentage is 10%; however, you may change this percentage at the Library Manager console.

Reclamation enablement

To minimize any impact on VTS activity, the storage management software monitors resource utilization in the VTS and schedules the reclamation when the Reclaim Threshold Percentage is reached. You can optionally prevent reclamation activity at specific times of day by specifying an Inhibit Reclaim Schedule on the Library Manager Console. However, the IBM VTS determines whether reclamation is to be enabled or disabled once an hour.

A panel on the Library Manager allows you to monitor the amount of active data on stacked volumes and help you plan for a reasonable and effective Reclamation Threshold Percentage. The navigation to this panel is: **Status** —>**VTS2** —>**Active Data Distribution**; an example of which is provided in Figure 7-10 on page 257.

Even though reclamation is enabled, there may not always be stacked volumes going through the process all the time. Other conditions must be met, such as stacked volumes that are below the threshold and drives available to mount the stacked volumes.

Reclaim is stopped by the internal management functions if a tape drive is needed for a recall or copy (as 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. If this happens reclaim is stopped after the current logical volume move is complete.

The information in Table 4-6 shows the reclaim settings that are available. The default value is 10%.

Table 4-6 VTS space reclamation algorithm

Number of scratch cartridges	Reclaim is enabled if:	Reclaim threshold percentage
>50	Non-invasive of VTS activity and time is outside Inhibit Reclaim Schedule.	10% or customer setting
10 to 50	Two or more drives available and time is outside Inhibit Reclaim Schedule.	10% or customer setting
<10	Always	As required for at least 15 free cartridges

Reclamation is enabled or disabled according to a new set of rules. These rules are required to enable the VTS to compensate for the new pooling structure as well as changes to scratch thresholds. Reclamation can now occur on multiple volume pools at the same time as well as processing multiple tasks for the same pool. As before, Reclamation will select 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 VTS would select all the stacked volumes from 0% to 29%. The reclaim tasks would then process the volumes from least full (0%) to most full (29%) up to the defined reclaim threshold.

Individual pools can also have different reclaim thresholds set. The number of pools can also influence the reclamation process, as the LM will always evaluate the stacked media starting with pool one. Details of these settings can be found in 4.2.6, “Define VTS management policies” on page 161.

Reclamation is also affected by Pvol Scratch Counts. It assesses the Scratch State for Pools as follows:

1. A pool enters a Low Scratch State when it has *access to* less than 50 but more than 2 stacked volumes.
2. A pool enters a Panic Scratch State when it has *access to* less than 2 empty Pvols.

Access to includes any borrowing capability. Borrowing is described in “Borrowing and returning” on page 43. The Common Scratch pool (VTS stacked volume pool 0) is allowed to be overcommitted.

Pools in either Scratch State (low or panic state) get priority for Reclamation. Table 4-7 summarizes the new thresholds.

Table 4-7 Reclamation priority table

Priority	Condition	Reclaim schedule honored	Active data threshold % honored	Number of concurrent reclaims
1	Priority move that benefits a pool in Panic Scratch State	No	No	1, regardless of idle drives
2	Pool in Panic Scratch State	No	No	1, regardless of idle drives
3	Priority move	Operator choice	No	1, regardless of idle drives
4	Pool in Low Scratch State	Yes	Yes	1, regardless of idle drives
5	Normal reclaim	Yes	Yes, pick from all eligible pools	(# idle drives / 2) minus 1 6 drv: 2 max 12 drv: 5 max

Note: A 3590 drive is said to be IDLE when there has been no activity for the previous ten minutes

Inhibit Reclaim Schedule

The Inhibit Reclaim Schedule defines when the VTS 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 may 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.

Figure 4-17 shows the Set VTS Management Policies window.

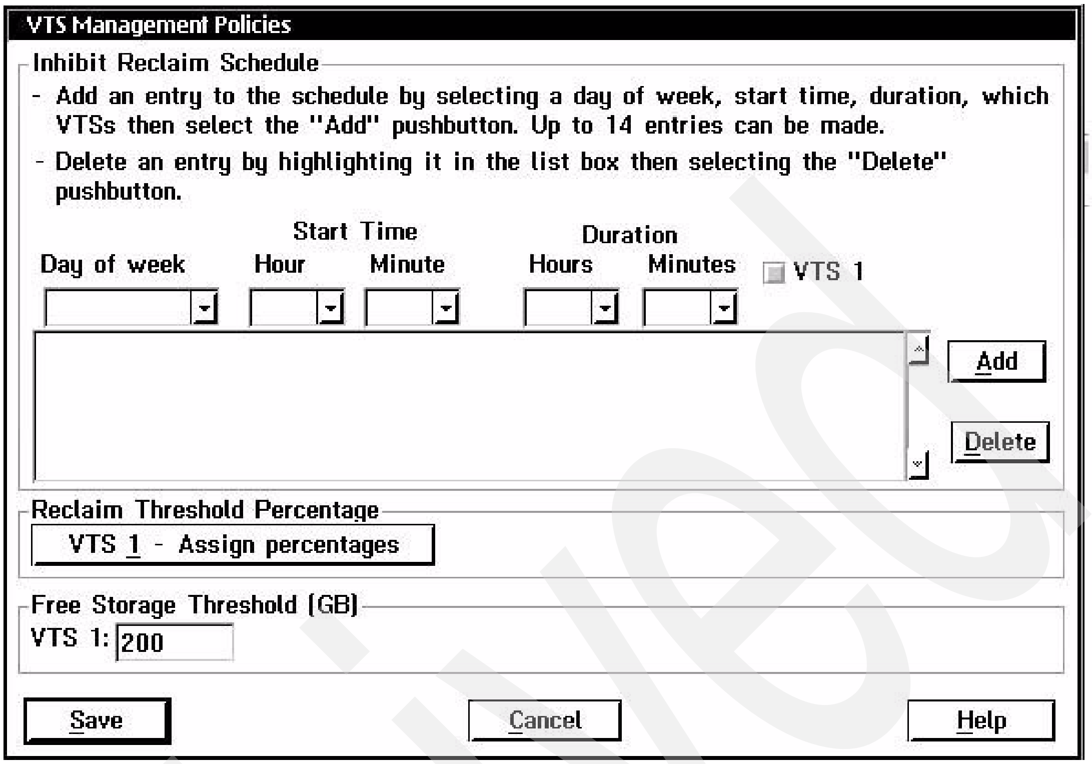

Figure 4-17 Set VTS Management Policies pop-up window

The following push buttons are shown in Figure 4-17:

Add Adds an entry to the inhibit reclaim schedule.

Delete Deletes an entry from the inhibit reclaim schedule.

VTS x - Assign percentages If the VTS is Advanced Policy Management enabled, then use this push button to define the reclamation thresholds for individual pools. Clicking this push button will open the window as shown in Figure 4-18. This is only available at the 527 or higher LIC level with APM enabled.

Save Closes the panel and saves all the changes made to the inhibit reclaim schedule and the free-storage threshold.

Cancel Closes the panel without saving any of the changes.

Help Displays the Set VTS Management Policies help panel.

Reclaim Threshold Percentage

The Reclaim Threshold Percentage (Figure 4-18) is used to identify when a Magstar 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 (FF04 post 527 LIC). 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 VTS to reclaim a stacked volume, because more data must be copied from one Magstar cartridge to the other Magstar cartridge.

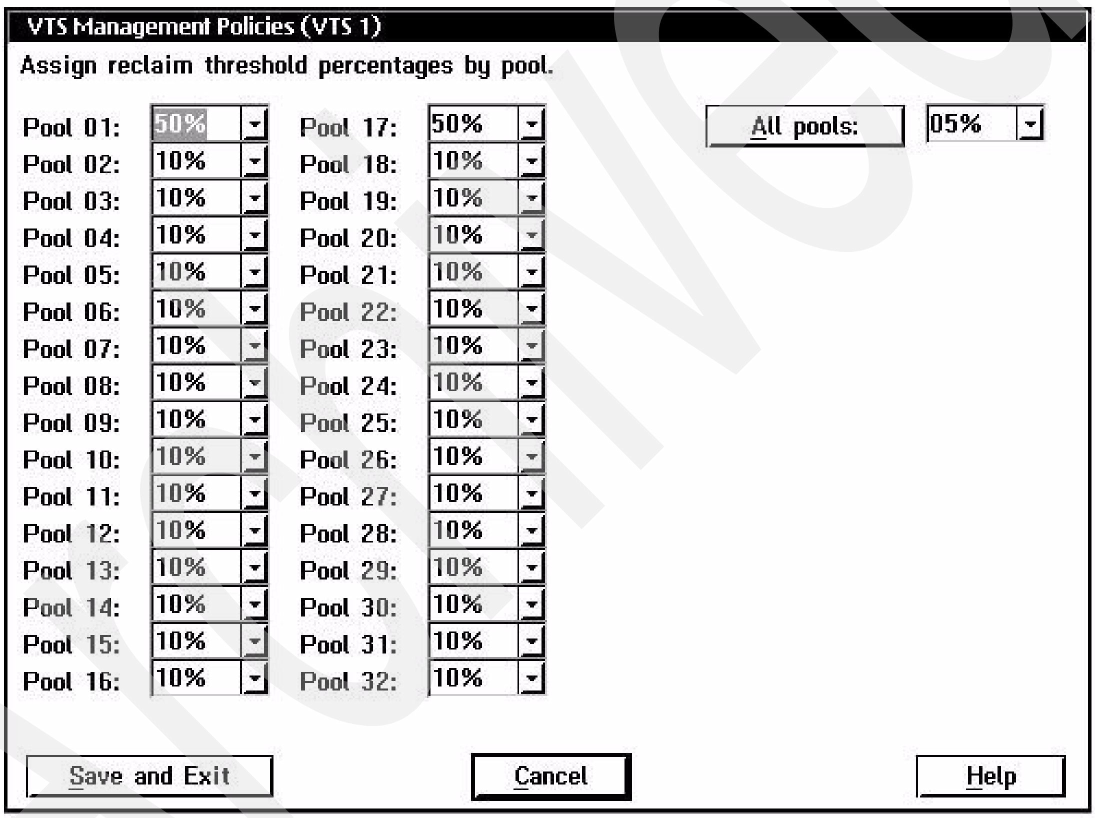

Figure 4-18 VTS reclaim thresholds by pool window with APM

Be aware that a multiple of two drives are involved in the reclamation process and because of this resource usage, you should not specify too high percentages. The *Active Data Distribution* bar graph will assist you in setting this number. See “Active data distribution” on page 256 for information about displaying the panel.

Free Storage Threshold percentage

The *Free Storage Threshold (GB)* provides a warning when the VTS is running low on free storage, the capacity of all the empty stacked volumes in the VTS. A threshold is provided for each VTS 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 3494 or 3953 Operator's Guide section under "*VTS Management Policies*" for additional information and Free Storage Threshold tables for 3590/3592 tables.

Note: You can see the effect reconciliation or reclamation has when they run by looking at the VTSSTATS report. See Figure 7-21 on page 268 for an example report. The field labeled SMF94VLA in the SMF record type 94 indicates the number of logical volumes managed by VTS. This actually includes all the logical volumes that have ever been used by the VTS, plus every instance of every volume that has been used by the VTS since the VTS has last performed a reconciliation. Therefore, when all the volumes have been used at least once, SMF94VLA value should be at least as large as the total number of logical volumes defined. The difference between SMF94VLA and the total number of defined logical volumes is the number of instances that logical volumes have been used since the last reconciliation. Therefore, the number would show an increase as logical volumes are being used until the next reconciliation.

4.2.7 Cleaning schedule pop-up window

As part of the VTS installation process, you should ensure that the Library Manager cleaning schedule for the VTS-managed IBM tape drives is set correctly. For the 3590 and the 3592 tape drives, no definition is required. The drives will request cleaning if necessary.

4.3 Defining Storage Groups and Pools

Storage Groups and Pools are only supported with APM. If APM is not installed on your VTS, you cannot use multiple physical volume pools: All scratch stacked volumes are in Pool 0 and all private stacked volumes are in Pool 1.

In order to determine whether your VTS is APM capable, bring up the VTS Status panel from the LM and check whether Advanced Policy Management Capable is set to **YES** (Figure 4-19). This figure shows an example of the Library Manager Status screen.

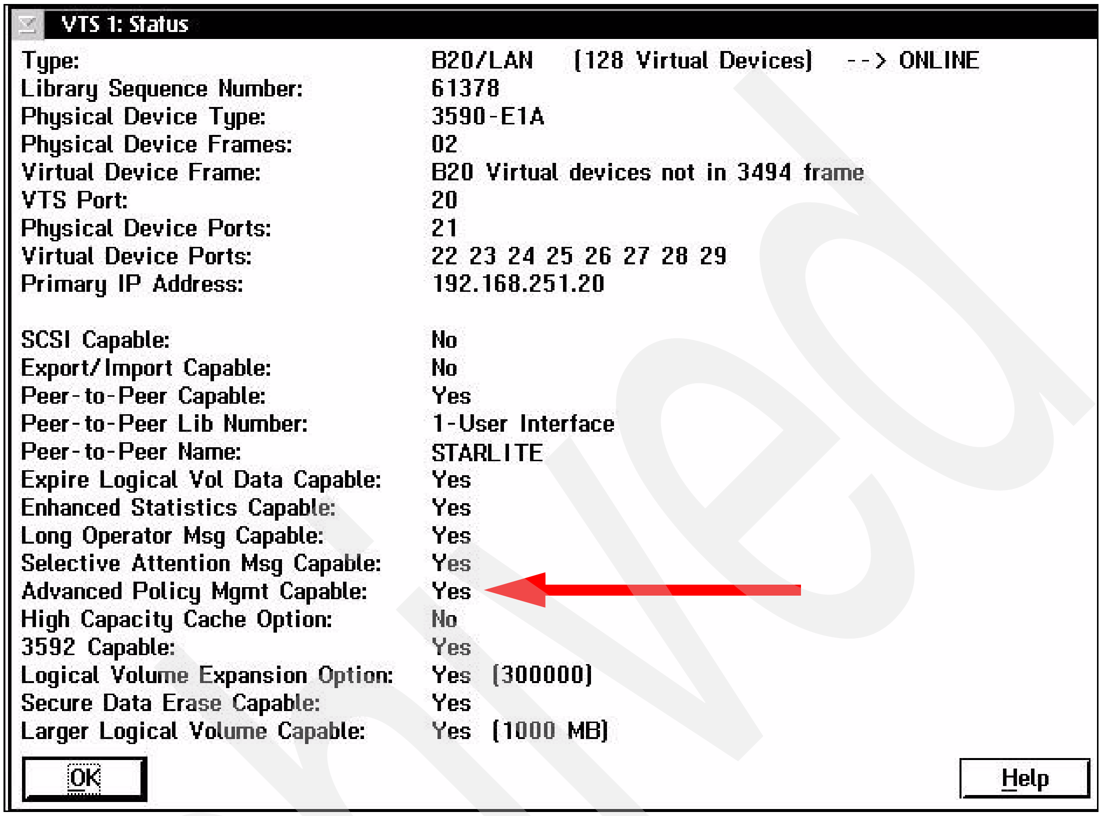

Figure 4-19 VTS status window

4.3.1 Library Manager definition panels

Definitions for the LM constructs are done through the *Commands Pull-Down Menu* (Figure 4-12 on page 152). Select **Manage Constructs and Pools** (Figure 4-20).

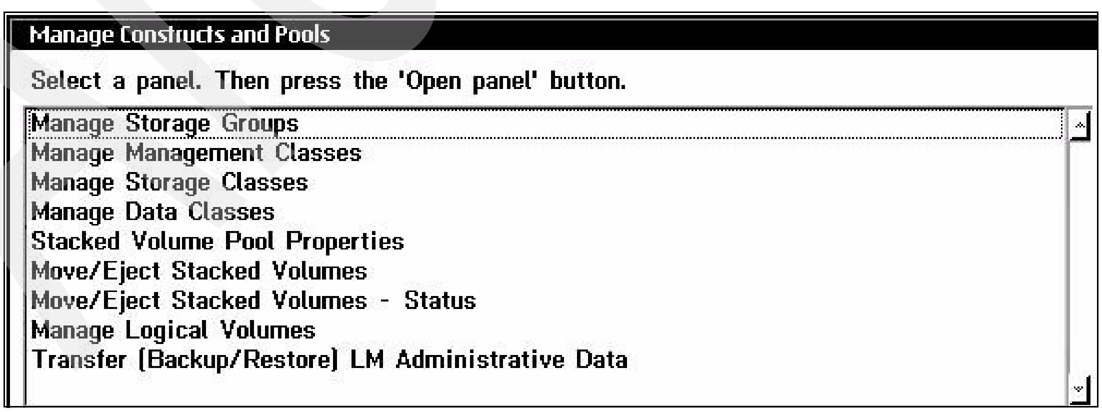

Figure 4-20 Manage Constructs and Pools window

The Manage Constructs and Pools window provides access to multiple panels that allow you to manage storage management constructs, stacked volume pool properties, move/eject stacked volumes and manage logical volumes.

To define Storage Groups (Figure 4-21), perform the following operations:

1. Select the **Commands - System Management** pull-down menu.
2. Click **Manage Constructs and Pools** (Figure 4-20 on page 167).
3. Click **Manage Storage Groups**.

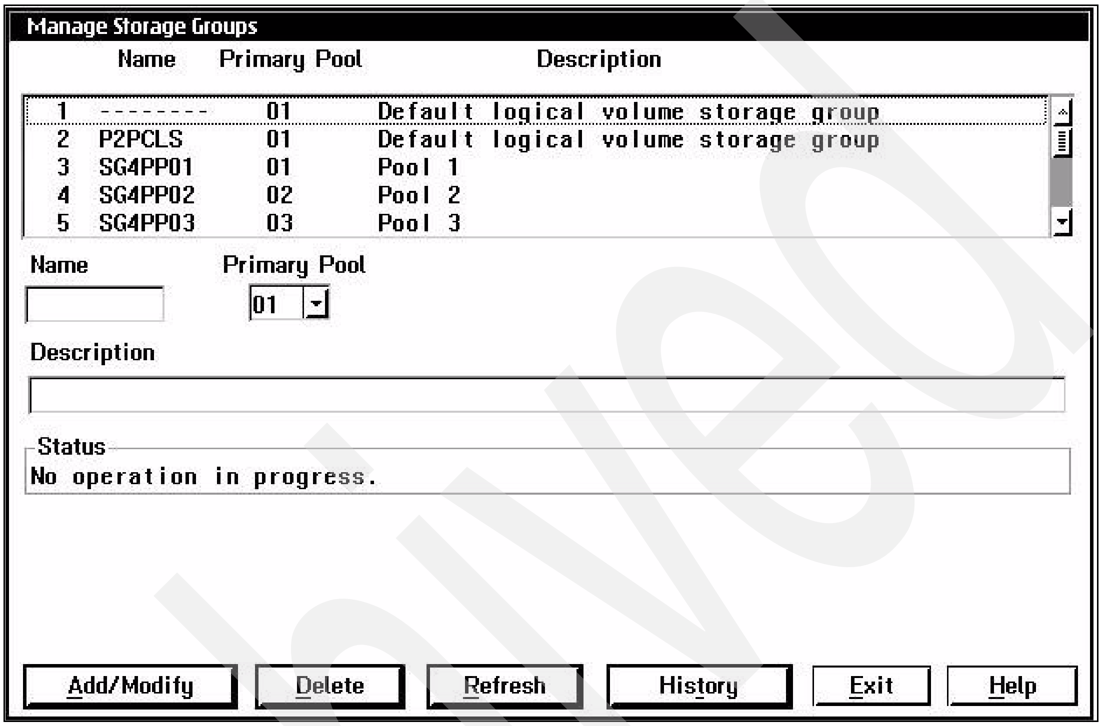

Figure 4-21 Manage Storage Groups window

Through the Manage Storage Groups Window (Figure 4-21), you will be able to:

- Add a Storage Group.
- Modify an existing Storage Group.
- Delete a Storage Group.

Creating a Storage Group

To create a Storage Group, proceed as follows:

1. Enter a one- to eight-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. Select a **Primary Pool** and enter a short description.
2. Select the **Add/Modify** button.

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 and/or description. Select the **Add/Modify** button.

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. Select the **Delete** button.

Add/Modify: Adds the entered Storage Group or modifies the selected Storage Group.

Delete: Deletes a Storage Group.

> **Note:** 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 VTS 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 VTS will query the logical volumes primary pool and the LM will return the primary pool associated with the default Storage Group.

Refresh: Refreshes the Manage Storage Groups window.

Exit: Closes the Manage Storage Groups window.

Help: Provides help about the Manage Storage Groups window.

The purpose of the *history table* is to:

- Indicate actions that are currently in progress or have already completed.
- Coordinate remote users (3494 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 a similar action.

> **Note:** The above example defines primary pool one to Storage Group TEST1. The definitions for primary pool one are completed through the Stacked Volume Pool Properties menu (see Figure 4-22).
>
> - This panel becomes available after upgrade to LM 527 and VTS 2.26. However, it provides no function without FC4001 installed.
> - The default Storage Group, identified by eight dashes (--------), cannot be deleted.
> - Up to 256 Storage Groups, including the default, can be defined.

> **Attention:** The use of multiple pools may have an impact on overall throughput in your VTS. Prior to utilizing this function of Advanced Policy Management , ensure that your VTS is configured correctly. To help identify any impact, use the modelling tool, Batch Magic, in consultation with your IBM Storage Specialist. Batch Magic has been updated to accommodate the new functions.

4.3.2 Defining Stacked Volume Pool Properties

To define Stacked Volume Pool Properties (Figure 4-22), perform the following actions:

1. Select the **Commands - System Management** pull-down menu.
2. Click **Manage Constructs and Pools** (Figure 4-20 on page 167).
3. Click **Stacked Volume Pool Properties**.

Figure 4-22 shows the Library Manager Panel.

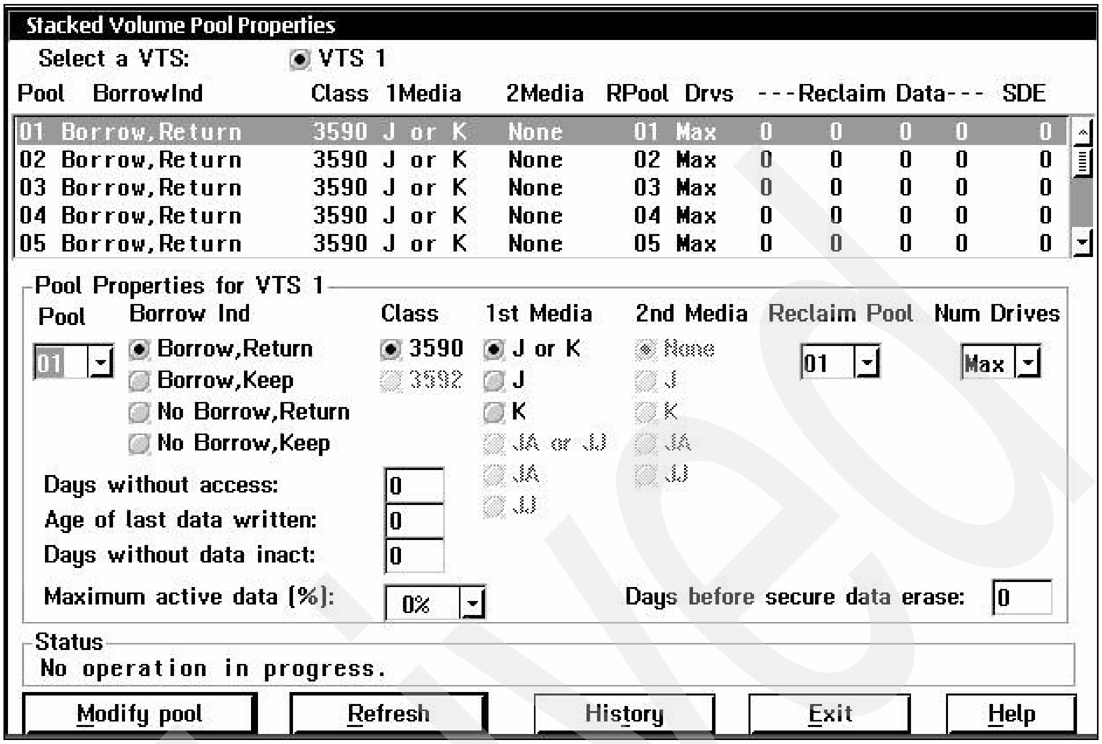

Figure 4-22 Stacked Volume Pool Properties window

The Stacked Volume Pool Properties window (Figure 4-22) allows you to modify pool properties. To modify pool properties, do the following steps:

1. In a multiple VTS environment, **select** the appropriate VTS.
2. A list of 32 General Purpose stacked volume pools for the selected VTS will be presented in the list box. Select one by using the mouse or keyboard to highlight the stacked volume pool you want to modify.
3. In the Pool Properties for VTS section, modify the fields as needed:

Pool
Contains the number of the pool from 1 through 32. The Common Scratch Pool (CSP), also referred to as Pool 00 is not contained in this list, because you cannot change its settings.

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.

Class
Shows the device type 3590 or 3592. In a B20 VTS, both 3590 and 3592 tape drives can be installed. However, only one device type is supported per pool and its supported media.

1st Media Shows the preferred media for the Media Class specified for this pool. For a Media Class of 3590 it can contain *J - HPCT*, *K - EHPCT*, or *J or K*. For a Media Class of 3592 it can contain *JA*, *JJ*, or *JA or JJ*.

2nd Contains the secondary media for this pool. If both medias types, for example *J or K* are specified for the First Media, the Second Media can only be set to *None*.

Reclaim Pool Lists the pool to which active logical volumes will be assigned when stacked volumes of this pool are reclaimed by the VTS. The stacked volume itself is treated according to the Borrow Indicator.

Num Drives Defines the maximum number of physical tape drive that this pool can use for pre-migration.

Days Before Secure Data Erase Can be 0 to 365 days and defines the number of days after which a cartridge is eligible for reclamation and, subsequently, must be physically erased. The count of the days begins, when the first byte of data on a cartridge has been invalidated after the cartridge was full.

A value of zero deactivates this policy. A non-zero setting defines that any stacked volume in this pool must be physically erased before it is returned to scratch. You can check whether your VTS is enabled for Secure Data Erase on the VTS Status panel shown in Figure 4-19 on page 167.

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 this policies.

4. Select the **Modify pool** button to modify the selected pool properties.

Refresh: Refreshes the Stacked Volume Pool Properties window.
Exit: Closes the Stacked Volume Pool Properties window.
Help: Provides help about the Stacked Volume Pool Properties window.

You can use one or 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 panel shown in Figure 4-22 is taken from the IBM 3494 Library Manager console. The panel of the IBM 3953 Library Manager looks similar and is shown in Figure 4-23.

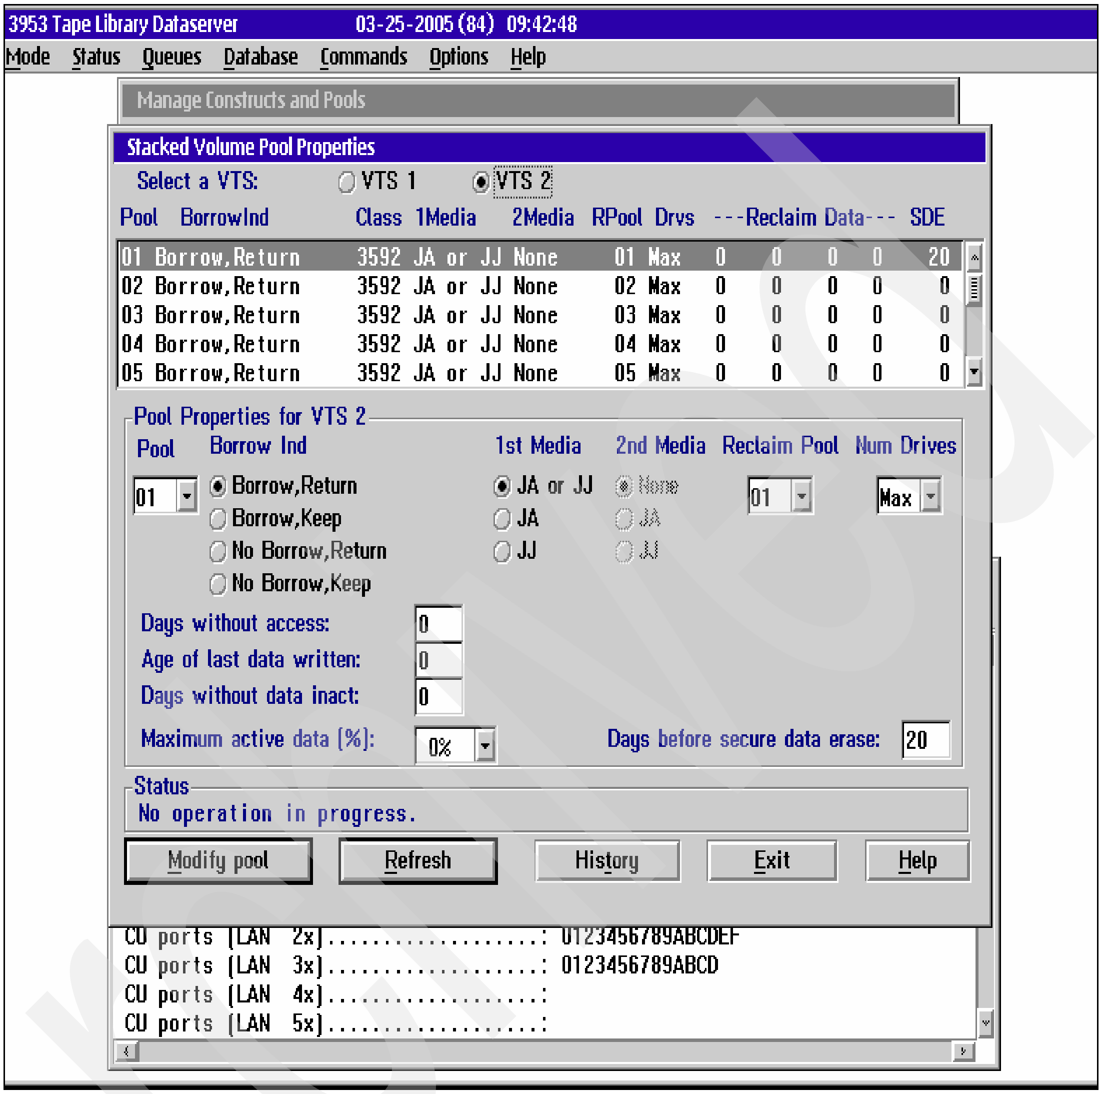

Figure 4-23 IBM 3953 Library Manager

Note that these definitions can all be done from the ETL Specialist panels as well. For more information on ETL Specialist panels, refer to the examples in the *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring,* SG24-2229.

4.3.3 Creating Management Classes

To set up logical volume dual copy in your environment (Figure 4-24), perform the following steps:

1. Select the **Commands - System Management** pull-down menu.
2. Click **Manage Constructs and Pools** (Figure 4-20 on page 167).
3. Click **Manage Management Classes**.

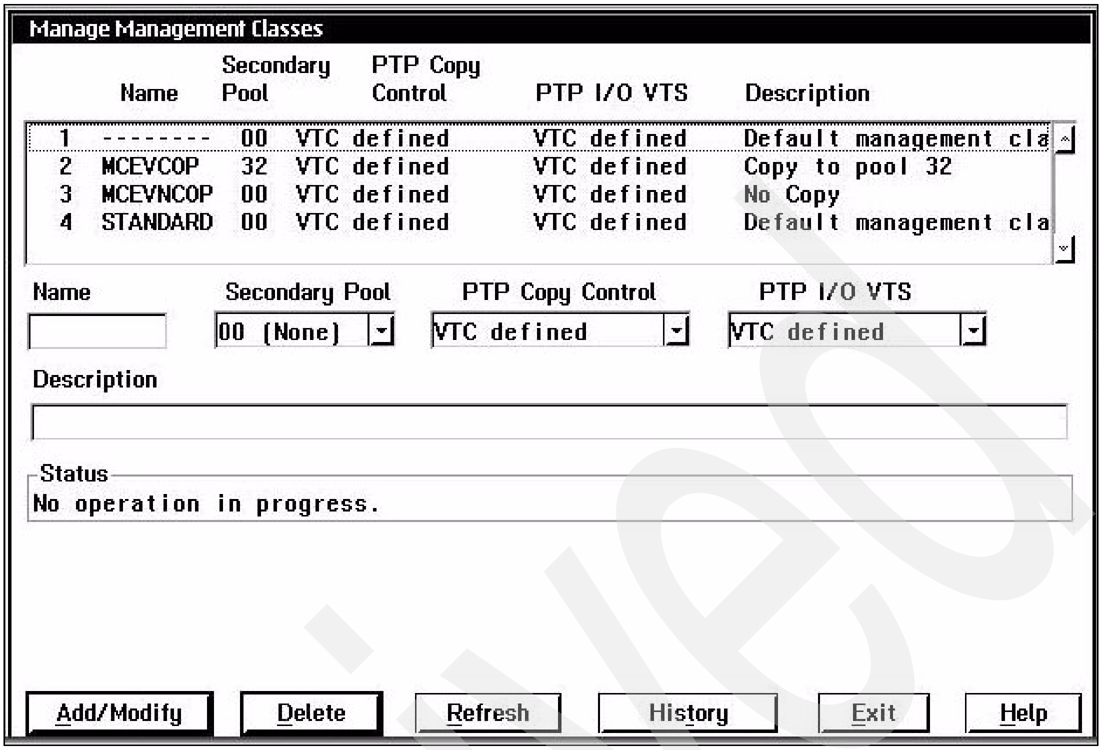

Figure 4-24 Manage Management Classes

The Manage Management Classes window (Figure 4-24) allows you view and manage Management Classes. Through this panel, you will be able to:

- Add a Management Class.
- Modify an existing Management Class.
- Delete a Management Class.

Adding a Management Class

To add a Management Class, proceed as follows:

1. Enter a one- to eight-character name in the **Name field**. This name must be unique within the Management Class construct names.
2. Determine dual copy option. 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. Specify the Peer-to-Peer Control:
   – Select **Immediate** to create the secondary copy at the same time as the first copy. This setting will over-ride the IBM service representative setting at the AX0 or VTC.
   – Select **Deferred** to create the secondary copy at a time after the first copy was made. This setting will over-ride the IBM service representative setting at the AX0 or VTC.
   – Select **No Copy** if you want the logical volume only to be written to a single VTS of the PtP VTS configuration.
   – Select **VTC Defined** to use the VTC setting as defined by the IBM service representative.

4. Specify the PtP I/O VTS:
   - Select **VTC Defined** to use the AX0/VTC setting as defined by the IBM service representative.
   - Select **Distributed Library 0** to use VTS0.
   - Select **Distributed Library 1** to use VTS1.

   Note that, if you specify that a specific VTS within a PTP VTS is to be used for the I/O VTS and the specified VTS is not available, any scratch mount request that specifies this Management Class construct will fail.
5. Enter a short description in the **Description** field.
6. Select the **Add/Modify** button.

To modify a Management Class, select 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, PtP copy options and/or description. Select the **Add/Modify** button.

To delete a Management Class, select from the current Management Classes presented in the list box. Use the mouse or keyboard to highlight the Data Class you want to delete. Select the **Delete** button.

Add/Modify: Adds the entered Management Class or modifies the selected Management Class.

Delete: Deletes the selected Management Class

Refresh: Refreshes the Manage Management Class window

Exit: Closes the Manage Management Class window

Help: Provides help about the Manage Management window

The purpose of the *history table* is to:

- Indicate actions that are currently in progress or have already completed.
- Coordinate remote users (3494 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.

> **Note:** This panel becomes available after upgrade to LM 527 and VTS 2.26. However, it provides no function without FC4001 installed.
>
> - The default Management Class, identified by eight dashes (--------), cannot be deleted.
> - Up to 256 Management Classes, including the default, can be defined.

> **Attention:** Depending on its use, dual copy may have an impact on overall throughput in your VTS. Prior to utilizing this function of Advanced Policy Management , ensure that your VTS is configured correctly. To help identify any impact, use the latest version of the modelling tool Batch Magic which has been updated to accommodate the new functions. Note that we do not recommend the use of Dual Copy in a PtP VTS environment.

4.3.4 Creating Storage Classes

To set up each Storage Class (Figure 4-25), perform the following steps:

1. Select the **Commands - System Management** pull-down menu.
2. Click **Manage Constructs and Pools** (Figure 4-20 on page 167).
3. Click **Manage Storage Classes**.

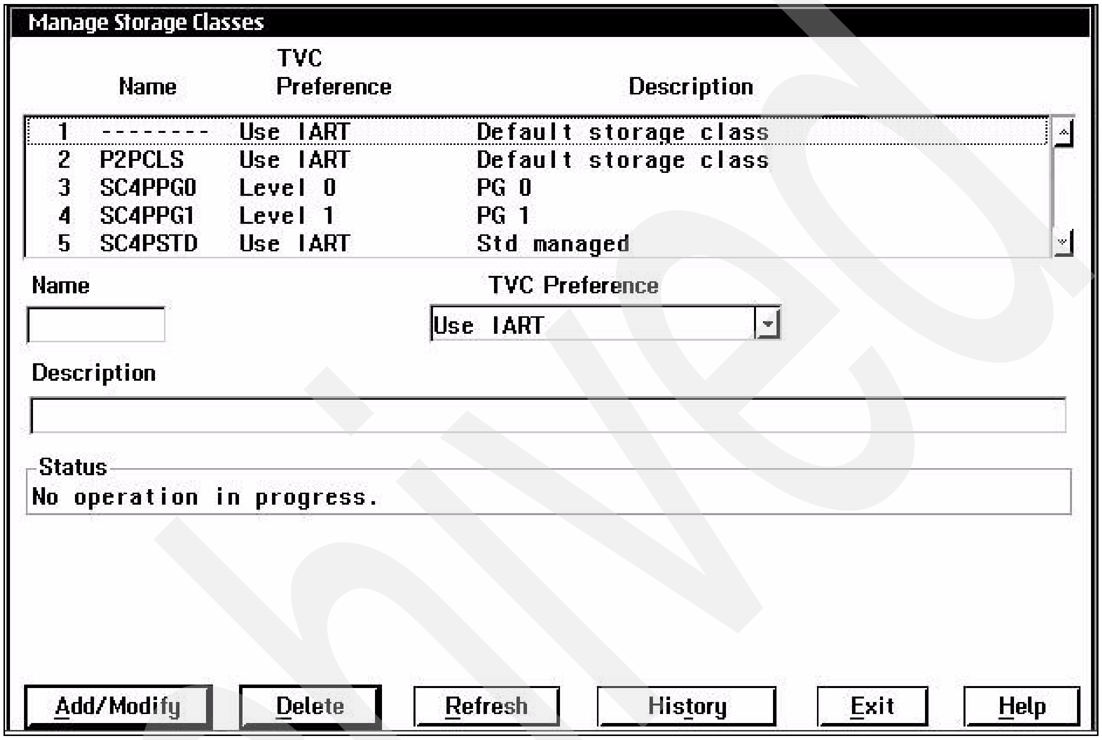

Figure 4-25 Manage Storage Classes window

The Manage Storage Classes window (Figure 4-25) allows you view and manage Storage Classes. Through this panel, you will be able to:

- Add a Storage Class.
- Modify an existing Storage Class.
- Delete a Storage Class.

Adding a Storage Class

To add a Storage Class, proceed as follows:

1. Enter a one- to eight-character name in the **Name field**. Define the same name as host defined DFSMS Storage Class construct. Name must be unique to within the Storage Class construct names.
2. Determine the **Tape Volume Cache Preference**. Three options are available:

IART	Instructs VTS to honor host supplied Initial Access Response Time (IART) value
Level 0	Instructs VTS to remove volumes from Tape Volume Cache (TVC) as soon as they are copied to tape. Remove largest first
Level 1	Instructs VTS to remove volumes from TVC after the copy has been made but only if space is needed in the TVC. Use Least Recently Used (LRU) algorithm.

3. Enter a short description of the Storage Class in the **Description** field.
4. Select the **Add/Modify** button.

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 and/or description. Select the **Add/Modify** button.

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 the **Delete** button.

Add/Modify: Adds the entered Storage Class or modifies the selected Storage Class

Delete: Deletes the selected Storage Class

Refresh: Refreshes the Manage Storage Class window

Exit: Closes the Manage Storage Class window

Help: Provides help about the Storage Class window

The purpose of the *history table* is to:

- Indicate actions that are currently in progress or have already completed.
- Coordinate remote users (3494 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.

> **Note:** This panel becomes available after upgrade to LM 527 and VTS 2.26. However, it provides no function without FC4001 installed:
>
> - The default Storage Class, identified by eight dashes (--------), cannot be deleted.
> - Up to 256 Storage Classes, including the default, can be defined.

4.3.5 Creating Data Classes

From a z/OS or an OS/390 environment perspective, for SMS managed tape, the DFSMS Data Class defines:

- Media type parameters
- Recording technology parameters
- Compaction parameters

In regards to Outboard Policy management, it currently serves no purpose. However, to keep your DFSMS environment consistent with your Outboard Policy Management constructs, these should be defined.

To set up each Data Class (Figure 4-26), perform the following steps:

1. Select the **Commands - System Management** pull-down menu.
2. Click **Manage Constructs and Pools** (Figure 4-20 on page 167).
3. Click **Manage Data Classes**.

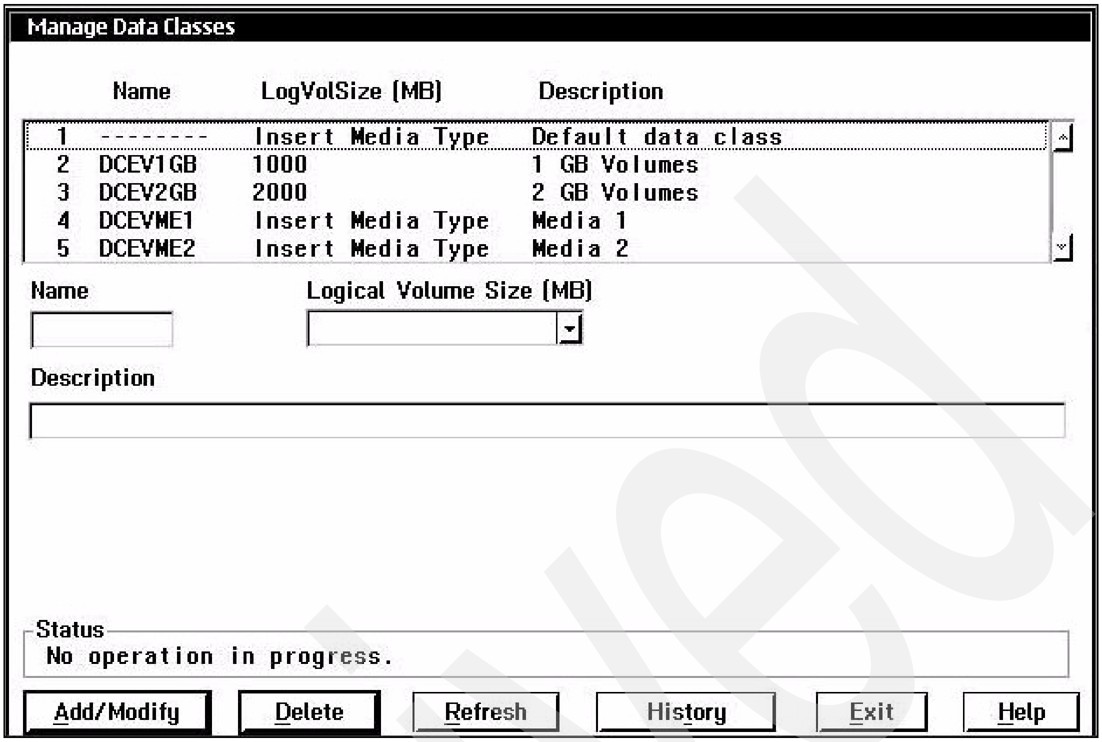

Figure 4-26 Manage Data Classes window

The Manage Data Classes window (Figure 4-26) allows you view and manage Data Classes. Through this panel, you will be able to:

- Add a Data Class.
- Modify an existing Data Class.
- Delete a Data Class.

Adding a Data Class

To add a Data Class, proceed as follows:

1. Enter a one- to eight-character name in the **Name field**. The name must be unique within the Data Class construct names.
2. Enter the **Logical Volume Size in MB** you want to use with this Data Class. Depending on the VTS hardware configuration, 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) which had been defined at the time of logical volume insert.

   Note that the hardware configuration of the VTS used in this sample does not support 4000 MB logical volume sizes. The IBM 3953 Library Manager panel shown in Figure 4-27 shows this selection.
3. Enter a short description in the **Description** field.
4. Select the **Add/Modify** button.

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. Select the **Add/Modify** button.

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. Select the **Delete** button.

Add/Modify: Adds the entered Data Class or modifies the selected Data Class

Delete: Deletes the selected Data Class

Refresh: Refreshes the Data Class window

Exit: Closes the Data Class window

Help: provides help about the Data Class window

The purpose of the *history table* is to:

- Indicate actions that are currently in progress or have already completed.
- Coordinate remote users (3494 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 a similar action.

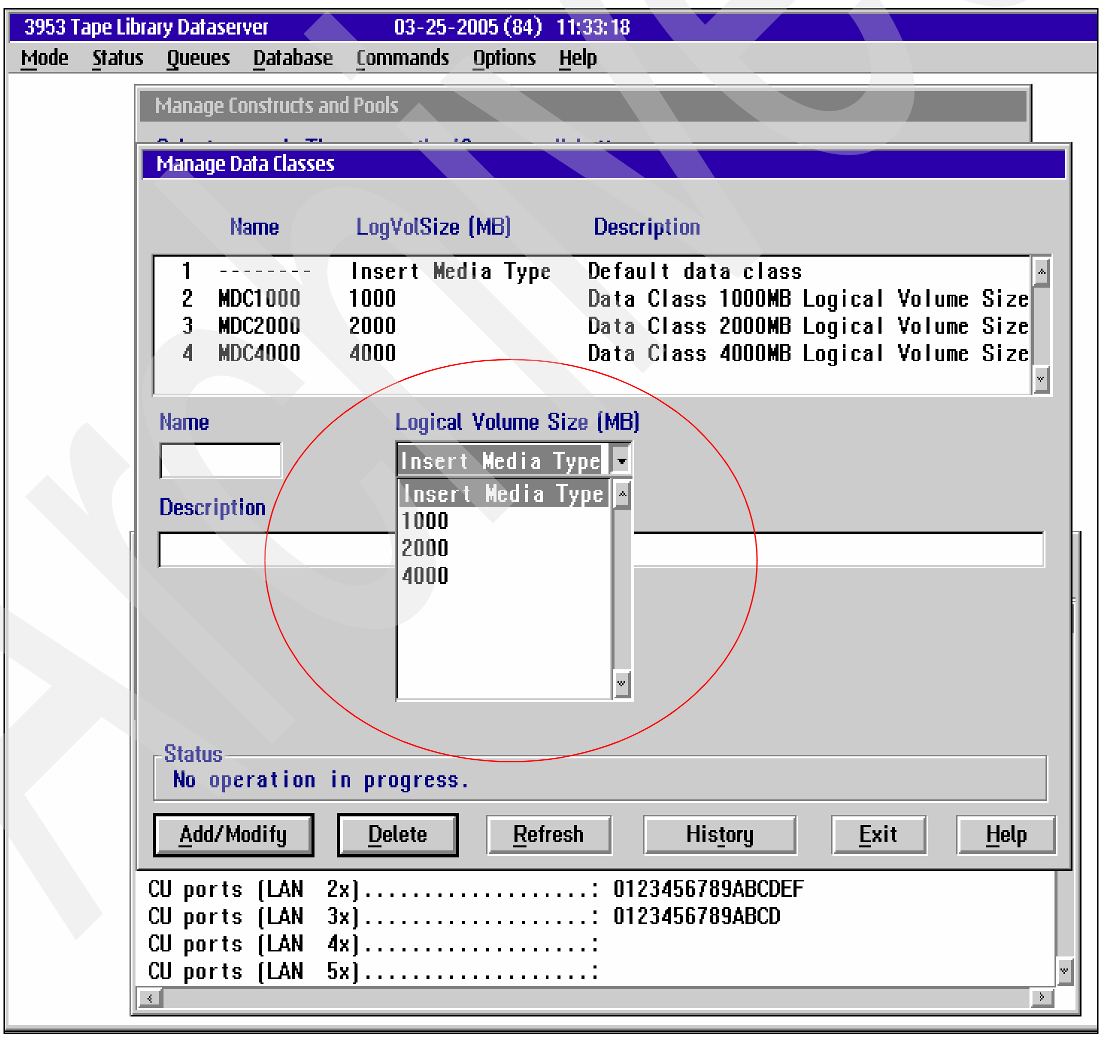

Figure 4-27 Definition of large logical volumes

4.4 Specific PtP VTS definitions and considerations

In the previous sections, we explained in detail how to define APM constructs. In the following section, we focus on PtP VTS specific implementation considerations of APM functions.

4.4.1 Construct definitions and transfer

It is essential to understand that the APM construct definitions which you make at the Library Manager are on a Library Manager Level only. Because both VTSs of a PtP VTS 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 one. This function cannot be performed through the ETL Specialist: you must do it from the Library Manager console. On the Library Manager console, select **Commands** —> **Systems Management** —> **Manage Constructs and Pools** —> **Transfer (Backup/Restore) LM Administrative Data.** The screen shown in Figure 4-28 is displayed, from where you can select which data 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.

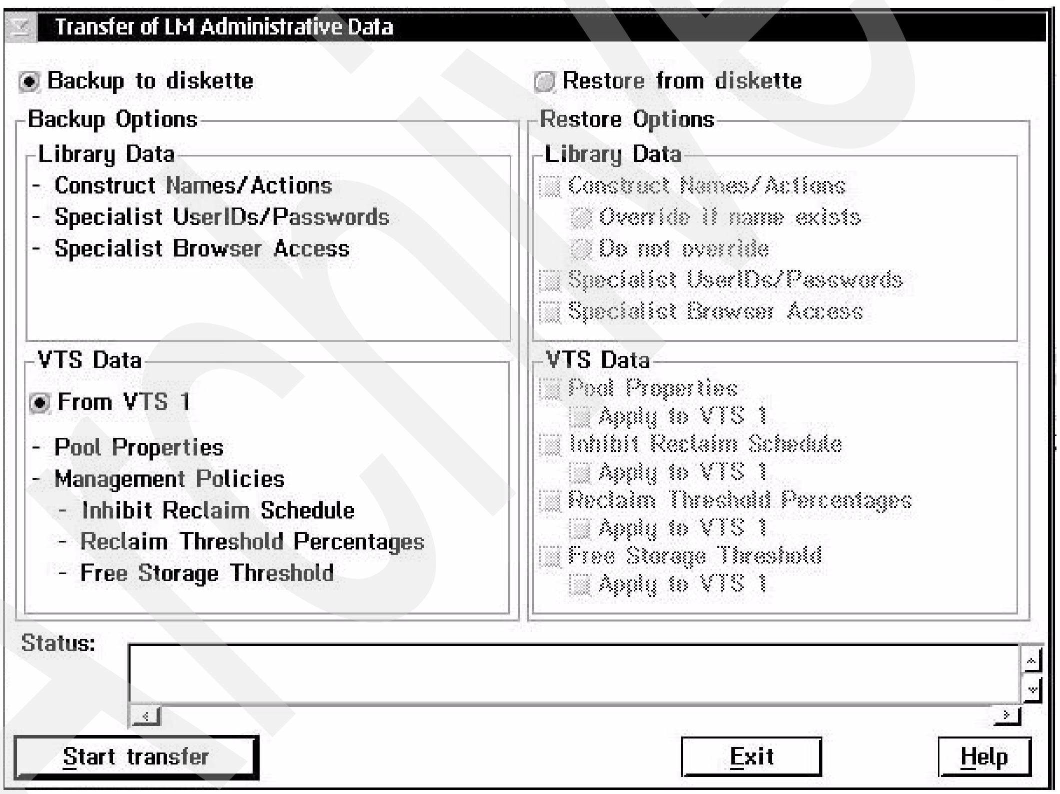

Figure 4-28 Transfer LM Administrative Data

In the example shown in Figure 4-28, **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 Restore Options.

> **Important:** In a PtP VTS configuration, you must define APM constructs and pool definitions on both Library Managers.

4.4.2 PtP Copy Mode Control

With the Advanced Policy Management features (FC4001 - FC4004) the copy mode (Immediate or Deferred) for individual logical volumes can be specified via the DFSMS Management Class construct. The copy mode is determined at logical volume mount completion, and will override the global CE setting or GDPS control setting. You should ensure that the PtP Copy Mode Control definitions are the same for both distributed VTSs.

Immediate Copy Mode

In Immediate Copy Mode the copy of the new or modified logical volume to the other VTS is made prior to the completion of the rewind/unload of the original logical volume. When the volume completes close processing, the PtP VTS has completed performing the copy of the logical volume. No further peer-to-peer copy processing is required.

Deferred Copy Mode

In Deferred Copy Mode the peer-to-peer copy of the new or modified logical volume is queued, and occurs some time after the volume completes close processing. When the peer-to-peer copy is performed depends on several factors, including current VTS workload, the length of time since the original volume was created or modified, and the amount of uncopied data in the Tape Volume Cache of the VTS which contains the original volume. It is important to note that, regardless of the operational mode, the PtP VTS will perform as much of the background work (peer-to-peer copies and copies to physical tape) as possible with the resources available at the time. In Deferred Copy Mode, although the peer-to-peer copy is initially placed on a lower priority queue than host work, the PtP VTS will perform the copy as soon as resources are available.

Deferred Copy Priority Threshold

The Deferred Copy Priority Threshold, or hours threshold, is selected by the customer and set by the service representative when the PtP VTS is installed. Once this hours age has been reached for a logical volume in the TVC and it has not yet been copied to the other VTS, its priority for being copied will be increased. The default threshold is 8 hours.

Deferred Copy Mode Peak Rate Duration

Deferred Copy Mode provides a higher write peak rate than immediate copy mode. However, the duration of the peak rate is not unlimited. The duration will primarily depend on TVC size, with larger caches affording longer peak rate duration, but at some point, the uncopied data will have to be copied. When this occurs, the host bandwidth capability will be significantly reduced.

Because of the need for “recovery” time, it is imperative that Deferred Copy Mode be used only in workload environments that have a definite peak versus average data rate profile (that is, a peak rate two or three times the average data rate), so that the PtP VTS has sufficient time to make a dual copy all the logical volumes before the next workload peak.

PtP Selective Dual Copy

Prior to the addition of the PTP Selective Dual Copy function, when a scratch logical volume was written by a host application in a PTP VTS configuration, it was first written to one of the VTSs in the configuration and then copied to the other one. That first VTS is typically referred to as the I/O VTS. The I/O VTS is determined by the Virtual Tape Controller (VTC) based on such factors as availability of both VTSs, workload on each, and customer specified defaults set during the installation of the system. The mode by which the copy is made, immediate or deferred, is determined by the your specified default, but the default mode can be overridden by an action specified for the Management Class construct assigned to the volume.

With the PTP Selective Dual Copy function, additional actions can be specified during the definition of a Management Class construct to allow:

- The Management Class to specify which VTS is to be used in writing data for a scratch mount.
- The Management Class to specify that no copy is to be performed.

These are in addition to the existing action that specifies a second copy of the data is to be performed in the same VTS.

Using the PTP Selective Dual Copy function

Applications are sometimes upgraded or new applications are created. These updated applications need to be tested and typically use a copy of production data. Once the testing has been completed, the data created is of no use. If the test data is written to a PTP VTS, it is also copied to the remote site and since the data is not used at all in any kind of recovery, it is a waste of the inter-site bandwidth to make the copy. The data copied to the other VTS will also be migrated to physical tapes In addition, testing for disaster recovery at the remote site also generates data that would waste inter-site bandwidth and the storage resources at the local site. The PTP Selective Dual Copy function provides a way to eliminate the copy and save the inter-site bandwidth.

There are several steps to take before the PTP Selective Dual Copy function can be used:

1. **Review the PTP configuration(s) that are installed.** The two VTSs that are part of a PTP configuration are labeled distributed VTS0 and distributed VTS1. You will need to understand this relationship when you define the Management Class constructs later on. If the library has two VTSs associated with it, the two VTSs are labeled as VTS1 and VTS2. However, that labeling has no relationship to the distributed VTS labeling.
2. **Decide on the Management Class construct names to be used** to control the function. We suggest that the names relate to the desired function. For example, if the desired action is to have the local VTS be the selected I/O VTS for scratch mounts and the data is not to be copied, a name like “MCLOCLNC” would be a good choice. “MC” for Management Class, “LOCL” for local VTS and “NC” for no copy. Likewise, if the desired action is to have the remote VTS be the I/O VTS for scratch mounts and the data is not to be copied, a name like “MCRMOTNC” would be a good choice. It is also suggested that a Management Class construct name be selected for the default action of letting the VTC determine the I/O VTS and to make the copy. An example could be “MCVTCCPY”.
3. **Define the Management Class construct names** and their associated actions through the Library Manager console or Specialist (WEB) interface. See 4.3.3, “Creating Management Classes” on page 172.
4. Define the Management Class construct names through the DFSMS panels at the host. Only the names are defined, there are no actions that need to be specified at the host.
5. Add selection logic to the Management Class Automatic Class Selection (ACS) routine to assign the new Management Class names as appropriate.
6. Activate the new SMS control data set with the updated ACS routines.
7. As new allocations are requested, the Management Class ACS routines will assign the new Management Class construct names as needed.

If the actions specified for PTP selective dual copy during the definition of the Management Classes are not identical on both libraries, the actions defined on the Master VTS will be executed and the following host console message will be generated:

- **VT0101 Warning:** PTP distributed library Stg Mgmnt Class volume XXXXXX copy control setting discrepancy, policy chosen: <VTC default | immediate | deferred | no copy>.

If the PTP VTS is operating in a GDPS environment, all I/O activity is through only one VTS (Primary). If a request is made to specifically mount a volume that when created specified that it was not to be copied and the volume is only valid on the other VTS (secondary), the volume will be copied from the other VTS to the VTS for I/O. In addition, the following host console message will be generated:

- **VT0103 Warning:** Primary VTS volume XXXXXX copied from Secondary VTS despite Stg Mgmnt Class specification to inhibit volume copy.

  If the Management Class specified in the mount request was defined indicating that a specific distributed VTS was to be used as the I/O VTS and that VTS is unavailable, the mount request is failed. The failure results in the following new or modified console messages (depending on whether the appropriate software PTFs have been applied):

4.4.3 Preferred versus no-preference selection for VTCs

In remote PtP VTS configurations, the customer may wish to direct the primary copies of virtual volumes from a host at one site to the distributed VTS at either the same site, or to one at the other site. Although all VTCs of one PtP VTS must be operating in the same copy mode of operations, they can be set up differently for VTS I/O preferencing. Each VTC can be set up by the service representative to operate with no-preference or with preference of either VTS0 or VTS1 as the I/O VTS.

If all components of the PtP VTS are available, and if both distributed VTSs have a valid copy of the virtual volume, and if none of them or if both of them have a copy in the Tape Volume Cache, the I/O VTS will be selected based on the VTS preferencing setup selected for the VTC. If no-preference is selected, no specific distributed VTS is chosen. The no-preference mode allows the spreading of the workload across both distributed VTSs and thus allows the maximum performance for the PtP VTS.

- In the case where all the PtP VTS components are located at the local site, no-preference is the recommended mode.
- In PtP VTS configurations where host(s) and one VTS are located at one site, and other host(s) and the other VTS are located at a remote site, the decision of whether the VTCs will be configured for preference of a particular VTS or for no-preference is more complex.

See the PtP VTS performance white papers, which are available on the TECHDOCS Web site, for guidance.

4.4.4 Delete Expired Volume Data PtP considerations

With the 526.07 level of 3494 Library Manager code and the 2.23.22.4 level of VTS code, a new function, “Delete Expired Volume Data” was added. With this optional function, the data associated with logical volumes assigned to a category with the Fast-Ready attribute set (scratch categories), are deleted from the VTS after a specified time period.

The function is provided to help you manage the amount of physical cartridge space occupied by expired data.

Tape Management System considerations

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 VTS will generate an image of a newly initialized volume. Although most tape management systems will 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.

IBM PtP VTS considerations

In using the Delete Expired Volume Data function with a PtP VTS, there are additional operational considerations that must be included in determining the expire time settings for the function because of the manner in which the VTCs select which VTS will be used in processing the I/O for a scratch mount.

The I/O VTS selection criteria that applies here are:

- The VTC will always select the VTS that has a valid version of a logical volume, independent of whether it has been expired by the host (assigned to a scratch category).
- If both VTSs have a valid version, then the preferred VTS is selected, if that mode of operation is specified.
- If a logical volume has been deleted from a VTS, it is no longer valid and that VTS cannot be selected for the I/O VTS.
  - But, if the volume has been deleted from both VTSs, then the preferred VTS is selected, if specified, and the VTC creates the volume as a newly initialized volume.

In setting up a PtP VTS ***and*** when using the delete expired volume data function, the following considerations apply, depending on the I/O selection mode of operation:

- **For no-preference:** Ensure that the non-zero expire time on BOTH Library Managers is set to the same value. If not, the VTS that expires the data first will not be used for scratch mounts.
- **For preferred VTS:** Ensure that a non-zero expire time is set on ***both*** Library Managers ***and*** that the non-preferred VTS/Library is set to delete the data ***before*** the preferred VTS/Library. Setting a non-zero expire time on only the preferred library or setting it to delete data before deletion on the non-preferred one can cause scratch mounts to use the non-preferred VTS for I/O. We recommend that the data on the non-preferred VTS/Library be set to be expired a minimum of 48 hours ahead of the data on the preferred VTS. An additional amount should be added to account for the amount of time the VTSs are to be disconnected for a disaster recovery test. For example, if the maximum disconnected time for a test is 24 hours, then the setting on the preferred VTS/Library should be 72 hours more than the setting on the non-preferred VTS/Library.

4.5 VTS software definitions

In this section, we describe topics related to the software setup of the Peer-to-Peer VTS. We provide specific information about the definition of an IBM Peer-to-Peer VTS to DFSMS/MVS SMS-managed tape and VM. For more detailed information about the software definitions for implementing an IBM VTS subsystem, please refer to *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

Tape management systems

From the host perspective, a PtP VTS is a single VTS subsystem. 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 Peer-to-Peer VTS and a stand-alone VTS. Therefore, there is no special setup for the Peer-to-Peer VTS. For basic considerations for tape management systems and the IBM VTS, please refer to *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

4.5.1 z/OS and DFSMS/MVS SMS-managed tape

In this section, we describe the DFSMS definitions for the PtP VTS. To define the Peer-to-Peer VTS, you use the ISMF panels to create a new definition of the VTS logical tape library to be recognized from the host. This definition is done in the same way as for a new installation of a stand-alone VTS except for making definitions for *three* tape libraries: One composite library and two distributed libraries. To use the Peer-to-Peer VTS, a Storage Group should be created to allow the VTS logical tape library virtual drives to be allocated by the ACS routines. Since 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. Refer to *IBM Magstar Tape Products Family: A Practical Guide*, SG24-4632, and to *z/OS DFSMS OAM Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC35-0427, for a complete discussion of the host software implementation tasks for IBM tape libraries.

Here is a summary of the basic definition procedures for a tape library installation:

1. Modify the SYS1.PARMLIB member (such as IEFSSNxx, SCHEDxx, IGDSMSxx, LOADxx, DEVSUPxx, and COMMANDxx).
2. Create VOLCAT.
3. IDCAMS IMPORT CONNECT 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 two distributed libraries. *Remember that library names may not start with a V.* Figure 4-29 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 ===> Peer-to-Peer VTS 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 4-29 Composite library definition

Figure 4-30 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 ===> Peer-to-Peer VTS 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 4-30 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 Peer-to-Peer VTS. Make sure that these defined construct names are the same as those you have defined at the LMs when Advanced Policy Management is being used.
7. Create the Storage Group(s) (SG) for the Peer-to-Peer VTS. Make sure that these defined construct names are the same as those you have defined at the LMs when Advanced Policy Management is being used.

   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 to assign VTS Storage Class. Translate, test, and validate ACS routines.
9. Activate the new SCDS.
10. Bring the composite library and distributed libraries online.
11. Bring the VTS virtual drives online.

For more detailed information about defining a VTS subsystem in a DFSMS environment, please refer to *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

JES3 initialization deck definition

In this section, we provide an example of a JES3 initialization deck for a four-VTC Peer-to-Peer VTS configuration with 64 logical drives. The INISH deck is comparable to the definition of a single IBM 3494 Tape Library containing four IBM 3490E tape control units, each of them with 16 device addresses. For a detailed description about JES3 in a SMS managed tape environment, please refer to *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

To balance the workload across all VTCs, the tape drives must be specified in a certain order without allowing JES3 to sort them. Therefore it is required that you specify `ADDRSORT=NO`. This specification should not cause any problems in your JES3 environment. For DASD devices, if they are still defined in the JES3 INISH deck, ADDRSORT does not matter at all. Even without DEVICE statements, the JES3-Data Set Awareness support will be kept, so most installations do not have DEVICE statements for DASD devices defined any more.

The following example is based on the Peer-to-Peer VTS outline shown in Figure 2-18 on page 33. The device addresses in this example range from 0110 to 014F. The library ID for the Peer-to-Peer VTS is 70299.

You must specify the library ID of the composite library (see Example 4-7).

Example 4-7 Specifying the library ID of the composite library

```
/*
/*SPECIFY ADDRSORT=NO
/*
ADDRSORT=NO
/*
/* Devices (Peer-to-Peer VTS virtual 3490E) 0110 to 014F
/*
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(0110,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(0120,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(0130,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(0140,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(0111,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(0121,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(0131,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(0141,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(0112,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(013F,SY1,,ON)
DEVICE,XTYPE=(PTPVTS,CA),XUNIT=(014F,SY1,,ON)
/*
/* This SETNAME is for Peer-to-Peer VTS virtual 3490E drives
/*
SETNAME,XTYPE=PTPVTS,NAMES=(LDGW3495,LDG70299,LDG3490E,LDE70299)
/*
/* Sysplex Wide Library Name
/*
HWSNAME,TYPE=(LDGW3495,LDG70299,LDG3490E,LDE70299
/*
/*
/* Library Specific Library Name
/*
HWSNAME,TYPE=(LDG70299,LDG3490E,LDE70299,LDGW3495
/*
/*
/* Sysplex Wide Device Name
/*
```

```
HWSNAME,TYPE=(LDG3490E,LDE70299,LDGW3495,LDG70299)
/*
/*
/* Library Specific Device Name
/*
HWSNAME,TYPE=(LDE70299,LDGW3495,LDG70299,LDG3490E)
/*
```

If you use the definition scheme, as strongly recommended for Peer-to-Peer VTS, you will not be able to benefit from the enhancements in JES3 R9 allowing simpler DEVICE definitions. With JES3 R9, ascending device addresses (max. 0001 - FFFF, usually 64 addresses on one statement) can be defined. However, you will still be able to benefit from much more important R9 enhancements, such as activation of new definitions per HOTSTART/REFRESH without IPL.

4.5.2 VM/ESA, z/VM, and guests running under VM

From the host perspective, a Peer-to-Peer VTS is a single VTS subsystem. VM supports the PtP only in VTS compatibility mode. There is no difference from the VM host’s point of view between a Peer-to-Peer VTS and a stand-alone VTS. Therefore, there is no special setup for the Peer-to-Peer VTS. This is also true for VM running with MVS or VSE guests. While the VSE native environment currently does not support the IBM PtP VTS, VSE running as a guest under VM is supported.

Note that there is one additional requirement for attaching a IBM PtP VTS to the VM host. VM does require that you also have the same IBM PtP VTS library attached to an MVS host. This is due to the fact that VM only knows about the composite library. The MVS host system and OAM are required in order to monitor and display information about the distributed libraries. For example, unsolicited attentions such as intervention required and CBR3750I hardware messages that the VTSs send to the host.

Here is a summary of the basic VM host environments and their software requirements for attaching an IBM VTS. These topics are discussed in much more detail in the *IBM Magstar Tape Products Family: A Practical Guide*, SG24-4632 and the VM-specific redbook, *Lights Out! Advanced Tape Automation Using VM/ESA*, GG24-4347.

In this section we briefly describe the following VM and VSE host environments:

- VM/ESA and z/VM native support using DFSMS/VM
- VM/ESA or z/VM
- z/OS guests
- VSE/ESA guests
- VSE/ESA as a VM/ESA guest using a VSE guest server (VGS)

VM/ESA and z/VM native support using DFSMS/VM

DFSMS/VM Function Level 221 (FL221) is the only way for a VM/ESA system to communicate with an IBM VTS. DFSMS/VM FL221 is part of VM/ESA and z/VM. The RMS function of DFSMS/VM FL221 provides VTS 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. Additional information on third-party tape management systems that support the IBM VTS in the VM/ESA and z/VM environment can be found in the *IBM Magstar Tape Products Family: A Practical Guide*, SG24-4632 and *Lights Out! Advanced Tape Automation Using VM/ESA*, GG24-4347.

Figure 4-31 shows the VM/ESA and z/VM native support for the IBM VTS.

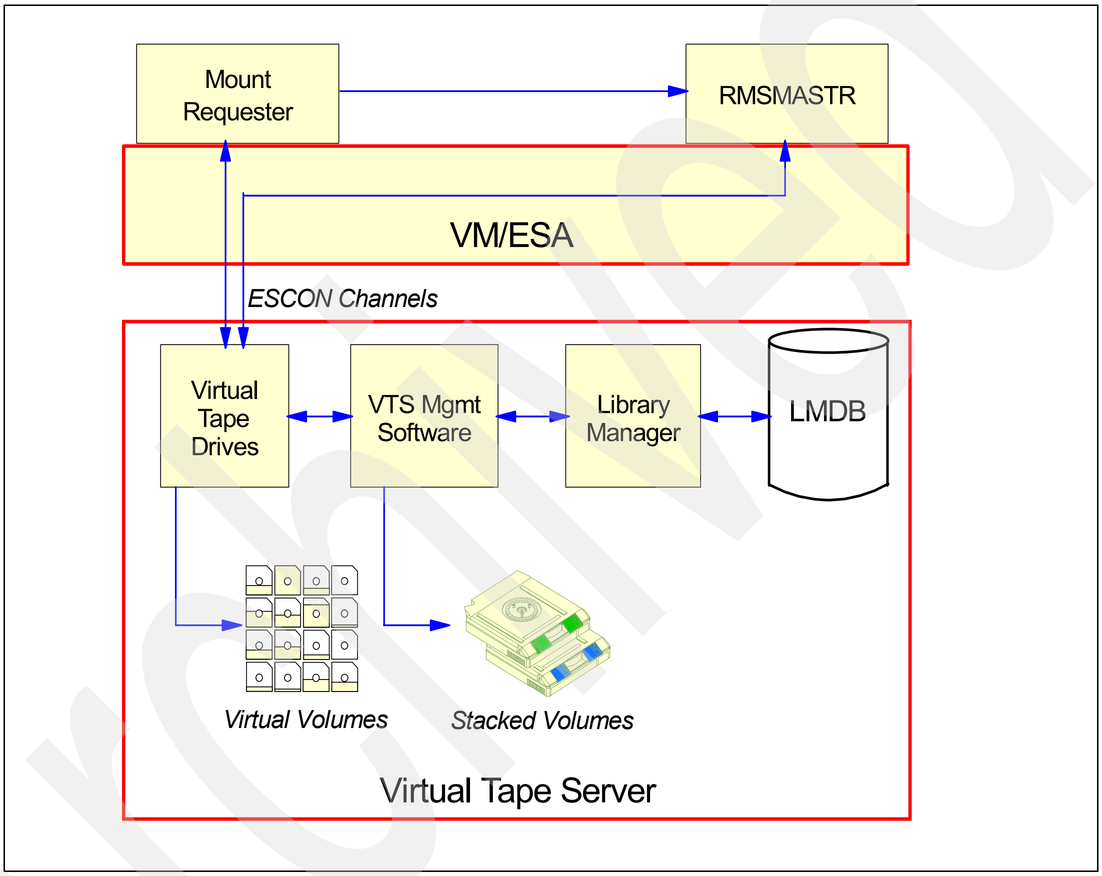

Figure 4-31 VTS in a native VM/ESA environment using DFSMS/VM

When you use the IBM VTS 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 VTS 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 VTS is not the best candidate for this data.

DFSMS/VM

After you have defined the new VTS tape library through HCD, you must define the VTS to DFSMS/VM if the VM system is to use the VTS directly. You define the VTS 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 the 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 library sequence number that is assigned to the VTS during hardware installation.

> **Note:** The Advanced Function feature and APM are currently not supported with VM/ESA.

For details on implementing of DFSMS/VM and RMS, refer to the *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141. If the VTS is shared by your VM system and other systems, for example in the case of an IBM PtP VTS you have a sharing MVS host, then additional considerations apply. Refer to the *Guide to Sharing and Partitioning IBM Tape Library Dataservers,* SG24-4409, for further information.

Running as a guest under VM/ESA or z/VM

Both z/OS hosts and VSE/ESA hosts allow you to use an IBM PtP VTS while running as a guest host system under VM/ESA or z/VM.

z/OS guests

It is possible for the environments described in 4.5.1, "z/OS and DFSMS/MVS SMS-managed tape" on page 184 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.

> **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`

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 on the STDEVOPT statement, refer to *VM/ESA Planning and Administration Guide and VM/ESA Running Guest Operating Systems*.

VSE/ESA and z/VSE guests

z/VSE and VSE/ESA support native attachment of a standalone VTS. However, to support a PtP VTS, z/VSE and VSE/ESA must run as a guest system under z/VM.

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 VTS library. If the VGS is required, define the LIBCONFIG file on the VGS service machine's A disk. This file simply cross-references the VSE/ESA guest's tape library names with the names that DFSMS/VM uses. To enable VSE/ESA 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 VSE/ESA 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 "VSE/ESA as a VM/ESA guest using a VSE guest server (VGS)" on page 190.

For further information, please refer to 7.6, “VSE Guest Server Considerations” in the *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409.

The CA DYNAM/TM-VSE does not use the VGS machine.

VSE/ESA as a VM/ESA guest using a VSE guest server (VGS)

When a VSE/ESA guest machine uses a tape drive in the IBM VTS, 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, VSE/ESA 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. VSE/ESA cannot use RMSMASTR directly, however, because RMS functions run only in (CMS) mode.

Therefore some VSE/ESA 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 3494 and the virtual drives of the IBM VTS. In Figure 4-32 you can see the flow and connections of a IBM VTS in a VSE/ESA environment under VM.

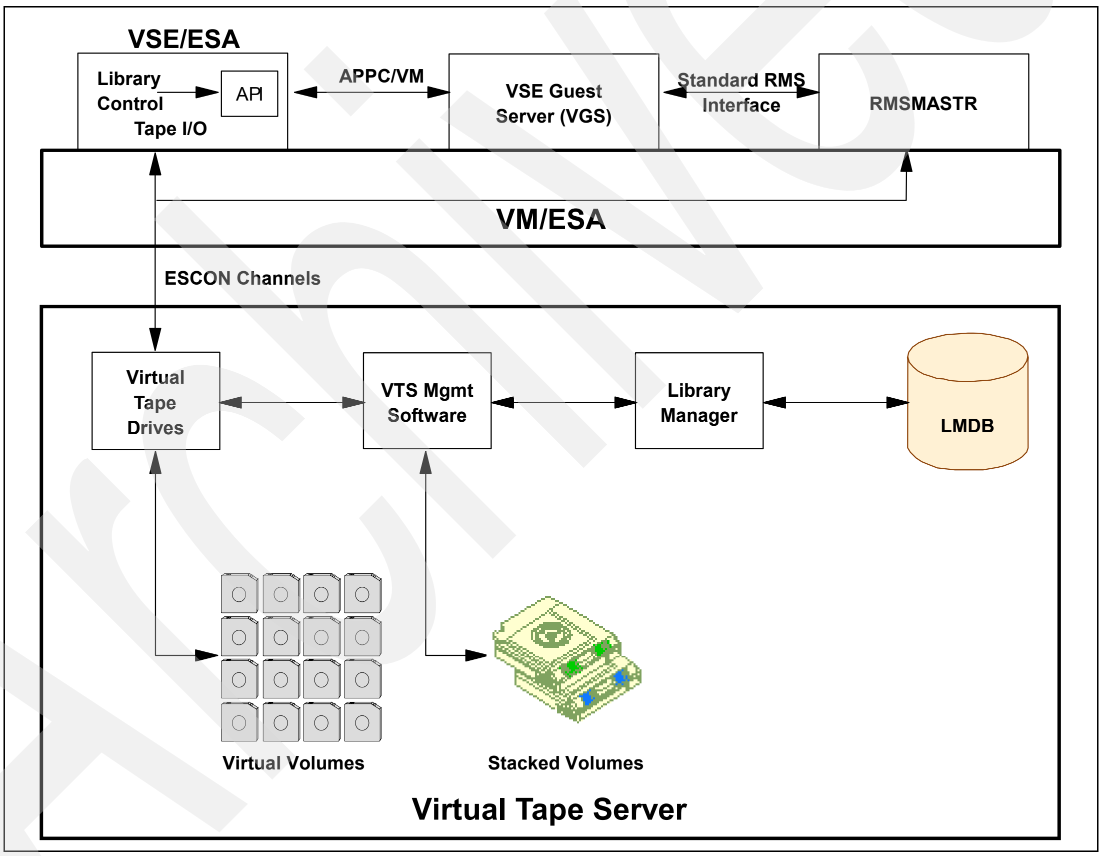

Figure 4-32 VTS in a VSE/ESA environment as a VM guest

Tape management

As with the VM/ESA native environment (see “VM/ESA and z/VM native support using DFSMS/VM” on page 187), the tape management system is responsible for keeping an inventory of volumes in the IBM VTS that belong to VSE/ESA. Some vendor tape management support scenarios do not use VGS. Instead they communicate directly with RMSMASTR through CSL calls. Figure 4-33 shows the case of CA-DYNAM/T VSE.

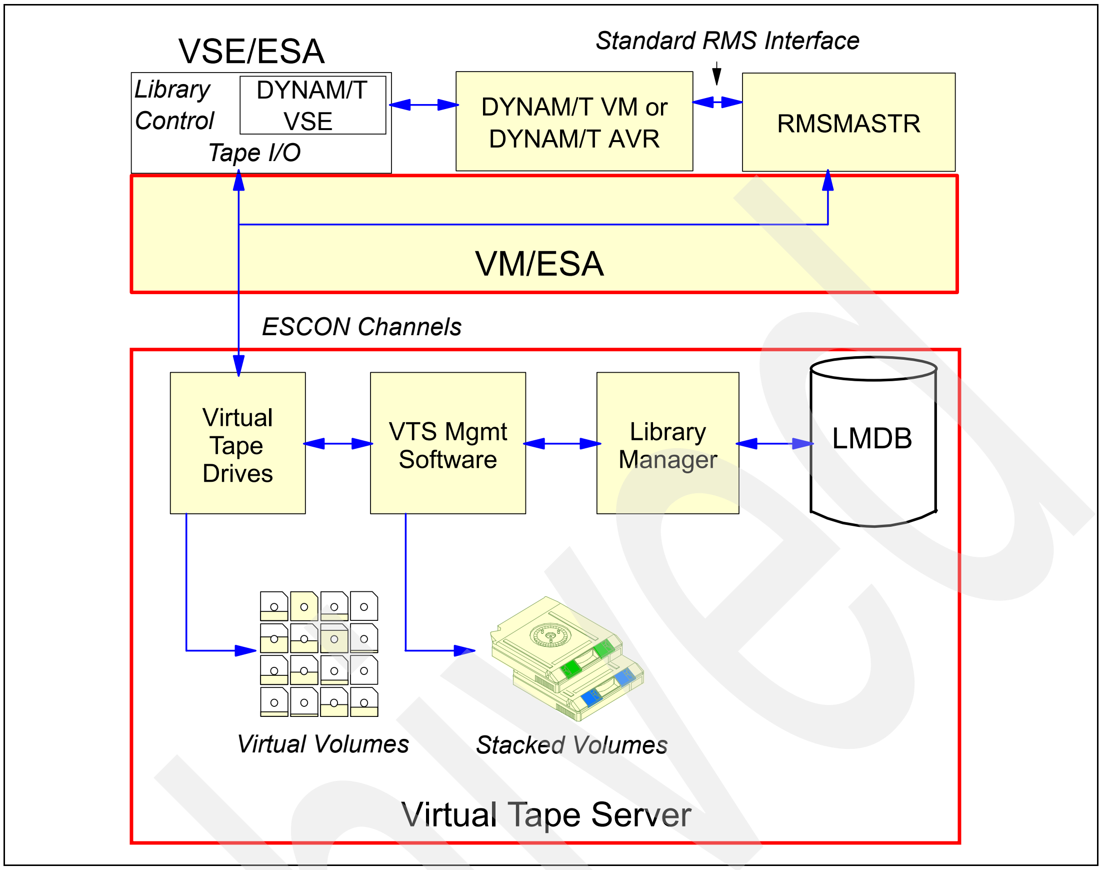

Figure 4-33 VTS in a VSE/ESA environment as a VM guest (no VGS)

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.

4.5.3 Other host operating systems

This section summarizes the software support for additional host systems.

Native z/VSE and VSE/ESA

Native z/VSE does not support the PtP VTS. Standalone VTS support is provided with:

- VSE/ESA V2.7 with LCDD and VTAM® LU6.2
- z/VSE V3.1, when the VTS attaches through ESCON channels to the host.

Open Systems

Open Systems do not support the IBM PtP VTS. They do, however, support the stand-alone VTS. For additional information on Open Systems and the IBM stand-alone VTS, refer to *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring*, SG24-2229.

TPF

TPF does not support the IBM PtP VTS. They do, however, support the stand-alone VTS. For additional information on TPF and the IBM stand-alone VTS, refer to *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring*, SG24-2229.

4.6 Hints and tips

In this section we present some hints and tips for installing the PtP VTS.

4.6.1 Sharing VTS within 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 PtP VTS, looks like a separate library to the host. A PtP VTS can be shared by multiple System z systems, VM and VSE guest systems in the same way that a physical tape library or stand-alone VTS can be shared.

> **Note:** If you are sharing a PtP VTS with the Advanced Function feature or APM installed between multiple hosts, you have to make sure that the proper toleration maintenance is installed in all attaching systems. See 3.2, "Planning for Advanced Policy Management exploitation" on page 82.

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 set, tape management system inventory and the catalog information, that is 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.

4.6.2 Partitioning the VTS 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 on sharing and partitioning a library, refer to Chapter 5, "Implementation in a DFSMS/MVS Environment" in the *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409.

4.6.3 VTS teach operation

The teach operation is performed by an IBM service representative as part of the installation process for the 3494 with a VTS and the definition of the PtP VTS. Therefore, you need not be concerned with teach operations. A teach must be performed as a part of the PtP VTS installation.

4.6.4 Library inventory

The IBM service representative typically performs an inventory of the 3494 as part of the PtP VTS installation. Either a reinventory complete system or an inventory new storage is performed, depending on whether the VTSs are part of a whole new 3494 library, or are installed in an existing 3494 library.

Before the inventory is performed, you can insert your stacked volumes in empty cartridge cells in the library frames to be inventoried. If you have many cartridges to insert, this is a good time to do so. Note that with inventory new storage, you should only insert cartridges in the newly installed library frames, as those are the only ones to be inventoried. During the inventory, the IBM service representative is presented with the option of defining stacked volser ranges. You should inform the service representative of the ranges that need to be specified. Later, you can add and modify the range definitions and insert more stacked volumes.

During the inventory, the IBM service representative is also presented with the option of defining the cleaner volume masks. At least one cleaner mask must be defined. Inform the service representative of your cleaner cartridge volsers.

For the IBM TS3500/3953, there is no Inventory function that the CE can initiate. The TS3500 library performs an automatic inventory whenever after a frame door is closed. See 4.2.5, “Physical cartridge insertion” on page 158 for more details on inserting physical VTS cartridges into an IBM TS3500 library.

With the 3494, logical and physical volsers could be deleted at the beginning of a *Re-inventory Complete* operation. However, with the 3953 Library Manager attached to the TS3500, the re-inventory complete function is removed, since the TS3500 performs its own inventory and the 3953 Library Manager electronically uploads the inventory.

To compensate for the removal of this delete capability, a new menu item was added under the *Utilities* pull-down called *Delete physical/logical volumes...* This function is accessible by the CE and provides the same delete capability that is found on the 3494. After the delete, an upload from the TS3500 of the physical volumes is automatically performed. The volsers are placed in the appropriate Insert category. If the logical volsers were deleted, they will need to be reinserted from the LM panel. The host will then need to upload the volsers and set them to the appropriate category.

5

PtP VTS migration and relocation

In this chapter, we describe the step-by-step migration approach to upgrade from IBM VTS to IBM Peer-to-Peer VTS. We assume that you already have a basic understanding of the IBM VTS and IBM PtP concepts discussed in earlier chapters and the *IBM Total Storage Virtual Tape Server: Planning, Implementing, and Monitoring*, SG24-2229.

Whenever one or more IBM VTSs are already installed and functional, if the existing hardware has to be reused to build a Peer-to-Peer VTS configuration, you will need to consider a *migration* approach to move to a Peer-to-Peer environment.

This means that, based on the hardware installed, some or all the components that are part of your 3494 or TS3500/3953 and VTS configuration, need to be upgraded and therefore temporarily stopped.

This migration section is divided into five main topics:

- Planning for migration
- Upgrading an existing VTS to a PtP VTS
- Merging two existing VTSs in to a PtP VTS
- Relocating an existing PtP VTS
- Migrating an existing 3494 PtP VTS to a new TS3500/3953 PtP VTS

The first topic, planning for migration, describes general migration considerations. The second topic has a detailed task list for the upgrade of an existing VTS in to a PtP VTS configuration. The third topic describes the tasks for the merge of two existing VTSs in to a single PtP VTS. The final two topics provide a step-by-step task list for relocating an existing PtP VTS, as well as the possibility of migrating to a TS3500/3953 when 3592 cartridges are used.

5.1 Planning for migration

In this section, we describe general migration considerations that apply to all four migration scenarios discussed later in this section.

5.1.1 The Peer-to-Peer VTS hardware configuration

You need to consider the difference in the data processing capability between the stand-alone existing VTSs and a Peer-to-Peer VTS with the copy processing overhead. The performance of the Peer-to-Peer VTS depends on selection of the copy mode, the I/O preference setting, and the level of Advanced Policy Management (APM) exploitation on each individual VTS within the PtP VTS. You may need to add TVC capacity and 3592 tape drives. For more information, please refer to Chapter 2, "Peer-to-Peer VTS Architecture" on page 7, and Chapter 3, "Preinstallation planning" on page 63.

5.1.2 The existing VTS subsystem configuration

To upgrade an existing VTS subsystem, there are prerequisite features for the VTS configured in a Peer-to-Peer environment. You have to check the existing VTS subsystem configuration and add the features, if required.

For more information about the prerequisite features, please refer to 3.1.2, "VTS configuration requirements" on page 65.

Note: If your current VTS configuration is a model B18, then you must first upgrade to a model B20 VTS before moving to a PtP VTS configuration. Refer to the *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring*, SG24-2229 for additional information on upgrading a stand-alone VTS from a model B18 to a model B20.

5.1.3 Stacked volumes

Additional stacked volumes may be required to hold the copies of logical volumes. Each distributed library must have enough stacked cartridges to hold all the logical volumes.

Note: Advanced Policy Management (APM) may require additional stacked cartridges as well, depending on your use of the Dual Copy in one or both VTSs. We do not recommend Dual Copy in a PtP VTS environment, with the exception of the use of Selective PtP Copy. Refer to 3.2, "Planning for Advanced Policy Management exploitation" on page 82.

5.1.4 Software implementation

If you are converting an existing VTS subsystem in to a PtP subsystem, you need to install the required PTFs on your system to use a Peer-to-Peer VTS. For more information about the software requirements of a Peer-to-Peer VTS, please refer to Section 3.3, "Host software requirements" on page 88.

I/O definition

You need to modify the I/O definition in IOCDS and HCD for each of the independent logical tape subsystems of 16 drives that is configured in a Peer-to-Peer VTS. For more information about the I/O definition of a Peer-to-Peer VTS, see 4.1, "Hardware I/O configuration definition" on page 128.

SMS tape library definition

When upgrading to a PtP VTS, the library name of an existing VTS is used for the composite library; therefore, the existing tape library definition can be used for the composite library. This causes all logical volumes in the existing VTS subsystem to be associated with the composite library, and thus no manual changes to the Tape Configuration Data Base (TCDB) entries are required.

In addition, you need to add two tape library definitions for the distributed libraries, remembering that a library name cannot start with the letter “V”. One of the two distributed libraries is the existing VTS subsystem that now has a new library name. Figure 5-1 shows a sample modification of the SMS tape library definition. Note that while the host library name associated with the original VTS changes, the library ID of the original VTS remains the same and is associated with the first distributed library in the PtP VTS configuration.

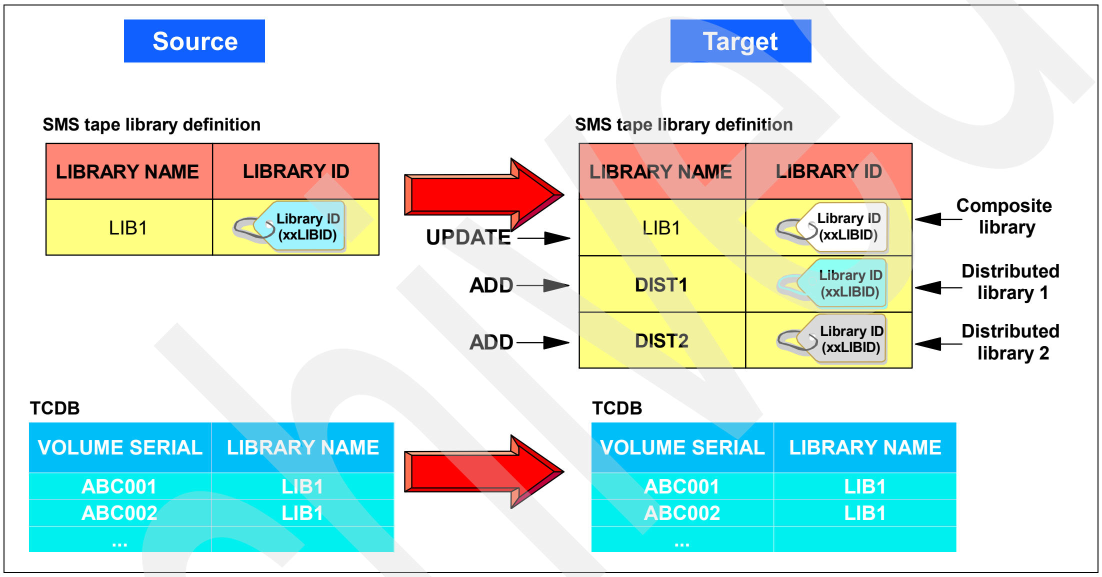

Figure 5-1 SMS tape library definition for upgrading from VTS to PtP VTS

> **Important:** The library ID field (also known as the library sequence number at the hardware level) defined in the SMS Tape Library Define panel, as well as in the HCD LIBRARY-ID parameter, must ***exactly*** match the library sequence number defined to the library by your hardware service representative

For more information about the SMS tape library definition of a Peer-to-Peer VTS, please refer to 4.5.1, “z/OS and DFSMS/MVS SMS-managed tape” on page 184.

JES3 environments

In a JES3 system environment, you need to modify the INISH deck to provide load balancing across all the VTCs.

For more information about the JES3 definition of a Peer-to-Peer VTS, please refer to “JES3 initialization deck definition” on page 186.

5.1.5 Application changes

In a Peer-to-Peer VTS, the hardware automatically creates dual copies of all logical volumes. Therefore, you may no longer need to use software functions like DFSMShsm TAPECOPY or DFSMSdss duplexing to create dual backup copies.

5.1.6 Volumes and data migration

Migrating from a stand-alone VTS configuration to a Peer-to-Peer VTS does not imply data movement; this differs from the migration process for a physical tape library to a VTS. You do not need to use a tape copy tool to move data from the stand-alone VTS to the Peer-to-Peer library.

The logical volumes already present in the stand-alone VTS are associated with the new composite library during the migration process without the need to copy the data into the Peer-to-Peer VTS.

Remember that you will need to provide an additional quantity of scratch stacked volumes equal to the original number used before the upgrade to hold the second copy of each volume.

The dual copy of a logical volume will only be made after the volume has been accessed (read or written) again by the host system. You may wish to force the copy of *critical volumes* to be done by reading volumes that have not been copied. You can use your tape library management system to identify volumes with a last reference date prior to the upgrade and build jobs to read from them.

For example, if you are using DFSMSrmm, you can get a list of volumes that have not been referenced since a Peer-to-Peer upgrade by issuing the command shown in Figure 5-2.

```
RMM SEARCHVOLUME VOLUME(*) -
   LOCATION(library name) -
   STATUS(NOTSCRATCH) -
   LIMIT(*) -
   TYPE(LOGICAL)
   OWNER(*) -
   CLIST
```

Figure 5-2 Finding volumes without a copy: Step 1

This command gives you a list of volumes that are not in scratch status. You now need to issue another command, as shown in Figure 5-3, to find out which volumes have not been copied.

```
RMM LISTVOLUME VOLUME(volser) -
   STATS -
   LIMIT(*) -
   OWNER(*) -
   CLIST
```

Figure 5-3 Finding volumes without a copy: Step 2

This command will tell you the date last read and the date last written. If either of these is earlier than the date of the upgrade to a Peer-to-Peer configuration, there is no second copy.

In addition, OAM has an SMS display command that retrieves volume information from the TCDB volume record and hardware Library Manager. This single volume display will retrieve and display whether the hardware is reporting that a "dual copy" currently exists for the logical volume displayed. Refer to Figure 6-18 on page 246 for an example of this display.

> **Important:** Do *not* try to copy every volume that does not have a second copy.
>
> The upgrade process does not include a step to copy these uncopied volumes, because the performance impact would be severe. The process of creating a second copy will require a specific mount for a recall to cache, cache space, and a write to a stacked volume. This will take resources away from normal activity and would impact performance.
>
> This technique should *only* be used to identify a small number of volumes for which you must have copies.

Once you have decided whether you need to copy any of these volumes, you could build jobs (such as IEBGENER to read the volume label, requiring BLP authority) or issue mount commands.

Please note that it is much better to allow the Peer-to-Peer VTS to create copies, since volumes are accessed as part of your normal workload. The technique discussed in this section should only be used for very small numbers of volumes because of the performance impact.

5.2 Upgrading an existing VTS to a Peer-to-Peer VTS

In this section, we give you a step-by-step approach to upgrade an existing IBM 3494 with a VTS to a Peer-to-Peer VTS configuration. For the purposes of this example we assume that the existing VTS is a 3494. The process would be the same if the existing VTS library is a TS3500/3953, with the exception that the TS3500/3953 only supports 3592 drives.

This phased procedure is intended to minimize the impact of the migration, reusing the previous VTS attributes and definitions as much as possible to reduce migration outages and impact on your operations.

Migration of the existing VTS subsystem is performed by your service representatives. It involves the physical replacement of essential parts of your 3494 Tape Library and VTS subsystem, so you must plan for the appropriate outage of your tape library and the VTS.

The hardware upgrade consists of the addition of:

- An IBM 3494 or TS3500/3953 Tape Library
- Additional 3590 or 3592 cartridges to hold the logical volume copies
- A new VTS with the copy feature
- The CXn auxiliary frame or frames
- The virtual tape controllers

The existing VTS must be upgraded with the copy features, together with any appropriate cache upgrade. If the current stand-alone VTS is a model B18, then it must first be upgraded to a model B20 before attempting the upgrade to a PtP VTS configuration. (Refer to the *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring*, SG24-2229 for additional information on upgrading from a model B18 to a B20.) Much of this work can and should be done as preparation work in advance of the Peer-to-Peer feature installation. The tasks that you can perform and test early are shown in "Updating the existing 3494 VTS" on page 200.

The hardware MES can be completed in two parts. First, we have the installation of the new hardware components: 3494 or TS3500/3953, VTS, CXn frame, and VTCs. This can be done without affecting the operations of the existing VTS. These tasks are shown in “Installing new hardware” on page 200.

The second part requires that the existing VTS is down, because we need to install the Copy Base feature. These tasks are shown in “Installing the Peer-to-Peer upgrade” on page 201.

5.2.1 Planning

The preparation steps can and should be done well in advance of the delivery of the new hardware. You should decide upon the configuration and collect data for the installation of the StorWatch Specialists. The *IBM 3494 Tape Library Introduction and Planning Guide*, GA32-0448, contains some worksheets that you may find helpful in documenting what will be done. After you have completed these worksheets, review them with your service representative. Similarly, information on the TS3500/3953 is found in the *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Introduction and Planning Guide*, GA32-0472 and the *IBM System Storage TS3500 Tape Library Introduction and Planning Guide,* GA32-0559.

See 3.1.2, “VTS configuration requirements” on page 65, and 3.1.14, “TotalStorage Specialist configuration” on page 78, to see what information you will need.

5.2.2 Hardware migration steps

The hardware migration is performed by your service representatives, who will complete the following sets of tasks.

Updating the existing 3494 VTS

To update the existing 3494 VTS:

1. Move any logical volumes accessed by SCSI hosts to another library.
2. Vary existing VTS offline.
3. Remove SCSI host attachment feature if installed.
4. Increase cache size if necessary.
5. Install LAN card in Library Manager if not already installed.
6. Install additional 3590 or 3592 drives if needed.
7. Install Library Manager licensed internal code supporting Peer-to-Peer.
8. Install the VTS licensed internal code supporting Peer-to-Peer.
9. Install Advanced Policy Management feature if desired.
10. Vary VTS online.
11. Configure and enable Web access to the Library Manager.
12. Install the DFSMS/MVS support and Toleration code for Peer-to-Peer.

Installing new hardware

To install new hardware:

1. Install CXn frame or frames.
2. Install four or eight VTCs.
3. Configure the VTCs.

4. Install new hardware:
   a. If 3494, install new 3494 Tape Library, 3590 or 3592 drives, and new stacked volumes.
   b. If TS3500/3953, then install the new TS3500/3953, 3592 drives, and new stacked volumes.
5. Install new VTS including the copy features.
6. Install Advanced Policy Management feature (both VTSs must match).
   a. Copy APM definitions from the original VTS over to the new VTS.
7. Configure the new Library Manager for Peer-to-Peer.
8. Install new ESCON/FICON directors if required.
9. Connect ESCON/FICON links from the new VTS to the VTCs.
10. Install and configure RSF for the new VTS and VTCs.
11. Configure and enable Web access to the VTCs.
12. Configure and enable Web access to the new Library Manager.

Installing the Peer-to-Peer upgrade

To install the Peer-to-Peer upgrade:

1. Vary original 3494 virtual devices and VTS offline.
2. Install the Peer-to-Peer features on the original VTS.
   a. The service representative must disable the Import/Export part of the Advanced Function feature.
   b. Installation creates the token database.
3. Confirm that the code levels on both Library Managers (3494 and/or TS3500/3953) and VTSs match.
4. Connect links from original (now upgraded) Bxx to VTCs.
5. Connect host ESCON/FICON channels to VTCs.
6. Back up the VTS partition of the original Library Manager database.
7. Restore the Library Manager database to the new Library Manager.
8. Configure the original Library Manager for Peer-to-Peer.
9. Designate one library as the user interface library.
10. Bring each VTS and the VTCs online.
11. Initiate the creation of tokens in the new VTS. The Peer-to-Peer VTS will initialize as available.
12. Vary online virtual devices at host.

5.2.3 Software migration steps

These are the steps that you must perform to define the Peer-to-Peer VTS and make the appropriate SMS changes:

1. Use HCD to delete the drive addresses for the existing VTS[1].
2. Use HCD to define device addresses for the Peer-to-Peer VTS as eight, or sixteen strings of 16 devices each.

[1] You could also use HCD to alter the definitions for the existing VTS; for completeness, we will show this approach in the section that describes merging two existing VTSs in 5.3, “Merging two existing VTSs into a Peer-to-Peer VTS” on page 205.

3. Use the ISMF library definition of the existing library as the definition for the composite library to avoid TCDB changes.
   – You must change the library ID to match the new library ID defined by your hardware service representative to the composite library. (The library ID of the original existing library will become the library ID of the distributed library.)
4. Use ISMF to define each of the distributed libraries.
5. If applicable, update Jes3 definitions
6. Activate the IODF with the device address for the Peer-to-Peer VTS.
7. Activate the SCDS to use changed SMS definitions.
8. Verify that the composite library and both distributed libraries come online.

In the rest of this section we discuss these software migration steps in more detail through the use of a specific example.

Let us assume that the existing VTS is configured as in Figure 5-4, installed and active, using 128 virtual drives with:

- Virtual drive addresses ranging from 100 to 17F
- A library name of LIBVTS1
- A library sequence number (LIBID) of 11111

The TCDB contains thousands of volumes in the range from V00001 to V99999. These virtual volumes have LIBVTS1 as the library name in the volume record, as shown in Figure 5-4.

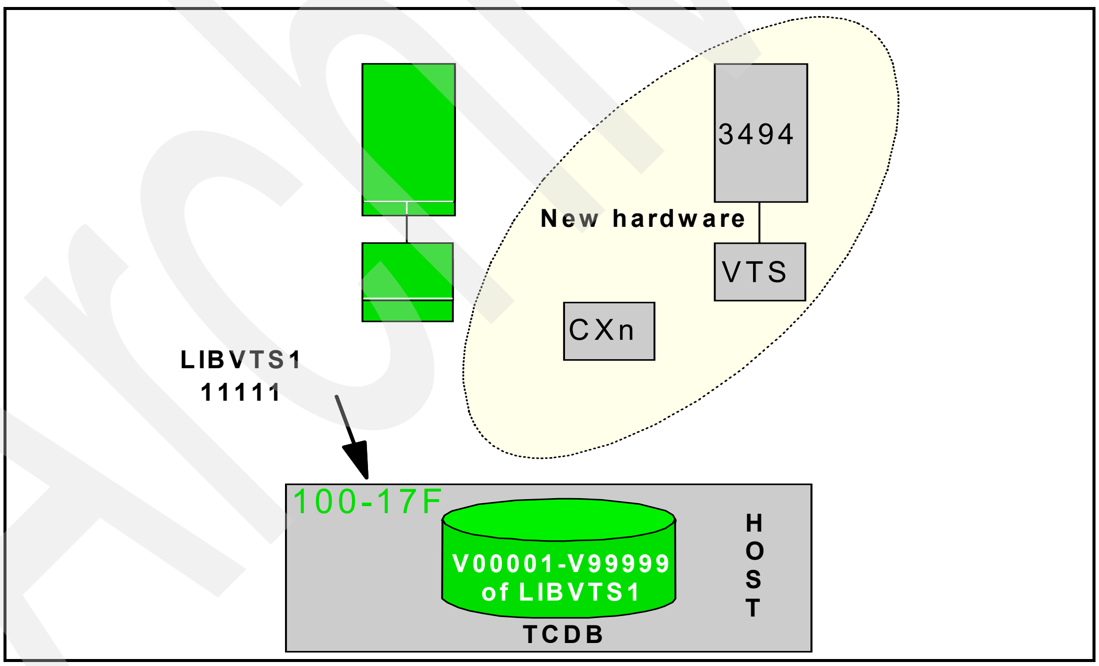

Figure 5-4 Upgrading an existing VTS: MES installation adding new hardware

Once the hardware installation is completed, you can start the software implementation as detailed in the following steps.

Deleting old device addresses

Using HCD, delete the range of 128 addresses used by the original VTS, 100 to 17F.

The new Peer-to-Peer VTS will use the 128 virtual devices of the composite library.

HCD define of new composite library devices

Using HCD, define a new range of 128 devices (addresses 200-27F) as documented in 4.1.3, “Defining devices through HCD panels” on page 140. These will be the new 128 virtual devices of the composite library.

Composite library definition

We will reuse the library name, LIBVTS1, of the original VTS so that we do not need to make any TCDB changes. The LIBPORT IDs will be 1, 2, 3, 4, 5, 6, 7, 8, respectively, for the eight emulated 3490 CUs. The device address range is 200-27F in this example. See the panels in 4.1.3, “Defining devices through HCD panels” on page 140.

We will use ISMF to alter LIBVTS1 to point to library ID 33333, which we will assume is the newly defined library ID of the composite library. If you do not supply a library ID, the convention is to use the last five digits of the CXn frame.

ISMF define of distributed libraries

Use ISMF to define the distributed libraries. In this example, we use the VTS library sequence number as the Library ID, for example, 11111 and 22222, and put in a name for the distributed VTSs, for example, DISTLIB1 and DISTLIB2, as you can see in Figure 5-5.

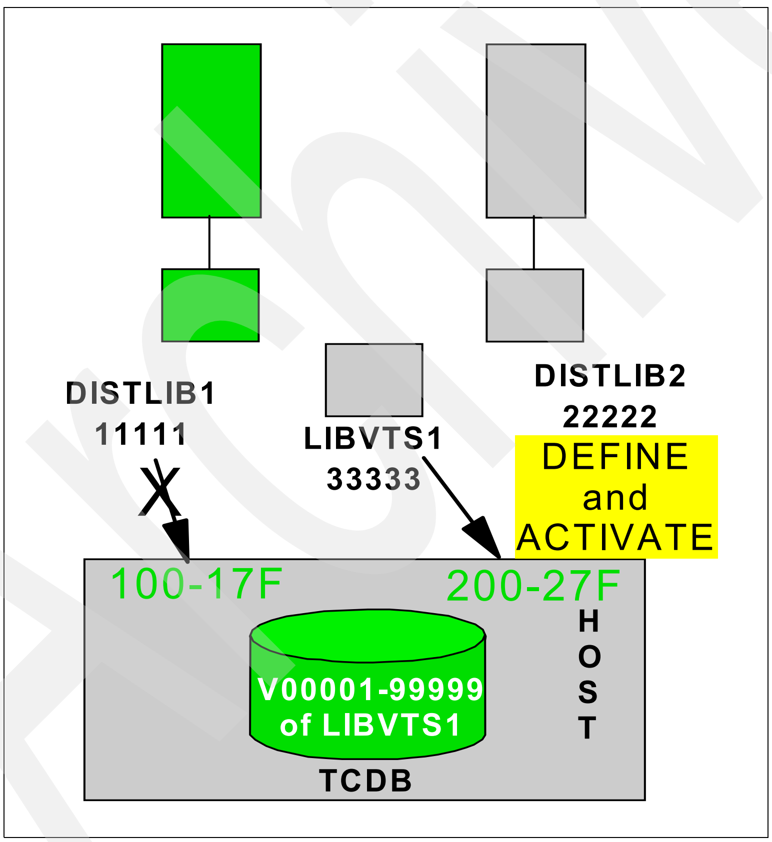

Figure 5-5 HCD and ISMF definitions for the newly installed hardware

Activating the HCD definitions

Activate the new definitions using the ACTIVATE IODF command from the MVS console for the HCD definitions, and using the SETSMS SCDS command for the SMS definitions.

After activating the IODF, the new 3490 addresses should be available and can be varied online to MVS.

Activating the SMS definitions

If your ISMF Library definitions for the three libraries have "Initial Online Status" set to "YES", activating the SCDS should bring the composite library online (LIBVTS1) and the two distributed libraries online (DISTLIB1 and DISTLIB2). If OAM is not automatically restarted upon SCDS activation (`RESTART=NO` is specified in the OAM started procedure), you need to restart it manually, using the **`F OAM,RESTART`** command.

Then you can submit jobs to LIBVTS1 as you did before the upgrade, because the TCDB records remain the same.

Verifying the definitions

If the new libraries go online, this means that the Distributed VTSs and the VTCs are successfully communicating with the host. The final Peer-to-Peer VTS configuration is shown in Figure 5-6.

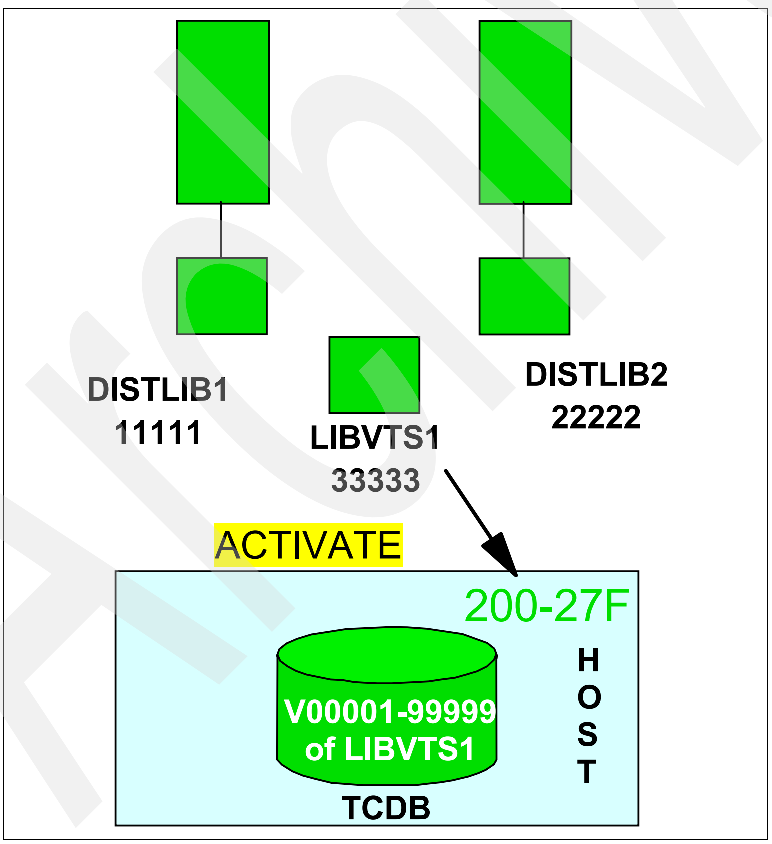

Figure 5-6 Peer-to-Peer VTS configuration after the MES is completed

From a host perspective, both the original stand-alone VTS and the Peer-to-Peer VTS can be seen as a single VTS. As in the original stand-alone VTS configuration, the host sees an address range of 128 virtual 3490 devices. Therefore, after the migration process, the host system still sees a single VTS image, even though it is now a Peer-to-Peer VTS.

5.3 Merging two existing VTSs into a Peer-to-Peer VTS

From a host system point of view, merging two existing stand-alone 3494 and/or TS3500/3953 VTS systems into one Peer-to-Peer configuration means that you are consolidating two libraries into one. This means that you should consider these points:

- You must make sure that there are no duplicate volume serial numbers for the logical volumes across the two source VTSs.
- A Peer-to-Peer VTS may have up to 500,000 logical volumes, as may one stand-alone VTS; so, the combined total of logical volumes in the two source stand-alone VTSs must be less than 500,000.
- Each stand-alone VTS can have up to 256 device addresses for a total of 512 addresses; a Peer-to-Peer configuration has a maximum of 256 addresses.
- Two existing VTSs may be used for different workloads, which would have to be combined for the Peer-to-Peer configuration.
- If the two existing VTSs are attached to the same 3494 or TS3500/3953 library, one must be moved to a different library.

In the previous upgrade, where we added a new VTS to an existing VTS, we were able to reuse the library name of the original VTS for the composite and hence avoid any changes to TCDB and DFSMSrmm entries. We can no longer avoid this, so there will be additional tasks to reconcile library names in control records. However, it is still wise to use the library name of one of the two source libraries as the library name for the composite library.

Figure 5-7 illustrates what we will be changing.

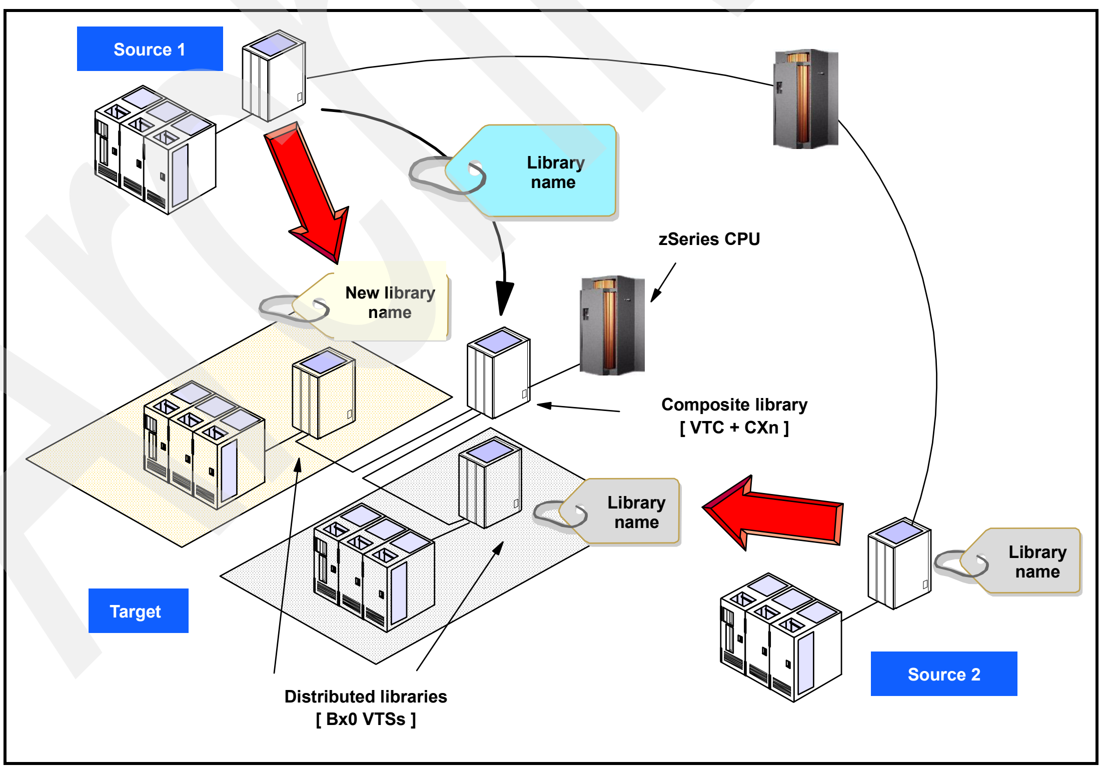

Figure 5-7 Merging two existing VTSs into a Peer-to-Peer VTS

5.3.1 Migration overview

The new hardware consists of:

- Any necessary upgrades to the existing VTSs, so that they have all the prerequisites to be able to install the copy features on each
- The VTCs and CX0 frames, including possible replacement to CX1 frames
- The copy features on each VTS

The migration can and should be broken down into separate stages to minimize outages:

- Planning
- Hardware and software preparation
- System definitions
- Merge

We describe each of these steps in more detail in the following sections.

5.3.2 Planning

The preparation steps can and should be taken well in advance of the delivery of the new hardware. You should decide upon the configuration and collect data for the installation of the StorWatch Specialists. The *IBM 3494 Tape Library Introduction and Planning Guide*, GA32-0448, contains worksheets that you may find helpful in documenting what will be done. After you have completed these worksheets, review them with your service representative. Also, you can find information on the TS3500/3953 in the *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Introduction and Planning Guide*.

See 3.1.2, “VTS configuration requirements” on page 65, and 3.1.14, “TotalStorage Specialist configuration” on page 78, to determine what information you will need.

Merging an existing 3494 VTS with a TS3500/3953 in to a single PtP VTS configuration is supported. However, the 3494 VTS must be a B20. If your 3494 VTS is currently a model B18, you must first upgrade it to a model B20 before merging it with a TS3500/3953 to form a PtP VTS. Refer to the *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring,* SG24-2229 for additional information on upgrading a stand-alone model B18 VTS to a model B20.

Consider what the channel types are, ESCON/FICON, for your two VTSs. In order to merge, they must be the same. If they do not match, you must consider the additional time needed to upgrade to FICON on one of the VTSs.

At this stage, you should also estimate how many logical volumes are defined across the two VTSs that you will be merging. If the total is less than 500,000, you need take no action. If the total exceeds 500,000, you must remove logical volumes. Once you are sure that the combined total of logical volumes is below 500,000, it would be wise not to insert further logical volumes until the merge has been completed to avoid the need to repeat the check.

Since the merged Peer-to-Peer VTS will create copies of logical volumes changed after the merge, you should review the number of scratch stacked 3590 or 3592 volumes and order sufficient extra volumes to insert into the libraries to hold the copies. If one VTS currently has 3590 cartridges and the other has 3592 cartridges, we recommend that, for performance reasons, you consider migrating both libraries to 3592 after the conversion to PtP. Refer to 3.5.3, “Performance considerations for 3592 drives” on page 101 for additional information.

You will need to check the logical volume serial number ranges defined to the two VTSs and resolve any volume serial number duplication. You should also check the stacked volume serial numbers on the two VTSs for any duplication with another stacked volume, as well as, any duplication with the logical volume serial ranges defined to the two VTSs.

5.3.3 Hardware and software preparation

The hardware preparation is performed by your IBM service representatives, who will complete the following set of tasks.

Updating existing VTSs in preparation for a merge

To update the existing 3494 and/or TS3500/3953 VTSs:

1. Move any logical volumes accessed by SCSI hosts to another library
2. Vary existing VTS offline.
3. Remove SCSI host attachment feature if installed.
4. Increase cache size if necessary.
5. Install LAN card in Library Manager if not already installed.
6. Install additional 3590 or 3592 drives if needed.
7. Install the Library Manager licensed internal code supporting Peer-to-Peer.
8. Install the VTS licensed internal code supporting Peer-to-Peer.
9. Install Advanced Policy Management feature (both VTSs must match).
10. Vary VTS online.
11. Re-check for duplicate logical volume serial numbers or duplicate stacked volsers.
12. Configure and enable Web access to the Library Manager.
13. If both VTSs are attached to the same 3494 Tape Library, you must move one of them to a different 3494 library.
14. Build a REXX exec, DFSORT™ job, or something similar to assist in the TCDB updates.

The level of licensed internal code that supports Peer-to-Peer configurations includes tools that your service representative should use to check for duplicate volume serial numbers.

In addition to the tasks done by your service representative, you need to:

15. Plan system definition changes (see 5.3.4, “System definitions” on page 207).
16. Install the DFSMS/MVS support and Toleration code for Peer-to-Peer.
17. Insert additional scratch stacked 3590 or 3592 volumes. Ensure there are no duplicate volume serial numbers across the VTSs before you enter any new stacked volumes.

Installing new hardware

The virtual tape controllers should be installed and configured before the VTSs are merged.

1. Install CX1 frame or frames.
2. Install four or eight VTCs.
3. Configure the VTCs.
4. Install new ESCON/FICON directors if required.
5. Configure and enable Web access to the VTCs.
6. Configure and enable RSF for the VTCs.

5.3.4 System definitions

Changes are needed to the device addresses and to the SMS definitions. You should prepare the changes before you merge the VTSs by taking copies of your existing definitions and making the changes to the copies. You cannot change the SMS definitions, TCDB records, or DFSMSrmm volume records at this stage; this is best done during the physical merge. But you can build the jobs to do the changes in preparation for running them during the merge.

Device addresses

The HCD definitions must be changed. The definitions for one of the libraries should be changed to reflect a configuration of four or eight independent control units, each with 16 addresses, for use as the definitions for the composite library. The other set of addresses can be deleted. Refer to 4.1.3, "Defining devices through HCD panels" on page 140 for HCD definition examples.

SMS definitions

The library name of one of the source VTSs will be used for the composite library so you can reuse one set of library definitions within ISMF with only a minor change to point to the new library ID. An additional benefit of reusing at least one of the library names is that at least some of the TCDB records will not have to be changed. You can reuse the ISMF library definition of the other source library for one of the distributed libraries, and you will need to provide a new set of definitions for the second distributed library. Figure 5-8 shows an example of how the SMS library definitions are changed. While you do reuse the definition of the second VTS for one of the distributed libraries, we did not use the second library's name. In this example, we elect to use a library name of DISTLIB2 instead so that the name itself makes it clear that it is a distributed library.

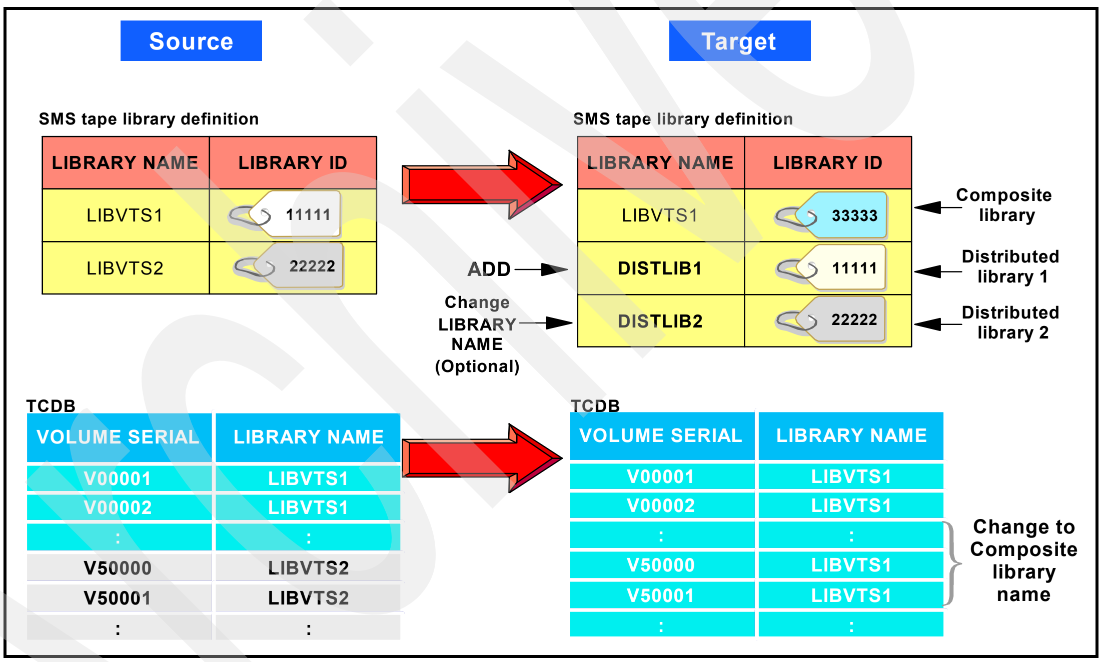

Figure 5-8 Changing definitions for merge

ACS routines

The original configuration started with two independent stand-alone VTSs and SMS tape using tape Storage Groups to direct allocations. You should review the tape Storage Groups and decide whether you want to keep them all while just using the new composite library as the allocation target, or redirect some of the work to other VTSs that you have. If you decide to keep the Storage Groups that were originally defined to the second VTS, LIBVTS2 in our example, then you must update that Storage Group definition to list the library name of LIBVTS1 (instead of LIBVTS2).

After reviewing tape Storage Group assignment to VTSs, you may need to prepare and test changes to your ACS routines. Those changes will then be activated following the completion of the merge.

JES3 definitions

In a JES3 environment, you should modify the INISH deck to balance the load across the virtual tape controllers. You should also remove the device definitions of the library whose name is not becoming the composite library name.

Definition changes

For the definitions described in this section, we assume that your existing VTSs are configured as shown in Table 5-1, installed and in production, and defined as follows.

Table 5-1 Definitions before merge

VTS	Library ID	Volser range	Device addresses
LIBVTS1	11111	V00001-V49999	100-17F
LIBVTS2	22222	V50000-V99999	200-27F

The TCDB contains thousands of volumes in the range V00001 to V99999. Volumes V00001 to V49999 are in LIBVTS1 while V50000 to V99999 are in LIBVTS2, as shown in Table 5-1.

In this example, we elect to reuse library name LIBVTS1 as the composite library name. The library name for the second stand-alone VTS, LIBVTS2, will disappear.

Table 5-2 shows the target configuration.

Table 5-2 Target definitions after merge

VTS	Library ID	Volser range	Device addresses
LIBVTS1	333333	V00000-V99999	100-17F
DISTLIB1	111111	Not applicable	Not applicable
DISTLIB2	222222	Not applicable	Not applicable

Using this example, you will need to complete the following definition changes at the host for the new PtP configuration:

1. Using the HCD panels, change the definitions for addresses 100–17F.
2. Delete the HCD definitions for 200-27F; they will no longer be used
3. Using ISMF, alter the definition of LIBVTS1 to have a new library ID of 33333.
4. Using ISMF, delete the definition for LIBVTS2.
5. In ISMF, define the distributed library DISTLIB1 with library ID 11111.
6. Define the distributed library DISTLIB2 with library ID 22222.

HCD

You must change the HCD definitions for addresses 100–17F so that you have eight control units with LIBPORT IDs from one to eight emulated control units. Ensure each logical control unit (CUADD) has a control unit address of 0. You need a new library ID, and in this example we use 33333. We recommend that you always list the LIBRARY-ID and LIBPORT-ID in your IOCP and HCD. It reduces having to reactivate the IODF when the library is not available at IPL as well as providing enhanced error recovery in some cases.

ISMF

In ISMF, you alter the library definition for LIBVTS1 by changing the library ID from 11111 to 33333.

By doing this, LIBVTS1 will be the composite library. This means that virtual volume entries in the TCDB and DFSMSrmm control data set, which refer to that stand-alone library LIBVTS1, now refer to the composite library and do not need to be changed.

Next, you define a new library (called DISTLIB1 in this example) and use the existing library ID of 11111. As the library ID is unchanged, there is no need for a hardware teach to define the library ID number of 11111.

Finally, you need a definition for the second distributed library. In this case, we define a new library, DISTLIB2, with a library ID of 22222. Again, by keeping the same library ID we avoid a hardware teach.

Figure 5-9 shows the effect of the HCD and ISMF changes.

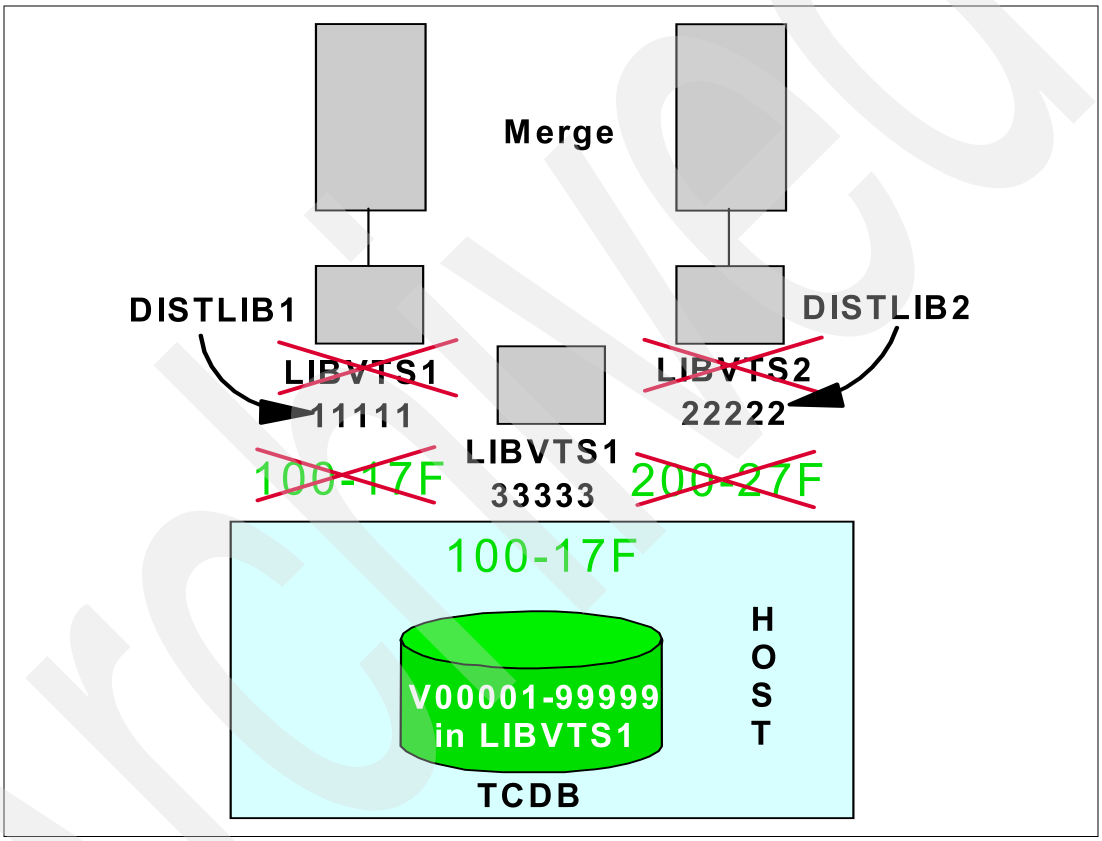

Figure 5-9 Summary of definition changes for merge

5.3.5 Merge the VTSs

The following steps outline the merge of two existing VTSs in to a Peer-to-Peer VTS configuration. The time required for the merge will depend on your current VTS configurations and should be planned and discussed with your IBM service representative well in advance.

Install copy features and merge

For each VTS:

1. Vary VTS offline.
2. Remove host ESCON/FICON cables.
3. Establish the VTS as a Peer-to-Peer distributed VTS.

4. Select the UI library.
5. Define the composite library name.
6. Vary the VTS online.
7. Back up database.
8. Install copy feature.
9. Merge database records from the other VTS being upgraded[2].
   – This step will also confirm that there are no duplicate volume serial numbers. If there are, the merge cannot be completed.
10. If APM is being used, copy the APM definitions from each VTS in to the other:
   – If only one VTS is currently using APM, copy the definitions to the other
   – If both VTSs currently use APM and neither is using the defaults, then you must manually combine the APM definitions.

For example, when you have individual stand-alone VTSs, it is often the case that the ACS routines assign unique Storage Groups for each individual VTS. When you merge them in to a single PtP configuration, you now have all the logical volumes from both libraries in each VTS. Hence, you must account for APM definitions for both Storage Groups in both VTSs. The same is true for any Management Class (MC), Storage Class (SC), and Data Class (DC) definitions.

Complete the physical merge

Next, complete the physical merge by performing these tasks:

1. Verify that both Library Managers have the same category attributes defined.
2. Cable VTCs to VTS in UI library.
3. Cable VTCs to other VTS.
4. Synchronize token databases.
5. Cable VTCs to host channels.

While the configuration is unavailable during the merge, you can correct volume references in the TCDB and, if applicable, the DFSMSrmm control data set. These must all refer to the composite library name.

TCDB changes

The virtual volume TCDB entries for volumes in LIBVTS2 must be changed to refer to the composite library, LIBVTS1. To do so, first issue the IDCAMS command:

```
LISTC VOLUMEENTRIES(V*) LIBRARY(LIBVTS2)
```

This will list all the TCDB volume entries referring to the old library name, LIBVTS2. At this stage, ensure that you have a valid backup of your TCDB in the event of any problems with the updates that you will do next.

For each volume, issue this IDCAMS command to correct the library name to refer to the composite library:

```
ALTER volser VOLUMEENTRY LIBRARYNAME(LIBVTS1)
```

You may choose to use a REXX exec or DFSORT to build a list of commands that you can review before submitting. A sample REXX exec is included in 5.3.7, “REXX exec for altering TCDB volume entry library names” on page 213.

[2] If the VTSs are physically remote, consider how to exchange the database backups quickly.

To verify your TCDB changes, you can either display the library's mountable tape volume list through ISMF or, individual volume displays can be done through OAM with the command, **D SMS,VOLUME(volser)** as shown in Figure 6-18 on page 246.

Updating your tape management system

We are not aware of any tape management systems, other than DFSMSrmm, that hold library names in control data set records. At this stage, for DFSMSrmm, you will need to update the DFSMSrmm volume entries. It would be wise to ensure that you have a current backup of both the DFSMSrmm control data set and the journal just in case of any problems.

Issue the following command:

```
RMM SEARCHVOLUME VOLUME(*)
   OWNER(*) LIMIT(*) -
   LOCATION(LIBVTS2) -
   CLIST('RMM CHANGEVOLUME ',' LOCATION(LIBVTS1) HOME(LIBVTS1)')
```

This will build a data set containing a list of DFSMSrmm **changevolume** commands. After ensuring that you have a valid backup of the DFSMSrmm CDS, you can execute this list of commands by issuing:

```
EXEC hlq.EXEC.RMM.CLIST
```

Where *hlq* is the High Level Qualifier.

This will change both the location and the home location name in the DFSMSrmm volume records to refer to the composite library name.

To verify your RMM changes, run an RMM volume report, or to display individual volumes, you may also use the RMM "LISTVOLUME" command. Refer to the *DFSMSrmm Reporting*, SC26-7406 manual for additional information on generating and viewing RMM volume reports.

5.3.6 Activating the definitions

When the upgrade has been completed, you can activate the new definitions.

Use the **ACTIVATE IODF** operator command to change the IOCDS and use the SETSMS SCDS operator command to activate the SMS definitions. (Refer to "Activating the HCD definitions" on page 204 and "Activating the SMS definitions" on page 204).

To verify the Device Services library control blocks built as a result of the activation, use the "QUERY LIBRARY" command. This function has multiple options and helps you view the existing library control blocks, as well as assisting in correcting any problems with the current library control blocks. Refer to Appendix C, " on page 407 for additional information and details on this function.

5.3.7 REXX exec for altering TCDB volume entry library names

The REXX exec in Figure 5-10 shows how you might correct the TCDB entries for one of the VTSs merged to form a Peer-to-Peer VTS. You would use this to alter the TCDB entries for the library whose library name is *not* reused for the Peer-to-Peer composite library.

```
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
```

Figure 5-10 ALTERVOL exec

5.4 Relocating an existing PtP VTS

In this section, we give you a step-by-step approach to moving an existing IBM Peer-to-Peer VTS configuration (3494 and TS3500/3953) to a new location with matching IBM hardware. This phased procedure is intended to minimize the impact of the migration, reusing the previous VTS attributes, and definitions as much as possible to reduce migration outages and impact on your operations.

5.4.1 Planning

As the source configuration in this scenario is already a PtP VTS, this migration will involve less time and fewer changes to both the hardware and the software host than were seen in the last section, 5.3, “Merging two existing VTSs into a Peer-to-Peer VTS” on page 205. The source PtP configuration allows host access to the data from the first VTS; at the same time, the second VTS and cartridges are being moved to the target location. This dramatically reduces your host outage time. As the target configuration is the same, all cartridges can be reused.

The preparation steps can and should be taken well in advance of the delivery of the new hardware. Migration of the existing PtP subsystem ***must*** be performed by your IBM service representatives. It involves:

- Installing a new PtP at the target location (with at least one TS3500/3953 or 3494, VTS, and two AX0)
- Target host HCD and software definitions/changes assuming that the target is a new host or the library name or library ID has changed:
  - *Remember*, if you change the composite library name, you will have additional changes required to update the host TCDB and DFSMSrmm volume records before you can bring production up at the target site. We discuss this in more detail in 5.3, “Merging two existing VTSs into a Peer-to-Peer VTS” on page 205.
- Copying of one of the source LM and VTS databases.
- Planning and special handling for the physical move of the stacked cartridges.
- Physical move of one of the VTSs.

You must work closely with your IBM service representatives so that you can plan for the appropriate outage of your tape library and the VTS.

Available PFE Services

There is a billable service, called *Two-Step Data Migration*, available from Hardware Product Field Engineering (PFE) to assist you in the planning and implementation of the moving of a PtP VTS to a new location. The migration steps discussed in this section are based on that service and satisfy the following requirements:

- Fast migration, determined by the time it takes to transport the cartridges to the new location
- Safe migration, through separate transfer of first and second VTS cartridges
- Fast and safe fallback, ensured through the following actions:
  - Single source VTS remains read accessible until the checkpoint at the target location
  - Checkpoint built in to test data accessibility from the target location before disconnecting and moving the second VTS
  - Ability to fall back to the source VTS in case of a problem at the target location
- Ability to start production with only one VTS available on the PtP through:
  - Read-only mode initially while the LM and VTS data bases and cartridges are being moved to the second site.
  - After that, Read/Write Disconnect mode (at target site) until the second VTS is moved.

Attention: We highly recommend that your service representative inquire with PFE regarding this “Two-Step PtP Data Migration” service.

5.4.2 Migration steps

Figure 5-11 shows a diagram of the Two-Step Data Migration service offering.

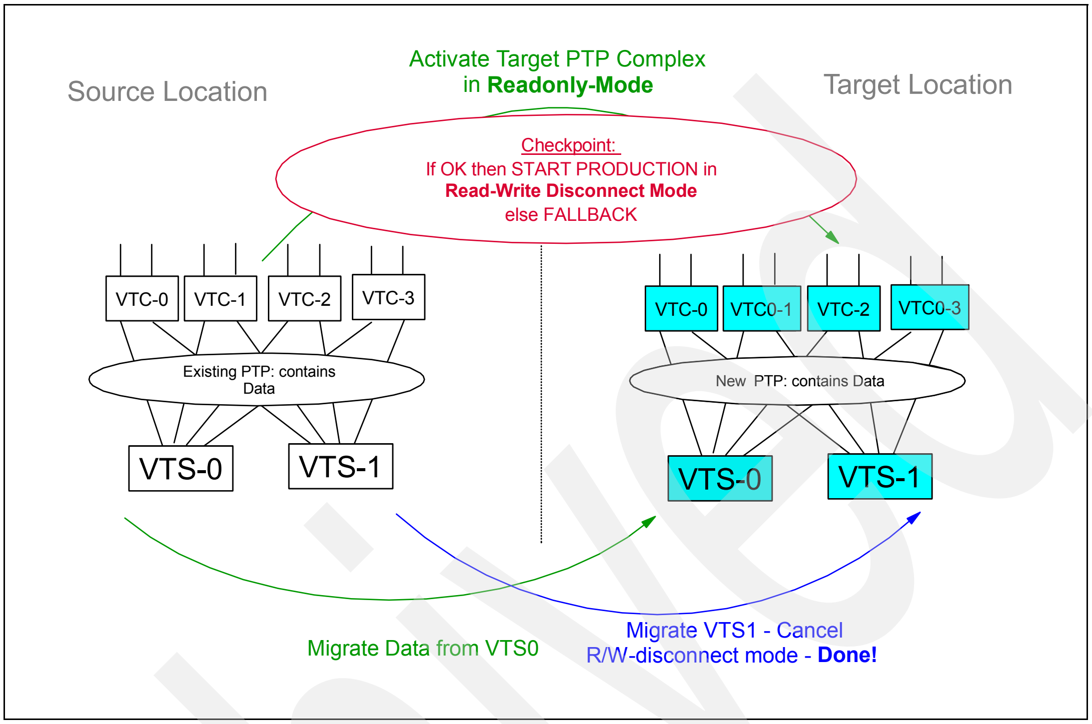

Figure 5-11 Two Step Data Migration

The following list guides you through the hardware and software implementation steps required to migrate your PtP VTS:

1. Install and configure a PtP VTS complex at the target site:

   – Include at least one TS3500 (with 3953) or 3494 Tape Library, one VTS, and two VTCs (tested and empty).

   – Place the source VTS that will be running production at the source site, while the second VTSs data is being migrated, in to read-only mode. Production cannot be allowed to write until after the checkpoint is completed at the target site, and the target VTS is brought up in Read/Write Disconnect mode. Production would then be allowed to continue, read, and write, on the single target VTS until the second VTS move is completed.

   > **Important:** The single VTS that remains running at the source site (while the second VTS is having its databases and physical cartridges moved to the target site) ***must*** be placed in read-only mode. If writes were allowed at the source site, before the switchover to run production at the target site's single VTS, then those additional volumes that had been written at the source site would not be available to the target host. You would have to wait until the second VTS was moved and brought online at the target location before accessing those volumes.

   – Migrate data and cartridges from *one* source VTS to the newly installed target VTS. You must move the source TS3500/3953 to the target TS3500/3953. (Migration from a VTS attached to a TS3500/3953 to a VTS attached to a 3494 is not currently supported.) Hence, if the new VTS installed at the target site is attached to a TS3500 with 3953, then migrate the data and cartridges from the source TS3500 first. If the new target VTS is 3494, then migrate from the source 3494 first.

This single VTS migration will include:

- LM and VTS databases (including APM if applicable)
- Physical stacked cartridges following manufacturer's media transport and handling recommendations

Refer to the *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring,* SG24-2229 for the complete VTS Data Migration procedure.

– Ensure that the appropriate host definitions and configurations for the new PtP VTS are completed, including:

  - HCD definitions (see 4.1.3, "Defining devices through HCD panels" on page 140)
  - ISMF library definition panels are created for both the composite library and distributed libraries

  Important: The library ID fields defined in the SMS Tape Library Define panels, as well as in the HCD LIBRARY-ID parameters, must ***exactly*** match the library sequence numbers defined to the individual VTS libraries by your hardware service representative.

  - Define/update the SMS constructs, especially Storage Groups
  - Prepare/update ACS routines
  - Ensure appropriate activates done to connect the device to the target host system
  - If applicable, update Jes3 definitions

  Refer to 4.5.1, "z/OS and DFSMS/MVS SMS-managed tape" on page 184 for additional information on defining the composite and deferred libraries to DFSMS and Jes3.

  As this example opted to use an identical composite library name at the target site, no TCDB or RMM updates are required.

2. Bring up the target PtP VTS with one VTS in *Read-Only* mode. Never allow both the source and the target site to write at the same time or data inconsistency will occur!!

   Attention: Initially bringing the target PtP up in read-only mode is important to prevent writes to the target site until the PtP is fully operational. If production is started at the target location before the checkpoint, there is no safe fallback to the source VTS.

3. Checkpoint:

   – Vary the drives online to z/OS and bring the library online to the target host

   – Verify data can be read from the target host

   – Verify Device Services library control blocks using "QUERY LIBRARY" command[3]

   – Checkpoint determines if fallback is required or not.

   – If no fallback is needed, then:

     - Stop production at the source location, vary the remaining source VTS offline, etc.

     - Set readwrite disconnect mode at the target site. Production read and write can begin from the target location with access to a single VTS.

   – If fallback is required, just continue to work at the source location with one VTS

[3] Refer to Appendix C, " on page 407 for additional information and examples on how to use this function.

- Move the second source VTS to the target location. This includes the physical move of the VTS and all of the stacked cartridges. While the second VTS is being moved, the target location will be running on a single VTS. It is important that this second VTS move be done ASAP to limit the risk of a single point of failure.
- Once the second VTS is at the target site and installed, start normal PtP operations.

5.5 Migrating a 3494 PtP VTS to a new TS3500/3953 PtP VTS

As in the previous section, 5.4, "Relocating an existing PtP VTS" on page 213, we give you a step-by-step approach to moving an existing IBM Peer-to-Peer VTS configuration (3494 in this case) to a new location. However, in this example we will also be upgrading the hardware to an all TS3500/3953 PtP VTS at the same time. Only the LM and VTS databases and the physical cartridges will be moved into the new TS3500/3953 PtP VTS with this procedure. The VTS hardware itself will not be moved. This procedure of course assumes that the source 3494 PtP contains 3592 drives, as the targeted TS3500/3953 only supports 3592 cartridges.

This phased procedure is intended to minimize the impact of the migration, reusing the previous VTS attributes, and definitions as much as possible to reduce migration outages and impact on your operations. Again, as the source configuration is already a PtP VTS, the host continues to have read access to the data from the first VTS, at the same time the library and VTS databases and cartridges of the second VTS are being moved. The "write" outage time is limited to the movement of the first VTS's databases and physical cartridges.

Migration of the existing PtP subsystem *must* be performed by your IBM service representatives. It involves both the copying of each VTS and LM database in to the new VTSs, as well as the physical movement of the stacked 3592 cartridges. You must work closely with your IBM service representatives so that you can plan for the appropriate outage of your tape library and the VTS.

5.5.1 Planning

As the source configuration in this scenario is already a PtP VTS, and there is no physical movement of the VTS itself (only databases and cartridges), the migration time of the second VTS will take less time than it did in the last section, 5.4, "Relocating an existing PtP VTS" on page 213. The source PtP configuration allows host read access to the data from the first VTS, at the same time, the second VTS and cartridges are being moved to the target location. This dramatically reduces your host outage time. As the target configuration is all TS3500/3953, the source 3494 configuration must already be using 3592, or have plans to upgrade and migrate to 3592, before this move is attempted. The TS3500/3953 does not support 3590 drives/cartridges.

The preparation steps can and should be taken well in advance of the delivery of the new hardware. Migration of the existing PtP subsystem *must* be performed by your IBM service representatives. It involves:

- Install of new PtP at the target location (with two TS3500 Tape Libraries with an IBM 3953 Tape System each, VTS, and four or eight VTCs)

- Target host HCD and software definitions/changes assuming that the target is a new host or the library name or library ID has changed
  - **Remember**, if you change the composite library name, you will have additional changes required to update the host TCDB and DFSMSrmm volume records before you can bring production up at the target site. We discuss this in more detail in 5.3, "Merging two existing VTSs into a Peer-to-Peer VTS" on page 205
- Copying of each of the source LM and VTS databases (one initially, the second later)
- Planning and special handling for the physical move of the stacked cartridges.

You must work closely with your IBM service representatives so that you can plan for the appropriate outage of your tape library and the VTS.

Available PFE Services

The billable service, "Two Step PtP Data Migration", available from Hardware Product Field Engineering (PFE), which we described in "Available PFE Services" on page 214, can also assist you in the planning and implementation of the moving of a PtP VTS to a new location and the upgrade to TS3500/3953 PtP VTS. In this case, only the LM and VTS databases and physical cartridges are being moved. Hence, the time in Read/Write Disconnect mode is limited to the second library's stacked tape transport time.

5.5.2 Migration steps

The hardware and software implementation steps required are:

1. Install and configure a PtP VTS complex at the target site
   - Including two TS3500/3953 tape library configurations, two VTS, and four or eight VTCs (tested and empty)
   - Place the source VTS that will be running production at the source site, while the second VTSs data is being migrated, in to read-only mode. Production cannot be allowed to write until after the checkpoint is completed at the target site, and the target VTS is brought up in Read/Write Disconnect mode. Production would then be allowed to continue, read and write, on the single target VTS until the second VTS move is completed.

     > **Important:** The single VTS that remains running at the source site (while the second VTS is having its databases and physical cartridges moved to the target site) **MUST** be placed in read-only mode. If writes were allowed at the source site, before the switchover to run production at the target site's single VTS, then those additional volumes that had been written at the source site would not be available to the target host. You would have to wait until the second VTS was moved and brought online at the target location before accessing those volumes.

   - Migrate data and cartridges from <u>one</u> source VTS to the newly installed target VTS. You must move the source TS3500/3953 to the target TS3500/3953. (Migration from a TS3500/3953 VTS to a 3494 VTS is not currently supported.) Hence, if the new VTS installed at the target site is 3953/3584, then migrate the data and cartridges from the source 3953/3584 first. If the new target VTS is 3494, then migrate from the source 3494 first. This single VTS migration will include:
     - LM and VTS databases (including APM if applicable)
     - Physical stacked cartridges following manufacturer's media transport and handling recommendations

Refer to the *IBM TotalStorage Virtual Tape Server Planning, Implementing, and Monitoring,* SG24-2229 for the complete VTS Data Migration procedure.

- Ensure that the appropriate host definitions and configurations for the new PtP VTS are completed:
  - Ensure that HCD definitions are done (see 4.1.3, “Defining devices through HCD panels” on page 140).
  - Ensure that ISMF library definition panels are created for both the composite library and distributed libraries.

  > **Important:** The library ID fields defined in the SMS Tape Library Define panels, as well as in the HCD LIBRARY-ID parameters, must ***exactly*** match the library sequence numbers defined to the individual VTS libraries by your hardware service representative.

  - Define/update the SMS constructs, especially Storage Groups.
  - Prepare/update ACS routines.
  - Ensure that appropriate activates are done to connect the device to the target host system.
  - If applicable, update Jes3 definitions.

  Refer to 4.5.1, “z/OS and DFSMS/MVS SMS-managed tape” on page 184 for additional information on defining the composite and deferred libraries to DFSMS and Jes3.

  As this example opted to use an identical composite library name at the target site, no TCDB or RMM updates are required.

2. Bring up the target PtP VTS with one VTS in *read-only* mode. Never allow both the source and the target site to write at the same time or data inconsistency will occur!!

   > **Attention:** Initially bringing the target PtP up in read-only mode is important to prevent writes to the target site until the PtP is fully operational. If production is started at the target location before the checkpoint, there is no safe fallback to the source VTS.

3. You have now reached a checkpoint and should perform the following tasks::
   - Vary the drives online to z/OS and bring the library online to the target host.
   - Verify that data can be read from the target host.
   - Verify Device Services library control blocks using “QUERY LIBRARY” command. Refer to Appendix C, ” on page 407 for additional information and examples on how to use this function.
   - Checkpoint determines if fallback is required or not.
   - If no fallback is required, then:
     - Stop production at the source location, vary the remaining source VTS offline, etc.
     - Set Read/Write Disconnect mode at the target site. Production read and write can begin from the target location with access to a single VTS.
   - If fallback is required, just continue to work at the source location with one VTS.

- Move the second source VTS to the target location. This includes only the physical move of the stacked cartridges and copy of the LM and VTS databases. While the second VTS's data is being migrated, the target location will be running on a single VTS. It is important that this second VTS migration be done ASAP to limit the risk of a single point of failure.
- Once the second VTS's physical (stacked) cartridges are at the target site, and the LM and VTS databases copied, start normal PtP operations.

6

Operations

In this chapter, we describe operational considerations and usage guidelines unique to the IBM Peer-to-Peer VTS. For general guidance on how to operate the IBM 3494 Tape Library, the TS3500/3953 Tape Library, and the IBM 3494 Virtual Tape Server, refer to:

- *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-00449
- *IBM TotalStorage 3584 Tape Library for zSeries Hosts,* SG24- 6789
- *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring,* SG24-2229
- *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632
- *z/OS Object Access Method Planning, Installation and Storage Administration Guide for Tape Libraries*, SC35-0427

6.1 Overview

For successful operation of your Peer-to-Peer VTS, it is important that you understand its concepts and its components. They are described in detail in Chapter 2, “Peer-to-Peer VTS Architecture” on page 7. In this chapter, we combine the components and functions of the Peer-to-Peer VTS 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 256 virtual addresses for tape mounts. A composite library can run in immediate copy mode or in deferred copy mode. The host view is aware of the existence of the physical libraries. The term *distributed library* is used to denote the physical libraries that are part of the Peer-to-Peer VTS.

The physical view is directly related to the hardware components and their functions. The hardware is clearly part of the physical view: the VTCs, the channel links, the VTSs, and the 3494 or TS3500/Tape Library with 3953 Tape System. Also, the Master VTS and the I/O VTS can be considered elements of the physical view. The preferred VTS for I/O is in this view too.

The UI distributed library is part of the physical view. The host view currently does not know about the UI library, which is the library you must go to when defining new logical volumes.

When you operate the Peer-to-Peer VTS, you deal with one view at a time. There are three operator interfaces for providing information about the Peer-to-Peer VTS.

- The Library Manager panels deal with information in the Library Manager database. This is part of the physical view. There is limited information available on the VTSs in the 3494 or 3953 Library Manager. 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 TotalStorage 3494 Tape Library Operator’s Guide*, GA32-00449 or the *IBM TotalStorage 3953 Library Manager Model L05 Operator Guide,* GA32-0473.
- OAM commands are available from the host operator console for information regarding the Peer-to-Peer composite and distributed libraries. This information represents the host view of the components within the Peer-to-Peer VTS. Other z/OS commands can be used against the virtual addresses. These commands are not aware that the 3490E addresses are part of a Peer-to-Peer VTS configuration.

> **Note:** The “Specialist” is a server running on the specific hardware you want to interface with. For example, the ETL Specialist is the Web-based server running on the 3494 Library Manager or on the IBM 3953 Library Manager. The Peer-to-Peer Specialist is the Web-based server running on the VTC. The TS3500 Specialist is used for IBM TS3500 Library-specific functions such as defining Cartridge Entry Policy.

- ETL Specialist functions are available via a Web-based user interface. There are four Specialists available for tape library management. There are three Tape Library Specialists. The 3494, the 3953, and the TS3500 Specialist server enable you to monitor the Library Manager of the 3494 and 3953, plus any attached VTSs. The TS3500 Library Specialist allows for management (configuration and status) of the TS3500 Library. The Peer-to-Peer VTS Specialist is used for the Peer-to-Peer VTS and runs in the VTC. Because a Web interface is provided, they can be used from any suitable Web browser allowed to use the IP addresses. Panels display information related to the physical view. For Specialist installation requirements, see “Peer-to-Peer VTS Specialist” on page 79.

6.2 Hardware

The VTC and the Model CX1 frame are new hardware components in a Peer-to-Peer VTS compared to a standalone VTS.

6.2.1 VTC control and indicators

The only control directly accessible on a VTC or AX0 is the power switch. The AX0 model also includes a warm boot button which functions as in any System p server operation. These affect only the operation of the particular VTC involved and will not impact the other VTCs. Service access to each VTC is provided by modem or ASCII terminal connection to the com port of the VTC. This function is not offered as a user interface. Service access is also provided by the service LAN via the TotalStorage Master Console. Each VTC provides an additional Ethernet connection that provides a Total Storage PTP Specialist user interface to the VTC and the PTP VTS as a whole. The following section provides further details.

6.2.2 Model CXn control and indicators

The Model CX0 or CX1 frame has a Unit Emergency switch, which removes power only to its frame and installed VTCs. The Unit Emergency switch on the Model CX0 or CX1 frame should be set to the OFF (O) position only in an emergency and must be set to the ON (I) position to allow the installed VTCs to be powered on manually.

The 3494 or 3953/TS3500 Tape Library does not control power on and power off of the Model CX0 or CX1 remotely. Each VTC has a power switch under the media bay cover.

6.3 Monitoring the Peer-to-Peer VTS

As explained in the introduction, there are three interfaces for querying information on the Peer-to-Peer VTS.

- The Library Manager panels deal with information about the Library Manager, 3494 Tape Library, and the attached VTSs.
- OAM commands are available from the host operator console for information regarding the Peer-to-Peer composite and distributed libraries.
- Specialist display functions are available via a Web-based user interface.

Note: The Library Manager in the TS3500/3953 library configuration does not control or manage the TS3500 library.

6.3.1 ETL Specialists

There are four new Specialists now available for the tape library. The ETL Specialist enables you to monitor the Library Manager, the VTS, the 3494 and the TS3500 library from a remote location. The Peer-to-Peer VTS Specialist is used for the Peer-to-Peer VTS. You activate the Specialist for your Peer-to-Peer VTS by entering the correct IP address on your Web browser.

The ETL Specialist home panel directs you to all the Specialists panels available for the 3494 or TS3500/3953 Tape Library. Depending on the configuration of your system, you can select one of the following items from the home panel:

- Library Manager

- VTS
- Virtual tape controller (the VTCs); this leads you to the Peer-to-Peer VTS panels
- TS3500 Tape Library

From each main panel, you are directed to panels with more detailed information. A sample of the Library Manager Home panel is shown in Figure 6-1 and Figure 6-2, followed by the System Summary of the Library Manager in Figure 6-3.

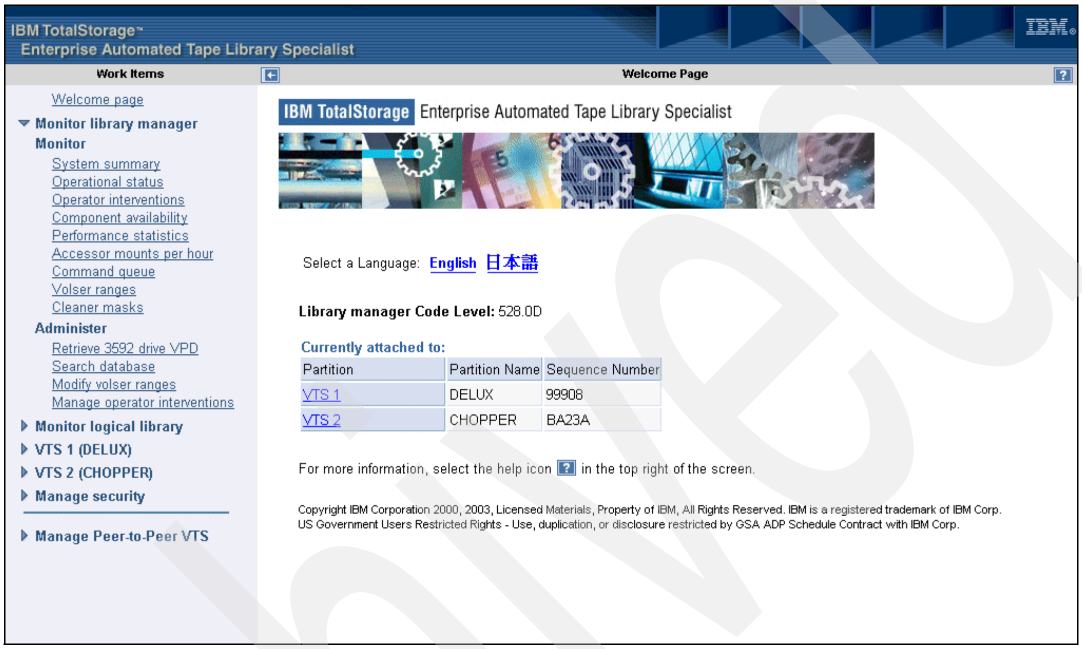

Figure 6-1 Enterprise Tape Library Specialist panel from the 3494 Library Manager

Figure 6-2 Enterprise Tape Library Specialist panel of the 3953 LIbrary Manager

Library Manager

The following detailed panels are available for the Library Manager of the 3494 or 3953. (Refer to Figure 6-1 for the 3494, Figure 6-2 for the 3953, and Figure 6-3 for a summary.)

These panelsinclude:

- System summary
- Operational status
- Operator interventions
- Component availability (Device availability on the 3953)
- Performance characteristics
- Command queue
- LAN hosts status
- LAN information
- Dual accessor zones
- Volser ranges
- Cleaner mask (not supported on the Library Manager Specialist of 3953)

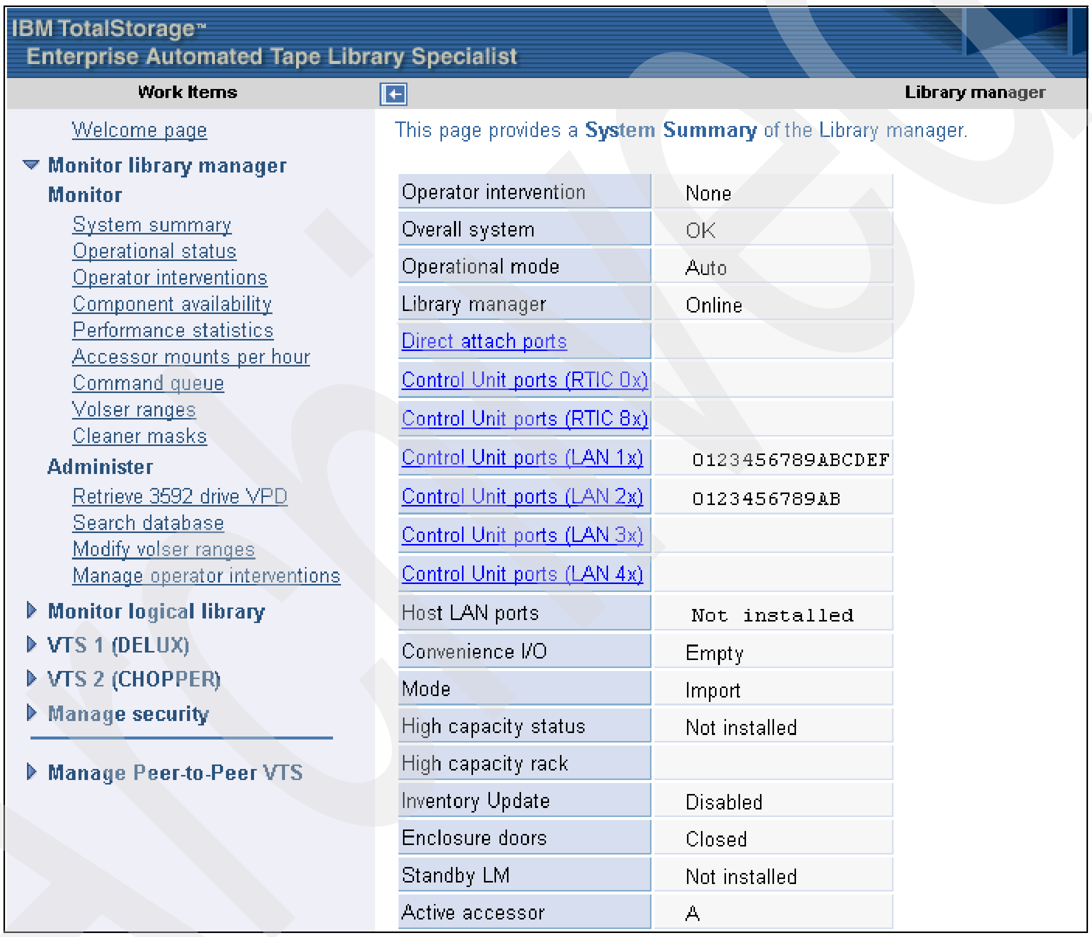

Figure 6-3 Tape Specialist Library Manager system summary

VTS

The following detailed panels (Figure 6-4) are available for each VTS in your configuration:

- Active data
- Active data distribution
- Data flow
- Logical mounts per hour
- Mount hit rate
- Physical device mount history
- Category attributes
- Management policies

- Status
- Volser ranges
- Real time statistics
- Find Logical Volume
- Manage Logical Volumes
- Request Stacked Volume Map
- Move/eject stacked volumes
- Security
- With Advanced Policy Management, you also receive:
  - Administration
  - Manage constructs
  - Select storage pool properties

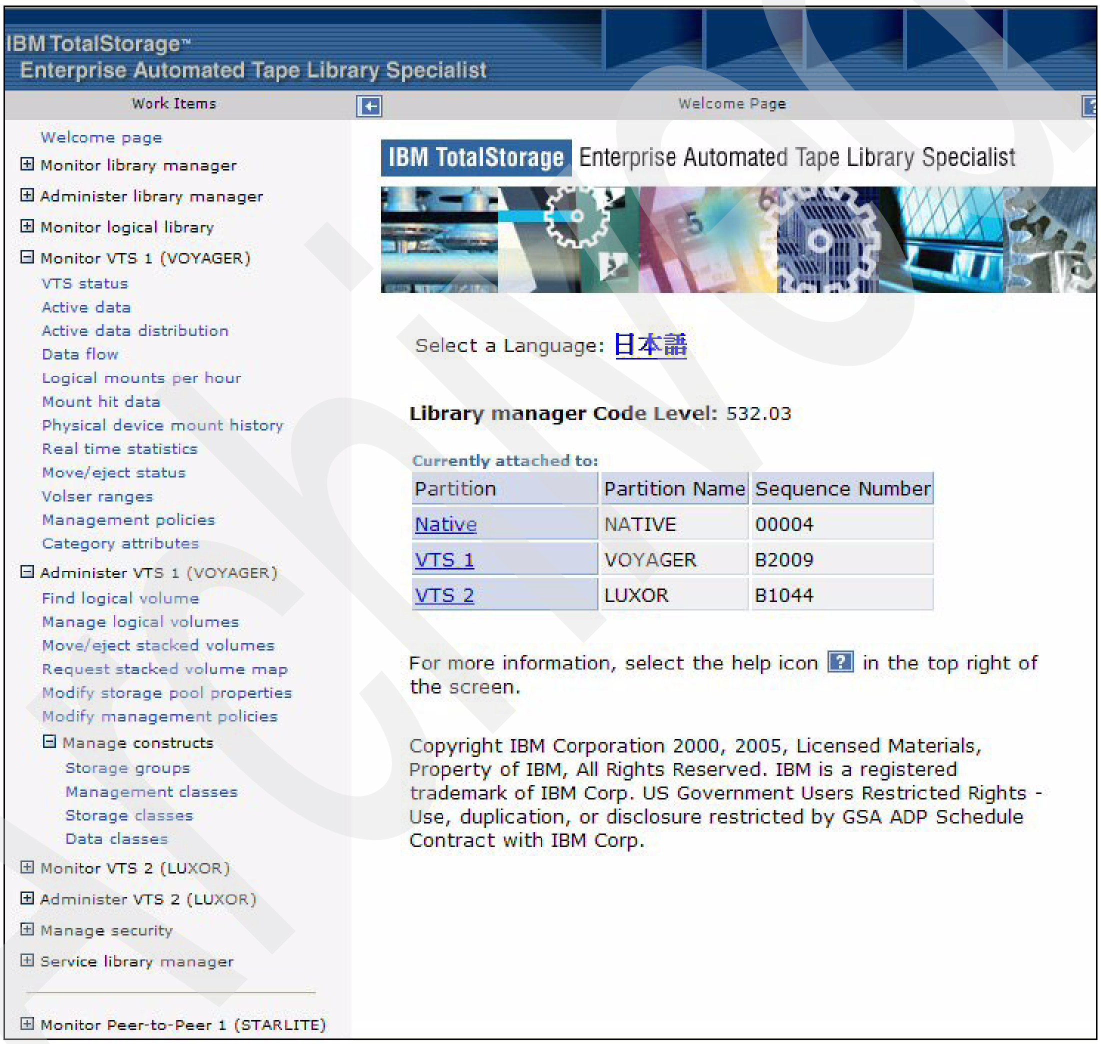

Figure 6-4 ETL Specialist VTS selection

Virtual tape controllers

These panels deal with the physical view of the Peer-to-Peer VTS. First, you must select the Peer-to-Peer VTS by name; this is the name of the composite library as defined to DFSMS/MVS using ISMF. The initial panel shows a list of the VTCs. You can then select one of them. They all present the same information because each controller maintains the complete view of the composite library. So, if a VTC is down or in service mode, you can select one of the remaining VTCs.

The following detailed panels are available for the Peer-to-Peer VTS:

- Home page
- System status
- System configuration
- Current device activity
- Logical volume status
- Current copy workload

6.3.2 Peer-to-Peer VTS configuration

You can control the hardware configuration and the installed features of the distributed libraries and the composite library in various ways. which we explain in this section.

Library Manager

On the VTS Status panel (Figure 6-5) of the Library Manager, information is added to show you if a VTS in the Tape Library is part of a Peer-to-Peer VTS. A VTS is Peer-to-Peer capable if FC4010 is installed on the VTS and it is configured as such on the Library Manager. The VTS name and number are presented on this panel.

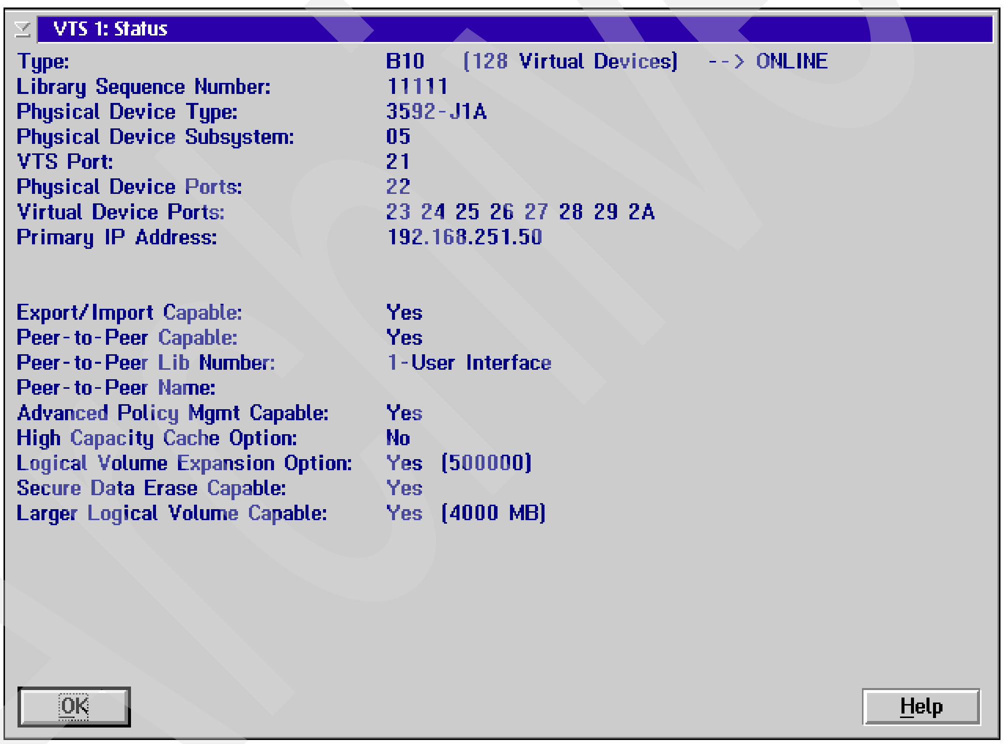

Figure 6-5 VTS status with VTS LIC Release 7.4

Host commands

From the host operator console, the following commands are enhanced to present information related to the Peer-to-Peer VTS:

```
DISPLAY SMS,OAM
```

The console message from D SMS,OAM now displays the number of composite and distributed libraries.

DISPLAY SMS,LIBRARY(*libname*),DETAIL

The messages produced by D SMS,LIBRARY have changed as follows:

- For a VTS composite library, `VCL` is displayed in the LIB TYP field.
- For a VTS distributed library, `VDL` is displayed in the LIB TYP field.
- Peer-to-Peer VTS library association is displayed. The display of the composite library provides the names of the distributed libraries, and the display of the distributed library provides the name of the related composite library.
- The display of the distributed library provides information on the private stacked volume count in addition to the already provided scratch stacked volume count.
- The following new operational states are displayed:
  - Copy operations disabled
  - VTS operations degraded
  - Immediate mode copy operations deferred
  - Service preparation occurring

An example of the display for a composite library is shown in Figure 6-6. The library type is a VCL (VTS Composite Library). There are logical volumes, but no stacked volumes and no slots are associated with the composite library. Note the distributed libraries information line and the status line. Also note:

- Scratch category = current defined scratch categories used by this host system
- Library support outboard policy = this is the host terminology of APM, in other words - this composite library is APM capable.

```
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY  SCRTCH  ON  OP
LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS   VOLS
ATLBA041  VCL  3494-L10   64   16    4      0      0    969    Y   Y
------------------------------------------------------------------
MEDIA        SCRATCH      SCRATCH     SCRATCH
TYPE           COUNT    THRESHOLD    CATEGORY
MEDIA1             0           10        0011
MEDIA2           969           10        0012
------------------------------------------------------------------
DISTRIBUTED LIBRARIES:  ATL70680 ATL70677
------------------------------------------------------------------
OPERATIONAL STATE:  AUTOMATED
ERROR CATEGORY SCRATCH COUNT:                0
------------------------------------------------------------------
Immediate Mode Copy operations deferred.
VTS operations degraded.
Bulk input/output not configured.
Library supports outboard policy management.
```

Figure 6-6 Composite library display

An example of a display of a distributed library is shown in Figure 6-7. The library type is a VTS distributed library (VDL). There are no drives or logical scratch volumes associated with a distributed library, but the display does show stacked volumes and cartridge slots.

```
TAPE       LIB  DEVICE     TOT  ONL  AVL  TOTAL  EMPTY SCRTCH  ON OP
LIBRARY    TYP  TYPE       DRV  DRV  DRV  SLOTS  SLOTS   VOLS
ATL70680   VDL  3494-L10     0    0    0    413    236      0  Y  Y
--------------------------------------------------------------------
COMPOSITE LIBRARY:      ATLBA041
--------------------------------------------------------------------
OPERATIONAL STATE:  AUTOMATED
SCRATCH STACKED VOLUME COUNT:                 95
PRIVATE STACKED VOLUME COUNT:                 81
--------------------------------------------------------------------
Bulk input/output not configured.
```

Figure 6-7 Distributed library display

6.3.3 Composite library

The composite library reflects the overall state of the Peer-to-Peer VTS. For some information it combines the states of the distributed libraries to generate the state of the composite library.

State condition

If the composite library is in a certain state based on a specific state of the distributed libraries, you must look at the status of the underlying distributed library to determine which tape library caused the composite library state condition.

Table 6-1 shows the relationship between the composite and the distributed states.

Table 6-1 Relationship between composite and distributed states

Condition	Composite value	Comments
Automated mode	If *any* distributed library is in that mode.	—
Paused mode	If *all* distributed libraries are in that mode.	—
Manual mode	If *all* distributed libraries are in that mode or if one is in manual mode and the rest are in paused mode.	—
Degraded operation	If *any* distributed library is in this state.	This indicates that some part of the Peer-to-Peer VTS is not functioning properly.
Vision system non-operational	If *any* distributed libraries are in this state.	The barcode scanner has failed.
Offline	If *all* distributed libraries are in this state.	The library is offline.
Library Manager check 1 condition	If *all* distributed libraries are in this state.	The Library Manager has entered the check 1 recovery state.
All storage cells full	If *all* distributed libraries are in this state.	Each storage cell has a volume assigned to it. No more physical inserts are allowed.

Condition	Composite value	Comments
Out of cleaner volume	If *all* distributed libraries are in this state.	There are no cleaner cartridges available.
Dual write disabled	If *all* distributed libraries are in this state.	The Library Manager with the 2nd hard drive feature or an HA1 configuration is not keeping the secondary database synchronized.
Environmental alert	Composite library is never in this state.	Smoke is detected in the library enclosure.
Insufficient resources for mounts	If *all* distributed libraries are in this state.	The Peer-to-Peer VTS is out of usable physical drives.
Library Manager switchover in progress	If *all* distributed libraries are in this state.	The Library Managers in an HA1 environment are in the process of switching states.
Out of empty stack volume	If *all* distributed libraries are in this state.	The Peer-to-Peer VTS has no more available physical stacked volumes to be used to copy data to (operations are halted).
Library Manager switchover in progress	If *all* distributed libraries are in this state.	The Library Managers in an HA1 environment are in the process of switching states.

Figure 6-8, Figure 6-9, and Figure 6-10 show examples of the detailed system status of the Peer-to-Peer VTS and VTC. A number of the Peer-to-Peer VTS conditions are presented in the library part of the screen.

Figure 6-8 Library status screen

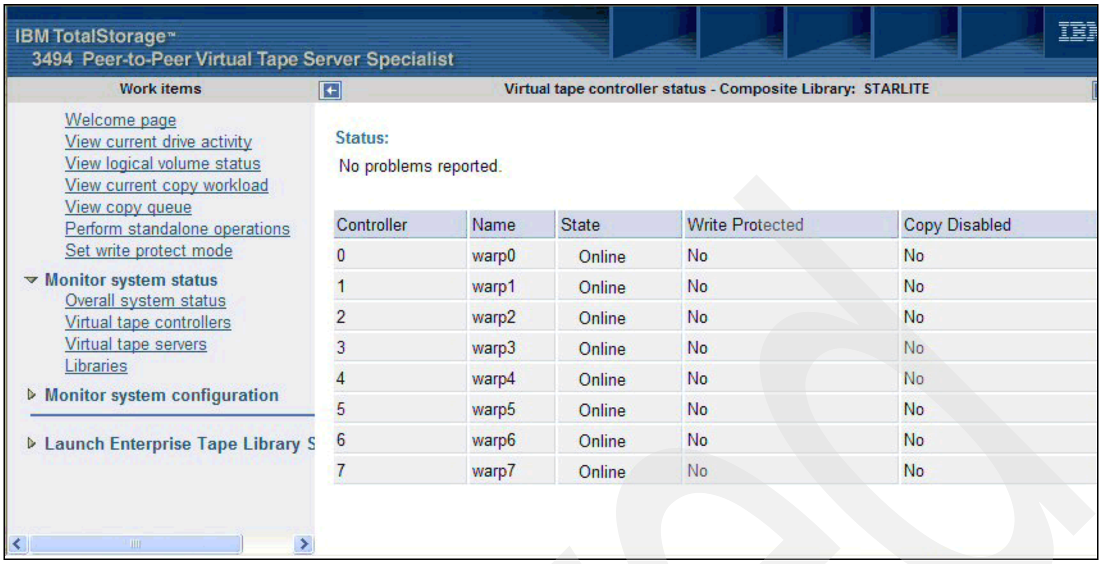

Figure 6-9 VTC Status screen

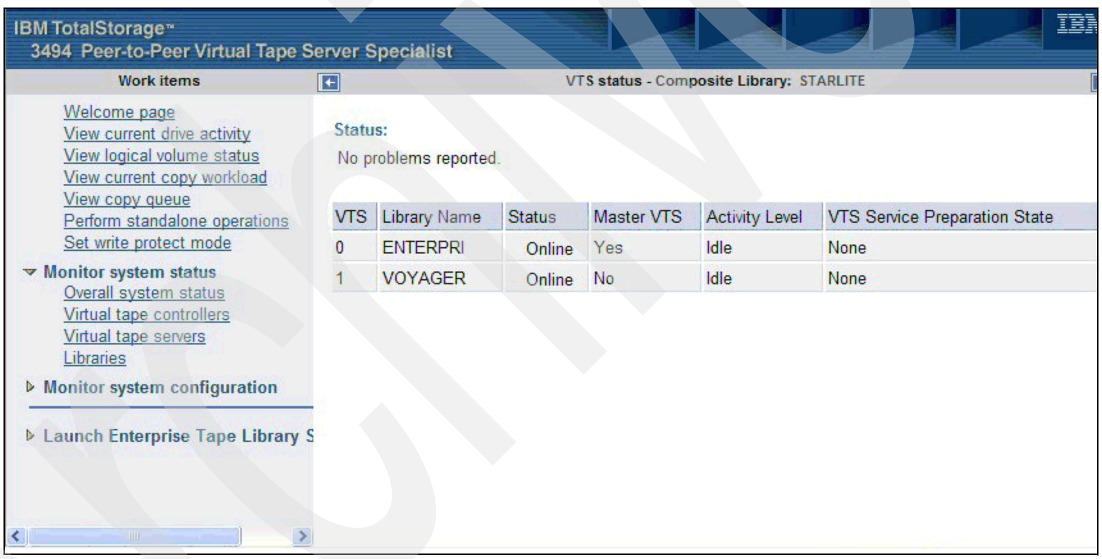

Figure 6-10 VTS Status screen

Copy mode

Prior to the addition of the PTP Selective Dual Copy function in VTS LIC 2.30.720.xx, control of when a scratch logical volume was written was set for the entire library at the initial install time. Now this control can be specified on the Management Class definition panel. This initial setting is still made and is now used as the default value. With the PTP Selective Dual Copy function, additional actions can be specified during the definition of a Management Class constructs to allow:

- The Management Class to specify which VTS is to be used in writing data for the initial scratch mount
- The Management Class to specify that no copy is to be performed

These new options are in addition to the existing actions specifying that a second copy of the logical volume is to be created in the same VTS.

Restriction: If the Peer-to-Peer VTS is part of a GDPS complex, copies are always created even if no copy is specified in the Management Class. Host message VT0103 will be presented when a forced copy overrides the Management Class parameter.

With Advanced Policy Management, these settings can also be defined at a Library Manager see (Figure 4-24 on page 173) or the Library Manager Specialist panel as part of a Management Class definition (see Figure 6-11 here).

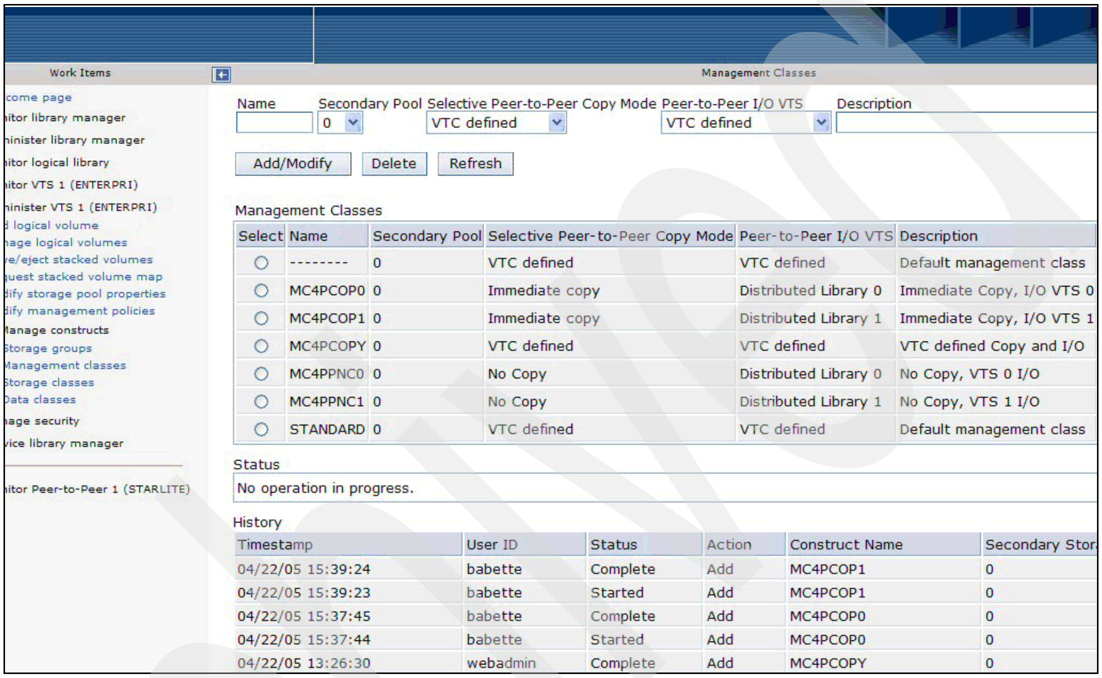

Figure 6-11 Management Class definition panel on the Library Manager Specialist

VTC mode control

The mode of all the VTCs under normal operations is Read/Write mode. This mode will be entered when the Peer-to-Peer system is initialized under normal circumstances. This mode should be changed under only two circumstances; a real disaster where the Master VTS in the Peer-to-Peer is inaccessible and any remaining VTCs could not establish a new Master or you are doing disaster recovery testing. With the microcode Release 7.4, control of the VTC mode is now selectable on the VTS Peer-to-Peer Specialist via the Write Protect Mode option. The mode you are able to select depends on the state of the Peer-to-Peer VTS. See “Peer-to-Peer Specialist” on page 347 for more information on how to change the VTC mode.

6.3.4 Distributed library

You can display the status of a distributed library just as you can today with stand-alone libraries. As described in 6.3.3, “Composite library” on page 229, the status of the distributed libraries determines the status of the composite library.

If, for instance, one of the distributed libraries is in paused mode and the other is in automated mode, the overall status of the Peer-to-Peer VTS will be in degraded mode. The composite library would still operate as if it were in automated mode and messages to the host would indicate which of the distributed libraries is paused. Also, if one of the distributed libraries is taken offline at the library, the state of the composite library will remain online and operational as long as one of the distributed libraries is operational.

Since all of the drives and volumes are defined to and associated with the composite library, the display of the distributed library will show that, from a host perspective, there are no drives associated with that library. The distributed libraries should be displayed for an accurate picture of the total slot count and empty slot counts.

The Peer-to-Peer VTS selects at initialization time which of the distributed VTSs becomes the Master VTS. But failures of some components of the Peer-to-Peer VTS afterwards may cause a master switchover.

The designation of a distributed library as a user interface (UI) distributed library is a Peer-to-Peer VTS installation task. It can be verified in the VTS summary panel of Enterprise Tape Library Specialist or in the VTS status window of the Library Manager.

6.3.5 Library Manager database

The Library Managers of each of the distributed libraries within the composite library has its own library management database. The information on the logical volumes in these databases is equal because synchronization of the status and the contents of the logical volumes is a primary task of the VTCs. The VTC communicates with both VTSs and the VTSs communicate with the Library Manager before a status change (for instance, a change from scratch to private) takes place.

The Peer-to-Peer VTS is designed to continue operation with the remaining configuration in case of failure of a component of one of the distributed libraries. This will result in temporary differences between the Library Manager databases of the two tape libraries. As soon as the Peer-to-Peer VTS is recovered, a synchronization process starts from the VTCs that results in synchronized distributed libraries and, as a consequence, also synchronized Library Manager databases.

6.3.6 Virtual drives

z/OS commands providing information on the status of your Peer-to-Peer virtual addresses returns information from the VTCs. there are two commands that give you information on the status of the library from the host’s perspective:

```
D U,,,address
```

```
LIBRARY DISPDRV,libname
```

In addition to the above two commands, you have the ability to display the current device activity panel from the Peer-to-Peer VTS Specialist.

When the host is not able to process logical volumes, it may be useful to find out whether communication is possible between your host and the VTC. If you are not able to communicate with the controller, your links may be broken. This can be investigated by issuing the following command:

```
DEVSERV QTAPE,(address),1,{RDC,DLD}
```

You will get information from the VTC in return; see Figure 6-12. The parameter RDC means Read Device Characteristics and DLD means Display Library Data. See Figure 6-13 for an example of the DLD parameter. If you have completely lost the connection, the command cannot get the data from the controller and is not able to fill in all the fields of the command output. It is an indication that the link may be causing problems and not your Peer-to-Peer VTS configuration. Normally, the command displays fields specific to the Peer-to-Peer VTS.

```
 DEVSERV QTAPE,2040,1,RDC
  IEE459I 20.10.55 DEVSERV QTAPE 953
  UNIT DTYPE  DSTATUS CUTYPE  DEVTYPE  CU-SERIAL  DEV-SERIAL ACL LIBID
  2040 3490L  ON-NRD  3494C2A 3590   * 0178-A1016 0178-A1016 I   04165
    READ DEVICE CHARACTERISTIC
  34905434905400E0 1FD88080B41B4168 0000080000000000 3494413590100002
  0041650500000000 4281000004000000 0400000000000000 0000000000000000
  ****      1 DEVICE(S) MET THE SELECTION CRITERIA
  ****      1 DEVICE(S) WITH DEVICE EMULATION ACTIVE
```

Figure 6-12 Example of DESERV QTAPE output with RDC data

Note the LIBID in the upper right of Figure 6-12. This corresponds to the first 3 bytes on second row (omitting the first half of the first byte) Also the fourth byte,'x05', is the Port ID of this particular control unit the command was issued to.

Figure 6-13 shows the same command but specifying the DLD data for the same VTC device.

```
DEVSERV QTAPE,2040,1,DLD
IEE459I 20.11.55 DEVSERV QTAPE 955
UNIT DTYPE  DSTATUS CUTYPE  DEVTYPE  CU-SERIAL  DEV-SERIAL ACL LIBID
2040 3490L  ON-NRD  3494C2A 3590   * 0178-A1016 0178-A1016 I   04165
  DISTRIBUTED LIBRARY DATA
010A101F04000DCF 0002000613780001 3494000000000000 0000000B20090002
3494000000000000 0000000000000000 0000000000000000 0000000000000000
****      1 DEVICE(S) MET THE SELECTION CRITERIA
****      1 DEVICE(S) WITH DEVICE EMULATION ACTIVE
```

Figure 6-13 Example of DESERV QTAPE output with DLD data

A Peer-to-Peer VTS (and, at suitable microcode levels, a stand-alone VTS) has a virtual device recovery function. If a virtual device is boxed, a reset allegiance channel command will allow a functioning alternate path to the boxed device to be used. The allegiance of the boxed device over the failed path is reset so that the alternate path can be used.

VTS virtual drive recovery enables the automatic recovery of virtual devices from boxed conditions that may be caused by link failures, director failures, or channel adapter failures.

No software support is needed for virtual device recovery.

6.3.7 Logical volume status

Information on logical volumes is available in the Library Manager database of both distributed tape libraries, in the Peer-to-Peer VTS, in the Tape Configuration DataBase (TCDB), and in the database of your tape management system.

A Peer-to-Peer VTS is transparent to your tape management system and to the TCDB. Those deal with the logical view of your data.

The logical volume information details in the Library Manager database are the same as today for a stand-alone VTS. The information for a Peer-to-Peer VTS logical volume is, however, available in two Library Manager databases. This information is the same in a fully operational Peer-to-Peer VTS. If there is a component failure in the Peer-to-Peer VTS, you may see temporary differences.

You can request additional volume-related information (see Figure 6-18 on page 246) by issuing the following command from the console:

```
DISPLAY SMS, VOLUME(volser)
```

Information on the dual copy status of your logical volume is also available on the logical volume status result panel of the Peer-to-Peer Specialist. You can submit the logical volume serial number on the logical volume status panel. The screen returned is shown in Figure 6-14. It summarizes all important information related to the logical volume.

Figure 6-14 Logical volume status screen

6.4 Peer-to-Peer VTS basic operations

In the next few sections, we explain some of the basic tasks that may be needed during the operation of a Peer-to-Peer VTS.

6.4.1 Power on/off

A complete IML of the VTC requires both VTSs to be functional prior to completing its logical subsystem IML. Any subsequent VTCs will IML successfully by communicating with any other successfully IML'ed VTC in the Peer-to-Peer system. When the VTC has completed its initialization by communicating with either both VTSs or other VTCs that have completed their initialization, the VTC enters a POI complete state and the control unit goes online. Both VTSs are not required as long as one VTC has been successfully IML'ed and has POI complete. An IML of a VTC is seen as an IML of a control unit and affects only the 16 or 32 devices defined to them. The rest of the Peer-to-Peer system remains the unchanged unless you IML the only VTC with a POI complete state. In that case the PtP VTS will be down.

The IML of the VTS does not affect the online status of the Peer-to-Peer VTS, but may temporarily limit access to certain volumes if that VTS has the only copy of that logical volume. Only after a successful IML of the Peer-to-Peer VTS can failures be tolerated. Upgrades and service action can be performed on individual components without making the whole Peer-to-Peer VTS unavailable.

Under certain circumstances, you might need to initialize the VTS Peer-to-Peer without all of the components it requires during a normal initialization. If access to one VTS in the PtP is not possible, because of broken links, building power failures, or catastrophic VTS hardware failures, or other conditions, the remaining components can be initialized into a state making the logical volumes available for use. If a VTC is powered on and it cannot communicate with one of the defined VTSs, it will not complete its IML because it cannot determine the state of all the components and logical volumes in the PtP VTS.

So it is possible to intervene and select a mode of operation you want the VTS PtP to enter. To do this, you can log on to the VTC via the service logon and use the service menus, usually an SSR activity, or connect to the VTS Peer-to-Peer Specialist and select one of the two modes available to you. You can allow the VTC to complete its initialization in either Read-Only mode or in Read/Write Disconnected mode.

Which mode you select would be based on your specific situation. A few things to consider are:

- Do you know how long the other VTS will be unavailable? Can you wait until the complete PtP VTS is available before powering on?
- Do you normally run in immediate copy mode? If you know that both libraries have a valid copy of every logical volume, then entering Read/Write Disconnected mode will not have any impact on any production jobs. Any jobs that are run will have full read/write access to the all the logical volumes. When access to the other half of the PtP VTS is re-established, all the logical volumes created or updated while in Read/Write Disconnected mode will be copied to the remote VTS. The VTCs will then enter Read/Write mode (normal mode) when the other VTS enters the Peer-to-Peer system.
- If you normally run in deferred copy mode or you know that valid copies exists in the inaccessible VTS, then coming up in Read/Write Disconnected mode could impact your production work. If a job is run and the most current copy of the data requested is located only in the inaccessible VTS, the job will fail because the token for this logical volume specifies that the only copy of the logical volume that it has access to is not the most current copy of the volume.

A scheduled power-off of one the components of the distributed libraries requires a service preparation to be performed in advance. Service actions on the VTC only require that the devices defined to the specific VTC be taken offline. For further details, see 2.7, “Service considerations” on page 55.

6.4.2 Clock setting

The Peer-to-Peer VTS adds some complexity to the clock setting of your configuration. In a stand-alone VTS, the clock of the VTS is synchronized with the clock of the Library Manager at VTS initialization time. As a consequence, you have to vary your VTS offline and change the tape library to pause mode before you are allowed to change the clock. An IBM System Service Representative (SSR) is responsible for performing this task.

With the Peer-to-Peer VTS, the clocks in the VTCs are synchronized with the clock of the UI distributed library. If this library is not available, the clock of the other distributed library will be used. Clock synchronization will take place at initialization time of the VTCs. Changing the clock to daylight savings time requires that you power off your Peer-to-Peer VTS.

The VTS clock time plays no role in the exchange of information between the host and the Peer-to-Peer VTS. Messages and SMF records related to the Peer-to-Peer VTS contain the time stamps of the host. Messages and statistics in the 3494 or 3953 Library Manager use the timestamp of its own internal clock. In case of problems with hardware components, it may be possible that Service Information Messages (SIMs) are sent to the host containing the time stamp based on the Library Manager clock.

Recommendation: We recommend that, for a VTS, you do *not* change your clocks twice a year, but rather, use your winter time or use GMT throughout the year.

You should realize that the time difference can be important in communication with your IBM SSR in case of problems with a hardware component. Problem records normally contain the host time or the real time of the occurrence of the problem. Your SSR must be aware of the clock time difference when performing the problem determination using information generated from the Peer-to-Peer VTS.

6.4.3 VARY online/offline

A composite VTS library can be varied online and offline like a stand-alone VTS. You can also use the VARY command for a distributed library, but VARYing a distributed library offline from the host really makes no difference, because it does not prevent the usage of the library as part of the Peer-to-Peer VTS. The warning message CBR3016I will be used when a distributed library is initialized or VARYed offline. A VTC can be taken offline by VARYing the associated addresses offline.

6.4.4 Library in pause mode

A distributed library can enter the pause mode as is possible today with a stand-alone 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 one of the distributed libraries is in pause or manual mode:

```
CBR3757E Library library-name in {paused | manual mode} operational state
```

With the Peer-to-Peer VTS support, a new functional enhancement allows logical mounts for scratch volumes, logical mounts for private volumes where the volume is resident in the Tape Volume Cache, demounts, and ejects, to be executed while in pause mode. These mounts will use the active distributed library as the I/O VTS.

Other mounts will be started, but will be held up by the resulting physical mounts needed to execute a recall. Those mounts will be queued by the Library Manager for later processing when the library leaves the pause mode.

Since 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 distributed library.

Although the Peer-to-Peer VTS has enhanced support for library pause mode, for performance reasons, it is very important that you minimize the amount of time spent with the library in pause mode condition.

6.4.5 Inserting and deleting logical volumes

Inserting and ejecting logical volumes require the Library Manager's of both distributed libraries to be online to the host before issuing the any eject or inserting of logical volumes. A A `CBR3726I` message with an error-code of `X'0D'` will be issued indicating that one of the distributed libraries is not available when the command was issued. See "Incompatibility error message" on page 243 for more detail on the CBR3726I message.

Inserting logical volumes

You must use the Library Manager console of the 3494 or TS3500/3953 Tape Library where the UI distributed library resides for defining new logical volumes to the Peer-to-Peer VTS. This function is restricted on the Library Manager of the other distributed library.

When you have defined the new logical volumes, the host is notified and can request the Library Manager to change the status of the logical volumes to private or scratch. OAM communicates with your tape management software, if available, to determine the category for the logical volume. In response to the request, the Peer-to-Peer VTS synchronizes the logical volume information in both Library Manager databases of the distributed libraries. When one of the distributed libraries fails during this process, the synchronization will be suspended and resumed after the failing distributed library becomes available.

> **Attention:** For performance considerations, limit the number of logical volumes to insert at no more than 10,000 at any one time. Wait for the previous list of volumes to complete before adding more.

Deleting logical volumes

Due to the permanent nature of the EJECT, the VTS 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 will fail. 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>.
```

There are considerations to be aware of when ejecting large numbers of logical volumes. Deleting logical volumes from a library must be performed by executing an eject command against the logical volume. The Peer-to-Peer VTS ensures that the logical volumes are deleted in both tape libraries providing both Library Managers are online at the time. The maximum number of ejects in progress is limited to 1000 for any one Peer-to-Peer VTS. Since 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. Plus 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 hosts performance. So the recommended limit for the number of outstanding ejects 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
```

Restriction: For performance considerations the recommended limit for the number of outstanding ejects requests on the host is no more than a couple thousand per system.

6.4.6 Logical Volume availability

The Peer-to-Peer VTS has improved capability to read a logical volume:

- If a permanent media read error occurs when a logical volume is recalled from a stacked volume
- If incorrect data is detected when being transferred to the host

Recovery is achieved either transparently within of the Peer-to-Peer configuration or through use of DDR by the.

Unreadable recall from stacked volume

Recovery for a logical volume that has a stacked volume media error is provided by the VTS and the VTC. The following actions are performed:

- If a recall fails due to a permanent media read error on a 3590 or 3592 tape drive, the VTS will retry the recall on another tape drive. Failure on the second tape drive results in the stacked volume being placed in read-only status. This action is the same as for a stand-alone VTS.
- The VTC detects when a VTS has not been able to recall a logical volume due to a 3590 media error and changes the I/O VTS for this logical volume to the other VTS. Recall of the dual copy into the Tape Volume Cache of the second VTS makes the logical volume available.
- In the unlikely event of another media error on the second stacked volume, an equipment check error will occur on the first I/O operation after the mount completion and the job will fail.

Read-only status volume recovery enhancement

Once an hour, a VTS checks for stacked volumes that are in read-only status and recalls the logical volumes into the Tape Volume Cache, thus removing the logical volumes from the damaged stacked volume. The Peer-to-Peer VTS enhances this process by copying logical volumes that cannot be read from the other VTS in a transparent manner, thus restoring two available copies.

If a dual copy of the unreadable logical volume has not yet been created in the second VTS at the time of the read-only status volume recovery, the intervention required message

```
Logical volume xxxxxx was not fully recovered from damaged stacked volume yyyyyy
```

is posted to the attached hosts. The stacked volume is placed in Library Manager volume category X'FF08' and, for security reasons, is not ejected from the tape library. An IBM SSR can initiate action for logical volumes that must be recovered.

Tape Volume Cache and adapter errors

Data from the Tape Volume Cache is checked for errors as read commands are being processed through the adapters that provide host attachment.

The recovery process proceeds as follows:

1. A read command failure causes MVS hosts to initiate the DDR process.
2. When the logical volume is remounted and positioned, and the read command that failed is reissued, the Peer-to-Peer VTC being used will access the copy of the logical volume in the other VTS.
3. A successful completion of a rewind/unload command causes the Peer-to-Peer VTS configuration to make a new copy of the logical volume in the VTS, which originally had the failure, thereby replacing the erroneous copy and ensuring that two copies remain accessible.
4. Continuing failures to read the logical volume with use of the DDR process must be reported to the SSR.

6.4.7 Adding stacked volumes

Stacked volumes must be inserted for both distributed libraries in a Peer-to-Peer VTS as in a stand-alone VTS. We recommend that both distributed libraries have about the same number of stacked volumes. Before adding stacked volumes, you must define the volser ranges as has been described in 4.2.2, “Define stacked volume ranges” on page 153.

When one of the distributed libraries runs out of stacked volumes, many of the operations of that library cannot continue. To prevent the Tape Volume Cache from reaching an overflow state and throttling, all logical mount processing is suspended as soon as the number of available scratch physical volumes in the distributed library is zero. The library enters the VTS out of stacked volumes state. The Peer-to-Peer VTS continues to operate in degraded mode allowing all jobs to run normally on the VTS with sufficient stacked volumes. Only the copies of all the volumes written during this time are deferred until the Out-of-Stacked-Volumes condition is resolved on the other VTS.

6.4.8 Deferred/immediate copy mode

As stated in 2.3, “Creating the secondary copy of a virtual volume” on page 26, the default copy mode, deferred or immediate, which the Peer-to-Peer VTS is using, is determined at installation time. This state can change when failures of components of the Peer-to-Peer VTS enforce it. You are able to display it on the system configuration panel of the Peer-to-Peer VTS Specialist. With VTS Release 7.4, the default mode can be dynamically changed with a Data Class parameter, see Figure 6-11. If no parameter is selected, then the Peer-to-Peer VTS will use the default copy mode as soon as the availability of the components allows it. The Advanced Policy Management feature is a requirement for modifying the copy mode through Management Class definitions. This has been discussed in 4.4.2, “PtP Copy Mode Control” on page 180.

6.4.9 Stand-alone support

In the standalone VTS, the Library Manager provides a Stand-Alone Device facility that allows you to set a virtual tape device in a VTS into stand-alone mode. In the Peer-to-Peer VTS, all access to the logical devices and volumes is through the VTCs. So the facility to give you support for mounting logical volumes is provided by the VTS Peer-to-Peer Specialist. To the host, the device appears as if it was not in a tape library. This mode allows you to use devices in the VTS with stand-alone programs that cannot issue a command to mount and demount volumes in the library. It also allows you to IPL from a tape-library-resident drive. Examples of stand-alone programs are stand-alone dump and stand-alone restore.

Figure 6-15 Stand-alone mode panel on VTS Peer-to-Peer Specialist

Stand-alone mode support for native tape drives is still available on the tape library (3494 or TS3500/3953) through the Library Manager panels.

6.4.10 Distributed library aspects

The VTS management policy definitions for a distributed library must be entered on the Library Manager console and can be changed dynamically. These definitions apply only to the distributed library associated with the 3494 Library Manager or 3953 Library Manager are not propagated to the other distributed library. See 4.2, “VTS definitions from the Library Manager” on page 151, for further details.

As described in 2.6.1, “Initializing the Peer-to-Peer VTS configuration” on page 49, the Peer-to-Peer VTS selects the Master VTS associated with one of the distributed libraries based on a number of selection criteria. If a master switchover has taken place because of a failure of a component of the Peer-to-Peer VTS, you cannot reassign the Master VTS to the originally chosen VTS to be the master. The IBM System Service Representative may reassign the master after completion of repairs, or it will automatically be reassigned to the preferred Master VTS once conditions allow the switchover to occur

6.4.11 Preparing for service

When an element of the distributed libraries within the Peer-to-Peer VTS needs to be serviced, it must be prepared prior to taking the element away; otherwise, continued host access to data may not be possible. The service preparation task is a SSR responsibility, and will remove a whole distributed library from all Peer-to-Peer activity. More details on service preparation can be found in 2.7, “Service considerations” on page 55.

Preparing a VTS for service

When an operational VTS needs to be taken offline for service, the Peer-to-Peer VTS must first be prepared for the loss of the resources involved, to provide continued access to customer data. The controls to prepare a VTS for service (Service Prep) are provided through the service menu on a VTC.

Here is the message posted to all hosts when the Peer-to-Peer VTS is in this state:

```
CBR3788E Service preparation occurring in library library-name.
```

Preparing the tape library for service

If the 3494 or the TS3500 Tape Library in a Peer-to-Peer VTS requires service, the VTS associated with it must first be prepared for service. Once the VTS has completed service preparation, the normal procedures for servicing the tape library can continue.

Preparing a VTC for service

Preparing a VTC for service is the same as preparing any tape control unit for service. All tape devices addressable by the control unit must be varied offline, and all jobs in process must be completed or stopped. There is no need to do a Service Prep for any part of the VTS Peer-to-Peer.

6.5 Peer-to-Peer VTS messages and displays

This section describes the enhanced message support and the new messages related to the Peer-to-Peer VTS.

6.5.1 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 Peer-to-Peer VTS, 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.
```

6.5.2 Peer-to-Peer VTS messages

This section lists some of the Peer-to-Peer VTS specific messages that you may see. For a complete and current list, please see the appropriate volume of the *z/OS MVS System Messages* book.

Above Threshold Warning state

This alert message is introduced with VTS LIC 2.32.740.xx. The VTS 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 VTS leaves the Above Threshold Warning state when the amount of data to copy is below the threshold capacity for 30 consecutive sample periods. The consecutive sampling criteria is to prevent the excessive messages being created.

```
CBR3750I Message from library library-name:OP0160 Above threshold for uncopied data in
cache, throttling possible
```

Along with the above message comes an option to have the VTS suspend Read-Only Recovery (ROR) to reduce the internal housekeeping workload in order to alleviate the large number of volumes requiring copying to tape. This option, along with the message, is enabled or disabled by the SSR via service panels on each VTS.

You might see this message if the Peer-to-Peer was separated due to multiple link failures or in disaster recovery testing where any production work has been creating large numbers of logical volumes with out being able to copy them to the other VTS. When the Peer-to-Peer VTS connections are re-established and all of the deferred copies are sent to the other VTS, there may be temporarily too much data to fit into the VTS cache.

Incompatibility error message

In case of an uncompatible function error, you may 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 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'0D' The Peer-to-Peer VTS 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 Peer-to-Peer VTS 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 Peer-to-Peer VTS subsystem.

X'10' The AX0 in the Peer-to-Peer VTS subsystem that the command was issued to is in read-only mode and is not accepting eject or change use attribute requests. This mode of operation is provided to support disaster recovery operations in a Peer-to-Peer VTS configuration where the configuration is split between two physical sites.

6.5.3 Displaying the Peer-to-Peer VTS status

The following messages (see Figure 6-16) can be issued for the Peer-to-Peer VTS.

- `CBR1100I OAM status:`

  Explanation: The operator has entered the following command:

  `DISPLAY SMS,OAM`

```
CBR1100I OAM status: 321
TAPE TOT  ONL  TOT  TOT  TOT  TOT   TOT   ONL   AVL   TOTAL
     LIB  LIB  AL   VL   VCL  ML    DRV   DRV   DRV  SCRTCH
       6    5    1    2    1    1   180    41    41    2498
There are also 2 VTS distributed libraries defined.
CBRUXCUA processing ENABLED.
CBRUXEJC processing ENABLED.
CBRUXENT processing ENABLED.
CBRUXVNL processing ENABLED.
CBRUXSAE processing ENABLED.
Access Backup processing INACTIVE for UNREADABLE VOLUMES.
Access Backup processing INACTIVE for OFFLINE LIBRARIES.
Access Backup processing INACTIVE for NOT OPERATIONAL LIBRARIES.
```

Figure 6-16 Display OAM status

  If there are Peer-to-Peer VTS subsystems defined to the system, the following status line will be 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 may 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:`

  Explanation: The operator has entered the following command:

  DISPLAY SMS,LIBRARY(*library-name*),DETAIL

```
CBR1110I OAM library status: 318
TAPE      LIB  DEVICE    TOT  ONL  AVL  TOTAL  EMPTY SCRTCH  ON OP
LIBRARY   TYP  TYPE      DRV  DRV  DRV  SLOTS  SLOTS   VOLS
ATLBA041  VCL  3494-L10   64   16   16      0      0      0  Y  Y
----------------------------------------------------------------------
MEDIA       SCRATCH       SCRATCH       SCRATCH
TYPE          COUNT     THRESHOLD      CATEGORY
MEDIA1            0           200          0031
MEDIA2            0           300          0032
----------------------------------------------------------------------
DISTRIBUTED LIBRARIES:  ATL70680 ATL70677
----------------------------------------------------------------------
OPERATIONAL STATE:                            AUTOMATED
ERROR CATEGORY SCRATCH COUNT:                   0
SCRATCH STACKED VOLUME COUNT:                 150
PRIVATE STACKED VOLUME COUNT:                1200
CORRUPTED TOKEN VOLUME COUNT:                   0
HIGH CAPACITY INPUT STATION CAPACITY:          10
HIGH CAPACITY OUTPUT STATION CAPACITY:         10
----------------------------------------------------------------------
Library supports Outboard Policy Management.
```

Figure 6-17 Display SMS composite library command

Note: Library type value is VCL for composite as apposed to VDL for distributed.

A display of OAM library status (Figure 6-17) has been generated. When a library name is supplied, there is one data line describing the specified library. When ALL is supplied, there is one data line for each library in the configuration. If both optical and tape libraries have been defined in the SMS configuration, the sample display above will be generated. Otherwise, only the data for the library type defined will be generated.

- Our example includes:
  - Changed/added status lines.
  - The media type, scratch count, scratch threshold, and scratch category lines will only be displayed for media that have a threshold value or a scratch count greater than zero.
  - Service preparation occurring in distributed library *library-name*.
  - Library supports Import/Export.
  - Library supports Outbound Policy Management.

- `CBR1180I OAM tape volume status` (Figure 6-18).

  Explanation: The operator has entered the following command:

  `DISPLAY SMS,VOLUME(volser)`

```
DISPLAY SMS,VOLUME(SL0026)
VOLUME  MEDIA     STORAGE   LIBRARY   USE  W  C  SOFTWARE  LIBRARY
        TYPE      GROUP     NAME      ATR  P  P  ERR STAT  CATEGORY
SL0026  MEDIA2    P2PCLS    STARLITE  P    N     NOERROR   PRIVATE
-------------------------------------------------------------------
RECORDING TECH:     36 TRACK         COMPACTION:         YES
SPECIAL ATTRIBUTE:  NONE             ENTER/EJECT DATE:   2005-02-03
CREATION DATE:      2005-02-03       EXPIRATION DATE:    2007-01-24
LAST MOUNTED DATE:  2005-02-16       LAST WRITTEN DATE:  2005-02-03
SHELF LOCATION:
OWNER: IBMUSER
LM SG: P2PCLS    LM SC: P2PCLS    LM MC: STANDARD  LM DC:
-------------------------------------------------------------------
Logical volume.
Volume is cache resident.
Valid copy in each distributed library.
```

Figure 6-18 Display volume details

  A display of volume status has been requested. Status is reported for the requested tape volume.

- `CBR4125I Valid copy of volume` *`volser`* `in library` *`library-name`* `inaccessible.`

  Explanation: An error has been detected during Library Automation Communication Services (LACS) processing for the MOUNT function. Library *library-name* has indicated that a valid copy of volume *volser* is not currently available. The volume cannot be retrieved using normal library automated function; operator or SSR intervention is needed. The error is reported by a unit check when the mount order is sent to the library.

- `CBR4126I Library` *`library-name`* `drive is in read only mode.`

  Explanation: An error has been detected during Library Automation Communication Services (LACS) processing for the MOUNT function. The requested drive in the VTS Peer-to-Peer library *library-name* is in Read-Only mode, causing the scratch mount request to this drive to fail. Read-Only mode is provided at an AX0 level to prevent hosts attached to that AX0 from modifying the contents of a logical volume or its category assignment.

- `EDG0235E ERROR IN LOCDEF FOR LOCATION` *`location`* `- DISTRIBUTED LIBRARIES CANNOT BE SPECIFIED`

  Explanation: This message applies only to DFSMSrmm users. A LOCDEF parameter in the DFSMSrmm PARMLIB has specified a location name that matches the name of a system-managed library that is distributed. You cannot use the names of distributed libraries with DFSMSrmm.

7

PtP VTS performance and monitoring

In this chapter, we describe the monitoring and reporting tools available for the VTS and PtP VTS. We focus on the usage of the Library and VTS Specialists for realtime statistics, and explain the enhanced reporting based on SMF type 94 records. We also highlight performance-related considerations that are unique to the Peer-to-Peer VTS:

- We describe how the Specialists interact with each other, if you have a 3494 Tape Library and a TS3500 Tape Library attached to either VTS.
- We demonstrate the scope of the Specialists by showing sample screen displays.
- We explain the usage and reporting capability of the VTSSTATS.
- We give you a brief overview of the IBM Storwatch Expert.
- We describe other possible display and monitoring options.

7.1 Monitoring and reporting tools

The IBM TotalStorage Peer-to-Peer VTS Specialist, the IBM TotalStorage Enterprise Tape Library (ETL) Specialist, and the IBM System Storage TS3500 Tape Library Specialist are the state of the art tools for realtime monitoring. By displaying information using the ETL Specialist and Peer-to-Peer Specialist, the information obtained is for both the logical image (PtP Specialist) and for the underlying hardware components (ETL Specialist).

> **Note:** The IBM 3494 and 3953 Library Manager consoles only provide information about the distributed VTSs. The PtP VTS Specialist is the only source of information about the composite library. We therefore recommend that you always implement the PtP VTS Specialist when you install a PtP VTS.

VTSSTATS can be used to obtain information about the Peer-to-Peer VTS and its distributed libraries based on the SMF94 records. Its daily statistics may be used for long-term reporting and trend reporting; the hourly reports are helpful for performance evaluation and problem determination. The Specialists are shipped as microcode for the VTCs and for the 3494 and 3953 Library Manager.

VTSSTATS and other MVS tools are available from the following links:

Internet FTP site `ftp://ftp.software.ibm.com/storage/tapetool/`

IBM intranet `http://w3-1.ibm.com/sales/systems/ibmsm.nsf/docnames/tapetools`

The Library Manager console and MVS console commands may be used as well to get status information of the PtP VTS. We describe the Specialists and VTSSTATS in more detail in the following sections.

7.1.1 Monitoring using the IBM Specialists

The Specialists are a 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 Peer-to-Peer VTS, its components, and the distributed libraries. They also provide information on the copies that have not been completed yet and on the amount of data to be copied.

The *Peer-to-Peer VTS Specialist* is a Web server that is installed in each virtual tape controller. You need to perform some installation steps to access the Web pages from your workstation browser. Please refer to 3.1.14, "TotalStorage Specialist configuration" on page 78, for more information.

Although this Specialist is not required for monitoring the Peer-to-Peer VTS, we highly recommend that you install it because it is a single source of current information about the complete hardware and logical components of the Peer-to-Peer VTS.

In a Peer-to-Peer VTS configuration, you have up to three possible Specialists available: the ETL Specialist (3494 or 3953), the TS3500 Tape Library Specialist, and the Peer-to-Peer VTS Specialist. You can access the Peer-to-Peer Specialist Web pages from the ETL Specialist and vice versa, as there is a link between the products which enables you to switch between them seamlessly and hence easily find the required information.

The *ETL Specialist* is the Web interface to the Library Manager and the VTS of each distributed library.

Library Manager information includes:

- System summary
- Operational status and operator interventions
- Component availability
- Performance statistics
- Command queue
- LAN host status and LAN information
- Dual accessor zones (if High Availability Unit is installed)
- VOLSER ranges and cleaner masks

VTS information includes:

- Active data and active data distribution
- Data flow
- Logical mounts per hour and mount hit data
- Physical device mount history
- Category attributes
- Management policies
- Real time statistics
- Move/Eject Status
- Volser Ranges
- Find and Manage Logical Volumes
- Manage Constructs
- Move/Eject Stacked Volumes
- Request Stacked Volume Map
- Select Storage Pool Properties
- Modify Volser Ranges
- Security

7.1.2 Monitoring screens

The real-time statistics include status, throughput, and performance statistics, the same information that has been added in the SMF 94 record. If not explicitly outlined on the screen, data is displayed for the previous 24 hours.

Figure 7-1 shows the home page of the 3494 ETL Specialist.

Figure 7-1 3494 ETL Specialist home page

The Peer-to-Peer VTS Specialist gives you more Peer-to-Peer related specific information on virtual tape controllers, composite and distributed library configurations, copy workload, and other useful statistics. Figure 7-2 shows the home page of the Peer-to-Peer Specialist.

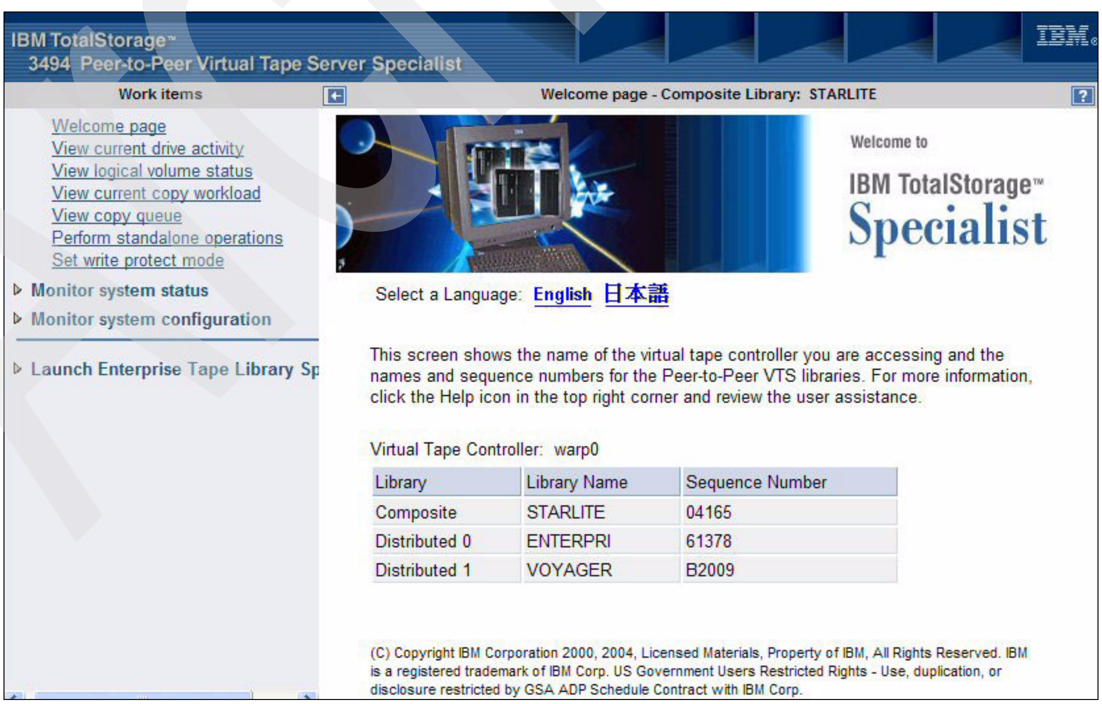

Figure 7-2 Home panel of the PtP VTS (Composite Library)

The page is divided into two main windows: the blue window at the left, where you can select the pull-down menus with the various information sources; and the right window, where the information is displayed.

From the home page, you can see the library name and sequence number of the composite and distributed libraries. Using the tabs in the left window, you can display the other information, detailed in the following sections.

System status

From this page, you can see the status of virtual tape controllers, distributed VTSs, distributed libraries, and system links. For each of these components, you have a status column; when this column is red, either the component is offline or the status is *not optimal* (for distributed libraries).

There is also a graphical view of the system status available and a detailed system status page, which gives you more detail on the status of each Peer-to-Peer VTS component.

If you need a better explanation of the fields shown on each page, click on the question mark in the upper-right corner of the page. Here you can find complete help text that explains the meaning of all the information displayed.

Figure 7-3 is an example of the graphical system status page. All components shown in red are not available because of a failure.

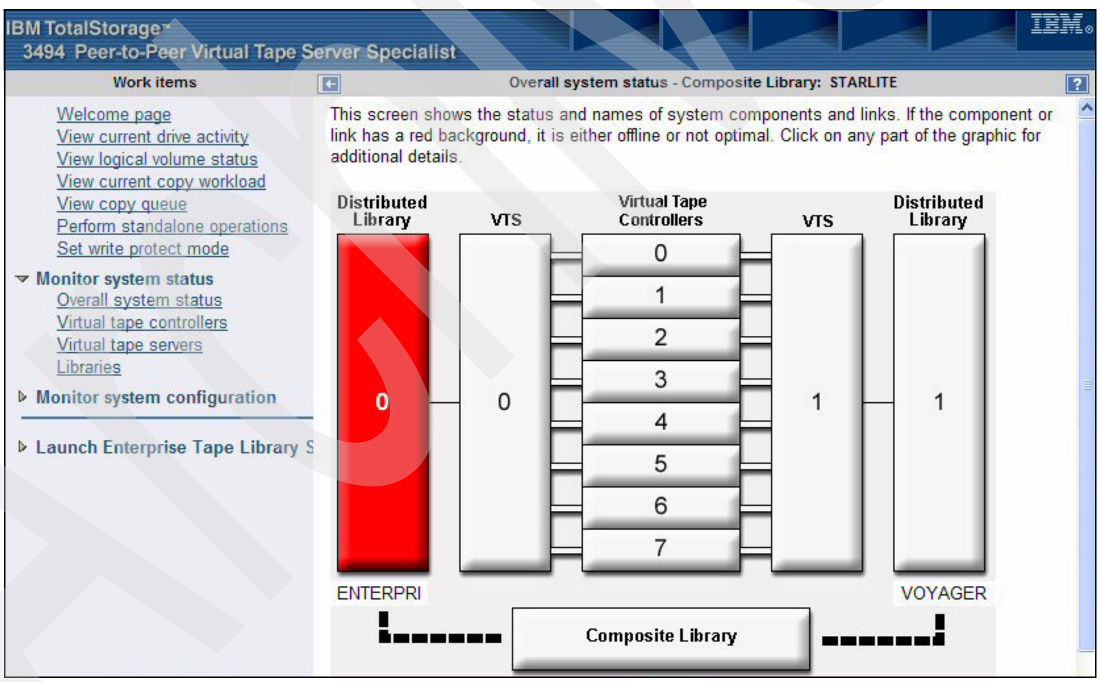

Figure 7-3 Peer-to-Peer VTS Specialist graphical system status

System configuration

This page contains the following information on the configuration of every component of the Peer-to-Peer VTS complex:

- VTC name, serial number, and network IP address
- Composite and distributed VTS library name, serial number, and sequence number
- Library Manager information
- Copy mode, deferred copy priority threshold, and UI library name

An example of the information included in this page is shown in Figure 7-4.

System Configuration

The following tables summarize the current Peer-to-Peer VTS configuration:

Virtual tape controllers:

Controller	Name	Serial Number	Network IP Address	Code Level	Number of Drives
0	warp0	0D00220	9.155.50.191	2.32.745.19	16
1	warp1	0D00221	9.155.50.192	2.32.745.19	16
2	warp2	0D00222	9.155.50.193	2.32.745.19	16

Controller	Name	Operational Mode	I/O VTS Selection Mode	Copy Mode
0	warp0	Normal	Balanced	Immediate
1	warp1	Normal	Balanced	Immediate
2	warp2	Normal	Balanced	Immediate

VTSs:

VTS	Library Name	Serial Number	Number of Drives	Code Level
0	ENTERPRI	0061378	128	******
1	VOYAGER	00B2009	128	******

Libraries:

Library	Library Name	Sequence Number	UI Library *
Composite	STARLITE	04165	-----
Distributed 0	ENTERPRI	61378	Yes
Distributed 1	VOYAGER	B2009	No

* Note: Only one of the distributed libraries can be the UI library.

Library	Library Name	Library Manager	Network IP Address	Code Level
Distributed 0	ENTERPRI	A	9.155.50.177	******
		B	9.155.50.178	******
Distributed 1	VOYAGER	A	9.155.49.12	******
		B	Not Configured	******

Figure 7-4 Peer-to-Peer VTS Specialist system configuration

7.1.3 VTS realtime statistics

The screens in Figure 7-5 to Figure 7-7 show a comprehensive overall status of the PtP VTS:

- Hardware related information such as Lib_ID, PtP name, operational mode, etc.
- Availability status of components such as channel adapter, power supplies, RAID array adapter, etc.
- Performance information such as reclamations in progress, read-only recovery in progress, overall throttling, etc.
- Physical device status such as number of cartridges filling and with active data, status of each single tape drive, etc.
- Hourly based information such as channel reads and writes, physical mounts for recall, migrate and reclaim, etc.

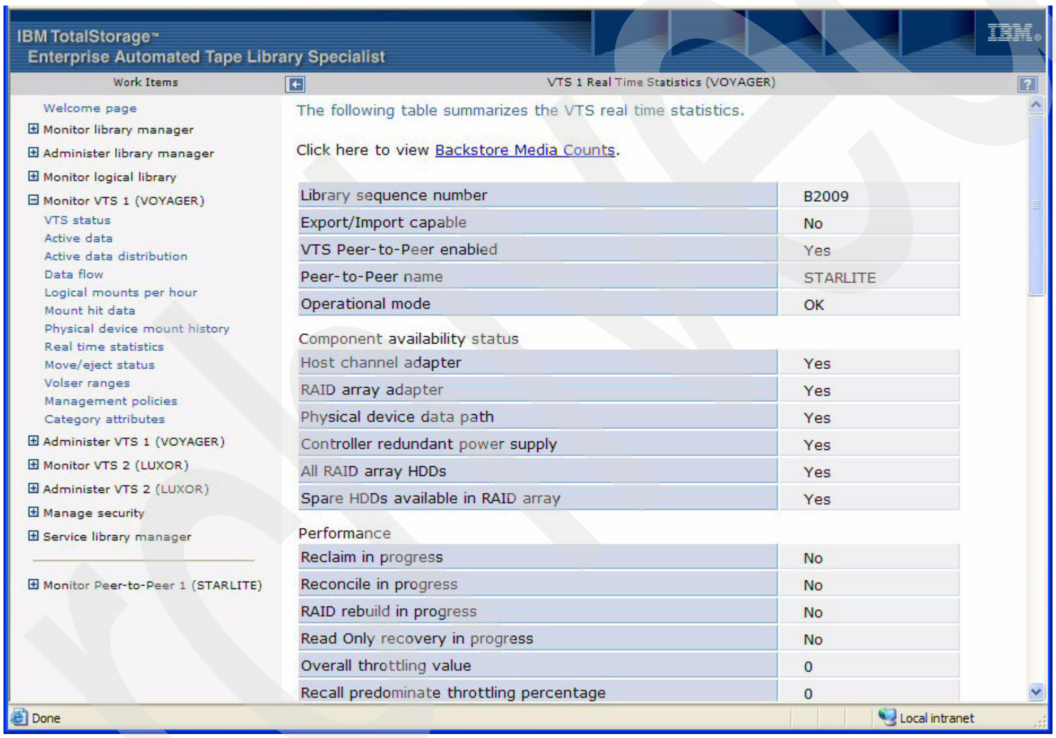

Figure 7-5 Realtime statistics page 1

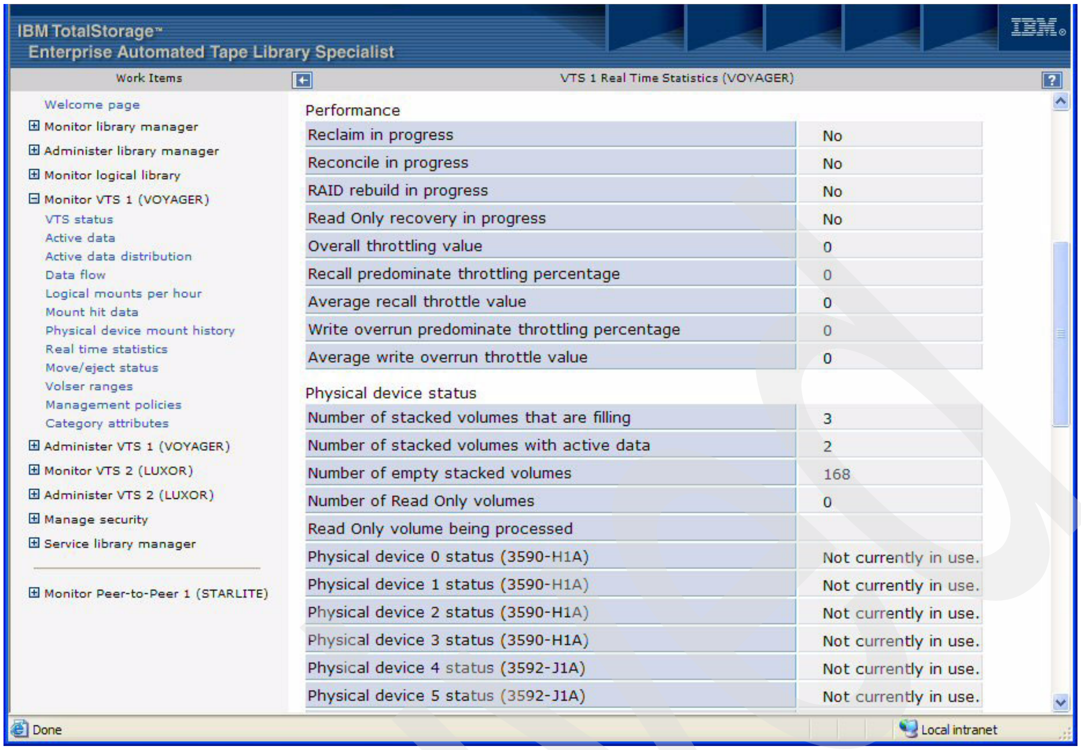

Figure 7-6 Realtime statistics continued

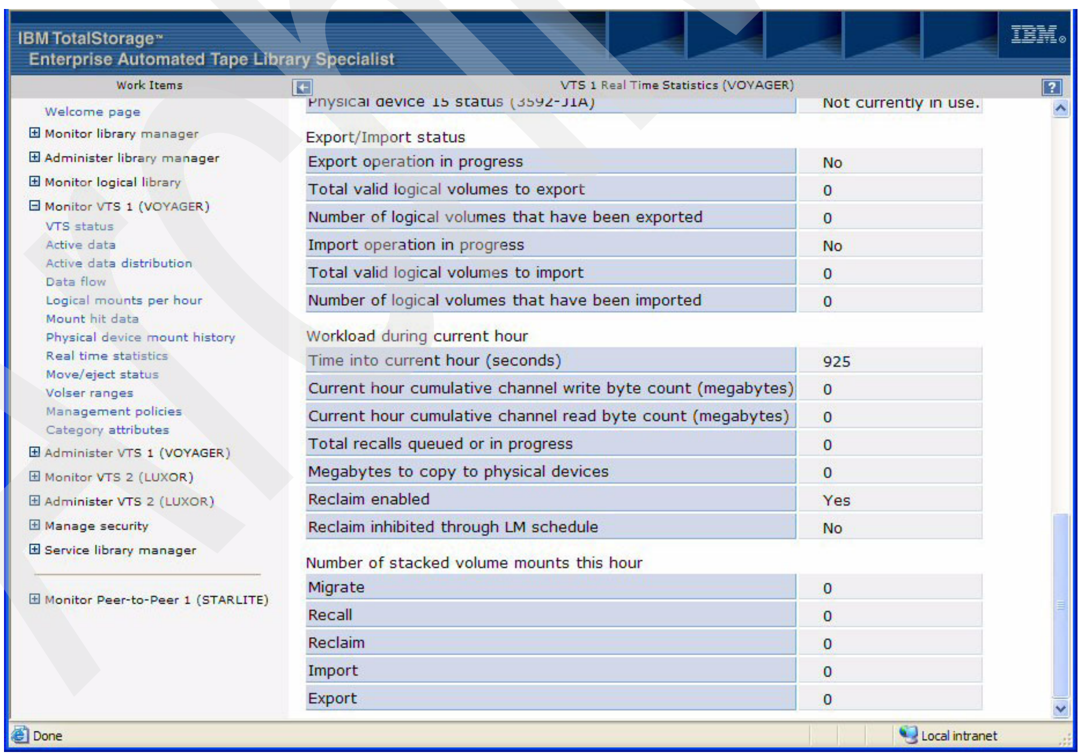

Figure 7-7 Realtime statistics continued

Status display

All attached VTCs are displayed in Figure 7-8 on page 255: their ON/OFF status, if one of them is in the *WriteProtect Mode*, and if *Copy is disabled*.

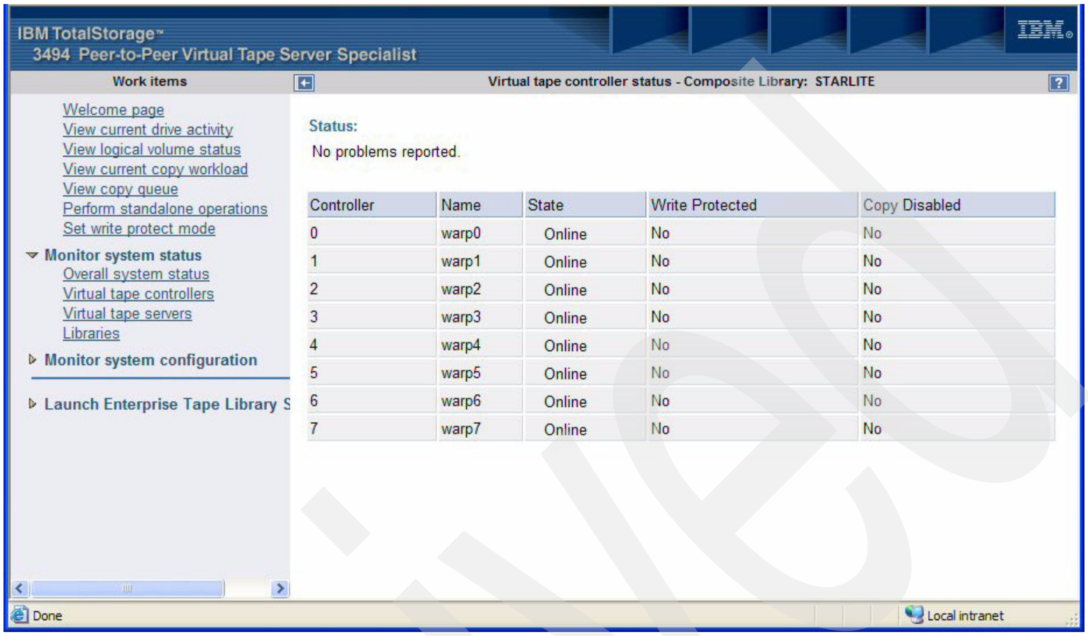

Figure 7-8 VTC status display

Current device activity

For each of the virtual tape controllers, this option shows the virtual device activity, with detailed information on:

- Mounted logical volume serial
- Written and read bytes
- Time to mount and time on drive for the mounted volume
- I/O VTS used for this logical volume

This information can be used to monitor Peer-to-Peer VTS activity, as well as evaluate performance constraints: The ratio between bytes written and read can indicate the balance between input and output, and the time to mount can give you an indication of a potential performance degradation.

Logical volume status

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 to obtain the desired information, as detailed in Figure 7-9.

Figure 7-9 Peer-to-Peer VTS Specialist logical volume status

Note that both the data level and the category level can be *current* or *downlevel*. This means that the second copy of a logical volume has already been completed for *current*, and has not yet completed for *downlevel*.

When the *in transition* keyword is present in the Stacked Volume, Compressed File Size, Last Modified Time and Last Access Time fields, this means that the specified volume is mounted on a virtual drive.

The Pre-Existing Volume field can be set to Yes or No. This flag indicates whether or not a logical volume contained valid data that existed on the VTS prior to conversion to a Peer-to-Peer VTS.

Active data distribution

The display shown in Figure 7-10 gives you the actual status of all physical cartridges and the percentage of active data they actually contain. You get the exact number of volumes with active data and the defined reclamation threshold.

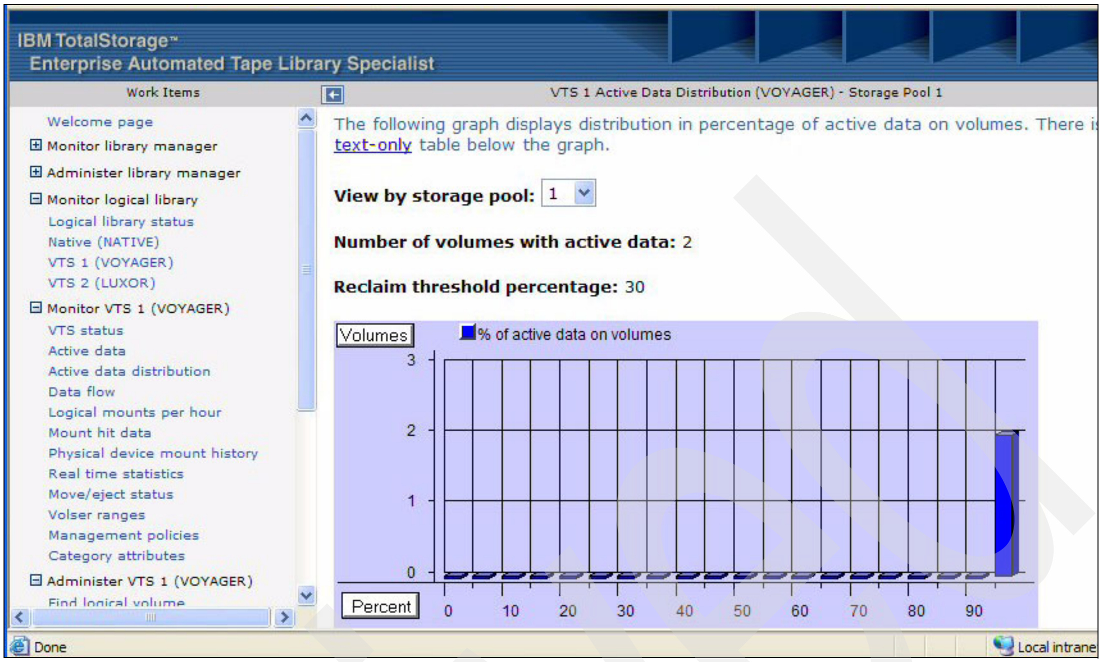

Figure 7-10 Active data distribution

VTC Copy queue display

The display shown in Figure 7-11 shows the content of the copy queue: the volser, source VTS where the first copy resides in, the number of bytes to be copied, and the age in the queue. It is obvious that this PtP VTS runs in deferred mode.

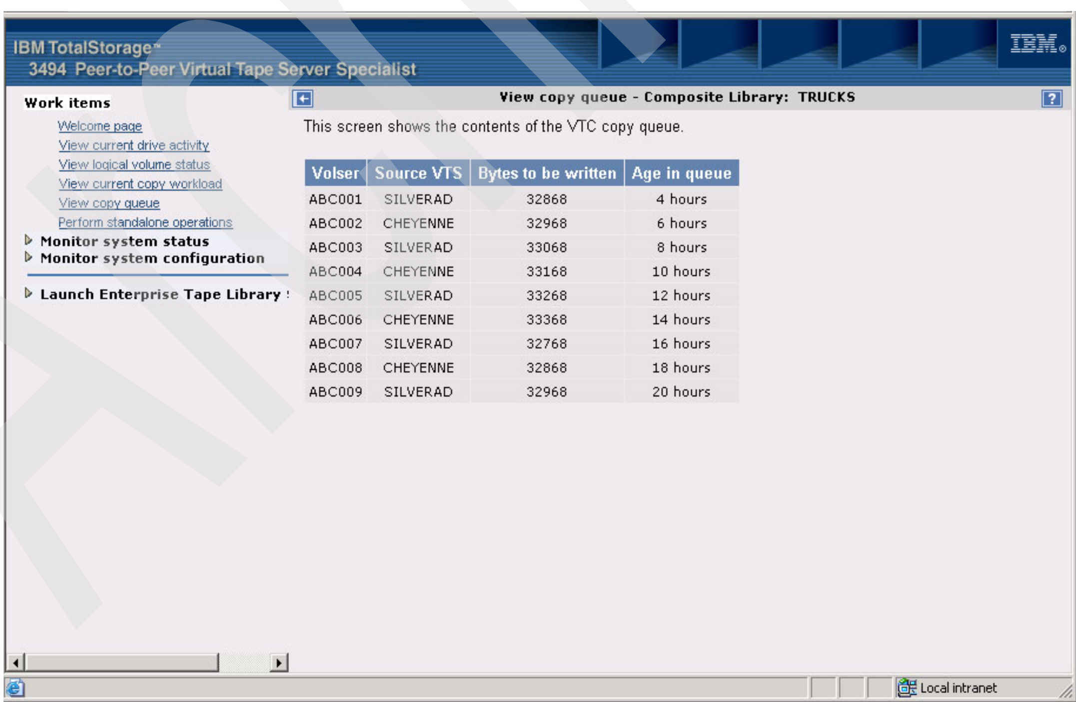

Figure 7-11 Copy queue

VTS Data flow

The graph in Figure 7-12 displays the amount of data written and read from the host channel.

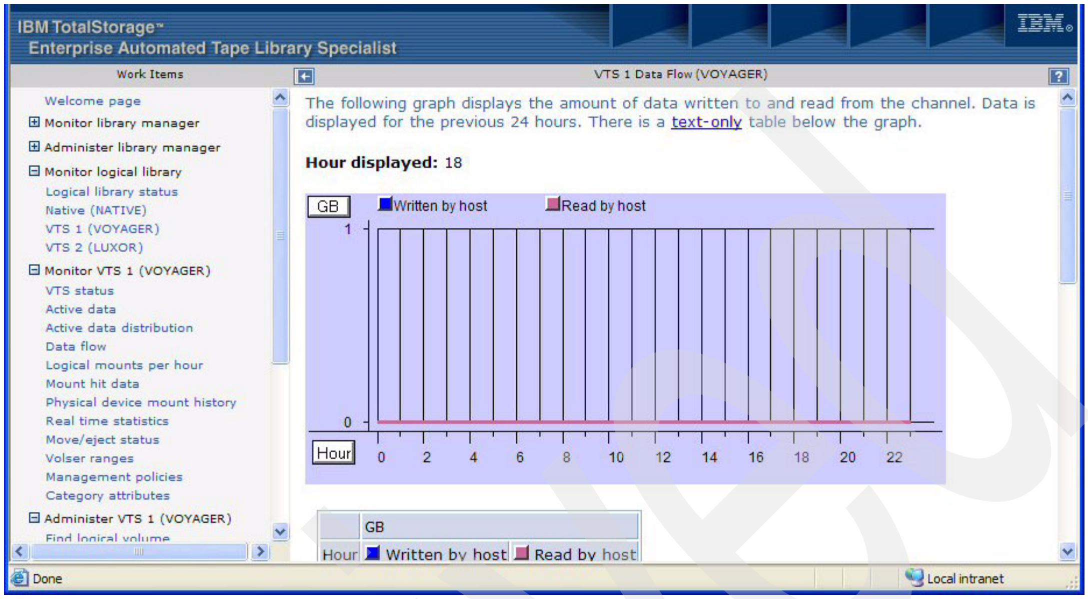

Figure 7-12 Data flow

The logical mounts display

Figure 7-13 shows the number of logical mounts per hour, which includes fast ready mounts, cache hits, and mounts for recalls .

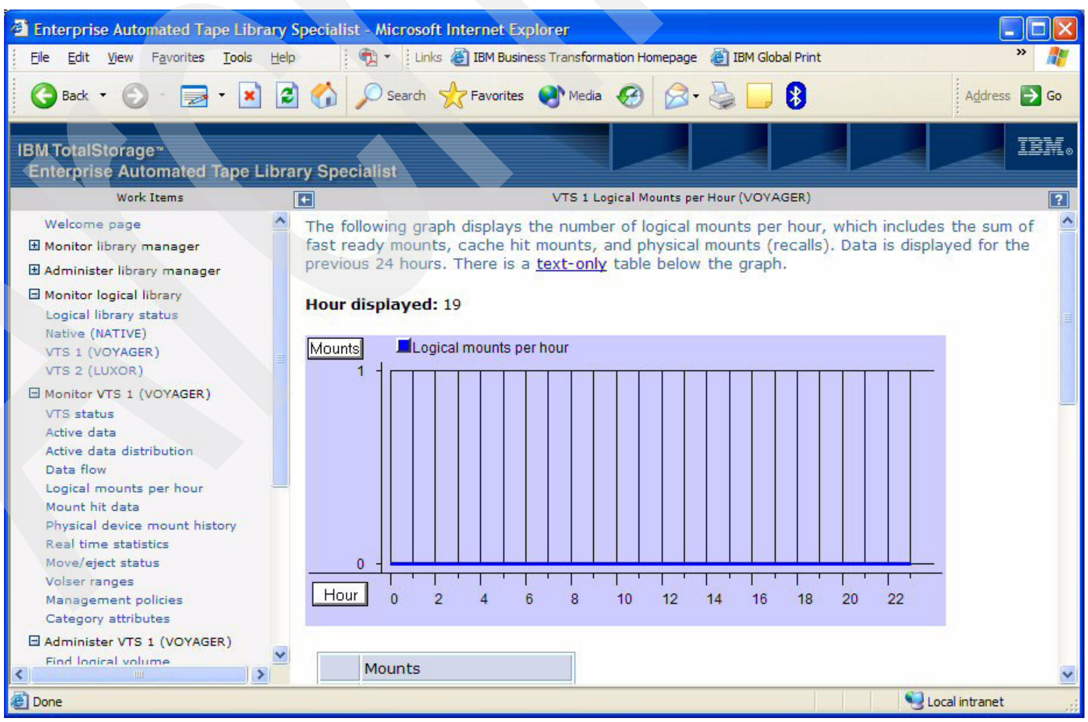

Figure 7-13 Logical mounts

Physical mounts

The graph in Figure 7-14 displays the maximum, average, and minimum numbers of physical drives and the time to mount stacked volumes.

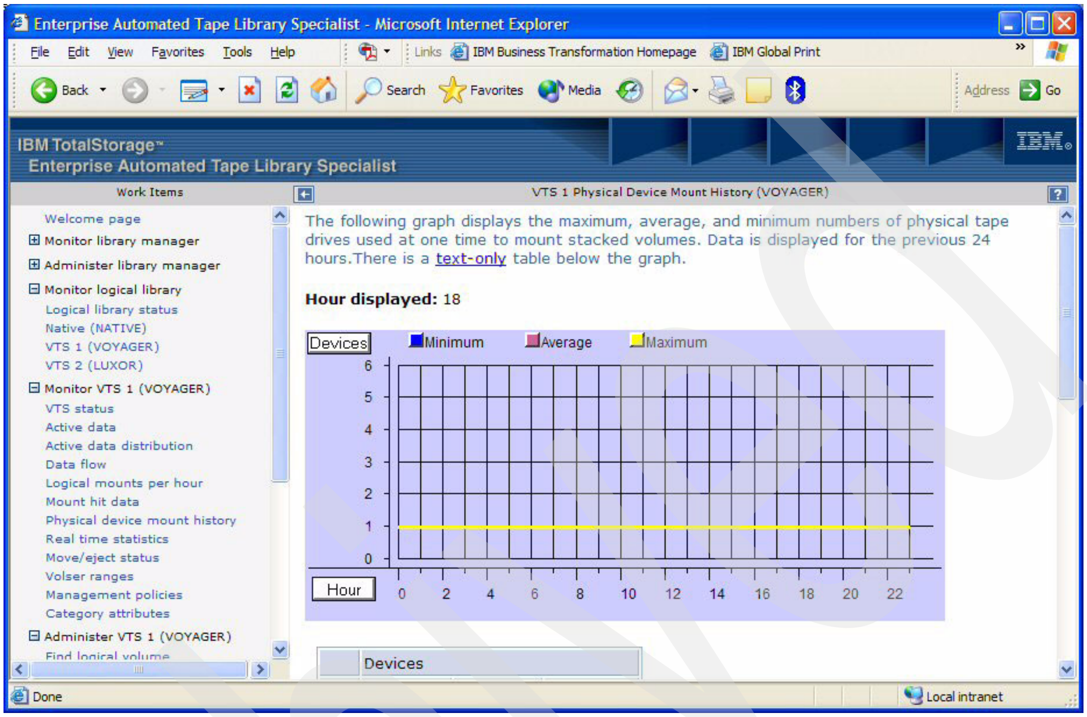

Figure 7-14 Physical mounts

> **Note:** There are many more information and status screens available that were not presented in this section.

7.2 Using VTSSTATS for performance evaluation and reporting

VTSSTATS is a program already used for performance evaluation and reporting of stand-alone VTSs, which has been enhanced to provide reports for the Peer-to-Peer VTS. With the links described in the following paragraph, you have access to the IBMTOOLS library, which contains different jobs to analyze the Library Manager logs if you have your VTS attached to a non-MVS system, and to the VTSSTATS tool, which uses SMF94 records for a comprehensive performance analysis. In this section we describe some examples based on the VTSSTATS reporting tool. The flat files called DAILY and HOURLY may be used as input to a graphic tool that generates multiple charts (see 7.2.10, “VTSSTATS graphical reporting” on page 276).

Every hour the 3494 Library Manager sends VTS statistics records to all attached hosts. In z/OS, these statistics are recorded as SMF type 94 records. See Appendix D, “SMF type 94 record layout” on page 415 for a detailed description of the contents of this record. VTSSTATS produces reports from SMF type 94 records and outputs flat files that are suitable for use with a spreadsheet program. You can find this tool as well as the spreadsheet program on IBMTOOLS.EXE.

7.2.1 VTSSTATS tool overview

The following sections provide a description of VTSSTATS.

What it is used for:

- Monitoring VTS drive and Tape Volume Cache (TVC) activity
- Doing trend analysis to see when upgrade is needed

What it does:

- Analyzes and reports on SMF 94 data
- Summarizes VTS activity on hourly and daily basis
- Reports virtual drive activity
- Number of drives allocated
- Number of mounts (Fast ready, read hits, read misses)
- Reports physical drive activity
- Number of drives allocated
- Mounts for recall, copy, reclaim
- Tape volume cache activity
- MBs read and written to TVC
- MBs read and written to physical tape
- Virtual volume size
- Logical volumes managed

What information in SMF 94 will be reported:

- 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%, etc. of active data
  - Reclaim threshold
- Min/Max/Avg mount times for fast-ready, read hits and read misses
- Compression ratio achieved in the ESCON card
- Compression ratio achieved between TVC and Drive
- Block sizes of data written to VTS
  - Average channel block size
  - Number of blocks written
  - Distribution of block sizes
    - 0 to 2K, 2K to 4K, 4K to 8K, 8K to 16K,
    - 16K to 32K, 32K to 64K, greater than 64K
- Reports by physical volume pool
- Reports by cache preference level

The following reports are available:

- Hourly active data and reclamation
- Hourly performance indicators
- Hourly import - export activity
- Hourly complib VTC activity (Peer-to-Peer ONLY)
- Hourly composite library activity (Peer-to-Peer ONLY)
- Hourly physical drive activity
- Hourly virtual drive activity
- Hourly Tape Volume Cache activity
- Hourly preference level virtual mounts
- Hourly cache IART mounts statistics
- Hourly pool active data and reclamation
- Hourly volume poolset statistics

- Hourly volume pool media statistics
- Hourly preference level virtual mounts
- Hourly preference level cache statistics

7.2.2 Tools download

The VTSSTATS tool used for VTS monitoring and the performance evaluation tool will be discussed in this chapter; they can be located at the following sites:

- Tools can be downloaded internally from the following intranet site:

  `ftp://submit.boulder.ibm.com/download/tapetools/`

- Business partner tools can be located at the following Internet site:

  `ftp://vtstools@service2.boulder.ibm.com`

- Customer tools can be located at the following Internet site:

  `ftp://ftp.software.ibm.com/storage/tapetool/`

The tools have an expiration mechanism, because there are so many updates for these tools to provide support for the newest hardware, or new functions. Therefore these tools expire to give you the opportunity to obtain the latest versions.

We recommend to examine the UPDATES.TXT file for updates to the tools library and to download the tools library at least every six months to prevent an unexpected expiration.

7.2.3 How to get started

After you have downloaded and installed the tools library, customize the VTSSTATS JCL and define the VTS or PtP VTS that you want to be reported (see Figure 7-15).

(VTSSTATS) - Control cards

Required for PTP

```
************************************************************************
*        FILL IN THE FOLLOWING RECORDS AS APPROPRIATE:                 *
************************************************************************
 CUSTOMER= THIS IS RELEASE THREE(3) OF VTSSTATS, BETA;<=CUSTOMER
*EUROFORMAT;
 REPORT=   ALL                                  ;<=REPORT TYPE
*SDATE=    04DEC2000                            ;<=START DATE OF INPUT
*SHOUR=    00                                   ;<=START HOUR ON INPUT
*EDATE=    04DEC2000                            ;<=END DATE OF INPUT
*EHOUR=    23                                   ;<=END HOUR ON INPUT
*LOGICALDAYSTART= 08                            ;<=2-DIGITS 00-23
**** INCLUDING VTSNUM= STMNT(S) WILL SELECT ONLY THOSE VTSS FOR RPRTNG
**** OMITTING VTSNUM= STMNTS WILL CAUSE ALL VTSS TO BE SELECTED
*VTSNUM=   059C2  059C8  21252 70677 70680      ;VTS'S
*VTSNUM=   70689  70695  70809 70810            ;VTS'S
*VTSNUM=   BA001  BA009  BA035 BA055            ;VTS'S
*VTSNUM=   BA008  BA040  BA043 BA044 BA045 BA071 ;COMPOSITE LIBS
*VTSNUM=   BA040  BA40A BA40B                   ;ONE COMPLIB CONFIG
*VTSNUM=   BA008  BA08A BA08B                   ;ONE COMPLIB CONFIG
**** GROUP= STMNTS CAUSE NO OPERATIONAL EFFECTS, FOR ID PURPOSE ONLY
*GROUP=  NAME   MODE  3494#  VTS1#   VTS2#      ;<=* PROTOTYPE
*GROUP= GRP001  NAT   10365                     ;<=*
*GROUP= GRP001  VTS   10365  30123  30124       ;<=*
*GROUP= GRP003  MIX   10397  60185              ;<==
*GROUP= GRP003  MIX   10397  60185  60191       ;<==
**** PTPVTS= AND &&PTPVTS= STMNTS ARE PAIRED VIA THE NAME ASSIGNED
*PTPVTS=    NAME    3494#                       ;PROTOTYPE
*&&PTPVTS=  NAME    3494#/VTS#    3494#/VTS#    ;PROTOTYPE
 PTPVTS=   JOEPTP   3301A                       ;<=COMPOSITE LIB ID
 &&PTPVTS= JOEPTP   71629/33010   71622/33011   ;<PEER2PEER PAIRS
/*
//*
//SMFDATA  DD  DISP=OLD,DSN=JBACCO.CUSTOMER.SMF94.DATA
```

Figure 7-15 VTSSTATS Control definitions

The PtP VTS name and library IDs have to be defined. Otherwise you do not get the VTC data reported.

7.2.4 Summary report

We recommend that you generate all reports which are provided by VTSSTATS. Based on the SMF94 input you get per default hourly reports and a summary report of each day and a total summary of all days.

The first report to be examined is the overall summary report, which summarizes all daily reports (see Figure 7-16 to Figure 7-19). You get a summary per day with the following topics:

- Physical drive activity
- Virtual drive activity
- Percentage of mounts not satisfied in cache, that is, cache miss
- Daily cache activity
- Daily performance indicators such as :
  - Blocksize distribution
  - Compression ratio
  - Throttling
- Daily VTC activity

Daily VTS Activity - page 1

```
DAYPHYS  GROUP: EXAMP1 -DAILY REPORT                 VTSNUM=D1055 *** Peer To Peer VTS***                   VTSSTATS(05145)
______________________EXAMPLE____________________    L1/V1= (17260,70277) L2/V2= (18454,B2339)        Run Date: 19:31 25MAY2005
Report Date 31MAR2005/00:42 to 06APR2005             Channels(ES+FI=16 SCSI=00)  DRIVES=12/12           Data timestamp: 00:42:13
Logical Day Start hour: 00            -PAGE 1 OF 4-  TVC size=1728/1728GB                                           LM ref hour: 23

                                           SMF94 FLD  31MAR05TH 01APR05FR@02APR05SA@03APR05SU 04APR05MO 05APR05TU 06APR05WE
VTS micro code level                     lvvcm/lvvcv
Library Manager micro code level         lvlmv/lvlmr
DAILY PHYSICAL DRIVE ACTIVITY (DEVTYP1/DEVTYP2)
    Maximum physical drives mounted............. maxcm    23/   0  24/   0  24/   0  24/   0  24/   0  24/   0  24/   0
    Average physical drives mounted............. avgcm   4.5/  .0 10.0/  .0  9.9/  .0  2.1/  .0  1.2/  .0  1.0/  .0  1.1/  .0
    Maximum time to mount physical drive (sec)...maxttm  954/   0  928/   0  785/   0  558/   0 1042/   0  482/   0  523/   0
    Average time to mount physical drive (sec)...avgttm   92/   0   93/   0   90/   0   73/   0   82/   0   78/   0   72/   0
    Physical mounts for recall..................stgmnts 1151/   0 1543/   0 1506/   0  121/   0  711/   0  665/   0  545/   0
    Physical mounts for copy....................migmnts  414/   0  726/   0  667/   0  281/   0  561/   0  574/   0  523/   0
    Physical mounts for reclaim.................recmnts  118/   0  212/   0  160/   0  143/   0  293/   0  268/   0  301/   0
    Total physical mounts for vts................  ...       1683      2481      2333       545      1565      1507      1369
    Total physical mounts for native.............  ...       1476      2287      2142       416      1309      1286      1110
    Virtual volumes premigrated..................  vmp       9559     15045     16427      7715     14372     16156     14310
DAILY VIRTUAL DRIVE ACTIVITY
    Maximum virtual volumes ready............... maxvdm       139       165       231        70       107       201       174
    Average virtual volumes ready............... avgvdm        45        51        70        18        35        35        27
    Maximum virtual mount time (sec)............   vrx       1525     28339     25838       812      1420      1364       896
    Average virtual mount time (sec)............   vra       71.6      45.9      31.5      11.7      33.4      27.8      22.2
    Total virtual mounts........................   ...       6034      8405     10874      4845      9750     10944      9490
    Fast ready - Number of mounts...............   vfr       4580      5782      7134      3876      7362      7971      7139
               - Maximum mount time (sec)........ maxfr       524       565       118        79       820       196       119
               - Minimum mount time (sec)........ minfr         1         1         1         1         1         1         1
               - Average mount time (sec)........ avgfr         7         5         5         4         6         5         4
    Cache hits - Number of mounts...............   vmh        868      2022      3347       841      1655      2270      1831
               - Maximum mount time (sec)........ maxch       258       462       542       134       476       136       188
               - Minimum mount time (sec)........ minch         1         1         1         1         1         1         1
               - Average mount time (sec)........ avgch        14        13        15         8         9         7        11
    Recalls    - Number of mounts...............   vms        586       601       393       128       733       703       520
               - Maximum mount time (sec)........ maxrm      1528      1458     26053       817      1421      1366       897
               - Minimum mount time (sec)........ minrm         5         5       138        61         2        59         8
               - Average mount time (sec)........ avgrm       661       548       660       265       363       352       314
    Percentage of mounts not satisfied in cache..  ...        9.7       7.1       3.6       2.6       7.5       6.4       5.4
    Virtual volume residency (sec) - Maximum.....  vvx      20465     31860     64735      9536     16076     10726     10654
    Virtual volume residency (sec) - Minimum.....  vvn          1         1         1         1         1         1         1
    Virtual volume residency (sec) - Average.....  vva      509.0     588.6     685.0     264.3     327.9     270.7     228.9
CONTINUED ON NEXT PAGE
```

Figure 7-16 VTS Daily Statistics Summary Report

Daily VTS Activity (cont'd)

```
DAILY TAPE VOLUME CACHE ACTIVITY
    Megabytes written to cache...................   vbw    3631628  5228437  6954647  3759467  7145464  8230666  6410594
    Megabytes read from cache....................   vbr    1453088  2102298  4650500   728785  2183697  2066449  1630745
    Maximum channel megabytes per second.........   ...     159.24   182.12   213.15   144.84   183.17   207.07   165.57
    Avg Max channel megabytes per second.........   ...        n/a      n/a      n/a      n/a      n/a      n/a      n/a
    Average channel megabytes per second.........   ...      88.27    84.84   134.31    54.20   107.97   119.17    93.06
    Ratio: max to avg channel megabytes per sec..   ...       1.80     2.14     1.58     2.67     1.69     1.73     1.77
    Megabytes written to physical 3590..........ppw+spw    4038277  5801477  6390059  3176469  6309725  7271113  6005838
    Megabytes read from physical 3590...........    n/a         na       na       na       na       na       na       na
    Weighted average hh:mm virt volume in cache..   vca      na:na    na:na    na:na    na:na    na:na    na:na    na:na
    Maximum hh:mm of virt vol in cache .......... mtvca    2698:27  2722:27  2745:27   306:36   277:10   295:23   319:23
    Average virt vol size in TVC (mb)............   vcz        382      404      471      510      610      582      559
    Average size all virt vols (mb)............ vba/vla        875      870      433      887      872      879      887
    Maximum number of virt vols in cache .......   vvic       9897    10483     9937     5372     5585     8171     7645
    Virtual volumes premigrated..................   vmp       9559    15045    16427     7715    14372    16156    14310
    Last hr # of virt vols managed by vts........   vla     195436   195087   195386   193790   199131   198299   193881
    Last hr private stacked volume count......... prict       4952     4961     3010     4859     4875     4926     4929
    Last hr scratch stacked volume count......... srtct        270      290      108      452      435      384      381
```

Figure 7-17 VTS Daily Statistics Summary Report sample, continued

Daily VTS Activity (cont'd)

```
DAYPHYS  GROUP: EXAMP1 -DAILY REPORT                VTSNUM=D1055 *** Peer To Peer VTS***                  VTSSTATS(05145)
____________________EXAMPLE___________________      L1/V1= (17260,70277) L2/V2= (18454,B2339)      Run Date: 19:31 25MAY2005
Report Date 31MAR2005/00:42 to 06APR2005            Channels(ES+FI=16 SCSI=00)  DRIVES=12/12        Data timestamp: 00:42:13
Logical Day Start hour: 00            -PAGE 2 OF 4- TVC size=1728/1728GB                                    IM ref hour: 23

    Last hr active data on stacked cartridges(gb)   vba     171125    169896     84705    172016    173765    174366    172049
    Last hr available stacked capacity (gb)......   vec      27133     29737      8131     40851     39134     34783     34211
DAILY IMPORT - EXPORT ACTIVITY
    Physical volumes imported...................    im1          0         0         0         0         0         0         0
    Logical volumes imported....................    im2          0         0         0         0         0         0         0
    Total megabytes imported....................    im3          0         0         0         0         0         0         0
    Megabytes moved from stacked to stacked vol..   im4          0         0         0         0         0         0         0
    Physical volumes exported...................    ex1          0         0         0         0         0         0         0
    Logical volumes exported....................    ex2          0         0         0         0         0         0         0
    Total megabytes exported....................    ex3          0         0         0         0         0         0         0
    Megabytes moved from stacked to stacked vol..   ex4          0         0         0         0         0         0         0
DAILY PERFORMANCE INDICATORS (*=hourly maximums)
      Percent of <= 2k blocks written.............    0kb        10*        7*        6*        8*       11*       10*       10*
    Percent of <= 4k blocks written.............    2kb         1*        1*        2*        2*        1*        1*        1*
    Percent of <= 8k blocks written.............    4kb         0*        0*        5*        0*        1*        1*        1*
    Percent of <=16k blocks written.............    8kb         4*        7*        6*        2*        1*        2*        4*
    Percent of <=32k blocks written.............   16kb        99*      100*      100*       96*       91*       65*       89*
    Percent of <=64k blocks written.............   32kb        83*       63*       48*       76*       86*       85*       81*
    Percent of > 64k blocks written.............   64kb         0*        0*        0*        0*        0*        0*        0*
    Total number of blocks written..............    ...  91649858K 43025644K  8706027K  1173322K 72912451K    27759K 63183344K
    Average blocksize written...................    ...     41000*    37000*    34000*    40000*    42000*    42000*    41000*
    Compression ratio into cache................ harat      4.30*     3.68*     4.01*     3.18*     3.35*     3.49*     4.75*
    Compression ratio to 3590................... bsrat      1.92*     1.92*     1.92*     1.92*     1.92*     1.92*     1.92*
    Weighted avg compression ratio into cache....   ...       2.3       2.7       OFL       2.4       2.2       OFL       2.3
    Maximum recall throttling percentage......... rcprt        0*        0*        0*        0*        0*        0*        0*
    Maximum write throttling percentage.......... wrovt        4*        5*        0*        0*        0*        0*        0*
    Average recall throttling value (msec)....... avrct        0*        0*        0*        0*        0*        0*        0*
    Average write throttling value (msec)........ avwot      847*     4495*        0*        0*        0*        0*        0*
    Average write throttling value (msec)........ avwot      847*     4495*        0*        0*        0*        0*        0*
    Overall throttling value (msec).............. totat       35*      112*        0*        0*        0*        0*        0*
CONTINUED ON NEXT PAGE
```

Figure 7-18   VTS Daily Statistics Summary Report sample, continued

Daily VTS Activity - (cont'd)

```
DAYPHYS  GROUP: EXAMP1 -DAILY REPORT                    VTSNUM=D1055 *** Peer To Peer VTS***                 VTSSTATS(05145)
______________________EXAMPLE_____________________      L1/V1= (17260,70277) L2/V2= (18454,B2339)       Run Date: 19:31 25MAY2005
Report Date 31MAR2005/00:42 to 06APR2005                Channels(ES+FI=16 SCSI=00)  DRIVES=12/12        Data timestamp: 00:42:13
Logical Day Start hour: 00            -PAGE 3 OF 4-     TVC size=1728/1728GB                                     LM ref hour: 23

    Reclaim threshold percentage.....see HOURPOLR thres         n/a       n/a       n/a       n/a       n/a       n/a       n/a
    Number of cartridges that can be reclaimed...   ...         n/a       n/a       n/a       n/a       n/a       n/a       n/a
DAILY VIRTUAL CONTROLLER ACTIVITY
    Comp lib...max logical volumes to be copied.. cllvc       3333*     2490*     1269*     1226*      379*     1547*      374*
    Comp lib...max mbs of data yet to be copied.. cldtc   15632540* 11728919*  6264857*  2513805*  1796887*  6062849*   766893*
    Comp lib...mounts completed for vts-0........ clmt0       5943      8328     10827      4845      9723     10929      9456
    Comp lib...mounts completed for vts-1........ clmt1         91        77        47         0        27        15        34
DAILY VIRT.CONT. VTC ACTIVITY
    Megabytes copied:.......................vtc-0 cldc0     285935    351780    426470    203897    403580    397273    426546
     Volumes copied:............................0 clvc0        659       995      1140       505       886       935       934
      Host read mbytes transferred:.............0 clrd0     157033    235339    500502    102884    276751    360914    173182
       Host write mbytes transferred:...........0 clwd0     492249    575758    616896    487733   1058554   1144128    764578
        Number of fast ready mounts:............0 clcm0        601       734       776       490       989      1000       839
         Number of cache hit mounts:............0 clsm0         88       245       356       118       225       315       245
          Number of recall mounts:..............0 clrm0         69        84        60        15       112       111        46
           Total mounts:........................0  ...        758      1063      1192       623      1326      1426      1130
    Megabytes copied:.......................vtc-1 cldc1     256486    350807    360151     82709    309860    357630    358235
     Volumes copied:............................1 clvc1        650       889       975       183       655       822       851
      Host read mbytes transferred:.............1 clrd1     257763    236166    746484     58681    268540    254598    222285
       Host write mbytes transferred:...........1 clwd1     474978    747276    848099    539894    880492   1243681    863913
        Number of fast ready mounts:............1 clcm1        627       757       876       478       897      1193       859
         Number of cache hit mounts:............1 clsm1        137       239       483        81       227       291       245
          Number of recall mounts:..............1 clrm1         88        77        47        14        76        83        81
           Total mounts:........................1  ...        852      1073      1406       573      1200      1567      1185
    Megabytes copied:.......................vtc-2 cldc2     307201    351496    449528    247998    505582    565965    382563
     Volumes copied:............................2 clvc2        750       893      1137       637      1080      1231       935
      Host read mbytes transferred:.............2 clrd2     176541    261395    563342     88025    250358    180082    221205
       Host write mbytes transferred:...........2 clwd2     315529    635398    837830    364275    565274    631177    603547
        Number of fast ready mounts:............2 clcm2        399       712       907       456       696       726       839
         Number of cache hit mounts:............2 clsm2         92       245       429       107       210       250       251
          Number of recall mounts:..............2 clrm2         73        75        37         2        65        73        69
           Total mounts:........................2  ...        564      1032      1373       565       971      1049      1159
    Megabytes copied:.......................vtc-3 cldc3     266587    305637    372720    193547    449657    492172    311962
     Volumes copied:............................3 clvc3        626       855      1033       459       966      1127       827
      Host read mbytes transferred:.............3 clrd3     141673    252788    513320     92703    264286    232441    162428
       Host write mbytes transferred:...........3 clwd3     510164    696414    890981    515432    778097   1043003    929798
        Number of fast ready mounts:............3 clcm3        631       779       975       486       848      1024       934
         Number of cache hit mounts:............3 clsm3         88       286       414       100       210       294       214
```

Figure 7-19 VTS Daily Statistics Summary Report sample, continued

7.2.5 Critical performance values

For reporting and statistical purposes, you may extract the key values, such as:

- Daily average virtual mount time
- Hourly maximum virtual mount time
- Hourly mounts not satisfied in cache
- Total MB/s throughput
- Scratch cartridge capacity left
- How many concurrent physical drives are used

If you have to do some performance evaluation or problem determination for your PtP VTS environment, then as with any other performance analysis, you have to follow the guideline to first look at the daily statistics of multiple days or of one day only and then, if there are any abnormalities, examine the hourly reports.

To be able to make the best use of the information provided in the reports and files, you must have an understanding of some of the key fields in the type 94 records. Throughout this section, we describe these key fields as necessary. In order to evaluate the performance of your VTSs, you should note the following information regarding your configuration and setup:

- Model - base or with FPAF

- Cache size
- Number of physical drives
- Number and type of channels attached to the VTS
- Distance between both sites
- Switches and/or DWDMs in this configuration

The following list of questions can be used to evaluate the performance of the VTS. These questions have been formulated based on some rules of thumb or guidelines that have been developed both from field experience and laboratory measurements and have proven valuable in the evaluation of VTS performance:

1. **Is your daily "Average virtual volume mount time" greater than 30 seconds?**

   Virtual mount response time will vary based on a number of factors, such as whether the mount request specifies a specific or non-specific volume to be mounted, whether a specific volume is in the Tape Volume Cache or must be recalled from a stacked volume and the overall workload of the library system.

   A daily average virtual mount time of 30 seconds or less is usually considered an acceptable level of mount performance.

2. **Do you have any hourly "Max virtual mounts" greater than 300 seconds or any hourly "Average virtual mounts" greater than 30 seconds?**

   Other useful information covers these questions:

   – How many concurrent virtual drives are being used?
   – How many mounts are being done?
   – What is the percentage of recalls?

   Figure 7-20 shows the VTSSTATS Hourly Virtual Drive Activity Report.

```
1GROUP: DUMMY  -  HOURLY VIRTUAL DRIVE ACTIVITY          VTSNUM=60514                                          VTSSTATS(12MAR1)
 ________SUPPLY YOUR TITLE INFORMATION HERE________      L1/V1= n/a           L2/V2= n/a                    Run Date: 13:41 05APR2001
 Report Date 28MAR2001/01:04                             Channels(ESCON=00 SCSI=00)  DRIVES= 6                Data timestamp: 01:04:59
 Logical Day Start hour: 00                              TVC size=0000GB                                           Libmgr ref hour: 03

            VIRTUAL      VIRTUAL
            VOLUMES       MOUNT     TOT  -----FAST READY-----  ------MOUNTS--------  ----MOUNTS NOT------  %MNTS  --VVOL RESIDENCY
             READY      TIME (SEC)  VMTS        MOUNTS          SATISFIED IN CACHE    SATISFIED IN CACHE   NOT    ------(SEC)-----
   HOUR    MAX   AVG    MAX   AVG          #   MAX  MIN  AVG    #   MAX  MIN  AVG    #   MAX  MIN  AVG   SAT      MAX   MIN   AVG
  smf94==>(vdx) (vda)  (vrx) (vra)       (vfr)  ---SECONDS---  (vmh)  ---SECONDS---  (vms)  ---SECONDS---         (vvx) (vvn) (vva)

    00     16     9    259    12    66    61    0    0    0     2    0    0    0     3    0    0    0     4    2799     9   431
    01     17    12    683    14    60    47    0    0    0    12    0    0    0     1    0    0    0     1    4311    12   589
    02     16    11     25     2    61    48    0    0    0    13    0    0    0     0    0    0    0     0    5311     7   728
    03     14     1     93     6    23    18    0    0    0     4    0    0    0     1    0    0    0     4    3710     5   658
    04      9     5     17     2    26    20    0    0    0     6    0    0    0     0    0    0    0     0    2897     4   701
    05      9     1      5     2    23    16    0    0    0     7    0    0    0     0    0    0    0     0    2459     5   320
    06      8     3    213    18    36    18    0    0    0    13    0    0    0     5    0    0    0    13    1773     5   306
    07      6     2    281    25    32    10    0    0    0    18    0    0    0     4    0    0    0    12     569     1   176
    08      5     2    498    65    24     0    0    0    0    19    0    0    0     5    0    0    0    20    2636    33   429
    09      2     1     73    11    18     1    0    0    0    13    0    0    0     4    0    0    0    22     515     2   128
    10      3     1    196    28    11     0    0    0    0     9    0    0    0     2    0    0    0    18     585     8   194
    11      3     1      2     1     5     1    0    0    0     4    0    0    0     0    0    0    0     0    2175    64   547
    12      3     1     36     3    42    20    0    0    0     6    0    0    0    16    0    0    0    38     476     1    78
    13      3     1     13     5     6     2    0    0    0     4    0    0    0     0    0    0    0     0    2939   167   666
    14      5     2    409    89    24     9    0    0    0     6    0    0    0     9    0    0    0    37     895     3   280
    15      6     4    502   210    26     6    0    0    0     0    0    0    0    20    0    0    0    76    2767    54   410
    16      8     4    637   158    46     6    0    0    0     4    0    0    0    36    0    0    0    78   10725    23   465
    17      5     3    589    51    71     7    0    0    0    38    0    0    0    26    0    0    0    36    2054     3   181
    18      5     3     38     4    15     4    0    0    0    11    0    0    0     0    0    0    0     0    1272    12   378
    19      2     1     16     4    16     4    0    0    0    12    0    0    0     0    0    0    0     0    4451    11   696
    20      5     2    134    34    13     5    0    0    0     0    0    0    0     8    0    0    0    61    4077    10   769
    21      4     1    235    77    41     2    0    0    0     2    0    0    0    37    0    0    0    90    3068     9   165
    22      6     1    198    46    63     9    0    0    0     4    0    0    0    50    0    0    0    79     735     1    45
    23     11     2    534    51    64    22    0    0    0    10    0    0    0    32    0    0    0    50    3169     1   193
                                ----- -----                -----                -----                ---
 DATE                             812   336                  217                  259                 31
```

Figure 7-20 Hourly Virtual Drive Activity Report

The fields in question are called SMF94VRX and SMF94VRA (Maximum and Average time in seconds used to complete a mount request on a virtual drive in the last hour).

If you have exceeded any of the two guidelines indicated for these records, it is likely to be caused by virtual volumes having to be recalled from the stacked volumes when drives are not available. This is an indication that the physical tape drives are over committed. See question 5 on page 269. In the sample report, you can see that the two fields in question do not have any values exceeding the guidelines.

Recall is driven by specific mounts that are cache misses. See question 3 on page 268.

During the recall process, the VTS subsystem needs to mount the physical stacked volume that holds the logical tape volume that is being requested by the host. The number of concurrent recall operations that are allowed is a function of several factors including the number of tape drives available and the amount of free space in the TVC. Queuing delays may occur if the VTS is recalling multiple logical volumes which are located on the same stacked volume. Another possible cause of long mount times could be accessor contention for non-VTS drives which are in the same 3494.

Please also be aware of these factors:

– If there are no drives available, the recall mount will be queued until a drive becomes available.

– If the virtual volume to be recalled resides on a stacked volume that is in use by a copy task or another recall task, the recall mount will be queued until the copy or other recall task completes.

– If the virtual volume to be recalled resides on a stacked volume that is in use by a reclaim task (target or source), then the mount will be queued until the reclamation completes processing the current logical volume, after which the reclamation task will be terminated and the recall mount will be processed.

Jobs requiring access to multi-volume datasets on virtual volumes, no longer in the TVC, are likely to see some delays in their mount times.

Other fields in the Hourly Virtual Drive report are explained here:

– **Number of Fast Ready mounts (SMF94VFR)** should be a nonzero value during a time period when scratch mounts are occurring. If a scratch mount occurs for a tape that is not in the Fast Ready category, the VTS has to mount the old stacked volume in order to recall the logical volume before it is reused as a scratch tape. If the category has fast-ready turned on, the old stacked volume will not be mounted.

– **Number of mounts satisfied in TVC (SMF94VMH)** shows the number of specific mounts for a logical volume which have been satisfied without having to mount the stacked volume (cache hits). These types of mounts will be satisfied very quickly.

– **Number of mounts not satisfied in TVC (SMF94VMS)** shows how many times stacked volumes were mounted to satisfy a request for a specific virtual mount (cache miss). This number may be different from the physical stage number because the VTS attempts to keep a stacked volume mounted. If another recall is attempted for a different logical volume on the same stacked volume, another mount is not needed, thereby, reducing accessor activity.

– **Virtual Volume residency time (SMF94VVX,VVN,VVA)** specifies the maximum, minimum and average time in seconds that a virtual volume was mounted on a virtual drive in this VTS during the last hour. Time is accrued from completion of the mount until a demount is issued.

3. **Are the hourly_% mounts not satisfied in TVC greater than 20%?**

   The percentage is shown in the report. It also can be calculated by dividing the number of virtual mounts not satisfied in TVC (SMF94VMS) by the total virtual mounts, times 100. If you are not using VTSSTATS, then the total virtual mounts is the sum of (SMF94VFR,SMF94VMH and SMF94VMS).

   Typically, if the number of cache misses is less than 20%, then the overall average virtual mount time should be about 30 seconds or less.

4. **How much data is my VTS managing and how much more data can it store before it becomes full?**

   Much more information is available as well:

   – How many MB/sec am I processing?
   – How many scratch stacked volumes are available?
   – How many volumes are in cache?
   – What is the average size of volumes in cache?
   – What is the total number of volumes being managed?

   This is measured by capacity used in MB (SMF94VBA) and capacity available in MB (SMF94VEC) of the defined and inserted stacked volumes. These volumes store virtual volumes and are either scratch, full, filling, or partially full (once full but now holding virtual volumes that may have expired).

   For a detailed analysis of the management of these volumes, refer to appropriate section.

   Figure 7-21 shows the VTSSTATS Hourly Tape Volume Cache Activity Report.

HOURLY TAPE VOLUME CACHE ACTIVITY (CHANGE)

```
HOURCACH PTPVTS=EXAMP1 - HOURLY TAPE VOLUME CACHE ACTIVITY VTSNUM=D1055 *** Peer To Peer VTS***            VTSSTATS(05145)
______________________EXAMPLE_____________________         L1/V1= (17260,70277) L2/V2= (18454,B2339)       Run Date: 17:14 25MAY2005
Report Date 01APR2005/00:42(FRI)                           Channels(ES+FI=16 SCSI=00)  DRIVES=12/12         Data timestamp: 00:42:33
Logical Day Start hour: 00                                 TVC size=1728/1728GB                                      LM ref hour: 00

                                                                                                      --STACKED--  -PHY TAPE CAPACITY-
      MBYTES   MBYTES   CHAN   MBYTES  MBYTES   MBYTES  AVG MIN  MAX MIN  ---VIRTUAL VOLUMES---  LOG VOLS  VOLUME
     WRITTEN  RD FROM  MB PER WRITTEN WRITTEN  RD FROM  VIRT VOL  VV IN  AVG SIZE  #IN  AVG SZ  MANAGED   COUNT          USED      AVAIL
 HOUR TO CACHE  CACHE   SECOND TO PRIM TO SECD PHYSICAL IN CACHE   CACHE   ALL(MB) CACHE IN CAC  BY VTS    PRI    SCR    (GB)      (GB)
smf94==>(vbw)  (vbr)          (ppwrit)(spwrit)  (vtr)    (vca)   (mtvca)   (calc) (vvic)  (vcz)   (vla)   (prict)(srtct)  (vba)     (vec)
ptpnote:  *5      *5            *4      *4       *2       *2        *2              *4            *1      *4     *4      *4        *4

 00    519266  136366 182.12  336647       0      na       na    161967      435   8971   389  196K/198K 4956   266    171406    26809
 01    470667  129820 166.80  355425       0      na       na    162027      434   9786   356  196K/199K 4960   262    171692    26418
 02    269568  111028 105.72  325949       0      na       na    162087      434   9959   351  196K/199K 4961   260    171931    26256
 03    288987   96504 107.08  301524       0      na       na    162147      434  10177   345  197K/200K 4967   254    172203    25749
 04    364946   67833 120.21  262083       0      na       na    162207      434  10483   338  194K/200K 4966   255    171490    25886
 05    341266   68667 113.87  314313       0      na       na    162267      435   9983   337  192K/201K 4971   250    171006    25310
 06    332047   63844 109.96  323248       0      na       na    162327      434   9884   340  193K/201K 4973   248    171288    25264
 07    128306   36912  45.89  299355       0      na       na    162387      434   9847   343  193K/202K 4978   243    171554    24780
 08     98504   49702  41.16  246169       0      na       na    162447      434   9868   344  194K/202K 4982   239    171760    24480
 09     82028   40065  33.91  210816       0      na       na    162507      434   9812   348  194K/202K 4985   236    171977    24271
 10    185857   61414  68.68  176436       0      na       na    162567      435   9663   355  194K/197K 4985   255    170438    26166
 11    282889   90500 103.71  222491       0      na       na    162627      435   9308   371  194K/193K 4980   271    169028    27758
 12    166818   75984  67.44  275837       0      na       na    162687      436   9113   378  195K/194K 4982   270    169293    27711
 13    226629   20704  68.70  234559       0      na       na    162747      435   8992   384  195K/194K 4981   269    169497    27664
 14    143650  118284  72.75  244038       0      na       na    162807      436   8836   392  194K/194K 4979   271    169233    27940
 15     35422   41723  21.42  207998       0      na       na    162867      436   8709   400  193K/194K 4981   269    169173    27755
 16     69560  142183  58.81  161277       0      na       na    162927      436   8241   398  194K/194K 4979   272    169290    28124
 17     35859  105207  39.18  165332       0      na       na    162987      436   8206   402  194K/193K 4976   275    168919    28492
 18     97673  194331  81.11  186023       0      na       na    163047      436   8254   397  194K/193K 4974   277    169024    28745
 19    165966  100875  74.12  164088       0      na       na    163107      436   8380   400  194K/194K 4966   285    169174    29207
 20    223879   55904  77.71  150644       0      na       na    163167      436   8580   396  194K/194K 4959   292    169302    29875
 21    230366   77260  85.45  190589       0      na       na    163227      435   8687   394  195K/194K 4957   294    169471    30084
 22    246513  115285 100.49  236211       0      na       na    163287      435   8529   403  195K/194K 4959   292    169691    29921
 23    221771  101903  89.90  210425       0      na       na    163347      435   8521   404  195K/195K 4961   290    169896    29737
      ------- -------         ------- ------- -------
DATE  5228437 2102298         5801477       0      na
peer-to-peer notes: *1=Vts1/Vts2, *2=max(V1,V2), *3=min(V1,V2), *4=V1+V2, *5=Complib
```

Figure 7-21 VTS Hourly TVC Activity Report

The appropriate SMF fields are explained here:

– **Number of Logical Volumes managed by VTS (SMF94VLA)** indicates the number of logical volumes managed by VTS. This actually includes all the logical volumes that have ever been used by the VTS, PLUS every instance of every volume that has been used by the VTS since the VTS has last performed reconciliation.
– **Capacity used in MB (SMF94VBA)** contains the sum of the compressed megabytes accumulated on the SMF94VLA logical volumes.
– **Capacity available in MB (SMF94VEC)** is calculated from the number of stacked volumes in the scratch category. Partially filled stacked volumes are not included in this calculation. **# Virtual Volumes premigrated (SMF94VMP)** refers to the number of copy to tape processes that have been completed. A high number here indicates high write activity, which is normal after heavy application write activity in prior intervals, where the VTS might not have been able to keep pace with the host data rate.
– **Average size of Virtual Volume (SMF94VCZ)** specifies the average size in MB of all logical volumes that are managed by the VTS. It does not include any volumes that are in cache. It is a calculated field, (SMF94VBA divided by SMF94VLA). The size is reported as the compressed size in case of an EHPO or PA configuration.
– **# Virtual Volumes in TVC (SMF94VNM)** specifies the number of logical volumes held in the TVC. If any of these volumes were requested by the host, it would result in a cache hit.

5. **Is performance constrained by the number of drives available?**

The additional relevant information covers:

– How many concurrent physical drives are being used?
– What is the average physical mount time?
– How many physical mounts am I doing?

Background tasks within the virtual tape server controller can also influence overall performance. Those background tasks include space reclamation, logical volume recalls (cache misses) and the copy process, all of which use the physical drives managed by the VTS. If your analysis proves that you are constrained by the total number of drives you have defined, then you may consider upgrading or adding more tape drives.

Figure 7-22 shows the VTSSTATS Hourly Physical Drive Activity.

HOURLY PHYSICAL DRIVE ACTIVITY

```
HOURPHYS PTPVTS=EXAMP2 - HOURLY PHYSICAL DRIVE ACTIVITY    VTSNUM=D1049 *** Peer To Peer VTS***                 VTSSTATS(05145)
______________________EXAMPLE_____________________         L1/V1= (13977,70774) L2/V2= (12022,B2433)      Run Date: 17:14 25MAY2005
Report Date 01APR2005/00:35(FRI)                            Channels(ES+FI=16 SCSI=00)  DRIVES=12/12        Data timestamp: 00:35:07
Logical Day Start hour: 00                                  TVC size=1728/1728GB                                     LM ref hour: 00
```

	--PHYSICAL--- ---DRIVES----		--PHYSICAL DRIVES-- --CONCURRENT MOUNTS			--SECONDS TO MOUNT-- ---PHYSICAL DRIVES--			ACCSSOR_MNTS (VTS+NATIVE)		TOTAL MNTS	---------------------- --PHYSICAL MOUNTS FOR--				VIRTUAL VOLUMES
HOUR	INSTALL	AVAIL	MAX	MIN	AVG	MAX	MIN	AVG	-A-	-B-	VTS	RECALL	COPY	RCLM	SDE	PREMIG
smf94==>	(vti)	(vta)	(vtx)	(vtn)	(vtv)	(vmx)	(vmn)	(vmv)	(aca) +	(acb)	(sum) =	(vps) +	(vpm) +	(vpr) +	(sde)	(vmp)
ptpnote:	*1	*1	*1	*1	*1	*2	*2	*2	*4	*4		*4	*4	*4	*4	*4
00	12/12	12/12	12/11	10/10	11/10	274	50	151	2	16	15	0	14	1	0	1498
01	12/12	12/12	12/11	10/11	11/11	482	52	126	11	19	20	3	17	0	0	1451
02	12/12	12/12	12/11	11/10	11/10	77	51	178	3	23	22	2	20	0	0	1086
03	12/12	12/12	12/11	12/10	12/10	313	46	107	1	34	34	11	21	2	0	854
04	12/12	12/12	12/11	12/ 9	12/10	88	51	67	1	25	25	4	13	8	0	424
05	12/12	12/12	12/11	11/ 6	11/10	68	46	215	9	23	25	0	18	7	0	957
06	12/12	12/12	12/11	11/11	11/11	345	50	115	3	30	30	10	16	4	0	1079
07	12/12	12/12	12/11	10/11	11/11	407	53	182	12	20	21	0	12	9	0	626
08	12/12	12/12	12/11	8/ 9	11/10	364	48	113	13	31	31	9	14	8	0	552
09	12/12	12/12	12/11	12/ 9	12/10	302	53	185	3	51	51	18	24	9	0	921
10	12/12	12/12	12/11	10/10	11/10	363	49	98	1	38	38	10	16	12	0	794
11	12/12	12/12	12/11	9/10	11/10	343	51	104	2	50	50	27	14	9	0	181
12	12/12	12/12	12/11	10/ 8	11/10	327	47	124	1	56	56	25	16	15	0	186
13	12/12	12/12	12/11	11/10	11/10	341	47	109	1	48	48	18	17	13	0	531
14	12/12	12/12	12/11	10/ 8	11/10	106	49	125	2	46	46	16	22	8	0	890
15	12/12	12/12	12/11	10/10	11/10	355	51	92	1	54	54	25	17	12	0	402
16	12/12	12/12	12/11	11/ 9	11/10	386	51	107	1	66	66	28	25	13	0	198
17	12/12	12/12	12/11	11/10	11/10	345	51	98	0	28	28	7	12	9	0	244
18	12/12	12/12	12/11	11/ 9	11/10	310	49	73	2	22	23	2	9	12	0	210
19	12/12	12/12	12/11	12/ 8	12/10	277	50	146	0	27	27	4	13	10	0	388
20	12/12	12/12	12/11	10/ 9	11/10	290	50	82	11	26	25	3	12	10	0	449
21	12/12	12/12	12/11	11/ 8	11/10	351	48	128	19	29	29	8	15	6	0	299
22	12/12	12/12	12/11	12/ 8	12/10	75	47	108	16	27	28	4	17	7	0	854
23	12/12	12/12	12/11	12/11	12/11	81	48	77	9	22	24	2	16	6	0	777
											-----	----	----	----	---	-----
DATE											816	236	390	190	0	10570

Figure 7-22 PtP VTS Hourly Physical Drive Activity

Here we describe the associated SMF records:

– **Mounted physical drive count (SMF94VTA,VTX,VTV)** numbers (actual, maximum and average) should indicate how busy the drives are. If the average numbers are close to the number of drives that you have configured, you may be constrained by the number you have. These values are established by periodic sampling. Tape drive residency time for stacked volumes tends to be high during periods of low host activity because of asynchronous copy processes to write logical volumes from TVC to stacked volumes. That is why you should use a daily summary to evaluate the physical drive usage.

  Alternatively, you may be constrained by the size of your TVC, because a constrained TVC tends to cause excessive movement of logical volumes to and from the stacked volumes.

– **Physical mount activity (SMF94VPS,VPM,VPR)** numbers (recall, copy, reclaim) are used to determine the usage of the physical tape drives.

  Recall indicates how many times stacked volumes were mounted to copy virtual volumes back into the TVC. The number of mounts for *copy* tells you how much physical mount activity is related to writing new data out to the stacked volumes. This number may not be very high because the VTS tries to keep the stacked volumes mounted when it can. *Reclaim* is used to move active data from one stacked volume to another once it has reached its reclamation threshold, therefore, returning the source to scratch. A high reclaim mount number could indicate that you have set the reclaim threshold at too high a value. See “Reclamation and reconciliation” on page 161 for more details.

Other possible reasons for high reclaim mount numbers include these:

- A previous time interval with high cache usage rate prevented reclamation.
- There has been a recent change of reclaim threshold on the Library Manager.
- A recent heavy use of virtual scratch volumes that had been previously used and, therefore, became eligible for reclamation from the stacked volumes
- The average expiration period is very small.
- You have crossed the stacked volumes threshold this is set to the minimum amount of volumes that the VTS allows before making reclaim a priority. You can issue the **`D SMS(libname),DETAIL`** command, which will show you the number of Scratch Stacked Volumes that you currently have available.

VTC statistics

In the composite library activity reports, we get the throughput and activity numbers of all VTCs reported. In the first report, shown in Figure 7-23, we get the list of logical volumes and MB to be copied and the completed logical mounts for VTS-0 and VTS-1. The high number of logical volumes and MB yet to be copied and the high number of mounts on one VTS only indicates that the VTS-1 in this example was not active for a certain amount of time.

Hourly Composite Library Activity For PtP

```
HOURP2P  PTPVTS=EXAMP1 - HOURLY COMPOSITE LIBRARY ACTIVITY VTSNUM=D1055 *** Peer To Peer VTS***            VTSSTATS(05145)
____________________EXAMPLE_________________________         L1/V1= (17260,70277) L2/V2= (18454,B2339)      Run Date: 17:14 25MAY2005
Report Date 01APR2005/00:42(FRI)                             Channels(ES+FI=16 SCSI=00)  DRIVES=12/12       Data timestamp: 00:42:33
Logical Day Start hour: 00                                   TVC size=1728/1728GB                                   LM ref hour: 00

H         LOGICAL_____VOLS       MBs__ OF___DATA          MOUNTS COMPLETED       MOUNTS COMPLETED
O         TO BE COPIED____       YET TO BE COPIED         FOR VTS-0_______       FOR VTS-1______
U              cllvc                  cldtc                    clmt0                  clmt1
R               *5                     *5                       *5                     *5

 00           2490                11728919                     997                      0
 01           2388                11478831                    1036                      0
 02           2295                11003097                     390                      2
 03           2266                10776697                     642                      2
 04           2170                10476911                     468                      0
 05           2077                10073664                     377                      0
 06           1995                 9695065                     393                      1
 07           1928                 9380732                     267                      9
 08           1886                 9176136                     151                      1
 09           1845                 8991362                     174                      2
 10           1801                 8847626                     204                      6
 11           1770                 8674593                     240                      3
 12           1727                 8424429                     329                      1
 13           1687                 8224173                     184                      3
 14           1641                 7983482                     189                      0
 15           1590                 7729998                     178                      1
 16           1556                 7603548                     175                     40
 17           1527                 7433889                     100                      2
 18           1488                 7257720                     276                      0
 19           1469                 7155704                     275                      1
 20           1430                 6960096                     409                      0
 21           1388                 6824711                     367                      0
 22           1349                 6664108                     254                      0
 23           1292                 6451771                     253                      3
peer-to-peer notes: *1=Vts1/Vts2, *2=max(V1,V2), *3=min(V1,V2), *4=V1+V2, *5=Complib
```

Figure 7-23   Composite library - hourly total VTC activity

The following VTC report shows the hourly activity of each VTC in the PtP VTS. This helps to identify any inbalances of activity that reduce the overall performance. Figure 7-24 displays, for each VTC, the volumes and megabytes copied, host read and write activity, and the mount characteristics (fast ready, cache hit or cache miss). The compression ratio for each VTC is displayed as well.

Hourly Complib VTC Activity For PtP

```
HOURVTC  PTPVTS=EXAMP1 - HOURLY COMPLIB VTC ACTIVITY       VTSNUM=D1055 *** Peer To Peer VTS***                 VTSSTATS(05145)
____________________EXAMPLE___________________             L1/V1= (17260,70277) L2/V2= (18454,B2339)       Run Date: 19:11 25MAY2005
Report Date 01APR2005/00:42(FRI)                           Channels(ES+FI=16 SCSI=00)  DRIVES=12/12          Data timestamp: 00:42:33
Logical Day Start hour: 00                                 TVC size=1728/1728GB                                       LM ref hour: 00

    _MBYTES_ VOLS__ HOST_READ HOST_WRT FAST_ CACHE RCALL  VTCn        _MBYTES_ VOLS__ HOST_READ HOST_WRT FAST_ CACHE RCALL  VTCn
    _COPIED_ COPIED MB__XFRRD MB_XFRRD READY _HIT_ MNTS_  COMP        _COPIED_ COPIED MB__XFRRD MB_XFRRD READY _HIT_ MNTS   COMP
                                       MNTS_ MNTS_        RATIO                                        MNTS_ MNTS_        RATIO
     cldc*   clvc*    clrd*     clwd*  clcm* clsm* clrm*  clcr*        cldc*   clvc*    clrd*     clwd*  clcm* clsm* clrm*  clcr*
     *5      *5       *5        *5     *5    *5    *5     *5           *5      *5       *5        *5     *5    *5    *5     *5
    ______________________hour 00__________________________           ______________________hour 01__________________________
VTC0   30726    114     10168    31249    71    35     0  3.98  VTC0   32196    157     10698    25736    69    50     0  3.15
VTC1   25685     58     18425    78508   107    23     5  2.90  VTC1   20212     91     24122    59984    78    40     8  2.96
VTC2   25482     68     21176   103587   103    25     1  2.84  VTC2   21281     94     14443    88222    98    58     0  2.90
VTC3   23072     81     23778    55900    80    42     0  3.10  VTC3   20116     96     22575    42641    57    62     0  3.84
VTC4   19410     70     20768    95055    97    45     4  3.49  VTC4   17205     85     15451    78376    88    39     9  4.34
VTC5   25818     87     11485    73063    98    32     1  2.68  VTC5   21491    107      2787    49516    83    31     0  3.10
VTC6   26289     83     19736    43658    90    33     0  3.13  VTC6   13116     93     12411    43550    71    48     2  3.40
VTC7   26779     77     10826    38242    69    34     2  3.54  VTC7   23044     90     27329    82640    85    48    12  3.52
    ______________________hour 02__________________________           ______________________hour 03__________________________
VTC0   22861     52     14265    17424    24    16     3  2.25  VTC0   19851     73      5039    20982    60    14     0  1.94
VTC1   16567     31     21655    64256    48    21     6  2.86  VTC1   20684     60     26152    28222    47    29     0  2.66
VTC2   18953     44      9416    51729    47    12     1  2.53  VTC2   20995     68     16339    34044    70    20     0  2.19
VTC3   14383     31      5930    27633    37    15     0  2.47  VTC3   18500     63     11922    75245    95     9     0  2.26
VTC4   23107     56     23409    29076    33    16     1  3.06  VTC4   23182     72      9611    42345    62    11     1  2.56
VTC5   20037     52      7413    26120    24    11     0  2.64  VTC5   13542     51     15650    30432    62    14     0  2.55
VTC6   21499     50     17857    27375    29    13     1  3.06  VTC6   15804     63      8031    25612    68    15     2  2.08
VTC7   16644     41     11079    25951    20    12     2  3.39  VTC7   20929     71      3757    32102    58     7     0  2.28
    ______________________hour 04__________________________           ______________________hour 05__________________________
VTC0   21633     55      9425    61218    65     5     5  1.91  VTC0   21416     50      7932    52454    44     7     0  2.59

UP to 24 HOURS
peer-to-peer notes: *1=Vts1/Vts2, *2=max(V1,V2), *3=min(V1,V2), *4=V1+V2, *5=Complib
```

Figure 7-24 Composite library - hourly single VTC activity

6. **Are you experiencing elongated or erratic job run times?**

   Elongated or erratic job run times can be caused by inbalanced load (read and writes) to the VTCs, not available VTCs, and long copy queues because one VTS was in service mode for example.

   If all components of the PtP VTS are available, these long job run times may be caused by a high level of throttling by the VTS. Throttling is directly reported by the VTS, in VTSSTATS. Also, RMF™ reports the VTS delay on the host write I/O as Disconnect Time.

   You can use RMF reports or monitors in conjunction with the VTSSTATS report to check the times when you suspect VTS throttling may have occurred. If you have 15 minute intervals when the disconnect time for VTS virtual drives exceed 500ms, then this is an indication that throttling is occurring.

7.2.6 More about throttling

The VTSSTATS data reports on two general types of throttling: recall and write. Recall throttling was something that happened for our original B16, but rarely, if ever, happens with B10, B18, or B20 models. This leaves write throttling, which involves slowing down host writes.

Write Throttling

There are three types or causes of write throttling, as described below.

Free-space throttling: This type of write throttling will occur if the amount of unused TVC space falls below a certain threshold, which is a few GB and varies by VTS model and channel type.

Pre-migrate throttling: This type of write throttling will occur when the pre-migrate threshold is reached. This threshold again varies depending on configuration, roughly between 60% and 80% 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 VTS 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.

Copy throttling: This type of write throttling will occur only for a PtP, and then for immediate copy mode when copying starts to take *too long*. (Write throttling can actually also occur if the amount of uncopied data in the TVC reaches a certain threshold, which would normally happen only for deferred copy mode.)

So, what is meant by “too long”? This is controlled by an algorithm that considers not only how long a copy is taking, but how much work is going on in the PtP. The idea is to try to recognize early on if a copy is in danger of taking more than 45 minutes, so that host write activity can be throttled sufficiently to allow the copy to finish in time. This might in many cases occur if a copy time exceeds twenty minutes, although this time can be shorter if the PtP is very busy.

Unfortunately, the SMF type 94 data does not tell us which type of write throttling is occurring. However, since copy throttling may usually not occur in deferred mode, you have the chance to isolate the cause of throttling (here it may be immediate mode with long copy times) by temporarily changing the mode to deferred.

If throttling went away when the customer switches from immediate to deferred copy mode, it is likely that the type of write throttling that had been occurring was copy throttling due to long copy times.

VTSSTATS reporting of throttling

VTSSTATS reprots throttling. Here is a brief explanation of the VTSSTATS values:

Maximum write throttling percentage (wrovt): This is the hour in the reported day with the highest percentage of write throttling. For your report, this means that there was one hour in the day where write throttling occurred 60% of the time. For all other hours, write throttling occurred less than or equal to 60% of the time.

Average write throttling value (msec) (avwot): This is the hour in the reported day with the highest average recall throttle value, which is the "average of the non-zero throttling values where write overrun was the predominant reason for throttling. The value is how long the VTS delays responding to the host."

Explanations of these and other SMF type 94 fields can be found in Appendix D, “SMF type 94 record layout” on page 415.

Immediate to Deferred switch

For immediate copy mode, the copy must complete before the VTS will give Rewind/Unload completion. This assures that the copy has finished before the job ends. However, this also requires that the copy complete in 45 minutes, which is the MIH time value set for VTS virtual drives. If the copy does not complete in 45 minutes, then the PtP reverts to deferred copy mode to avoid the time-out and resulting job failure. In this case the job completes, but the copy is not complete.

To avoid this situation, the VTS has implemented *copy throttling* for immediate copy mode. This will slow down incoming host traffic if a copy is starting to take too long, so that the copy will go faster and hopefully complete in the required 45 minutes. Whether copy throttling is invoked depends on the workload and size of the volumes being copied. In deferred copy mode, completion of the job is not tied to completion of the copy, so the copy can exceed 45 minutes and there is no need to invoke copy throttling. Going with TIME(0) means that copies will always run in *fast copy mode*, which tries to get the copy done as quickly as possible, even though there is no particular time requirement. If, for example, you specified TIME(4), then a copy would go into fast copy mode if it were not completed in four hours.

7.2.7 Warning and alerts

If a secondary copy operation cannot be completed immediately, it is very likely that this was caused by the failure of a component in the Peer-to-Peer VTS or internal thresholds have been met. Upon a component failure, the immediate copy mode may operate as if in deferred copy mode until the failing component has been brought back online. For details on the possible recovery scenarios, please refer to Chapter 9, “Disaster recovery scenarios” on page 321.

If a secondary copy operation is delayed, the host is notified with a warning message (CBR3787E) that indicates that immediate copy mode operations are deferred in a specific composite library. The warning message is retained until all the immediate mode copy operations that were deferred have completed.

Alert messages

We introduced new pre-migration threshold messages using the CBR3750I message format which have to be enabled by the CE after code activation.

When the VTS transitions into the “Above Threshold Warning” state, it provides a warning to the host operator that throttling may be started; the internal “Above Threshold Warning” state is reported, and this message will be sent once every hour while in this state.:

```
CBR3750I Message from library library-name: OP0160 Above threshold for uncopied data in
cache, throttling possible
```

When the VTS transitions out of the “Above Threshold Warning” state, the following message will be provided:

```
CBR3750I Message from library library-name: OP0161 Below threshold for uncopied data in
cache
```

The “Above Threshold Warning” state

The VTS enters the “Above Threshold Warning” state when the amount of data to copy exceeds the threshold for the installed cache capacity for five consecutive sample periods (amount of data to copy is sampled every 30 seconds). The VTS leaves this state when the amount of data to copy is below the threshold. This consecutive sampling criteria is to prevent excessive messages being generated on the host.

7.2.8 How to solve throughput bottlenecks

If the workload is appropriate, in any case unchangeable, or an increase in the number of MB/s written and read was ascertained and the VTS is indicating signs of stress, that is, throttling with the result that the batch window becomes too small for all batch processing, then it may be necessary to add system capacity.

Upgrading the TVC capacity can provide improvements in virtual mount times, throughput, and read cache hits and offer protection against throttling.

If the VTS subsystem is configured with less than the maximum allowable drives in the VTS and you are suffering from long virtual mount times in case of a cache miss, consider installing additional tape drives.

Also, especially in conjunction with the upgrade from 3590 to 3592 drives, consider the installation of the FICON Performance Accelerator Feature - FPAF.

Again, we recommend to analyze your environment using the BatchMagic tool in order to get a sound projection. You may compare the throughput numbers before the PtP VTS was installed and a certain amount of time later to examine any significant changes of the load profile.

7.2.9 Other VTSSTATS reports

Examples of the other reports that were added as part of the functionality update for Advanced Policy Management are in the "Performance and Monitoring" chapter of the redbook, *IBM TotalStorage VTS Planning, Implementation, and Monitoring*, SG24-2229.

New reports

Here is a list of new reports:

- Hourly preference level virtual mounts:
  - Shows the mounts in PG0
- Hourly preference level cache statistics:
  - For how long data is kept in cache for PG1
- Hourly cache IART mount statistics:
  - Shows these mounts
- Hourly pool active data and reclamation:
  - How many cartridges are available
  - Hows how many can be reclaimed
- Hourly volume poolset statistics:
  - Media usage for all pools
- Hourly volume pool media statistics:
  - Number of media used per pool

7.2.10 VTSSTATS graphical reporting

All daily and hourly reports may be graphically preprocessed and can be used for day-by-day reporting. Perceptible deviations become more transparent in a graph than as a table or a flat file.

For example, in Figure 7-25, Figure 7-26, Figure 7-27, and Figure 7-28, we present graphs for stacked volume capacity, daily megabytes transferred, virtual, and physical mounts by type, and VTC activities. You can find the tools under:

`ftp://ftp.software.ibm.com/storage/tapetool/`

In Figure 7-25, the chart on the right shows in the solid bars, the amount of data written; and in the striped bars, the amount of data read.

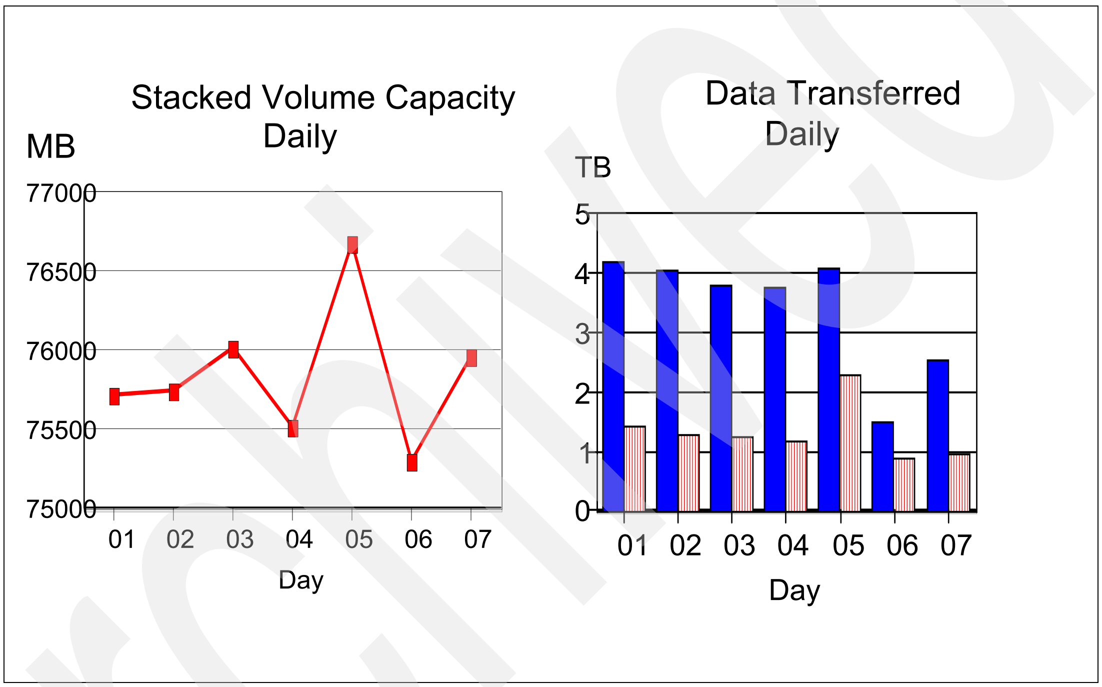

Figure 7-25 Graphics for Stacked Volume Capacity and for Data Transferred

The chart in Figure 7-26 illustrates how much physical mounts can be avoided by using a VTS: only the striped top portion of each of the bars (Misses) requires a physical mount. TVC Hits and Fast Ready mounts are satisfied from the TVS without requiring a physical mount.

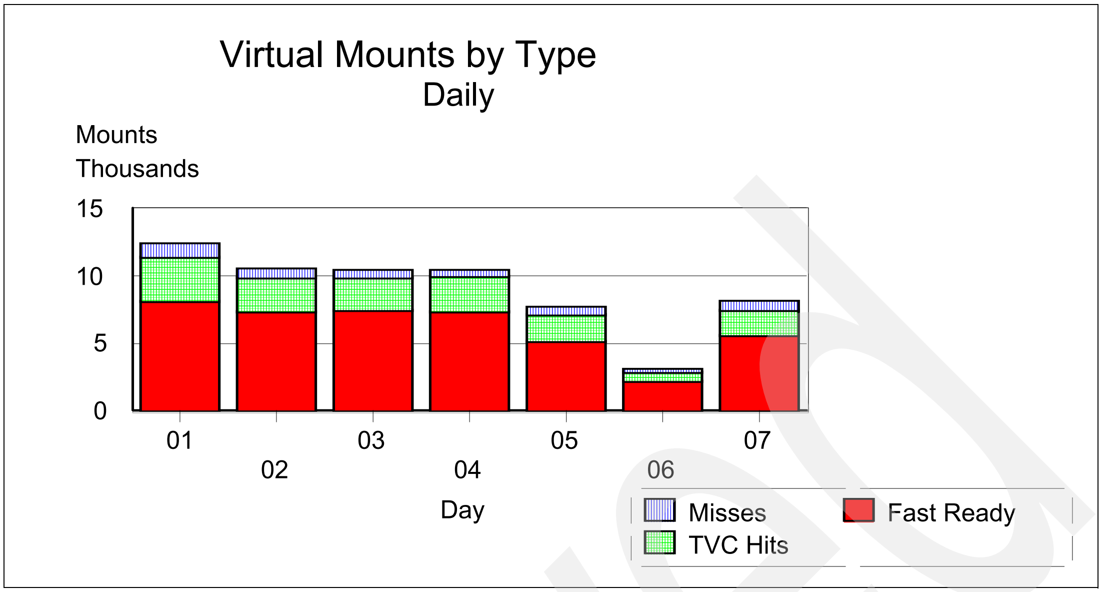

Figure 7-26 Graphics for virtual mounts

Figure 7-27 categorizes the physical mounts. The left, solid bar of each group *(Recall)* refers to the *Misses* bar shown in Figure 7-26.

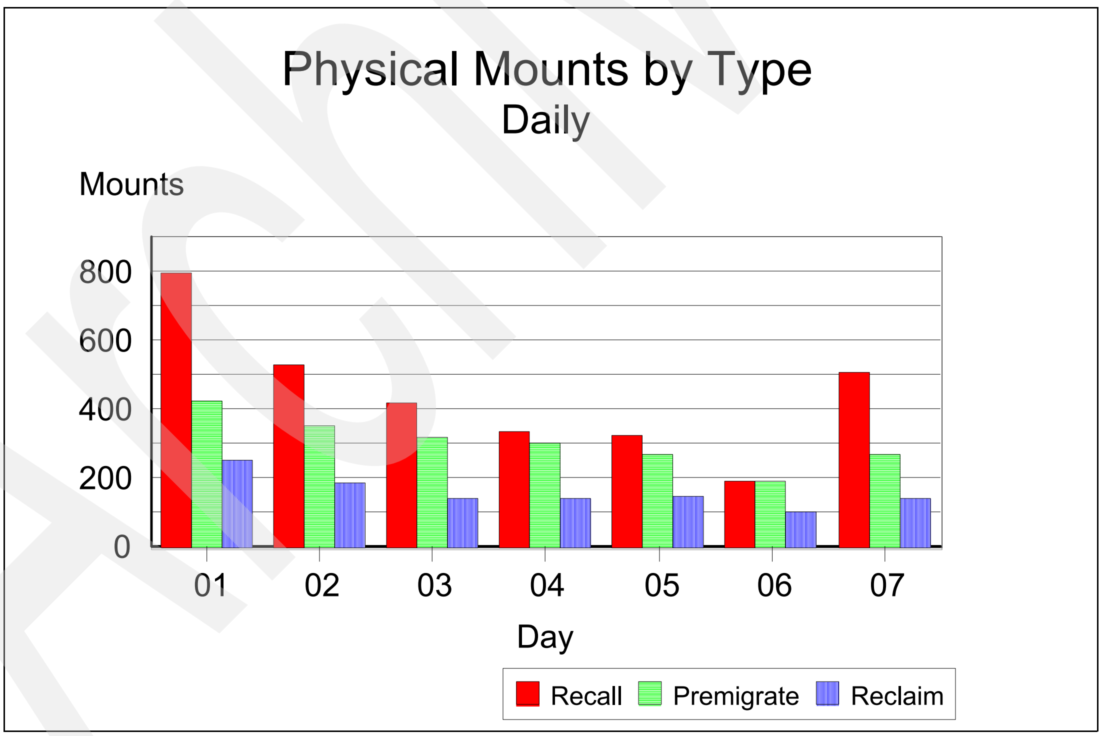

Figure 7-27 Graphics for physical mounts

While the charts shown in Figure 7-25 through Figure 7-27 are available for a standalone VTS and can therefore also be created for each distributed VTS in a PtP VTS separately, the charts in Figure 7-28 are only available for PtP VTS. They show statistics taken at the VTSC level.

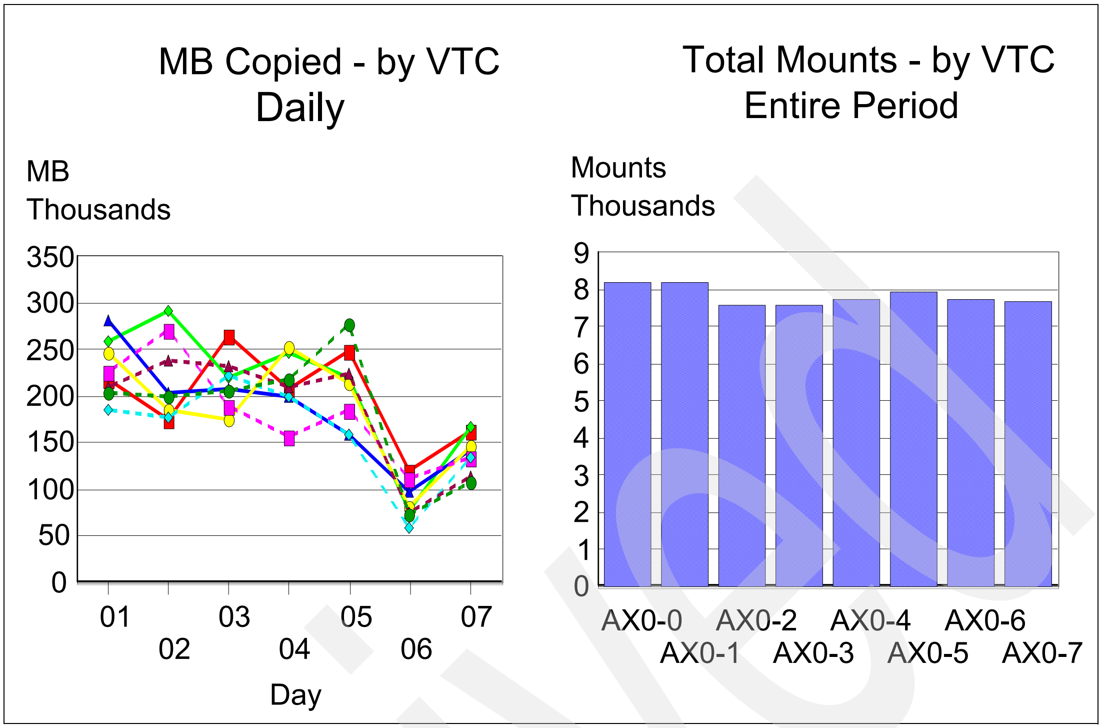

Figure 7-28 VTC mounts and copied data

7.3 Bulk volume information retrieval (BVIR)

As the capability of the base VTS and its high availability Peer-to-Peer (PTP VTS) configuration has grown, so has the need to efficiently manage the large number of logical volumes a VTS supports. The VTS architecture currently supports a host interface that allows an application to obtain information about a specific logical volume, but is not an efficient method to use when information for a large number of logical volumes is required. With BVIR you are now able to obtain information about all of the logical volumes a VTS manages:

- Physical volume to logical volume mapping information
- Cache content information
- Peer-to-Peer volume status information

7.3.1 Overview

With the potential to support hundreds of thousands of logical volumes in a VTS subsystem, 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 VTS subsystem that allows it to virtualize a tape volume, also allows for a simple and effective method to transfer the information to a requesting application. The VTS converts the format and storage conventions of a tape volume into a standard file managed by a file system within the VTS.

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 information is also presented in human readable form, so no special interpretation software is required either.

Attention: The BVIR function in a VTS is initially disabled. To use the function, an IBM Service Representative enables it through the service panels provided on the VTS.

If the function is disabled and a request formatted volume is created, the VTS will not report an error nor will it examine or modify the volume. A subsequent read of the volume will only result in the request records being read.

Once enabled, obtaining information from a VTS involves two steps.

First, a single data set with the information request is written to a logical volume. The logical volume can be any logical volume in the VTS 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 specifies 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 close of the volume, the VTS will recognize it as a request volume and 'prime' the VTS for the next step. Note that in a PTP VTS, the PTP Selective Dual Copy function must be used to direct the request volume to the specific VTS in the PTP configuration that the information is needed from.

Second, the request volume is again mounted, this time as a specific mount. Seeing that the volume was "primed" for a data request, the VTS 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. Once the VTS 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. Like the request records, the response records are also in human readable form, that is, lines of character data.

Note: 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.

Once the response data set has been written to the request logical volume, that logical volume functions identically to any other logical volume in the VTS. 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.

Important: In the PtP VTS environment, you must define from which distributed VTS you want to obtain the information. To do so, use the I/O VTS definition in the Data Class construct.

Figure 7-29 shows the process flow of BVIR.

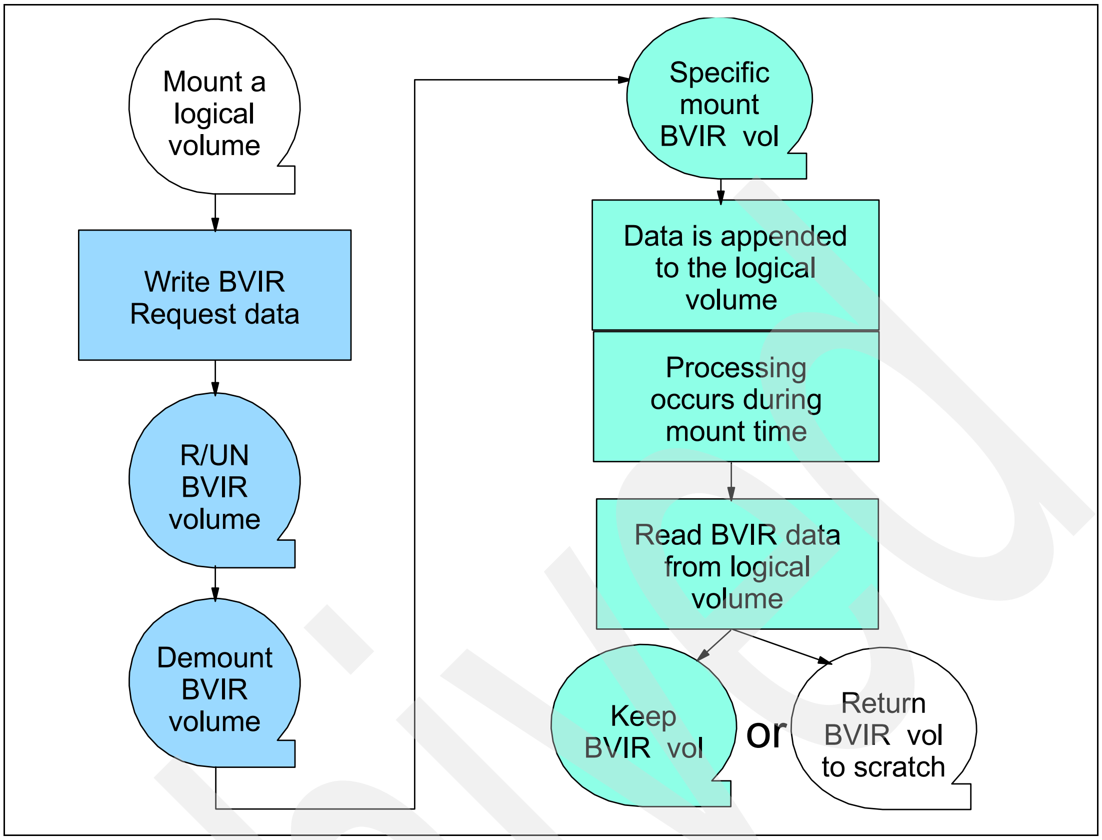

Figure 7-29 BVIR Process flow

The building of the response information does require a small amount of the resources of the VTS. 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. To minimize the use of the function from impacting normal VTS processing, the VTS automatically 'throttles' the number of requests it will handle by setting a minimum request handling time of approximately two minutes. Since host systems are tolerant of a wide range of mount response times, using a minimum handling time is the simplest way to control the impact on the VTS of the requests.

7.3.2 Prerequisite

BVIR was initially introduced with VTS code level 2.30.720.40 which was made available on 5/21/2004. Additional information is provided with VTS code level 2.3x.740.xx. Talk with your IBM representative about obtaining the necessary level. Although there are no library manager code changes to support the BVIR function, there may be other functions in a VTS code level that require a specific library manager code level. There are no host software updates required for this function.

7.3.3 Requesting a report

Any logical volume defined to a VTS can be used as the request/response volume. Logical volumes in a VTS 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. The 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.

We recommend normal scratch allocation methods for each request (DISP=(NEW,CATLG)). In this way, any of the available scratch logical volumes in the VTS 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).

Several types of data can be requested from the VTS. 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 FB 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.

The format for the request data set records is shown in Table 7-1.

Table 7-1 BVIR Request Records

Record 1: Request Identifier		
Bytes	**Name**	**Contents**
1 - 28	Request Identifier	VTS BULK VOLUME DATA REQUEST
29 - 80	Blanks	Blank padding
Record 2: Request		
1 - 40	Request	'PTP VOLUME STATUS' or 'CACHE CONTENTS' or 'VOLUME MAP' left justified, padded with blanks on the right
41 - 80	Blank	Blank padding
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 being failed.		
Note: In a PTP VTS, the PTP Selective Dual Copy function must be used to direct the request volume to the specific VTS in the PTP configuration that the information is needed from.		

7.3.4 Report output

When the request data set has been written to the volume and subsequently closed and demounted, when mounted again, the VTS will validate the contents of the request volume and append the requested data records to the data set. All appended records are 80 bytes in length. 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 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.

> **Note:** In the response records, the date and times presented are all based on the internal clock of the VTS handling the request. The internal clock of a VTS is not synchronized to the host or any other VTS.

The general format for the response data set is shown in Example 7-1.

Example 7-1 BVIR output format

```
VTS BULK VOLUME DATA REQUEST
VOLUME MAP
11/20/2003 12:27:00 VERSION 01
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
```

Records 1 through 5 are identical for all requests, and records 6 and following contain the requested output which differs depending on the request:

- **Records 1 and 2** contain the request data.
- **Record 3** contains data and time when the report was created and the version of BVIR, currently Version 1.
- **Record 4** contains the serial number and the Library ID of the VTS. The VTS sequence number field (S/N) is the 5 character identifier assigned to the VTS. The S/N is defined by the IBM CE during installation of the VTS. The Library ID field (LIB ID) is the 5 character identifier defined for the library associated with the VTS the request is issued to. The LIB ID 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** is contains all blanks.
- **Record 6** contains the title of the output columns, in our example for a Volume Map.
- **Records 7 and following** contain the requested data. The fields of these reports are described in the next sections.

PtP Volume Status Information

Within a PTP VTS, a database is maintained on each distributed VTS that contains information about the state and data validity of each logical volume on that VTS. The information from the databases on each of the VTSs in the PTP is used by the Virtual Tape Controllers and the Virtual Tape Servers to manage the copy and resynchronization processes in the PTP subsystem. They also contains other information related to the management of the logical volumes. The PTP Volume Status request can be issued to a non-PTP VTS, however, other than the logical volume serial number, the rest of the information fields will contain zeros.

For more information on the PtP VTS Status Information output, please refer to the *BM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115.

Cache Contents Information

Volumes accessed by a host are maintained in the tape volume cache in the VTS. The VTS controls the movement of logical volume out of the cache as space is needed for newly created or recalled volumes. The primary goal of the cache management algorithms in the VTS is to maximize the utilization of its cache for volumes that have some likelihood to be accessed again.

The cache management function of the VTS arranges the volumes in cache 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 premigrated (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. The ordering of the volumes in cache can be influenced through the use of Storage Class policies as part of the Advanced Policy Management function of the VTS.

Two policies are currently supported:

- **Preference Group 0 (PG0):** When space is needed in the cache, premigrated 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.
- **Preference Group 1 (PG1):** When space is needed in the cache and there are no premigrated preference group 0 volumes to remove, premigrated 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).

Example 7-2 shows a sample Cache report.

Example 7-2 Sample Cache report

```
ORDER VOLSER DATE/TIME IN CACHE PG
    1 VOL020 11/30/2003 11:57:00 0
    2 VOL019 11/29/2003 03:00:00 0
    3 VOL023 11/20/2003 09:57:00 1
    4 VOL016 11/20/2003 10:01:00 1
```

The fields displayed in this report are:

Order The order in which volumes are to be removed from the cache, right justified and filled with blanks.

Volser The six-character volser of the logical volume.

Date/Time in cache The time the logical volume was created or recalled into cache in the format of HH:MM:SS.

PG The preference group the volume is assigned to.

The contents of the cache typically are all private volumes; however, it is possible that some may have been returned to scratch status soon after being written. The VTS does not filter the cache contents based on the private or scratch status of a volume.

Physical Volume to Logical Volume Mapping Information

The VTS has databases that maintain the mapping between logical and physical volumes. Because of the multiple database design and that it can be up to 24 hours between reconciliation of the databases, it is possible that there are inconsistencies in the mapping information provided with this function. 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 are not presently located on.

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 a stale mapping relationship is recalled, it will recall correctly, but an additional mount of a different physical volume may be required.

Example 7-3 shows a sample Volume Map report.

Example 7-3 BVIR Volume Map

```
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
P00024     GK0040  P  000005 1 OF 1    1549.65 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
```

The fields contained in this report are:

Physical The six-character physical volser the logical volumes are located on.

Logical The six-character logical volser.

P/B The Pool Indicator: P indicates the logical volume is the primary copy, and B indicates that the volume is the backup copy.

Order The relative order of the logical volume on the physical volume.

Part Indicates whether the logical volume spans on.to another physical volume and if so, which part it is.

Size The size in MB.

Unknown or invalid request

If the request file does not contain the correct number of records or 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 the first record is correct but the second is not, the response file will indicate that the request is unknown as shown in Example 7-4.

Example 7-4 Invalid request

```
VTS BULK VOLUME DATA REQUEST
VOL MAP
11/20/2003 12:27:00 VERSION 01
S/N: 0F16F LIB ID: AB123

UNKNOWN REQUEST TYPE
```

7.3.5 Sample JCL

This section contains 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 JCL shown in Example 7-5 obtains a scratch volume to write the request data to.

Example 7-5 Obtain a scratch volume

```
//VTSQUERY 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 VTS THE REQUEST IS TO GO TO AS
//* A CHECK THAT IT IS GOING TO THE CORRECT VTS
//*****************************************************************
//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=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1 DD *
VTS BULK VOLUME DATA REQUEST
CACHE CONTENTS
/*
...
```

The sample JCL shown in Example 7-6 issues the mount request and then reads the response data which has been created during mount time and sends it to a printer.

Example 7-6 Read response data

```
//VTSRESP JOB ...
//***************************************************************** //* DO NOT USE
COMPACTION WHEN WRITING THE REQUEST FILE
//***************************************************************** //* 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=FB,LRECL=80,BLKSIZE=80)
```

The sample JCL shown in Example 7-7 only works for JS2for JES2 (this will not work for JES3 because it will not demount/remount the volume between steps) combines the 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 7-7 JES2 Sample JCL

```
//BVIRINFO JOB ...
//STEP0 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
DELETE TAPE.BVIR.B63M2N36
DELETE TAPE.BVIR.B63M2N36.OUTPUT
/*
//STEP1 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT2 DD DSN=TAPE.BVIR.B63M2N36,
// UNIT=B63M2N36,LABEL=(,SL),
// DISP=(NEW,CATLG),
// DCB=(RECFM=FB,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=FB,LRECL=80,BLKSIZE=80)
//SYSIN DD DUMMY
```

To achieve the same results with JES3, use the JCL shown in Example 7-8. It separates the create and read steps into two different jobs.

Example 7-8 JES3 Sample JCL

```
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=FB
```

For the latest information and additional sample JCL, refer to the *VTS Bulk Volume Information Retrieval Function User's Guide*:

http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100488

7.4 Other PtP VTS control and monitor options

The Library Manager console and MVS console commands may be used as well to get status information of the PtP VTS. Library Manager console

The Library Manager provides information about the distributed libraries and its attached VTSs. There is no reporting on the VTCs. Although you will have the ETL Specialists installed, the Library Manager PC must still be used to perform library related operations like PAUSE Mode settings.

7.4.1 LM Console displays

There are several screens available to display the status of the library and the VTS - not the VTCs. Task-related information can also be found in Chapter 6, "Operations" on page 221. For detailed information go to the *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-00449.

7.4.2 MVS Console commands

There are several display SMS commands available to display the OAM status, the composite library, and volume details. Task-related information can also be found in Chapter 6, "Operations" on page 221.

7.4.3 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 Peer-to-Peer VTS, IOSTATS can be used for different purposes:

- To evaluate the effect of the Peer-to-Peer VTS, compare job execution times before implementation of the Peer-to-Peer VTS to those after having migrated, especially if you are operating in immediate copy mode of operation.
- To evaluate the effect of hardware upgrades, compare job execution times before upgrading components of the Peer-to-Peer VTS, for example, if you are increasing the size of the Tape Volume Cache or the number of Magstar tape drives.
- To evaluate the effect of changing the copy mode of operation on elapsed job execution time.

7.4.4 IBM Enterprise Tape Library Expert

The ETL Expert is a chargeable software product with program number 5648-SWV. It complements the information that you can see from the Specialist, and it is integrated with other members of the TotalStorage Expert family, such as the Enterprise Storage Server® Expert.

The ETL Expert is accessed from a suitably configured Web browser. The browser must be a graphical (that is, not text-only) browser and must be able to support Java and JavaScript. It provides you with information about stand-alone VTSs and Peer-to-Peer VTSs. The ETL Expert server runs on either AIX® 4.3.3 or Windows® NT 4 (with service pack 6) and maintains historical data in a database so that you can identify trends. Data is captured from the Specialists.

For the Peer-to-Peer VTS composite library, you can view the same data that is written in the SMF type 94 records without needing to post-process the SMF data. For a stand-alone VTS and for Peer-to-Peer distributed libraries, you can see information about the VTSs in a report form, and you can also use a health monitor that gives you an automatically updated view with a warning if there are unexpected conditions that need your attention. In some cases, you can also view graphs.

Restriction: The ETL Expert does not support IBM 3592 drives.

8

Upgrade scenarios

In this chapter, we describe the various upgrade scenarios you might consider for your existing Peer-to-Peer VTS. We discuss why you might decide to upgrade, and the steps in how the upgrade is performed. We describe:

- B10 VTS to Model B20 VTS upgrade
- ESCON to FICON conversion
- Eight VTCs to four VTCs reduction
- IBM 3590 to 3592 upgrades

For more information about tape library upgrades, refer to *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632 for IBM 3494 related upgrades, and to *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning and Implementation*, SG24-6789 for upgrades of an IBM TS3500/3953 configuration.

For additional information on VTS feature upgrades, see the *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring,* SG24-2229.

8.1 General considerations

Here we summarize general considerations for the upgrade of a PtP VTS:

- **Availability:**

  The nature of a Peer-to-Peer VTS provides the ability for concurrent maintenance and upgrades in almost all situations. Upgrades and maintenance are performed on the distributed libraries by placing one VTS into a Service Preparation mode (Service Prep) on the VTC. Once selected on one VTC, it sends the request to place that VTS into Service Prep mode to all other VTCs in the Peer-to-Peer. By putting one side of the Peer-to-Peer into this mode it allows the host jobs to run as normal to the active VTS but without the copies being created. Any maintenance or upgrade can be performed at this time to the VTS and its associated distributed library. Once the required operations are completed, the Service Prep mode is canceled and the distributed library is brought back into the Peer-to-Peer. All deferred copies of the volumes are sent over from the other VTS.

- **Prerequisite code and features:**

  There are some code level and feature restrictions when upgrading one VTS model to another within a Peer-to-Peer or stand-alone. See the *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring,* SG24-2229 for more details on the upgrade scenarios.

- **User interface library:**

  If the VTS being upgraded is attached to the user interface (UI) library, then you will not be allowed to perform any UI function, like inserting of logical volumes, until the upgrade is completed and the library is back in the Peer-to-Peer.

- **General upgrade steps:**

  a. The VTS getting the upgrade must be put into Service Prep mode. If this VTS is the Master, a switchover is performed to make the other VTS the Master.

  b. The VTS upgrade is performed.

  c. After the upgrade is completed and the VTS is back online to its distributed library, Service Prep is canceled on the VTC and the active VTS send the copies of all volumes that were created during the upgrade.

- **PtP VTS resynchronization:**

  Depending on the length of time it takes to complete the upgrade and the amount of activity the active VTS is performing, you could have hundreds or thousands of logical volumes to resync after the upgrade. This process is automatic and could take many hours to complete. All logical volume access to any non-copied volumes will have to be done on the non-upgraded VTS.

  If you are planning on upgrading both VTSs, you must wait until all the copies are completed and no deferred volumes are left before you start the upgrade on the second VTS in the Peer-to-Peer.

8.2 VTS Model B10 to a VTS Model B20

Upgrading the model of VTS within a Peer-to-Peer is very similar to the process of upgrading when the library and VTS are stand-alone. However, you have to consider the channel connections of the VTS. If you currently have an ESCON VTS and your new VTS will be FICON, you will have to upgrade the VTCs to support FICON as well in order to have the new VTS communicate with the host and the other VTS.

There is no way to have a single VTC with ESCON and FICON channels at the same time. So a change to the channels on the VTS or VTC will affect both sides of the Peer-to-Peer. Therefore, a non-disruptive PtP VTS model upgrade that includes a channel type upgrade (ESCON to FICON) will need to be done in two steps. First upgrade the VTS with ESCON to a VTS Model B20 with ESCON at both sites. This is shown as Step 1 and Step 2 in Figure 8-1 and is discussed in this section.

Then upgrade both VTSs and all the VTCs to FICON. This is shown as Step 3 in Figure 8-1. See 8.3, “ESCON to FICON conversion” on page 296 for information on the channel upgrade.

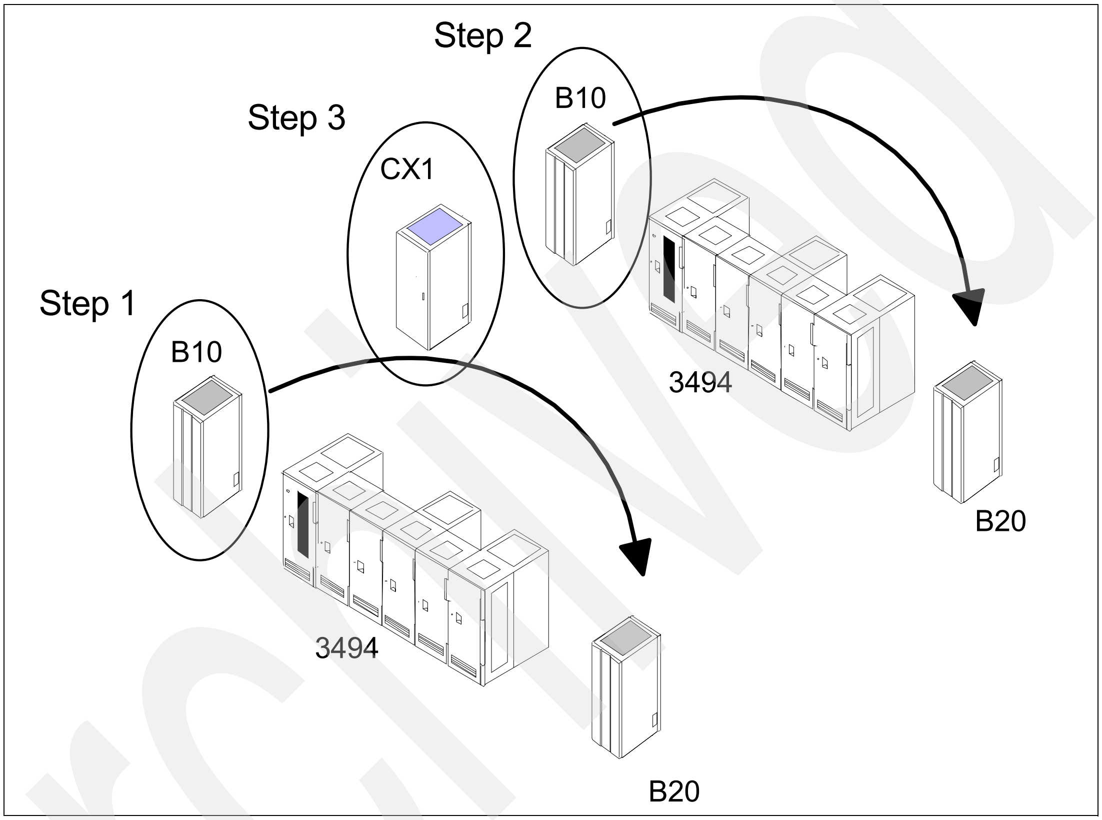

Figure 8-1 Upgrading VTS models

You may also select to change the number of VTCs from 8 to 4 if you are choosing to use FICON in your final configuration. See the Pre-installation Planning chapter for information on the performance of ESCON and FICON VTCs and see Figure 8-2 on page 293 for possible VTC configuration options you might consider. Details on the upgrade path to move from 8 VTCs to 4 VTCs are provided in 8.4, “Migration from 8 ESCON VTCs to 4 FICON VTCs” on page 306.

8.2.1 Virtual device considerations

A PtP VTS Model B18 or B10 can have 64 virtual devices on 4 VTCs. With FC5624, a VTS Model B18 can have 128 virtual devices spread over 8 VTCs. A PtP VTS Model B20 can have 128 virtual devices on 4 or 8 VTCs or 256 virtual devices on 8 VTCs. Table 8-1summarizes valid VTC and virtual device configurations.

Table 8-1 Supported VTC configurations (new installation)

VTS Models	# of VTCs	# of CUs per VTC	# of Dev/ VTC	Total Devices to Host	Total Channels (ESCON or FICON)	Number of Logical Paths [a]	Number of FC5265 installed[b]
B10	4	1	16	64	8	512/1024	NA
B20	8	1	16	128	16	1024/2048	zero
B20	4	2	32	128	8	512/1024	two
B20	8	2	32	256	16	1024/2048	four

a. ESCON/FICON: ESCON channels provide 64 logical paths each, FICON channels provide 128 logical paths each.
b. If the B20 is upgraded from a B18 which already had two FC5264 installed, then the target B20 can have four FC5264 instead of FC5265 installed.

Having 32 devices supported on a single VTC began with VTS microcode Release 7.3. FC5265 allows you to increase the number of logical devices to a maximum of 256 on the B20 VTS. This can be a concurrent upgrade to the PtP VTS, providing the B20 already has two of these features installed. Defining and activating the additional devices on the host may or may not be concurrent, depending on the host environment.

For more details on upgrading the number of virtual devices, see "Number of virtual devices" on page 294.

These combinations also come in consideration when deciding to convert from 8 VTSs to 4 VTCs. See 8.4, "Migration from 8 ESCON VTCs to 4 FICON VTCs" on page 306.

8.2.2 VTS microcode requirements

When you are upgrading a Peer-to-Peer VTS with two Model B18s to two Model B20s, you have to consider the supported levels of microcode on the current VTSs and VTCs to that of the new B20. A B18 cannot support microcode levels above 2.29.710.xx. A new VTS model B20 ship with a higher level than this. The first step in all model upgrades is to ensure all existing hardware have the minimum required code levels and if not upgrade the current hardware before doing any model upgrades. This will enable the selected upgrade to be non-disruptive to your Peer-to-Peer. The minimum code levels are:

- Library Manager - 530.xx
- VTS (B18/B10/B20) and VTCs - 2.29.710.30

8.2.3 Upgrade steps

The model upgrade is done in three steps:

1. **Step 1:**

   One side of the Peer-to-Peer gets a Service Prep so the current VTS can be upgraded to a VTS Model B20. Once the upgrade is complete, the Service Prep status is canceled and the Peer-to-Peer VTS is complete again. Wait until all deferred logical volumes have completed copying to the new VTS before beginning to upgrade the other side.

2. **Step 2:**

   The processes in Step 1 are repeated for the second VTS. Service Prep is selected for the remaining B18 VTS and the upgrade is started. Once complete, Service Prep is canceled again and all deferred logical volumes are copied to the second new VTS Model B20.

3. **Step 3:**

   The last step is to upgrade the VTC microcode to the level the two new Model B20 VTSs were shipped with. One at a time each VTC is taken offline (the 16 or 32 devices defined to the VTC) and the microcode is installed. Once complete the VTC is put back online and then the next VTC can be upgraded. It may be possible to upgrade more than one VTC at a time, if workload requirements permit.

8.2.4 Upgrade scenarios

Table 8-2 summarizes the possible upgrade scenarios for a PtP VTS configuration which we cover subsequently in this section.

Table 8-2 Possible PtP VTS model upgrades

	From	To
1.	B10/B10 with 4 VTCs and 64 virtual devices	B20/B20 with 4 VTCs and 128 virtual devices
		B20/B20 with 8 VTCs and 128 virtual devices
2.	B20/B20 with 4 VTCs and 128 virtual devices	B20/B20 with 8 VTCs and 256 virtual devices
3.	B20/B20 with 8 VTCs and 128 virtual devices	B20/B20 with 8 VTCs and 256 virtual devices

Scenario 1

When upgrading from a VTS Model B10 to a VTS Model B20 configuration, you have two options. If your current Model B10 has 64 devices, the B20 will have at least 128 devices. That gives you two options:

- A B20 PtP with 128 devices over 4 VTCs
- A B20 PtP with 128 devices over 8 VTCs

Each option requires HCD changes to increase the number of logical devices support by the host. The main decision you should make is whether to have four or eight VTCs. Note that eight VTCs are more costly but gives you twice the number of paths to VTS and therefore higher throughput and availability. Later upgrades from 128 devices over 4 or 8 VTCs to 256 devices are non-disruptive to the PtP. See 8.3, “ESCON to FICON conversion” on page 296 for information on changing the number of VTCs in your Peer-to-Peer VTS.

If you intend to for your final configuration to have FICON, you must upgrade to a B20 with ESCON and then convert the whole Peer-to-Peer to FICON> This will enable a non-disruptive upgrade. See 8.3, “ESCON to FICON conversion” on page 296 or 8.4, “Migration from 8 ESCON VTCs to 4 FICON VTCs” on page 306 for details on the ESCON to FICON part of your upgrade.

Upgrade steps for four VTCs

The upgrade steps are:

- Update HCD to support 2 CUs on each existing path (32 devices/VTC) and activate new IODF
- Upgrade each B10 with a B20 with 2x FC5265
- IML each VTC to pickup additional devices (32 instead of 16)

Upgrade steps for eight VTCs

If the target configuration has eight VTCs:

- Update HCD to support four additional VTC Control Units (16 devices/VTC) and activate new IODF.

- Upgrade each B18 with a B20 with no FC5265 and 8 ESCON channels.
- Install 4 additional VTCs.

Scenario 2

The upgrade steps are:

- Update HCD to support two CUs on each existing path (32 devices/VTC) and activate new IODF.
- Update HCD to support the additional four VTCs (2 CU per per path).
- Upgrade each B20 with 4x FC5265.
- Upgrade each B20 with four additional ESCON or FICON channels.
- IML original 4 VTCs to pickup additional devices (32 instead of 16).
- Install four additional VTCs.

Scenario 3

This upgrade has the least impact to your current configuration. The following upgrade steps are required:

- Update HCD to support 2 CUs on each existing path (32 devices/VTC) and activate new IODF.
- Upgrade each B20 with 4x FC5265.
- IML each VTC to pickup additional devices (32 instead of 16).

Checklist

Here is a checklist of upgrade considerations:

- Minimum code levels - LM = 530.xx or above, VTS/VTC = 2.29.710.30 or above
- HCD changes — see 4.1, "Hardware I/O configuration definition" on page 128
- Decide on number of VTCs — see "Four versus eight VTCs" on page 95
- 8 to 4 VTCs — see 8.4, "Migration from 8 ESCON VTCs to 4 FICON VTCs" on page 306
- ESCON to FICON — 8.3, "ESCON to FICON conversion" on page 296 or 8.3.4, "VTS Model B10 upgrade from ESCON to FICON" on page 303

8.2.5 Additional VTS upgrades

In this section we discuss the other feature upgrades that are available to the VTS and show where the upgrade considerations are different when the VTS is part of a Peer-to-Peer. Where there are no additional considerations for these feature upgrades, one side of a Peer-to-Peer VTS gets a Service Prep. Then upgrading the VTS is the same as if it were a stand-alone VTS. Please refer to the *IBM Virtual Tape Server Planning and Implementation* - SG24-2229 for information.

Number of virtual devices

FC5265 provides the ability to increase the number of virtual devices in pairs of 64 devices. This allows either 128 or 256 virtual devices on the Model B20 VTS. This feature replaces FC5264. The difference being FC5265 is concurrently upgradable, providing the first pair of FC5265 is already installed. Installing the first pair of FC5265 is disruptive to the VTS and tape library. The Model B10 VTS only supports 64 devices and does not support FC5265.

When the first pair of FC5265 are installed, each Library Manager is taught for 256 virtual devices and each VTS is configured for 256 devices, but only the first 128 are available for you for use. The VTCs must be IMLed one at a time to reconfigure themselves to support 32 devices instead of the normal 16 devices.

Prerequisites

- IBM 3494 Peer-to-Peer VTSs must both be B20.
- IBM 3494 Virtual Tape Server must be at microcode level 2.31.730.xx or higher.
- IBM 3494 Library Manager Microcode must be at microcode level 530.xx or higher.
- The B20 must have zero or two FC5265.
- If FC5265 is installed, verify the LM has been taught for 256 devices.

If the VTS is part of a 4–VTC Peer-to-Peer, the current host adapter configuration must be one of the following:

- 2 x FC3412 (ESCON) OR
- 4 x FC3415 (LW FICON) OR
- 4 x FC3416 (SW FICON)

Supported VTC configurations

Table 8-3 details the number of features FC5264 and FC5265 which can be installed on the two VTSs of a PtP VTS and indicates whether the combination is valid or not.

Table 8-3 VTS Model B20 Peer-to-Peer configurations based on FC5264 and/or FC5265

VTS0 installed features	VTS1 installed features	PtP configuration
none	none	8 x VTC x 16 devices
2 x FC5265	2 x FC5265	4 x VTC x 32 devices
2 x FC5264	2 x FC5264	8 x VTC x 32 devices
2 x FC5264*	4 x FC5265*	8 x VTC x 32 devices
4x FC5265	4 x FC5265	8 x VTC x 32 devices
2 x FC5264*	none*	not valid
2 x FC5265*	none*	not valid

*It does not matter which VTS in the PtP VTS has which features installed.

VTS cache upgrades

All the existing features for a stand-alone VTS cache upgrade apply to a PtP VTS too. There is no difference, and no additional processes are involved, when upgrading the cache size of a PTP VTS compared with a stand-alone VTS, other than the addition of putting the selected VTS into Service Prep before the upgrade is started. The only consideration is that, if you upgrade one VTS in the Peer-to-Peer, you should upgrade the other side to keep the capacity of both sides equal.

Upgrading to 3592 tape drives

It is always recommended, although not required, to have both sides of the Peer-to-Peer equal in capacity and performance. If you choose to upgrade the VTS Model B20 to support 3592 tape drives on one VTS, then you should plan to upgrade both VTSs. There are microcode level considerations when upgrading the B20 VTS to handle 3592 tape drives. Upgrading a VTS to support 3592 tape drives and migrating the stacked volumes from 3590 media to 3592 media is detailed in the stand-alone VTS Redbook, *IBM Virtual Tape Server Planning and Implementation* - SG24-2229.

Minimum prerequisites

- VTS and VTC microcode at LIC 2.31.730.xx or higher
- VTS Model B10 or B20 only (Not supported on the B18)

The basic upgrade steps would be:

- Upgrade microcode on all VTSs, VTCs, and Library manager (if required).
- Upgrade one VTS and distributed library with 3592 tape drives.
- Upgrade second VTS and distributed library with 3592 tape drives.
- Modify storage pools and dataclass definitions on the Library Manager of one distributed library to start using new tape drives and media.
- Copy these definitions to the Library Manager of the other distributed library. See the chapter on Implementation in this book for details on copying APM definitions from one Library Manager to another.
- Modify host definitions to point to the new data constructs.

8.3 ESCON to FICON conversion

This section describes the steps to follow for a non-disruptive full upgrade to a Model B20 FICON PtP VTS.

8.3.1 Considerations

A standalone B10 or B20 cannot be upgraded to a FICON PtP VTS without first being upgraded to an ESCON PtP. Figure 8-2 explains the upgrade path. As mentioned previously, all VTCs must be at the same level of feature code.

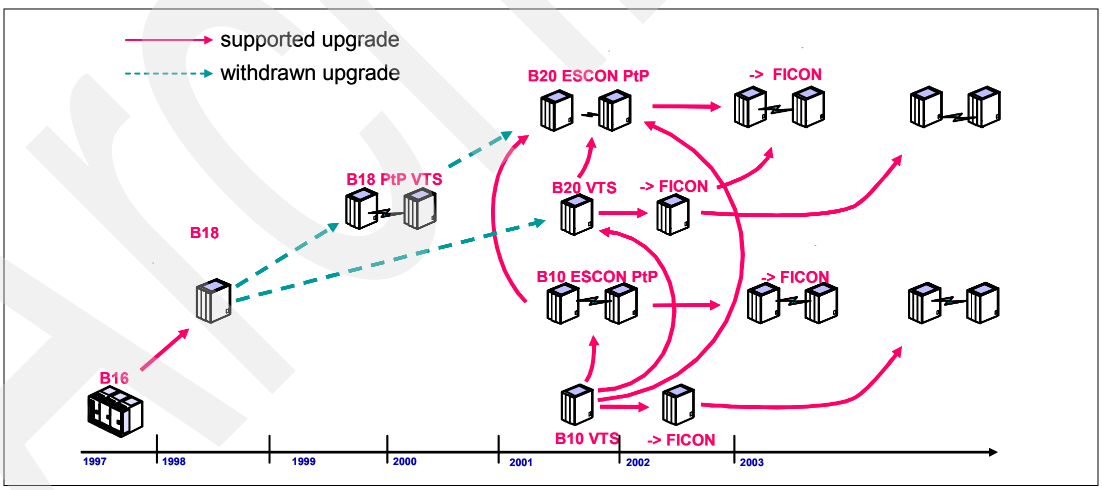

Figure 8-2 Upgrade paths for a VTS

Attention: Effective June 30, 2006, IBM will withdraw from marketing all remaining features of the 3494 Virtual Tape Server Model B18 and the 3494 Model B18 to B20 upgrade. Orders received before June 30, 2006 will be acknowledged on an as-available basis. On or after the effective date of withdrawal, you can no longer order these products directly from IBM.

8.3.2 Prerequisite activities

The following levels of code are the required minimum before commencing this upgrade:

- LM 528 code
- 2.28.700.X on all VTSs
- 2.28.700.X on all VTCs

HCD definitions for all the new FICON paths need to be created and activated on the host before the new FICON paths and devices with be available for use. Not having these definitions activated on all hosts before the upgrade is started could impact whether this upgrade is non-disruptive or not.

8.3.3 Upgrade steps

The following steps are required to upgrade from ESCON to FICON. Each step is described in more detail in Figure 8-3 on page 298 though Figure 8-8 on page 303.

1. Update VTS0 to Half ESCON/Half FICON (replace ESCON in primary I/O Drawer). New FICON will not be connected (No FICON VTCs available).
2. Replace CX0-1 with a FICON CX1-1. This allows VTS0 to use the FICON CX1-1 and ESCON CX0-0.
3. Update VTS1 to Half ESCON/Half FICON (replace ESCON in primary I/O Drawer). New FICON will not be connected to CX0-0 (no FICON VTCs available).
4. Update VTS1 to All FICON (Replace all ESCON in both I/O Drawers). Half of the FICON (Primary Drawer) will be connected to CX1-1. The other half of FICON will not be connected (2nd FICON CX1 is unavailable).
5. Replace CX0-0 with a FICON CX1-0. This allows VTS1 to use the both FICON CX1s.
6. Replace remaining ESCON with FICON in VTS0. This allows both FICON VTSs to communicate to all eight FICON VTCs.
7. MES is complete.

The diagram in Figure 8-3 represents the non-disruptive upgrade of a full ESCON B20 PtP VTS subsystem to a full FICON B20 PtP subsystem. It is a pictorial representation of the steps outlined previously.

It does not matter which VTS or which group of VTCs are selected to start the upgrade with. Your decision on which to select first would be based more on which VTS is more remote. If one VTS is connected through extended links, CNT for example, you might consider starting with the remote site. But each VTS will be unavailable at least twice during the full upgrade. In our example we are starting with VTS0 and VTCs 4-7.

Figure 8-3 shows the starting configuration for the FICON MES.

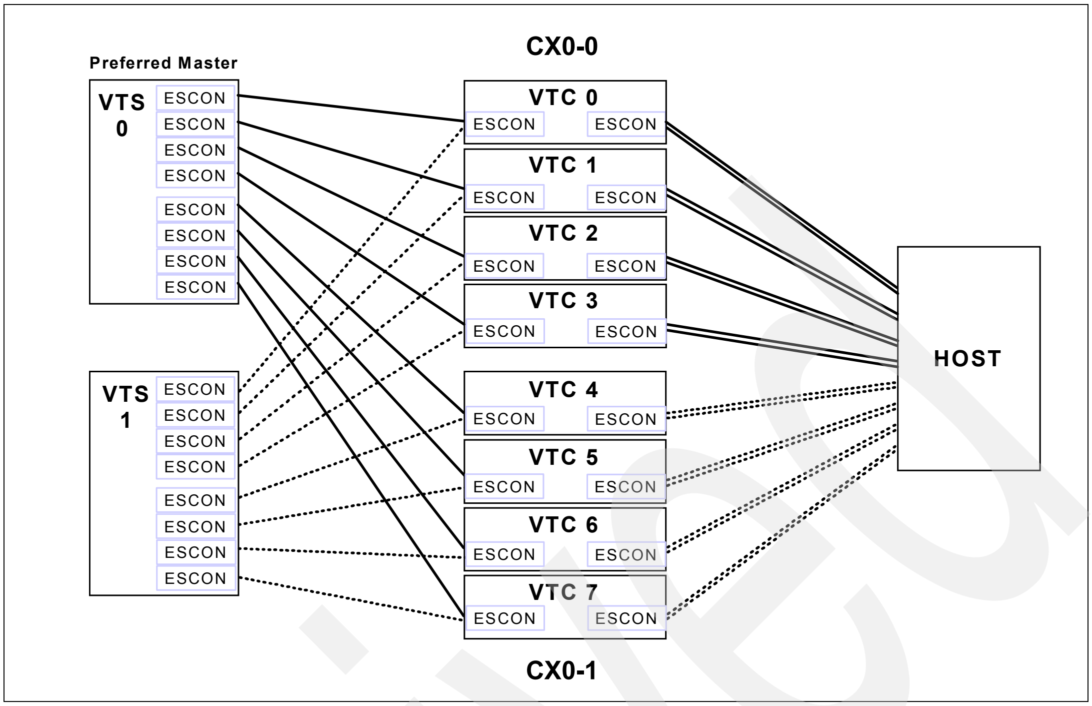

Figure 8-3 Starting configuration of a B20 VTS with 8 VTCs

Step 1

Proceed as follows:

- Perform a Master Switchover, if required.
- Place VTS0 in Service Preparation mode.
- Power off VTS0 and remove half of the ESCON adapters and replace them with FICON adapters.
- Power on VTS0
- Wait for deferred copy volumes to complete copying before continuing to the next step.

Figure 8-4 shows the result at the end of Step 1.

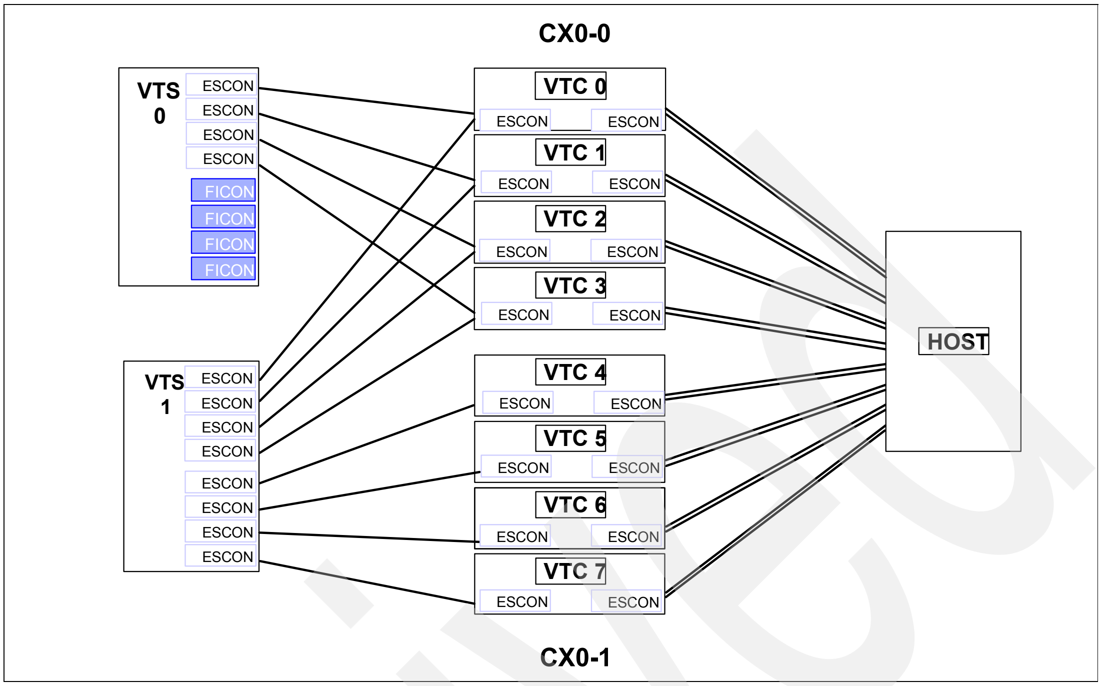

Figure 8-4 Step 1, upgrade VTS0 to 1/2 FICON and 1/2 ESCON

During this stage of the upgrade, all the defined logical devices and all the ESCON paths are online and can access all of the logical volumes. Since VTS0 is not available, all logical volume copies are deferred.

Step 2

Proceed as follows:

- Take the devices defined to VTC4-7 offline to all hosts.
- Take VTC4-7 offline and power off.
- Disconnect all the ESCON channels from both VTSs to these 4 VTCs.
- Upgrade VTC4-7 to FICON by either replacing the hardware or replacing the ESCON adapters within each VTC.
- Cable up the FICON VTCs to VTS0 and the host
- Power on VTC4-7 and then bring new FICON channels online.
- Vary devices defined to VTC4-7 back online.

Figure 8-5 shows the configuration at the end of Step 2.

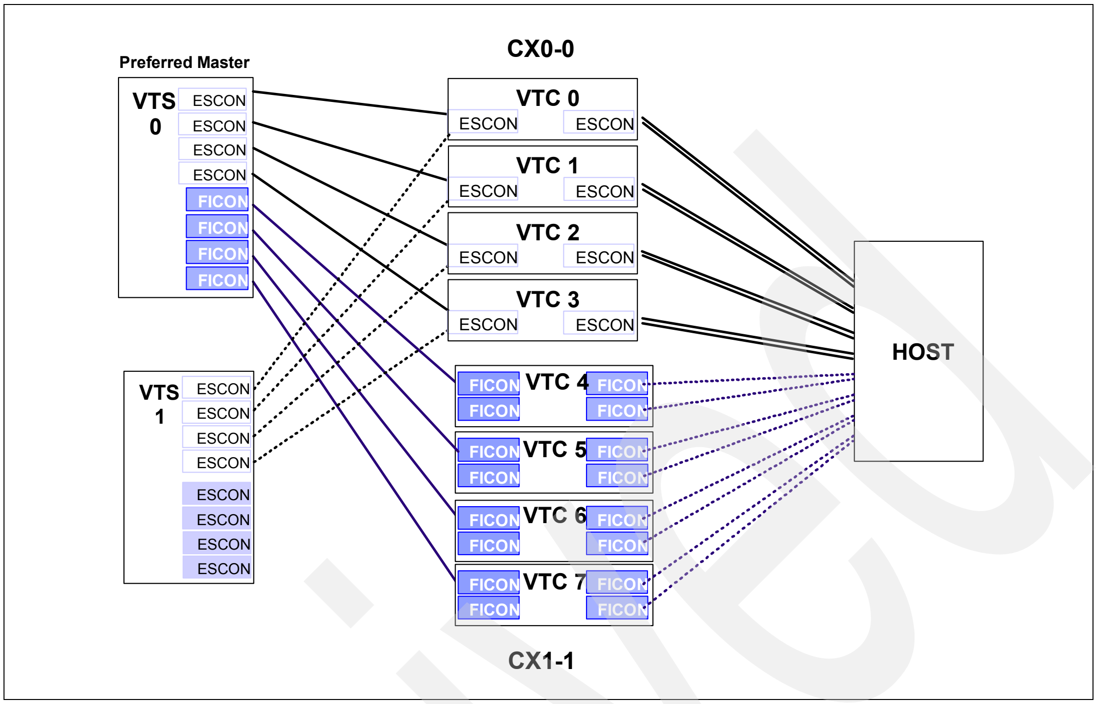

Figure 8-5 Step 2, upgrade VTC4-7 to FICON

During this stage, you have half of the logical devices offline and half of the physical paths offline. All remaining devices and paths have full access to all logical volumes in the PtP VTS. When this step is completed, you have 8 ESCON paths and 8 FICON paths online, each communicating with a different half of the total address range of the Peer-to-Peer VTS.

Step 3

Proceed as follows:

- Perform a Master Switchover.
- Place VTS1 in Service Preparation mode.
- Power off VTS1 and remove all of the ESCON adapters and replace them with FICON adapters.
- Cable up the VTS1 FICON connections to VTC4-7.
- Power on VTS1.
- Wait for deferred copy volumes to complete copying before continuing to the next step.

Figure 8-6 shows the configuration at the end of Step 3.

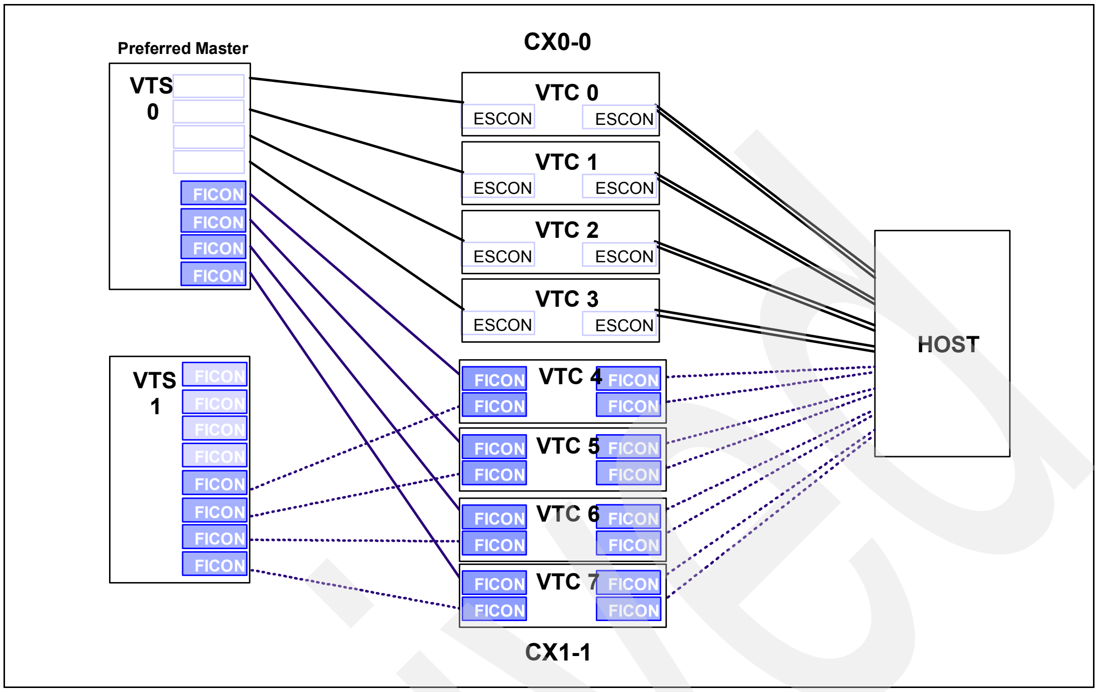

Figure 8-6 Step 3, upgrade VTS 1 to FICON

During this stage of the upgrade, VTS1 will be unavailable but full access to all logical devices will be available as will be access to all logical volumes. Only dual copies will be deferred until the PtP VTS is back together.

Step 4

Proceed as follows:

- Vary the devices defined to VTC0-3 offline to all hosts.
- Upgrade VTC0-3 to FICON by either replacing the hardware or replacing the adapter within each VTC.
- Cable up the FICON VTC0-3 to VTS1 and the host.
- Bring new FICON paths of VTC0-3 and all logical devices defined to VTC0-3 online.

Figure 8-7 shows the result at the end of Step 4.

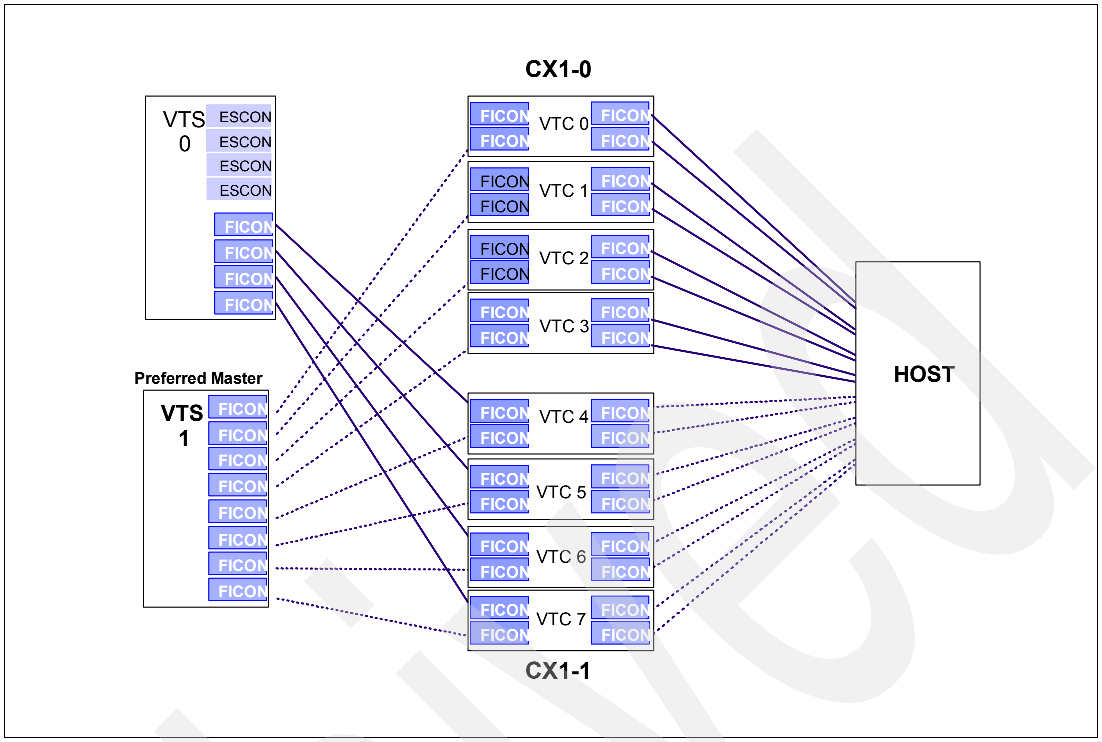

Figure 8-7 Step 4 - Upgrade VTCs0-3 to FICON

During this stage, you have half of the logical devices offline and half of the physical paths offline. All remaining devices and paths have full access to all logical volumes in the PtP VTS. When this step is completed, you will have 16 FICON paths online to host and all logical devices will be available with full access to all logical volumes.

Step5

Proceed as follows:

- Perform a Master Switchover and Service Prep VTS0
- Power off VTS0 and remove remaining ESCON adapters and replace with FICON adapters.
- Cable up these last 4 FICON adapters to VTC0-3.
- Power on VTS0.

Figure 8-8 shows the result at the end of Step 5, when the upgrade is complete.

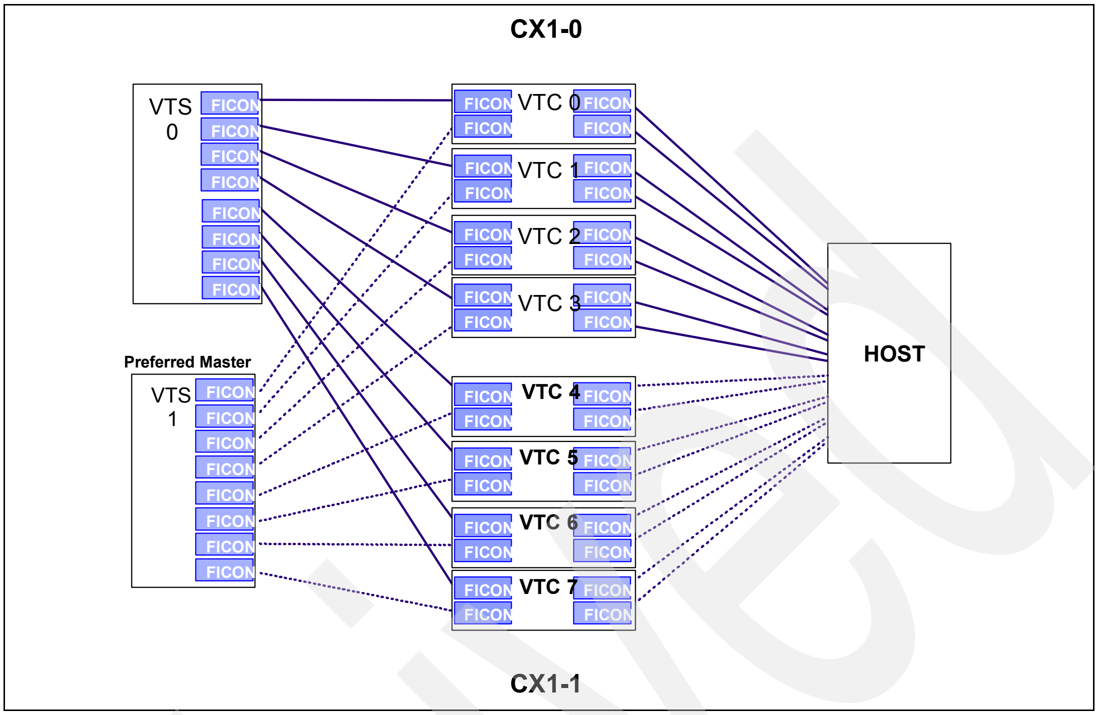

Figure 8-8 Step 5, upgrade VTS1 to all FICON - upgrade complete

During this stage of the upgrade, all logical devices have access to all logical volumes. Only copy volumes are deferred until VTS0 is back in the Peer-to-Peer.

8.3.4 VTS Model B10 upgrade from ESCON to FICON

Upgrading a PtP VTS Model B10 from ESCON to FICON is exactly the same as for upgrading a PtP VTS Model B20 with ESCON to FICON. The only difference is the number of VTCs and the number of channels involved. See 8.3, “ESCON to FICON conversion” on page 296 for a detailed upgrade path.

Prerequisite code

The following levels of code are the required minimum before commencing this upgrade:

- LM 528 code
- 2.28.700.X on all VTSs
- 2.28.700.X on all VTCs

Figure 8-9 shows a diagram of the start configuration of the Model B10 PtP VTS using all ESCON channels.

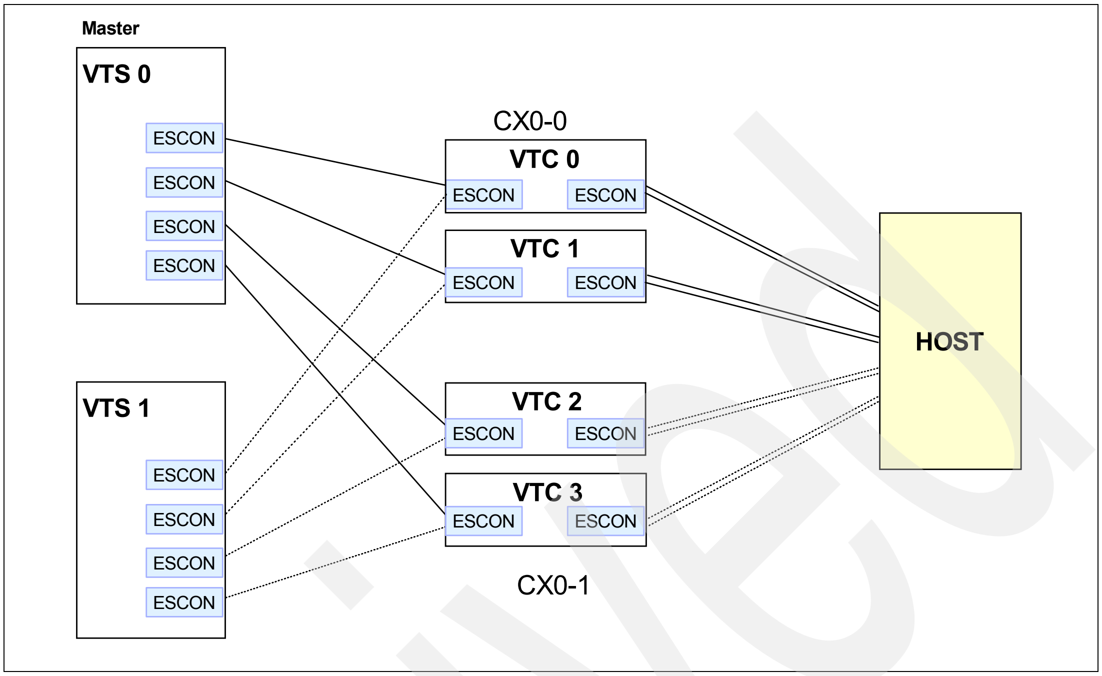

Figure 8-9 Initial B10 ESCON PtP System

Steps in this upgrade

The steps involved are as follows:

1. Service Prep the first VTS to be upgraded (VTS0 in this example).
2. Upgrade VTS 0 to Half ESCON/Half FICON (replace ESCON in slots 3 and 4).
3. Cancel Service Prep.
4. New FICON will not be connected (No FICON VTCs available, VTS0 has only 2 paths).
5. Take VTC0-1 offline and replace VTCs 0 and 1 with FICON VTCs, by either:
   – Installing a new CX1 with the FICON VTCs 0-1 already installed
     OR
   – Replacing the VTCs with a FICON VTCs

   Now VTS0 can use FICON on VTCs 0-1 and ESCON on VTCs 2-3.
6. Service Prep VTS1.
7. Upgrade VTS1 to All FICON.
8. Cancel Service Prep.
9. Half of the FICON (Slots 3 and 4) will be connected to VTCs0-1.
10. The other half of FICON will not be connected (FICON VTCs2-3 are unavailable).

    At this point both VTSs are running normal production with 2 FICON paths each.

11. Take VTC2-3 offline and replace with a FICON VTCs 2 and 3, by either:
   - Installing a new CX1 with the FICON VTCs2-3 already installed
     OR
   - Replacing the VTCs with a FICON VTCs

   Now VTS1 can use all 4 paths with the FICON VTCs.
12. Service Prep VTS 0.
13. Replace remaining ESCON with FICON in VTS 0.
14. Cancel Service Prep.

Now both VTSs can use all 4 paths with the FICON VTSs to communicate to the host. The MES is complete. Figure 8-10 shows a diagram of the final configuration.

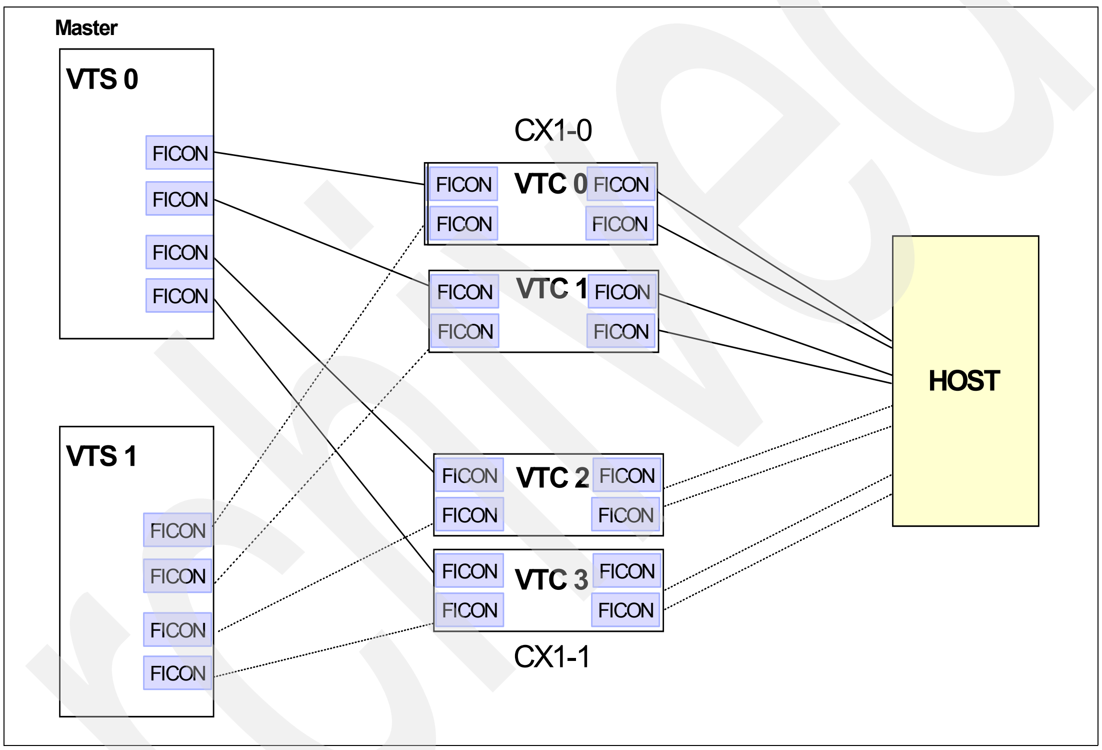

Figure 8-10 Final all FICON configuration

8.4 Migration from 8 ESCON VTCs to 4 FICON VTCs

This upgrade scenario incorporates several upgrade options in one change event. At the end of this upgrade, the PtP VTS will have the following upgrades installed:

- Model B18 to Model B20 VTS
- ESCON to FICON upgrade
- 8 VTCS to 4 VTC conversion

The final stages of this upgrade are not concurrent. The overall throughput will be very similar between 8 VTCs with ESCON and 4 VTC with FICON. See "Four versus eight VTCs" on page 95 for more information.

If you currently have two B20 VTSs in your Peer-to-Peer, then you would only need to replace the ESCON adapters with FICON and add/delete the appropriate FC5264 / FC5265 instead of replacing the B18 with a B20 as described in the example. The rest of the upgrade process is the same.

Attention: Effective June 30, 2006, IBM will withdraw from marketing all remaining features of the 3494 Virtual Tape Server Model B18 and the 3494 Model B18 to B20 upgrade. Orders received before June 30, 2006 will be acknowledged on an as-available basis. On or after the effective date of withdrawal, you can no longer order these products directly from IBM.

8.4.1 Advantages

There are several advantages in performing the 8 ESCON to 4 FICON VTC portion of this upgrade:

- Less hardware, 1/2 the number of VTCs
- 1/2 the number of channel connections at the host
- 1/2 the number of extended links if one VTS is at another site
- Reduced cost

8.4.2 Prerequisites

All VTSs, VTCs and Library Managers must be at the required microcode levels before starting the change from 8 VTCs to 4 VTCs:

- LM - 531.xx
- VTS and VTC - 2.32.740.xx

Attention: Any upgrades to microcode or features must be installed prior to starting this upgrade. You cannot combine upgrading microcode at the same time as you upgrade the hardware. The entire Peer-to-Peer system must be at the required prerequisite microcode and installed feature levels before starting this upgrade.

8.4.3 Upgrade sequence

In this section, we depict the sequence of events and the status of the Peer-to-Peer VTS during the upgrade of having eight ESCON attached VTCs to four FICON attached VTCs. Half of the devices and paths will be offline to all hosts at certain stages during this upgrade. Select a time for this upgrade when the PtP VTS activity is at a low point to prevent performance impacts to production work. Also, the last steps of this upgrade are disruptive, so schedule the time when taking the PtP VTS down.

You must create and activate definitions for the new FICON connections on the host. See "Model B20 PtP VTS with 4 VTCs and 128 virtual drives" on page 135, for example, IOCDS definition examples. If you do not have enough free addresses to activate the new FICON range at the same time the existing ESCON range is active, or you wish to reuse the existing address range with the new FICON VTCs, then you can activate the new IODF when the you are in Step 4 of the upgrade process and the PtP VTS is down.

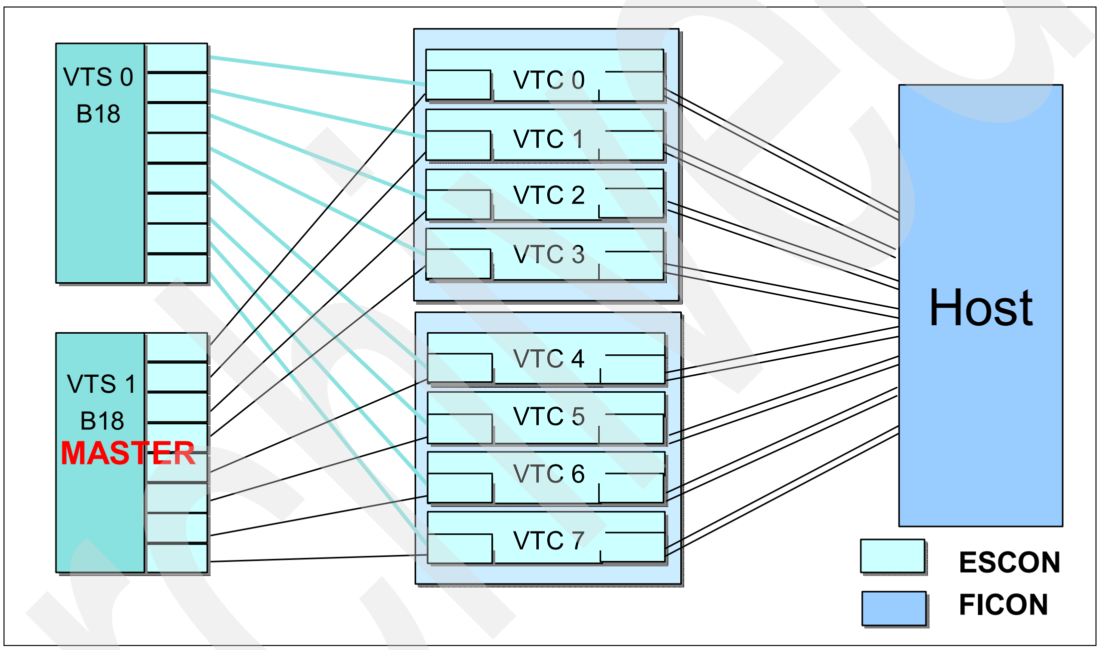

Figure 8-11 Starting configuration

In this case, you can upgrade from an ESCON B18 to a FICON B20 PtP VTS in one step, because the migration from eight to four VTCs disruptive already.

Note: In the following example, we start with VTS0 and VTC0-3. Alternatively, you can start with VTS1 and VTC4-7.

Step 1

Proceed as follows:

- Perform a Master Switchover to make VTS1 the Master VTS, if required.
- Place VTS0 into Service Prep mode.
- Vary VTS0 offline to the distributed library and power off the VTS.
- Migrate VTS to a new VTS Model B20 with 4 FICON adapters and 2x FC5265 installed.

Figure 8-12 shows the status after Step 1 has been completed.

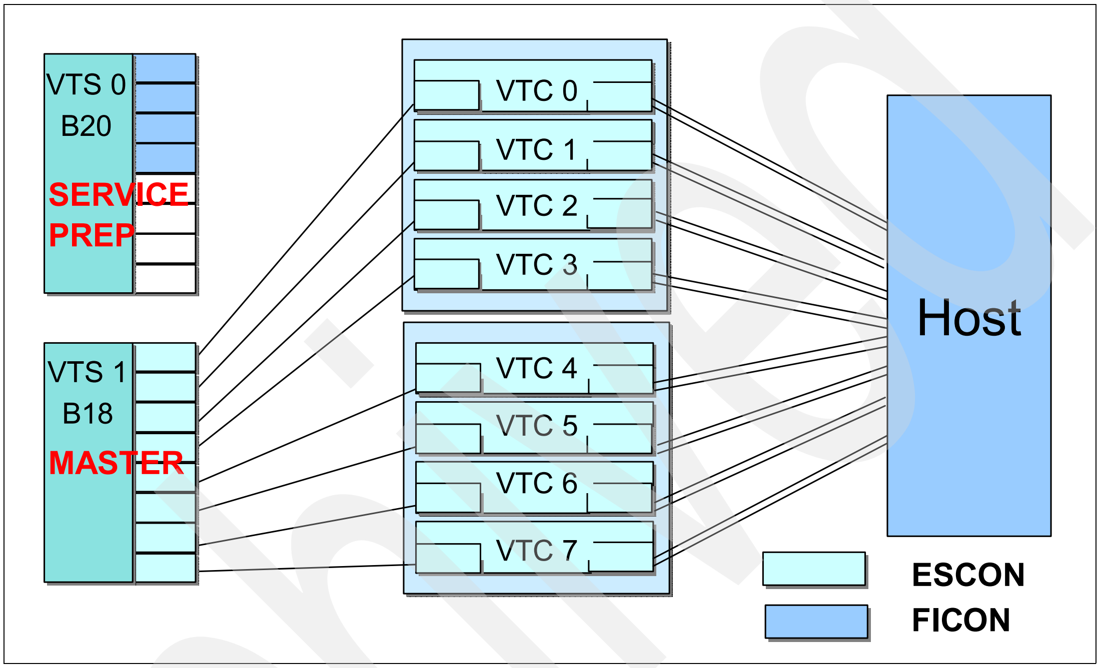

Figure 8-12 Step 1: Upgrade one VTS to B20 FICON

The current status of the Peer-to-Peer is degraded but operational. All logical volumes created in VTS1 will not have a copy on VTS0. All ESCON paths are online to all defined devices and have full access to all logical volumes.

Step 2

Proceed as follows:

- Vary the devices associated with VTC0-3 (64 devices) offline to all the hosts.
- Vary VTC0-3 offline and power off the VTCs.
- Disconnect VTS1's 4 ESCON connections to VTC0-3 (blocking ports if required).
- Upgrade VTC0-3 to FICON. Depending on the model of VTC you have, either replace the VTCs with new hardware or replace the adapters in the VTCs.

Figure 8-13 shows the PtP VTS configuration after Step 2 has been completed.

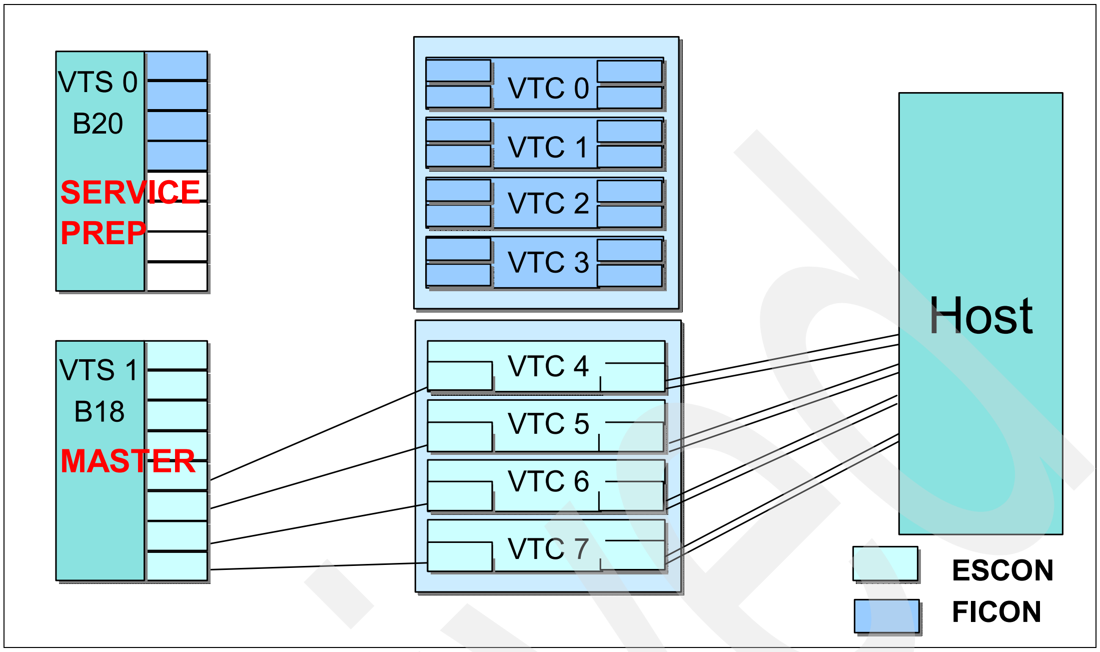

Figure 8-13 Step 2: Replace CX0 with ESCON AX0s with CX1 with FICON VTCs

The current state of the PtP VTS at this step is that:

- Half of the logical devices (64) are offline.
- Half of the channel paths are offline to the host.

Therefore, the performance of the PtP will be degraded by half at this time, but all remaining online devices have full access to all logical volumes.

Step 3

Proceed as follows:

- Cable up new VTC0-3 FICON connections to VTS0 and the FICON connections at the host.
- Power on the VTC0-3 and verify connectivity with VTS0 and any hosts. This can be done by putting the VTC0-3 into Read-Only mode. This will enable the control units in VTC0-3 to come online and can in turn be varied online to any host. Verification can be done by host display commands detailed in Chapter 4., “Implementation” on page 127.

Restriction: In order to be able to vary these new FICON control units online when the VTC0-3 are in Read-Only mode, you must have already had your existing PtP VTS defined with staggered addresses. If not, then you can only bring on the devices for the first LCU in each VTC, the first 16 devices, as the original ESCON VTC was defined.

Figure 8-14 shows the configuration after completion of Step 3.

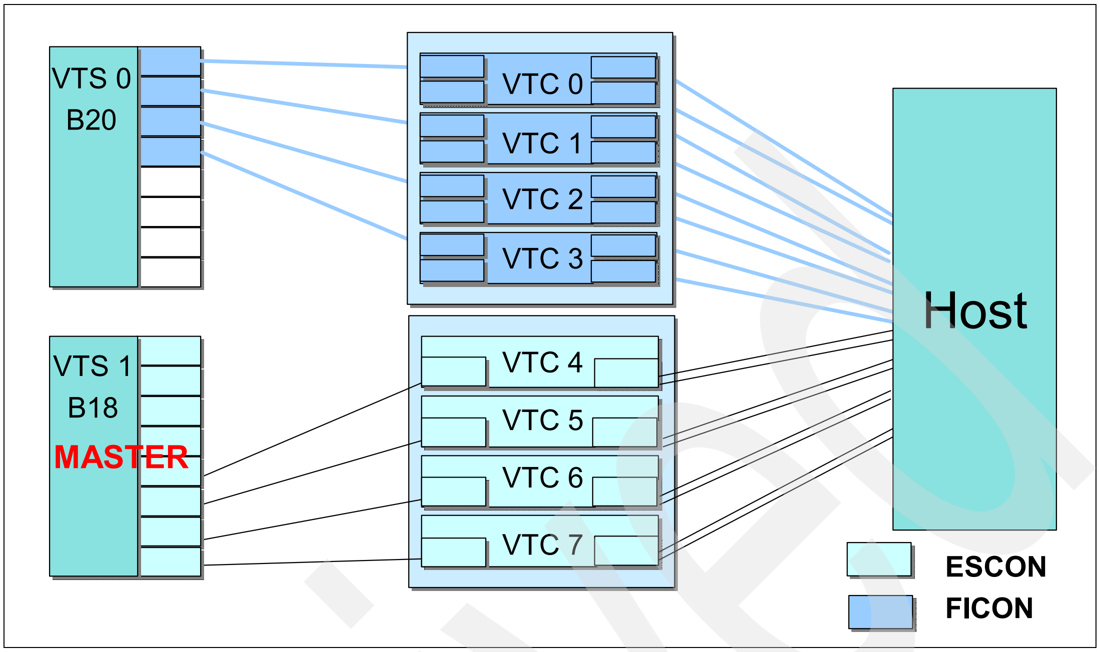

Figure 8-14 Step 3: Cable up new FICON VTCs

At this stage in the upgrade, VTS 0 and VTCs 0-3 are fully configured and could perform Read-Only work. The reason is that the FICON VTCs cannot come online in normal read/write mode because they cannot communicate with VTS1 or VTC4-7. If these VTCs are brought up in Read-Only mode, jobs can be run to verify that the first upgrade was completely successful before continuing. The test Read-Only jobs must be run to only the device range that is defined to the FICON VTCs.

Step 4

If your current host job requirements allow, it is possible to use VTS0 in Read-Only mode. No logical volume writes or Library Manager volume modifications can take place during this time. Any job specifying a write or modification to a logical volume will fail. Also, any logical volumes updated when VTS0 was down, during steps 1 to 3, will not be accessible by any host at this time. VTS1 is the only library with copies of these recently updated logical volumes until the PtP VTS is connected back together in Step 5.

If your environment will not tolerate job activity with only reading logical volumes, you must stop all job activity and take the PtP VTS down.

- Vary the devices associated with VTC4-7 offline.
- Vary PtP VTS offline to all hosts (if not running in Read-Only mode).
- Migrate VTS1 to the new B20 with FICON.

Figure 8-15 shows the configuration after completion of Step 4.

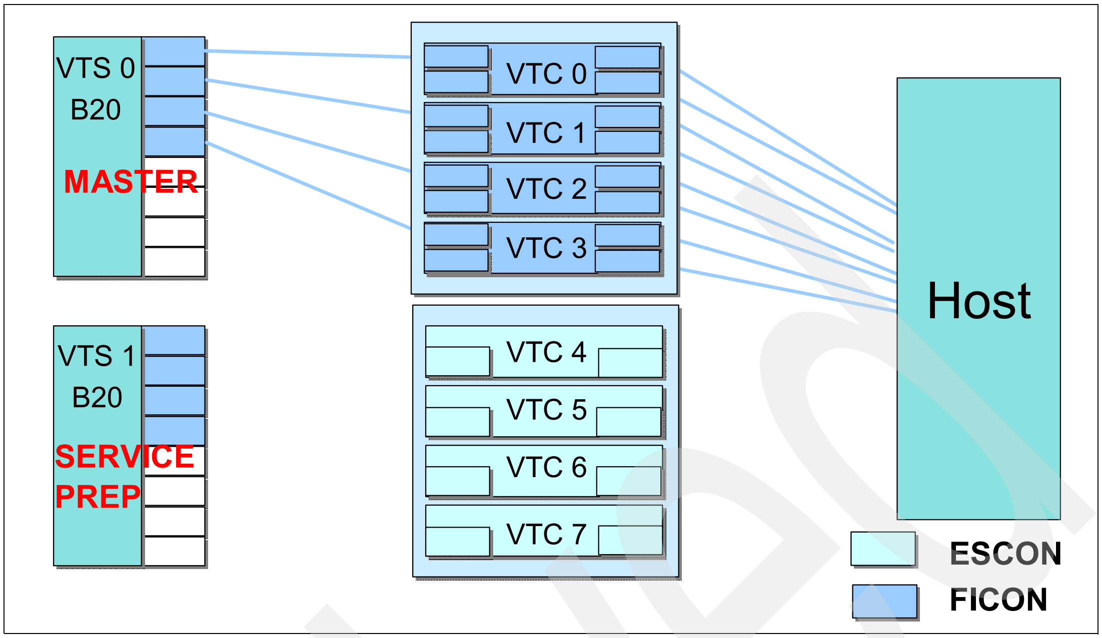

Figure 8-15 Step 4- Upgrade second VTS to FICON B20 - PtP is down

At this stage in the upgrade, the entire PtP VTS is offline to all hosts or, if your production workload will tolerate, only jobs reading logical volumes are allowed to be running.

Step 5

Proceed as follows:

- Cable up the new VTS1's FICON connections to VTC0-3.
- Once VTS1 is online, the VTCs will enter normal read/write mode.
- Vary any devices, distributed libraries, and composite library online, if not already online.
- Verify connectivity for all hosts connected to the PtP VTS.

Figure 8-16 shows the Model B20 PtP VTS after the upgrade has been completed.

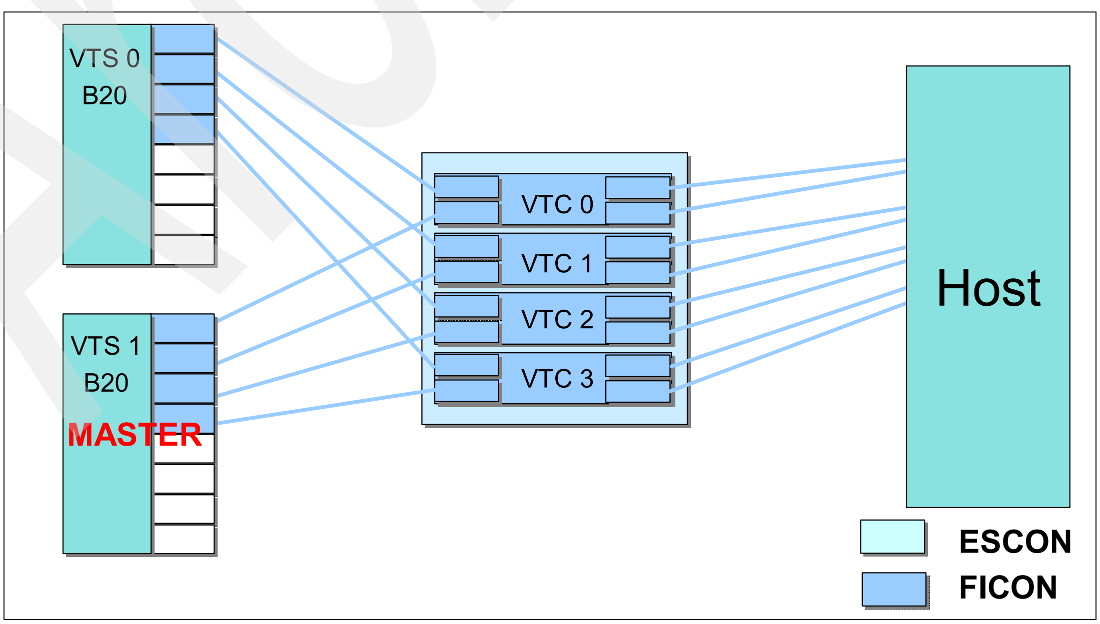

Figure 8-16 Upgrade completed and Peer-to-Peer fully operational

When this last step is completed and the VTCs can communicate to both VTSs, the token database from each VTS are synchronized and all the deferred logical volumes begin copying from VTS1 to VTS0. This may take many hours depending on your configuration and how long VTS1 was running with VTS0 in Service Prep mode. The deferred copy workload can be monitored from the VTS Peer-to-Peer Specialist. There may be reason to wait for the copies to complete or at least reduce in number before starting new work. All the copies are transferred in the background but use VTS resources. Normal production workload performance might be impacted if large numbers of logical volumes must be copied to the other VTS.

8.5 IBM 3590 to 3592 upgrade

The 3592 tape drive and media were introduced in September of 2003 for both direct System z as well as open systems attachment environments, but not attached to a VTS. The model J1A drive and Enterprise Tape Cartridge provide a non-compressed data storage capacity of 300GB and a maximum transfer data rate of 40MB/sec. In May of 2004 an additional cartridge, the Enterprise Economy Tape Cartridge, was made available, which provides a non-compressed storage capacity of 60GB. VTS support for the 3592 drive and the two cartridge capacities are provided with the October 2004 VTS R7.3 and associated library manager 530 code levels.

In November of 2005, a new version of the 3592 tape drive was released, the model E05. The drive, also known as the TS1120, is capable of a non-compressed data storage capacity of 500GB and a maximum transfer data rate of 100 MB/sec using the existing media. However, when installed in a VTS configuration, the drive emulates the model J1A from a data storage capacity perspective, limiting the capacity to 300GB and the data transfer rate to 50 MB/sec. Even though the VTS cannot take advantage of the improvement in the storage capacity of the new model, it can benefit from the improved data transfer rate and the other performance improvements such as load/thread time and search speed.

Because the model E05 drive emulates the model J1A drive, the setup and usage of the drive, as described in this white paper, are the same.

Restrictions:

- The functions to scale the 300 GB cartridge between a 300 and 60 GB capacity are not currently supported under the VTS.
- WORM 3592 media types are not supported in the VTS.

8.5.1 Upgrade prerequisites

The 3590 to 3592 upgrade is only supported in a Model B20 VTS. In a PtP VTS configuration, you can upgrade one VTS after the other, or you can only upgrade one VTS to 3592 and leave the other with 3590 drives. We do, however, recommend that both VTSs have the same hardware configuration.

Hardware requirements

Additional hardware is required to support the attachment of the 3592 drive to Model B10 or B20 VTSs. This hardware is added via feature codes. The 3494 model D22 frame also needs appropriate feature codes to support the installation and attachment of the 3592 drives to the VTS. Refer to the latest version of the *3494 Introduction and Planning Guide*, GA32-0448, for the required features codes and their descriptions.

If only 3592 drives are to be supported by the VTS, the VTS controller can be located up to 75 feet away from the D22 frame in the 3494 library. If both 3590 and 3592 drives are to be supported, the distance between the VTS controller and the 3494 library is determined by the 3590 attachment distance allowed (up to 46 feet). The D12 and D22 drive frames are not required to be adjacent in the 3494 library.

If the 3592 drives are in a TS3500 library, the VTS controller connects to a 3953 library manager and the 3953 library manager then connects to the drives in the library. The distances between the VTS controller and the 3953 and the 3953 and the 3592 drives can be up to 100 feet each with the supplied cables. Customer supplied cables can extend the distance between the 3953 and the 3592s to support greater distanced.

Intermixing 3590 and 3592 tape drives within the same D12 or D22 drive frame is not allowed. If 3590 and 3592 drives are to be intermixed, the VTS must have the Advanced Policy Management feature (FC4001) installed.

Field upgrade requirements

Field upgrade of a model B10 to support 3592 is not provided. An existing model B20 can be upgraded in the field to support the attachment of 3592 drives. In addition to the upgrade of the model B20, additional field upgrades are supported for the following combinations:

- Conversion of a 3494 D12 frame attached to a model B20 VTS to a D22 frame (applies if the VTS is attached to two D12 frames)
- Addition of a 3494 D22 frame to a model B20 VTS (applies if the VTS has a single D12 frame attached)
- Addition of a 3494 D12 frame to a model B20 VTS (applies if the VTS has a single D22 frame attached)
- Conversion of an open systems 3494 D22 frame to one that can be attached to a model B20 VTS (applies if the VTS has a single D12 frame attached)
- Attachment of 3592 drives within a TS3500 library.

Refer to the latest version of the 3494 and TS3500 *Introduction and Planning Guides*, GA32-0448 (3494), or GA32-0559 (TS3500), respectively, for further information about the necessary features and any prerequisites.

Microcode and software requirements

Support for 3592 drives and media was introduced with VTS code level 2.31.730 and Library Manager code level 530. The above listed code levels are included with feature code 5238 - Fibre Card Drive Attachment, that is required to attach 3592 drives to the VTS. Emulation support for the 3592 model E05 drive was introduced at VTS code level 2.32.740 and library manager code level 532.

Although existing host software levels can support VTS systems using the 3592, it is recommended that for z/OS V1 R3 or later, APAR be applied OA08777. These software levels support the reporting of additional statistics and pool control settings included in the SMF type 94 record.

8.5.2 Planning for the integration of 3592 drives and media

When planning for the installation of the 3592 drives and media in the VTS, whether the VTS is new or an existing one, the model of the VTS and whether existing data is to be retained are all key considerations.

If 3592 drives and media are to be installed in an existing VTS along with existing 3590 drives and media, there are several considerations:

- The VTS must be a model B20.
- The library must be a 3494.
- Decide whether existing data will be retained indefinitely or migrated to 3592 media allowing the 3590 drives and media to be removed from the VTS.
- Special consideration needs to be given to an existing VTS that has more than six 3590 drives, as only six of those 3590 drives are supported when both 3590 and 3592 are installed in the same VTS.

  Note: If the reason for having more than six drives is because of recall activity, some impact to recall times may be experienced until the data to be recalled has been migrated to 3592 media.

- If the VTS does not have the Advanced Policy Management feature installed, it must installed either prior to or with the installation of the 3592 drives and associated features. The volume pooling function of the APM feature is used control whether data is written to 3590 drives and media or 3592 drives and media. In addition, with the installation of the feature, all existing 3590 media with active data is assigned to pool 1 by default. You must take that into account as you set up the data management definitions and determine how you want to migrate data from 3590 to 3592 (described in subsequent sections).

If the VTS is part of a PTP VTS configuration, drives and media types do not have to match on each of the VTSs in the configuration.

The dual copy functionality of advanced policy management is available for use by both 3590 and 3592 drive types and media. There is no requirement for both the primary and secondary pools to be of the same drive/media type. For example, a pool used to store the secondary copy of logical volumes could be defined with 3592 drives and 300GB media, while the primary copy is written to a pool defined with 3590 drives and media.

8.5.3 Data migration considerations

Existing data in a VTS can be migrated from 3590 media to 3592 media by combining one or more of the following methods.

Direct new data to a pool with 3592 drives and media

With this method, all scratch logical volume allocations specify a Storage Group whose physical volume pool uses 3592 media. This could either be a new Storage Group defined for 3592 media or an existing one whose primary pool definition has been modified from 3590 media to 3592 media (Storage Groups and how to set them up are covered in the Setting Up Data Management Definitions section). Defining a new Storage Group is recommended as it allows data to be written to either drive type and, depending on naming conventions used for Storage Group names, may be less confusing. This method only deals with new allocations, so one or more of the other methods is also required.

Recall from 3590 and copy to a pool with 3592 drives and media

With this method, existing logical volumes are recalled and, when subsequently copied from the cache, are written to a physical volume pool that uses 3592 media. To accomplish this, the physical volume pool associated with the logical volumes must be changed to one that contains 3592 media. This is done by changing the primary pool defined for the existing Storage Group assigned to the volumes, or by assigning a different Storage Group to the volumes that is defined with a primary pool containing 3592 media (Storage Groups and how to set them up are covered in the Setting Up Data Management Definitions section).

Changing the primary pool for the existing assigned Storage Group is probably the simplest to do (one change through the manage Storage Groups window versus determining volser ranges for the volumes to change and then using the Manage Logical Volumes window). However, if the naming conventions used for the Storage Group names use the drive type (such as 3590) it may cause some confusion. Also, if a host continues to assign the original Storage Group name in its ACS routines, the logical volumes created will be directed to the pool with 3592 media.

Once the Storage Group changes have been made, old data is automatically migrated as it is accessed. Old data can also be forced to migrate. For example, if there is a set of logical volumes to migrate from 3590 to 3592 media, a job could be run that would cause them to be recalled. When they are subsequently unloaded and demounted, the VTS will see that the pool now associated with the volumes is different than the one the last time they were copied from the cache and it will copy them to the new pool. This method, combined with directing all new allocations as described above, is probably sufficient if the time period allowed for migration is long enough such that most of the old data has expired or recalled as part of normal job processing, leaving a relatively small amount of data to migrate by forcing a recall. If there are a large number of volumes to migrate, the effort to set up the host jobs to force the recalls becomes cumbersome and cache space will be taken up by the recalled volumes.

Consider that:

- If there are a large number of volumes to migrate and this method is desired, the efficiency for recalling the data from the 3590 tapes can be improved by using a map of the logical volumes on them. Rather than randomly recalling the logical volumes, the map can be used to structure the recall job such that the volumes are requested in the order they are located on the 3590 tape. The Bulk Volume Information Retrieval (BVIR) function can provide the physical volume to logical volume mapping information. Refer to the 3494 Bulk Volume Information Retrieval Function User's Guide for details on how to use the function.
- If a logical volume is in cache at the time its primary pool assignment is changed, the pool assignment will not effect where the data will be copied to until the volume is subsequently unloaded and demounted. It is during the rewind/unload command processing time that the primary pool destination for a logical volume is determined.
- In a VTS that does not have the Advanced Policy Management feature installed, logical volumes are managed as if they were all assigned a Storage Group that directed the data to pool 1. The 'under the covers' Storage Group assigned is the default, all dashes '--------', and it is associated with pool 1. When the 3592 drives and media are added and the Advanced Policy Management feature is installed, the existing logical volume retain their Storage Group assignment of '--------' and association with pool 1. To ensure that the recalled data is redirected to a pool with 3592 drives and media, you must either change the Storage Group assignment for the existing logical volumes or change the pool associated with the default Storage Group name of all dashes. If you do not make a change, even though you are using the reclaim setting to move an existing logical volume from 3590 to 3592 media, on a subsequent recall, the volume is written back to the pool with 3590 (pool 1).

Move logical volumes on a set of 3590 to a set of 3592 physical volumes

With this method, the valid logical volumes on one or more 3590 physical volumes are moved to a target pool that contains 3592 media. This is accomplished by using the move stacked volume function of the library. Through that function, the logical volumes resident on a range of stacked volumes are moved to a target pool immediately or as each stacked volume becomes eligible for reclaim. This method is a better alternative than the force recall method to move a small number of logical volumes. The VTS manages the migration internally so it does not require any host jobs to be run. Data moves from tape to tape so no cache space is needed.

Allow reclamation to move logical volumes to a pool with 3592

With this method, data from a pool with 3590 media is migrated to a pool with 3592 media based on a selected reclamation policy and associated parameters. Normally, the target pool for reclamation is the same as the source pool so that the active logical volumes remain in the pool specified by their associated Storage Group name. With this method, however, the target pool specified for a pool(s) with existing 3590 data is changed to a pool that contains 3592 media.

In addition to specifying a target pool with 3592 media, it is also recommended that the definition of the Storage Group name assigned to the volumes to be migrated be modified to reflect the new primary pool, or a new Storage Group name with the new primary pool be assigned to the volumes (see the discussion on the recall data method for the reason one method may be chosen over the other). If the primary pool is not changed, when a migrated logical volume is recalled and subsequently unloaded and demounted, the VTS will copy the data to the old 3590 pool.

With VTS release 7.3, additional reclamation policies are supported to allow better control over the migration time period. A reclamation policy for each pool is specified as part of the VTS management policies definitions. The following policies are supported:

- **Reclaim Percentage Threshold:**

  A physical volume is eligible for reclaim when the amount of active data on the volume falls below the threshold defined for the pool (Prior to VTS release 7.3, this was the only reclamation policy supported). A value of 5 to 95 percent can be specified for this field. This is the default policy.

- **Days Since Last Accessed:**

  A physical volume is eligible for reclaim when the number of days defined in the *days without access field* has elapsed since any data on the volume has been accessed because of a recall. A value from 1 to 365 days can be specified as a criteria for reclaim. A value of 0 disables this criteria for reclaim.

- **Days Since Last Written:**

  A physical volume is eligible for reclaim when the number of days defined in the *age of last data written* field has elapsed since any data was written to the volume. A value from 1 to 365 days can be specified as a criteria for reclaim. A value of 0 disables this criteria for reclaim.

- **Days Since Last Data Inactivation:**

  A physical volume is eligible for reclaim when the number of days defined in the *days without data inactivation* field has elapsed since any data was invalidated on the volume and the amount of active data on the volume falls below the threshold defined in the maximum active data field.

  A value from 1 to 365 days can be specified as a criteria for reclaim. A value of 0 disables this criteria for reclaim. A value of 5 to 95 percent can be specified for the maximum active data field.

  Note that the maximum active data field is only used in conjunction with the days since last data inactivation policy. It is independent of the reclaim percentage threshold field. A portion of the data on a physical volume is invalidated when the logical volume it represents has been modified/rewritten or deleted (as part of the delete expired volume data function). The remaining data on the physical volume is considered active data.

A portion of the data on a physical volume is invalidated when the logical volume it represents has been modified/rewritten or deleted (as part of the delete expired volume data function). The remaining data on the physical volume is considered active data.

Note: The reclamation workload for a VTS will increase while the migration of old data is taking place. The impact of that increased workload should be monitored and if it is impacting production use of the VTS, the inhibit reclamation policy setting for the VTS should be used to minimize that impact during peak production demand times.

8.5.4 Dealing with physical volumes during the data migration

One of the tasks to be performed during the data migration is the monitoring and management of the 3590 media that is emptied as part of the migration to 3592.

Important: It is very important that, during the migration process, the library does not run out of 3590 media in its common scratch pool. If one or more pools have borrowed 3590 media from the common scratch pool and the common scratch pool becomes empty of 3590 media, the library will enter the Out of Empty Stacked Volumes state and logical mounts will stop being processed regardless of the eventual physical pool and media the logical volumes would be written to. It is recommended that at least three 3590 volume remain in the common scratch pool until the migration has been completed and all of the borrowed 3590 media has been returned.

In case of an *Out of Empty Stacked Volumes* state in one of the VTSs, the operation of the PtP VTS is also inhibited. If no logical mounts can be performed, no copies can be created, even if host data can be still be written into the other VTS which has empty stacked volumes.

As part of preparing for the migration, there are steps to be taken with the pool(s) with 3590 media.

- First, change the borrow indicator to No Borrow/Return in the pool definition (this is done through the Stacked Volume Pool Properties panel on the library manager). This prevents any 3590 physical volumes from being moved into the pool from the common scratch pool and as the physical volumes are reclaimed, they will return to the common scratch pool so long as they were borrowed in the first place. Except for the last three, borrowed physical volumes that are already scratch will be returned to the common scratch pool. A pool must be dormant (no write activity or need for a scratch volume) for 72 hours before the last three borrowed cartridges are returned to the common scratch pool. Do this after making sure that there are at least three 3590 physical volumes in the common scratch pool.
- Next, if physical volumes were explicitly assigned to the pool, you must move or eject any scratch volumes using the Move Stacked Volume function of the library. Explicitly assigned volumes are not returned to the common scratch pool when scratch.
  - Specify the Move Number of Stacked Volumes selection or the Eject Number of Stacked Volumes selection.
  - Make sure that you specify Scratch, the source pool and the number to move/eject.
  - If you are moving the volume to the common scratch pool, indicate a target pool of 00, or if you are ejecting volumes, specify the I/O station type where they are to go. You can determine the number of volumes in the pool using the Search Database - Volser, Constructs, Pools, etc. function of the library. After the scratch physical volumes have been moved, only physical volumes with data on them remain in the pool.

As the migration proceeds and physical volumes become scratch, any borrowed volumes will be returned to the common scratch pool. Explicitly assigned volumes will not and must be moved/ejected using the Move Stacked Volume function as described above. You should occasionally monitor the number of cartridges remaining in the pool and eject volumes as needed. Use the Move/Eject Stacked Volumes function of the library manager to indicate how many scratch volumes you want to eject. The 3590 volumes can be ejected at any time during the migration, however, take care to leave a minimum of 3 volumes in the common scratch pool if any had been borrowed and not returned or the library will enter the Out of Empty Stacked Volumes state and mounts will stop being processed. Once all of the borrowed 3590 volumes have been returned, then the remaining 3590 volumes can be ejected from the library.

After all of the previously full volumes in a pool have been migrated (and assuming you redefined the pool to return scratch volumes), a few physical volumes (generally less than 10) will be left in the pool. Those volumes are ones that were partially filled when the migration began and are not considered for reclamation. You will need to use the move stacked volume method to complete the migration of the data from the pool. It will then take up to 72 hours for all of the borrowed physical volumes to return to the common scratch pool.

Once you have completed the migration, you may redefine the storage pools previously used by the 3590 for other uses.

Important: To prevent the *Out of Empty Stacked Volumes* state after you have ejected all 3590 volumes from the VTS, you need to change the media class for the default pool, typically Pool 1. We recommend that you re-define the class of all pools from 3590 to 3592.

See 4.3.2, "Defining Stacked Volume Pool Properties" on page 169 for details on how to do this.

8.5.5 Setting up data management definitions

With DFSMS, you control where data is stored and how it is managed through a set of names associated with a logical volume during allocation. The names are called storage constructs, and storage construct names have policies associated with them that define where and how data is stored and managed. One of the constructs, Storage Group, is used to manage where data is stored.

For the VTS, a group of physical data storage cartridges, referred to as a physical volume pool, can be associated with a Storage Group construct name. If both 3590 and 3592 drives and media are included in the configuration, physical volume pools must be defined for each to allow for system level control to direct data to the appropriate drive and media type.

With the 3592 comes the option of two cartridge types. A 60 GB capacity faster nominal access time cartridge and the high capacity 300 GB cartridge. You also need to define different volume pools if both 3592 media types are to be used.

For separation of application data or because of different reclamation policies, you may also want to define additional volume pools even if only the 3592 drives are installed and a single media type is used.

Note: If the VTS does not have the Advanced Policy Management feature installed, internally, all data is managed as if pool 1 was defined, but all definitions for the pool are hidden. What this means is that 'under the covers' all logical volumes are assigned a Storage Group name of all dashes '--------' and that Storage Group is associated with pool 1.

When the APM feature is installed, the definitions for pool 1 are now available to be modified. If you want logical volumes created prior to the integration of 3592 to be written to a pool with 3592 drives and media when they are recalled, you must change the pool associated with the default Storage Group name of all dashes or assign the volumes a Storage Group name that is associated with a pool containing 3592 media.

If you do not make a change, recalled data will be written to 3590 media in pool 1, even though you had set up the reclaim policy to move the data to a pool with 3592 media.

Prior to setting up the definition of a physical volume pool, you will need to decide how physical volumes become part of a pool and how they will be reclaimed.

Physical volumes become part of a physical volume pool either by 1) being explicitly assigned to the pool as the volumes are inserted into the library, or 2) by being borrowed from the common scratch pool when needed or 3) moved to the pool using the move stacked volumes function of the library. If volumes are to be explicitly assigned to a pool during insert, set up a volume serial number range for a specific pool using the Volser Range window on the library prior to inserting them.

A pool can “borrow” physical volumes from the common scratch pool when needed and optionally “return” them to the common scratch pool when they have been reclaimed. The borrow and return rules are specified as part of the definition for physical volume pool. With the move stacked volume function, one or more scratch physical volumes can be moved from one pool to another. A physical volume that has valid logical volumes on it can also be moved, however, the logical volumes are first copied to another physical volume in the originating pool.

There are six primary steps in setting up the data management definitions and controls:

1. **Decide on Storage Group construct names:**

   As part of the plan for integrating the 3592 drives and media into your VTS installation, you need to decide on the Storage Group construct names. It is suggested that the names be meaningful and perhaps relate to the type of data to reside on the pool or the drive type used in the pool. For example, if there is to be one general pool that uses 3590 drives and media, a name like “SG3590GP” could be used. “SG” for Storage Group, “3590” indicates the drive type for the pool and “GP” indicates that it is the general pool. Likewise, for the pool used to store large archive type data that uses 3592 drives and media, “SGLARGEA” would be a good choice.

2. **Define the Storage Group names to DFSMS and the library:**

   See 4.2, “VTS definitions from the Library Manager” on page 151 and 4.3, “Defining Storage Groups and Pools” on page 166 for more information.

3. **Define the volser ranges for the 3592 media:**

   See 4.2.2, “Define stacked volume ranges” on page 153 for details.

4. **Define the characteristics of the physical volume pools to the library:**

See 4.3.2, “Defining Stacked Volume Pool Properties” on page 169 on how to define the pools. Note that modifying the properties of an existing pool to change its drive class is not allowed if the pool contains media for an incompatible drive class.

There are two aspects that must be considered and action taken prior to modifying a pool’s drive class:

– A Volser range for a media type incompatible with the new drive class was previously established for the pool. For example, through the Volser range for stacked volumes function, a range was set up for a media type compatible with a drive class of 3590 and a home pool of 3. Pool 3's drive class cannot be changed to 3592 because it is not compatible with 3590 media types. Either the media type for the range or the home pool in the Volser range table must be modified before the pool property change is allowed.

– A pool already contains stacked volumes whose media type is not compatible with the new drive class for the pool. For example, pool 3 already contains one or more 3590 compatible stacked volumes. To change the drive class for pool 3 to 3592, the 3590 compatible stacked volumes and any associated active logical volumes must first be moved to another pool.

5. **Code/modify the Storage Group Automatic Class Selection (ACS) routine:**

   For additional information, see 4.5.1, “z/OS and DFSMS/MVS SMS-managed tape” on page 184.

6. **Activate the new construct names and ACS routines:**

   Before new allocations will be directed to the newly defined pools with 3592 drives and media, the source control data set (SCDS) with the new Storage Group definitions and ACS routines must be activated.

8.5.6 Additional considerations

If you are installing 3592 in a library that has two VTSs, there are some additional considerations when deciding on Storage Group names and pools. Storage Group names and the pools associated with them are a shared definition with both VTSs in a library. You cannot use the same Storage Group name to direct the data in one VTS to a pool and to a different pool in the other VTS. The pool properties can be different between VTSs, but the Storage Group name (and all of the other storage construct names) are common.

There is a special condition when one of the VTSs in the library is enabled for Advanced Policy Management (APM) and the other VTS is not. For the VTS that is not enabled for APM, all data is written to pool 1, regardless of the definitions for the default Storage Group. For example, assume that one VTS has only 3590 drives and one VTS has only 3592 drives. The VTS with 3592 drives is enabled for APM and if pool 2 is defined for the 3592 drives and media, the default Storage Group definition is changed to specify that pool 2 is to be used. Even though the Library Manager definition of the default Storage Group specifies pool 2, the VTS that is not enabled for APM will continue to write any data to its pool with 3590 drives and media, and since it can only have one pool, that is pool 1.

Additional information on the migration is available in the White Paper, *IBM TotalStorage Virtual Tape Server: Using 3592 In a VTS*, which is available from Techdocs at:

http://www.ibm.com/support/techdocs

9

Disaster recovery scenarios

In this chapter, we describe methods of testing the various components of the Peer-to-Peer VTS. We discuss the impact various component failure scenarios have on your PtP VTS and show how to test each scenario while your production work is still operating normally.

We cover the following topics:

- Peer-to-Peer VTS availability and testing considerations
- Failure scenarios
- DR testing with stand-by VTCs
- DR testing without stand-by VTCs
- GDPS considerations

9.1 Peer-to-Peer VTS availability and testing considerations

With a stand-alone VTS, when a failure occurs that causes the VTS subsystem to become unavailable, all the data stored in the VTS is unavailable until the repair is complete.

The Peer-to-Peer VTS is a solution to this problem. Its major design criterion is improving the data availability and recoverability characteristics of the VTS by duplicating all components of the VTS and by maintaining a dual copy of the data in the subsystem. It is also a solution for disaster recovery and electronic tape vaulting with transparent failover and failback.

The availability, recoverability, and performance characteristics will depend on the mode of operation selected. A performance enhancement is gained in the deferred copy mode of operation by scheduling the creation of the dual copy as Peer-to-Peer VTS activity permits. This, however, creates the risk that access to a successfully written volume may not be possible when there is a failure in one VTS. In this mode, the system as a whole continues to operate and provide access to most of the previous written virtual volumes. The amount of data that is inaccessible will depend on the particular workload and how much delay has occurred for the deferred copy operations.

It is, therefore, important that you understand the characteristics of the Peer-to-Peer VTS, and in particular, the consequences of the mode of operations prior to a failure. This section describes various availability scenarios. For each scenario, it summarizes what the failing component is, what happens in your system if a failure occurs, and how to regain normal operations. The theoretical number of failure scenarios, particularly with multiple component failures, is too great for us to address all of them. If you understand the consequences of a failure of a single component, you will be able to deal with multiple component failures.

But first, we will address the failure mode principles of the Peer-to-Peer VTS, which provides the availability service characteristics in case of a failure.

It is important that you check from time to time that your Peer-to-Peer VTS is in a fully operational state. This is not unique for the Peer-to-Peer VTS; it is good practice for all the equipment in your data center. Experience shows that you can avoid problems with your equipment if you verify your configuration at regular intervals. We identify the major items to verify in 9.1.2, “Verification procedure” on page 324. Some of the items can by automated if you use z/OS system automation and monitoring software.

As stated earlier, disaster recovery is a major reason for having a Peer-to-Peer VTS. The disaster recovery procedures, as performed by your SSR, are globally described together with changes in mount processing and disaster recovery alternatives.

The message support is enhanced for the Peer-to-Peer VTS. This includes new messages as well as improvements to the current message support.

9.1.1 Component availability principles

The Peer-to-Peer VTS provides the flexibility to separate hardware components for remote backup and recovery, which can improve data availability. Interconnected VTSs can also enable remote copy operations for remote vaulting applications. Each component in a Peer-to-Peer VTS is replicated one or more times to maximize data access and eliminate single points of failure. Access to volumes can then be maintained after component failures. In addition, service of one component of the Peer-to-Peer VTS can occur concurrently with operation of the rest of the Peer-to-Peer VTS. The Peer-to-Peer VTS provides the following availability and service characteristics:

- As long as one of the VTCs remains online, only one of the VTSs needs to be operational for the Peer-to-Peer VTS system to be operational. Data access when only one VTS is operational depends on the mode of operation prior to the failure.
  - Immediate copy mode: Access to data written by completed jobs (successful rewind/unload) prior to the failure is maintained. Access to data is lost for jobs that were in process at the time of the failure.
  - Deferred copy mode: In addition to the loss of access to data as in immediate copy mode, access to logical volumes written by previously completed jobs may not be possible until the failed VTS is returned to service.
- When only one VTC is operational, a second controller must be brought online before the first controller may be taken down for service.
- If all VTCs fail at the same time, both VTSs must be operational to bring one of the VTCs online and restore Peer-to-Peer VTS operation.
- If both VTSs fail at the same time, both must be brought up to restore normal Peer-to-Peer VTS operation.
- Changing logical volumes from the insert category to another category requires both VTSs and one VTC to be operational.
- A failure of either a VTCs or a VTS may cause jobs to abend. The jobs must be restarted.
- A failure of the I/O VTS where host jobs are performing writes will abend. As with VTC failures, the jobs must be restarted and the VTC will pick the other VTS for I/O operations.
- A failure of the I/O VTS where hosts jobs are performing reads will abend if there is no path from the VTC to the other VTS or the other VTS does not have a valid copy of the data. If there is a path to a VTS with a valid copy of the data, the processing for the job continues after a swap to another virtual drive and or path (subsystem initiated DDR). As part of the swap operation, the logical volume is demounted and re-mounted. As part of the re-mount operation, the VTC will select the remaining VTS with a valid copy of the data for I/O operations.
- When the VTS which was the Master VTS fails, and a switch to the other VTS as the master cannot be performed due to some other failure, processing of logical volumes that are using the remaining operational VTS as their I/O VTS can continue until they are closed, but no further mount operations can be completed. Mount operations can be re-started in one of two special modes after SSR action to enable the remaining VTS for read-only operation.
- Failure of a VTS which is the Slave VTS will allow jobs using the Master VTS for I/O operations to continue processing logical volumes until they are closed. Subsequent to the failure, all write type jobs can continue to be processed and read type jobs where a copy of the previously written data was made prior to the failure can also continue to be processed.
- In the event of a failure in the Peer-to-Peer VTS where one of the VTCs can no longer communicate with the Master VTS, the other VTS will become the Master VTS. The process to change the Master VTS is called *switchover*.
- To successfully complete a master switchover, at least three (for a four VTC configuration) or five (for an eight VTC configuration) VTCs must be fully operation and have a connection to the VTS proposed as the new master.
- If the current master VTS and more than one VTC fails, data operations with the Peer-to-Peer VTS cannot continue until the IBM SSR or you put the VTC in either Read/Write Disconnecteded mode or Read-Only mode with the remaining VTS. For more details, see 2.6.7, “Special VTC operational modes” on page 54.

- If you are in immediate copy mode, copy operations are deferred if:
  - One of the VTSs is not operational.
  - All links to a VTS are lost.
  - The tape library associated with a distributed VTS is not operational.
- If the PtP VTS was unable to create secondary copies during the time of a component failure, re-synchronization of the PtP VTS is given highest priority after the failing components have been repaired and the PtP VTS is fully operational again. This priority cannot be changed. You may encounter a performance degradation during re-synchronization, because the host cannot use the full bandwidth of the PtP VTS until the re-synchronization has completed.

9.1.2 Verification procedure

Here is a list of some helpful console commands to verify the status of a Peer-to-Peer VTS:

- `D SMS,LIBRARY(ALL),DETAIL` (check distributed and composite library; check state messages)
- `D M=DEV` (verify channel path)
- `LI DD` (verify virtual drives)
- `D R,L,KEY=OAM` (list outstanding replies)

9.1.3 Disaster recovery test enhancements

Disaster recovery (DR) testing in a PtP VTS environment requires some preparation and planning. Recent enhancements to the PtP VTS functionality can help you in planning and executing your disaster recovery testing. These enhancements include:

- **Selective PtP Dual Copy:**

  Selective PtP Copy allows you to write only one copy of a logical volume into a PtP VTS. When you use the No-Copy option during DR testing, you can write your logical volumes into the remote VTS only and minimize the impact on your production workload. No copy will be created, so you will not waste cartridge storage and bandwidth to copy logical volumes which need to be deleted again after the test.

  In addition to specifying No Copy in a Management Class for test data, you can also that you can also specify that only the VTS local to the test site is to have the volume by selecting it as the I/O VTS. If you use the concurrent testing method where no links are broken, it is a way to ensure that only the local VTS is used for the scratch mount.
- **Physical Volume Pooling:**

  Separating your logical volumes on cartridges within a dedicated pool during DR testing will help you to easily remove logical volumes no longer needed after the test.
- **Customer control of VTC modes:**

  Three VTC modes can defined through the PtP VTS Specialist as described in 9.1.4, “Setting VTC mode” on page 325. These modes are:
  - **Read-Only mode:**

    A VTS can be set in Read-Only mode, so that logical volumes can be read from the VTS, but no logical volumes can be created or modified as long as the VTS is set to this mode. You can use Read-Only mode during DR testing if you are testing at a remote site simulating that all connections to the local site are unavailable, and if you do not have to write data during the test.

– **Read/Write Disconnect mode:**

If the Master VTS fails, the remaining VTS can be set into Read/Write Disconnect mode, which allows write operation into the VTS to continue.

Attention: It is not allowed to operate both VTSs of a PtP-VTS in Read/Write Disconnected mode in parallel. If the VTSs cannot communicate with each other through the VTCs, only one VTS should operate in this mode; the other VTS can then operate in Read-Only mode.

Should both VTSs be inadvertently set to Read/Write Disconnect mode, the same virtual scratch volume could be used from different applications. When the VTSs once again are linked, part of the checking that is performed is to see if there are any volumes that have conflicting tokens, meaning that each was modified under Read/Write Disconnected mode. Any volumes volume with that state are blocked from access. Engineering intervention is needed to resolve these inconsistencies.

– **Write Protect mode:**

A single VTC can be set in Write Protect mode so that neither the contents nor the attributes of an existing volumes can be changed when mounted on a virtual drive of a VTC in this mode. YOu can define one or more VTCs to operate in Write Protect mode while the others are operating in normal read/write mode.

This function can be used to during DR testing, for example, by setting VTCs dedicated to the test system to this mode.

In 9.2, "Failure scenarios" on page 328, we have included information on how a component failure will affect the overall operation of the PTP VTS, and how to recover from a component failure. This information may support you in planning your disaster recovery procedures and in preparing DR testing.

9.1.4 Setting VTC mode

With VTS microcode Release 7.4, the selection of the VTC mode of operation can be selected to from the VTS Peer-to-Peer Specialist panels. By connecting to one of the VTCs, the Read-Only, Read/Write Disconnected, or Write Protect mode can be selected. The panel provides you with different selections based on the VTC state:

VTC online A VTC enters the online state at initialization when it can communicate with both VTSs or when it can communicate with only one VTS that is connected to another VTC already in the online state. When in the online state, the VTC can be set to Write Protect mode.

VTC offline A VTC enters the offline state at initialization when it can communicate with only one VTS that is not connected to another VTC in the online state. When in the offline state, the VTC can be set to Read-only or to Read/Write Disconnected mode.

Write Protect mode

The operator can only enable or disable Write Protect mode when the VTC is in the online state. Figure 9-1 shows the panel which is shown when Write Protect mode can be enabled. It contains the following informational items:

- The current status of Write Protect mode, either enabled or disabled
- A warning to the operator that enabling Write Protect mode causes pending write requests to fail, possibly resulting in job failures

- A caution to the operator to ensure that there are no active jobs being written to tape before enabling Write Protect mode

Before we discuss information on how a component failure will affect the overall operation of the PTP VTS, and how to recover from a component failure. This information may support you in planning your disaster recovery procedures and in preparing DR testing.

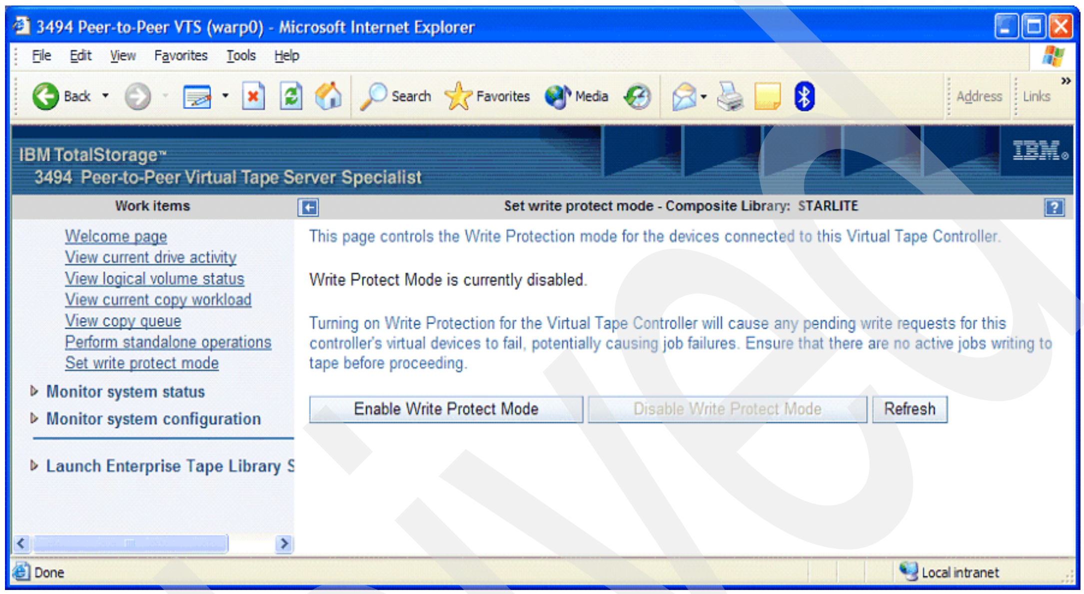

Figure 9-1 Write Protect mode on Specialist panel

Write Protect mode can be selected on a single VTC independent of any other in the library. This mode can also be enabled or disabled without taking any VTC offline. This will create an environment where one VTC and its defined devices will only operate in this mode while the rest of the VTCs run normal production work. While in this mode any of the following host operations will fail on any of the devices defined to the VTC in this mode:

- Any scratch mounts
- A host attempt to modify a logical volumes's contents that was mounted while in this mode
- A host attempt to change the attributes for the volume (category, constructs)

> **Note:** The Write Protect mode does not affect write type commands unless the volume is mounted while the mode is active.

In Figure 9-2, you see the Library Manager screen when Write Protect mode is enabled.

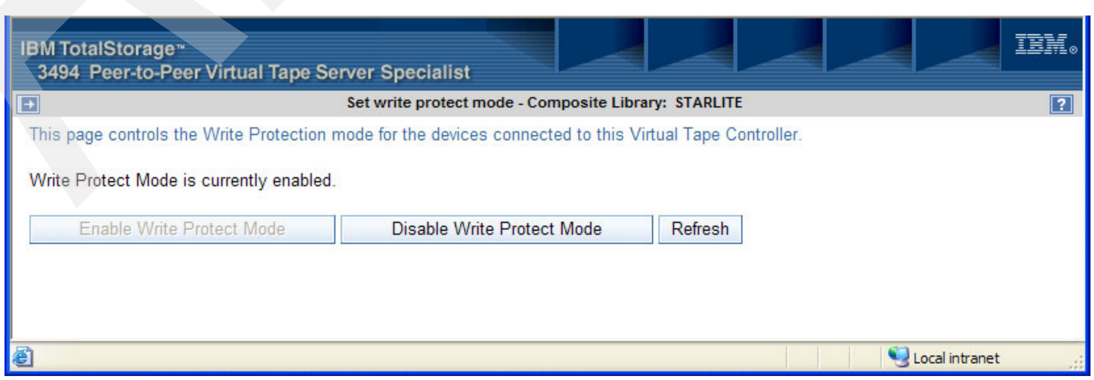

Figure 9-2 LM Panel with Write Protect mode enabled

When Write Protect Mode is enabled on all VTCs of a host, an attempt to modify a logical volume's constructs (VOLSER SL0186) or to insert logical volumes (VOLSER SL3000) results in the following host error messages listed in Example 9-1.

Example 9-1 Host error messages in Write Protect Mode

```
CBR1086I LIBRARY LMPOLICY results for volume SL0186
         FUNCTION NOT COMPATIBLE WITH THE LIBRARY.

CBR3606I Entry of volume SL3000 into library STARLITE failed. Unable to set the volume
category.
CBR3726I Function incompatible error code 10 from library STARLITE for volume SL3000.
```

Read/Write Disconnected and Read-Only modes

Read-Only and Read/Write Disconnected modes are only available when the VTC is in the offline state. The option to enable or disable Write Protect mode is not available then, as you can see in Figure 9-3. The page instead contains information about two other modes that were normally only available within the service panels of the VTC. These modes are still available to the SSR but are now available via the Specialist panels. The panel contains the following two push-buttons:

- **Activate Read/Write Disconnected Mode:**

  Pressing this push-button activates Read/Write Disconnected mode. In this mode, changes can be made to the contents and attributes of logical volumes; the VTC tracks these changes so that the VTSs can be synchronized when both are again available.

- **Activate Read-Only Mode:**

  Pressing this push-button activates Read-Only mode. In this mode, the data in logical volumes can be read but no changes can be made to their conte nts or attributes.

Figure 9-3 Read-Only and Read-Write Disconnected mode options on PtP Specialist.

Attention: Only one of these modes can be activated at a time. Once a mode is activated, both push-buttons disappear from the page and the mode cannot be deactivated until both VTSs are again available.

9.2 Failure scenarios

In this section, we discuss the impact and recovery procedures for failures in each of the Peer-to-Peer VTS components. Our examples are for a four-VTC configuration. For a complete list of failure scenarios, refer to the White Paper, *IBM Peer-to-Peer Virtual Tape Server Failover Test Scenarios in an OS/390 Environment*, which can be found on Techdocs at:

http://www.ibm.com/support/techdocs

9.2.1 Failure of host to VTC link

Figure 9-4 shows the failing component: Either one or both host connections to VTC 3 fail in this example.

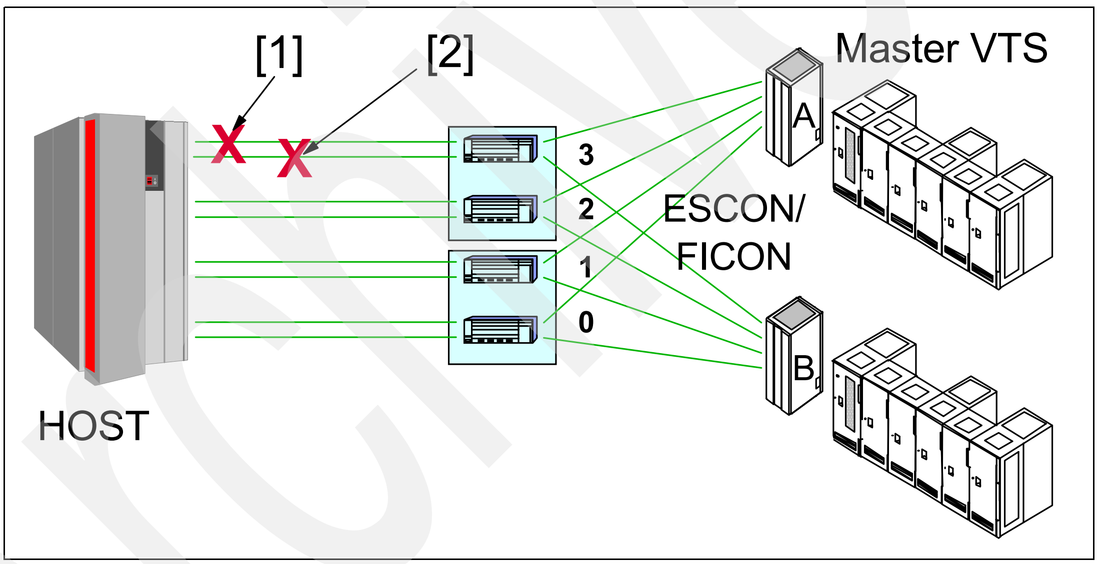

Figure 9-4 Host to VTC link failure

What fails?

One or more links between the host and the VTC fail. It may be that the intermediate directors, channel extenders, or remote channel extenders fail.

What happens?

If link [1] fails, only one link to VTC 3 remains available:

- All Peer-to-Peer VTS components continue to operate.
- Copied data remains available.
- All channel activity on the failing link is stopped.
- Channel errors are reported or error information becomes available from the intermediate equipment.
- Jobs that were using the 16 virtual drives on VTC 3 may fail.

If link [1] and link [2] fail, the host has no connection left to VTC 3:

- The virtual tape drive addresses for VTC 3 become unavailable, all other Peer-to-Peer components continue to operate.
- All channel activity on the failing host links is stopped.
- Host channel errors are reported or error information becomes available from the intermediate equipment.
- Copied data remains available.
- Jobs which were using the virtual tape drives of VTC 3 will fail.
- Logical volumes in use on VTC 3 are made available for subsequent use on other VTCs.

How to recover?

- Normal error recovery procedures and repair apply for the channel and the intermediate equipment. Contact your IBM Systems Service Representative (SSR) for repair of the failed components.
- Rerun the failed workload.

9.2.2 Failure of VTC or VTC links to both VTSs

Figure 9-5 shows the failing components. In this example, either VTC 3 or both links from VTC 3 to the VTSs fail.

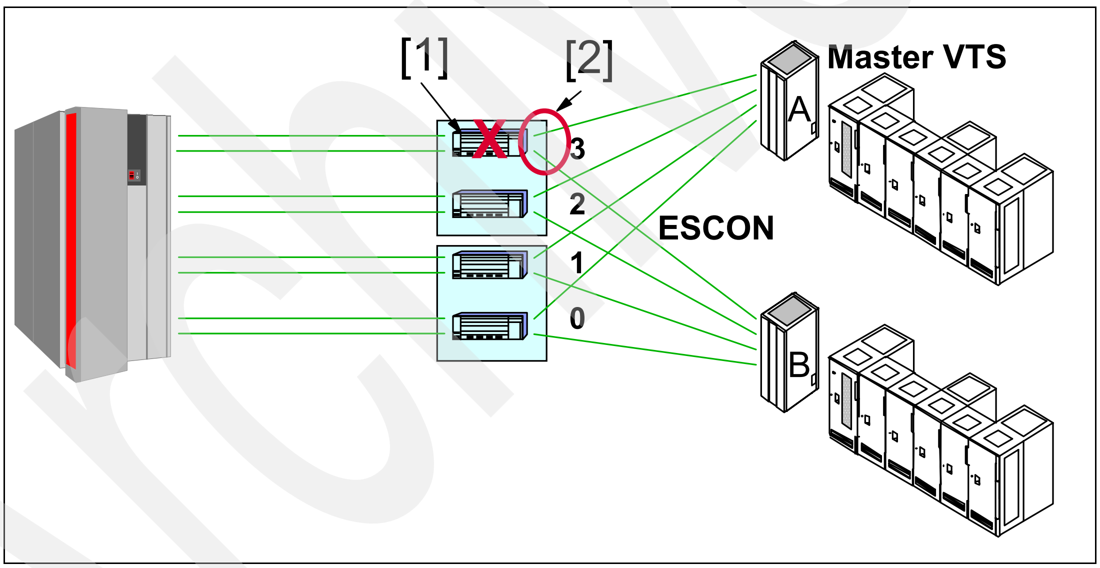

Figure 9-5 VTC failure or multiple link failure

What fails?

One or more of the VTCs stops processing or both links from one VTC to the VTSs fail. No communication is possible on the links to the host, nor on the links to the distributed VTSs.

What happens?

If the link failure between VTC 3 and VTS A occurs before the link failure between VTC 3 and VTS B, Master switchover may occur.

When [1] VTC 3 fails, or when [2] both VTS links fail from VTC 3:

- All remaining Peer-to-Peer VTS components continue to operate.
- Copied data remains available.
- The virtual drive addresses provided by VTC 3 are unavailable.

- The virtual addresses for the failing VTC may be boxed.
- Jobs using the virtual drives on the failing controller will fail.
- Logical volumes in use by the failing VTCs are made available for subsequent use by the other VTCs.
- Performance degradation is possible because the remaining VTCs must handle the total workload.
- Call home support is invoked by the VTC.

How to recover?

- Contact your SSR for the repair of the failing VTC or links.
- Rerun the failed workload.

9.2.3 Failure of VTC link to non-Master VTS

Figure 9-6 shows the failing component: one link between one VTC and the non-Master VTS.

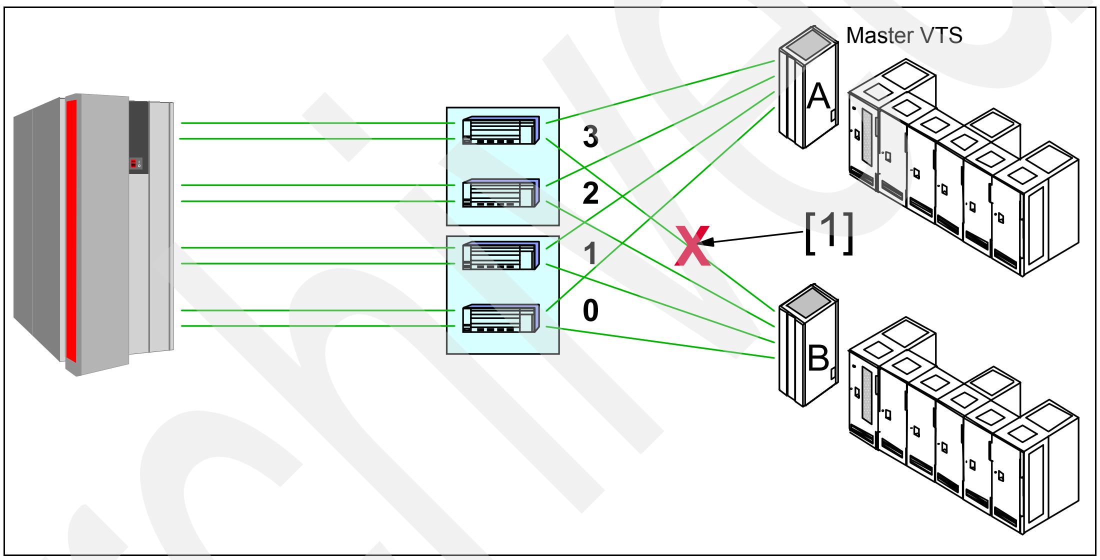

Figure 9-6 VTC to VTS link failure

What fails?

One link between a VTC 3 and VTS B fails. VTS A is the Master VTS.

What happens?

When connection [1] with VTS B fails:

- All Peer-to-Peer VTS components continue to operate.
- Copied data remains available.
- DDR is performed for the tape drive addresses for VTC 3 if they were reading logical volumes with VTS B as the I/O VTS. When DDR is performed and a tape drive address of another VTC is used, the job continues. If a tape drive address for the VTC with the failed link is used, the job continues if a copy of the logical volume is available in the connected VTS A.
- Any jobs writing logical volumes with the tape drive addresses for VTC 3 with VTS B as the I/O VTS will fail and must be rerun.

- In immediate copy mode, any dual copy of a logical volume in process at the time of the failure will be processed by other VTCs in the configuration as if the were deferred copies (as workload permits).
- Logical volumes that had failure of I/O operations in process due to the failed link are made available for subsequent use by other VTCs.
- Performance degradation is possible because only one VTS must handle the total workload.
- Call home support is invoked.

How to recover?

- Contact your SSR for repair of the failed connection.

> **Note:** After the failed link is repaired, immediate copy mode operation is automatically resumed by the VTC that had the failed link, if this was the copy mode before the failure.

9.2.4 Failure of VTC link to the Master VTS

Figure 9-7 shows the failing components.

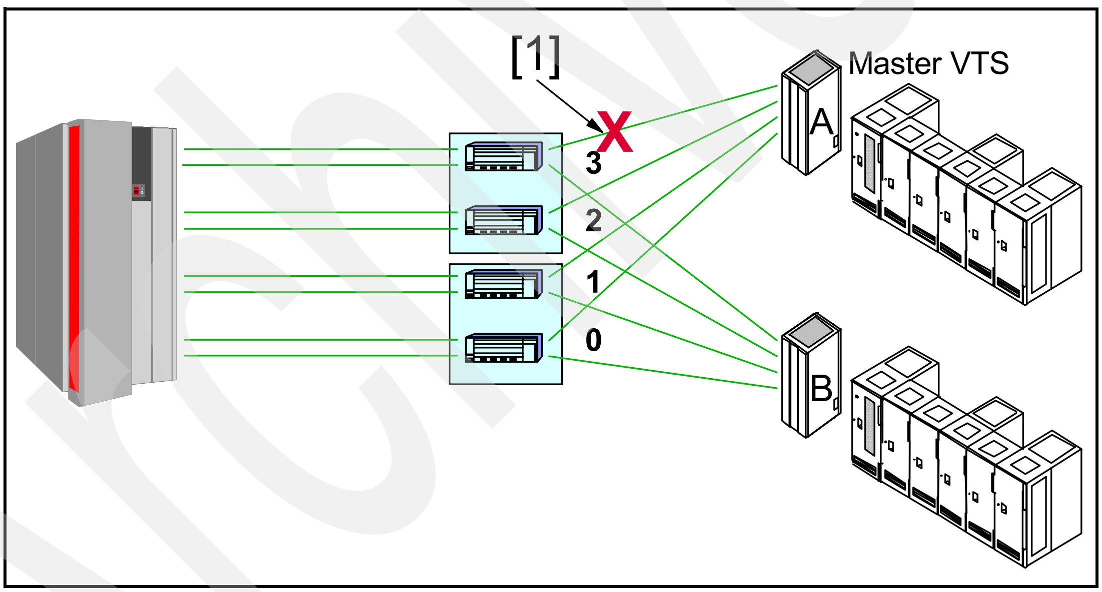

Figure 9-7 VTC to Master VTS link failure

What fails?

One link between a VTC and the VTS A fails. VTS A is the Master VTS.

What happens?

When connection [1] with VTS A fails:

- All Peer-to-Peer VTS components continue to operate.
- Copied data remains available.
- On the next mount request to VTC 3 (including a DDR mount), VTC 3 will initiate Master VTS switchover. Switchover occurs, and VTS B becomes the Master VTS.

- DDR is performed for the tape drive addresses for VTC 3 if they were reading logical volumes with VTS 3 as the I/O VTS. When DDR is performed, the job continues.
- Any jobs writing logical volumes with the tape drive addresses for VTC 3 with VTS A as the I/O VTS will fail and must be rerun.
- In immediate copy mode, any dual copy of a logical volume in process at the time of the failure will be processed by other VTCs in the configuration as if they were deferred copies (as workload permits).
- Logical volumes that had a failure of I/O operations in process due to the failed link are made available for subsequent use by other VTCs.
- Performance degradation is possible because the remaining VTCs must handle all workload.
- Call home support is invoked.

How to recover?

- Rerun any failed workload.
- Contact your SSR for repair of the failed connection.

After the failed link is repaired, immediate copy mode operation is automatically resumed by the VTC that had the failed link, if this was the copy mode before the failure.

If a Master VTS switchover has occurred, and if VTS A is the preferred Master, 30 minutes after restoring the failed link, the Master will be automatically switched back to VTS A.

9.2.5 Failure of VTC link to Master and non-Master VTSs

Figure 9-8 shows what has failed.

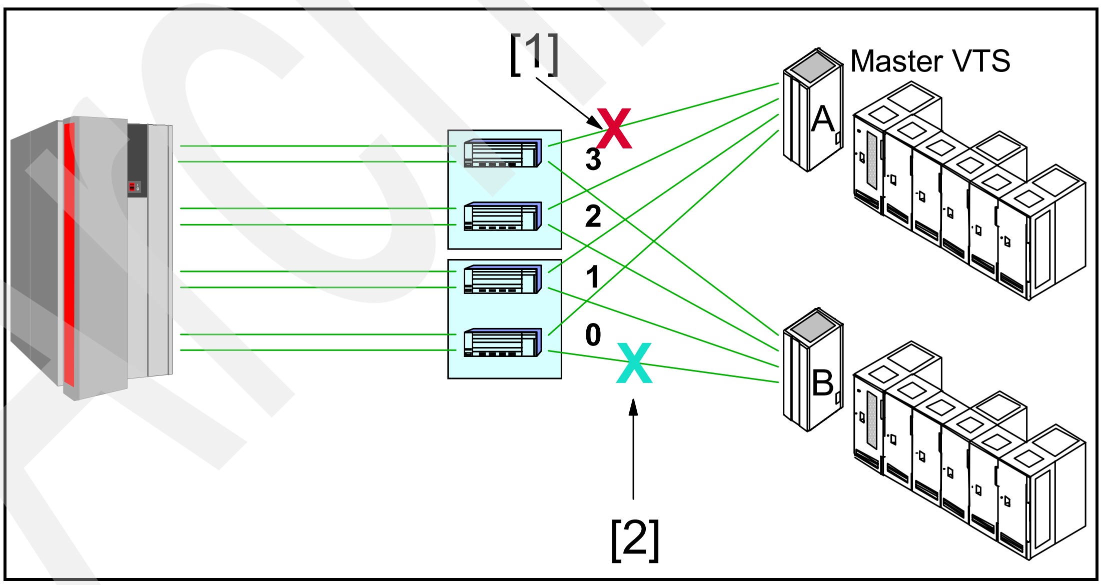

Figure 9-8 VTC to Master VTS link failure

What fails?

One link between a VTC and the Master VTS, and one link between another VTC and the non-Master VTS fail. VTS A is the Master VTS.

What happens?

Connections [1] and [2] with the VTSs fail.

- All Peer-to-Peer components continue to operate.
- Copied data remains available.
- On the next mount request to VTC-1 (including a DDR mount), VTC-1 initiates Master VTS switchover. Master VTS Switchover fails.
- DDR is performed for the tape drive addresses for VTC 3 if they were reading logical volumes with VTS A as the I/O VTS. If, when DDR is performed, a tape drive address for VTC 3 is used, the DDR will fail. If, when DDR is performed, a tape drive address for another VTC is used, the job continues.
- DDR is performed for the tape drive addresses for VTC 0 if they were reading logical volumes with VTS A as the I/O VTS. If, when DDR is performed, a tape drive address for VTC 3 is used, the DDR will fail. If, when DDR is performed, a tape drive address for VTC 1 or VTC 2 is used, the job continues. If, when DDR is performed, a tape drive address for VTC 0 is used, the job continues only if a copy of the logical volume is available in VTS A.
- Any jobs writing logical volumes with the tape drive addresses for VTC 3 with VTS A as the I/O VTS will fail and must be rerun.
- Any jobs writing logical volumes with the tape drive addresses for VTC 0 with VTS B as the I/O VTS will fail and must be rerun.
- In Immediate Copy Mode, any dual copy of a logical volume in process at the time of the failure will be processed by other VTCs in the configuration as if it were a deferred copy (as workload permits).
- Logical volumes which had failure of I/O operations in process due to the failed link are made available for subsequent use by other VTCs.
- Performance degradation is possible because the remaining VTCs must handle all workload.
- Call home support is invoked.

How to recover?

- Rerun all failed workloads.
- Contact your SSR for repair of the failing connections.
- After the failed links for VTCs 0 and 3 are repaired, vary the virtual tape drive addresses online from the host as required.

> **Note:** After the failed links are repaired, immediate copy mode operation is automatically resumed by the VTCs that had a failed link, if they were in copy mode before the failure.

9.2.6 Failure of all links to a VTS

Figure 9-9 shows that either all VTC to VTS A links or all VTC to VTS B links fail.

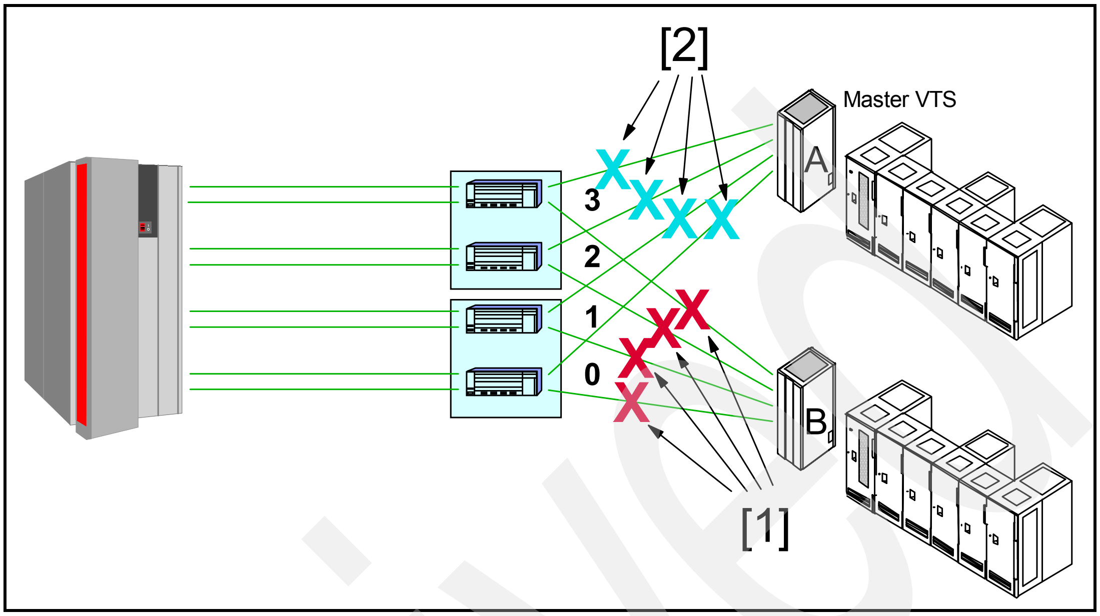

Figure 9-9 Failure of all links to a VTS

What fails?

All links between all VTCs and one VTS fail. VTS A is the Master VTS.

What happens?

When all connections [1] with **VTS B** fail:

- All Peer-to-Peer components continue to operate.
- Copied data remains available.
- Logical volumes in VTS B that have not been copied will not be available.
- If in immediate copy mode, the dual copy of logical volumes will be deferred.
- DDR is performed for tape drive addresses that were reading logical volumes with VTS B as the I/O VTS. When DDR is performed, the job will continue if a copy of the logical volume is available in the connected VTS A.
- Any jobs writing logical volumes with VTS B as the I/O VTS will fail and must be rerun.
- Logical volumes that had failure of I/O operations in process due to the failed links are made available for subsequent use by other controllers.
- Performance degradation is possible because only one VTS must handle the total workload.
- Insert or eject of logical volumes is not possible.
- Preferencing is bypassed for VTCs that were preferencing VTS B.
- Call home support is invoked.

When all connections [2] with **VTS A** fail:

- All Peer-to-Peer components continue to operate.
- Copied data remains available.
- On the next mount request to any VTC (including a DDR mount), that VTC initiates Master VTS switchover. Switchover occurs, and VTS A becomes the Master VTS.

- Logical volumes in VTS A that have not been copied will not be available.
- If in immediate copy mode, the dual copy of logical volumes will be deferred.
- DDR is performed for tape drive addresses that were reading logical volumes with VTS A as the I/O VTS. When DDR is performed, the job will continue if a copy of the logical volume is available in the connected VTS B.
- Any jobs writing logical volumes with VTS A as the I/O VTS will fail and must be rerun.
- Logical volumes that had a failure of I/O operations in process due to the failed links are made available for subsequent use by other controllers.
- Performance degradation is possible because only one VTS must handle the total workload.
- Insert of new logical volumes is not possible.
- Preferencing is bypassed for VTCs that were preferencing VTS A.
- Call home support is invoked.

How to recover?

- Rerun any failed workload.
- Contact your SSR for repair of the failed links.

> **Note:** After the failed links are repaired, immediate copy mode operation is automatically resumed by all VTCs, if they were in copy mode before the failure. This starts the copy operations for the deferred copies and will impact the bandwidth for data between the host and the Peer-to-Peer VTS if there is a large backlog of copies to be made.

If a Master VTS switchover has occurred, and if VTS A is the preferred Master, 30 minutes after restoring the failed link, the Master will be automatically switched back to VTS A.

9.2.7 VTS failure

Figure 9-10 shows the failing components.

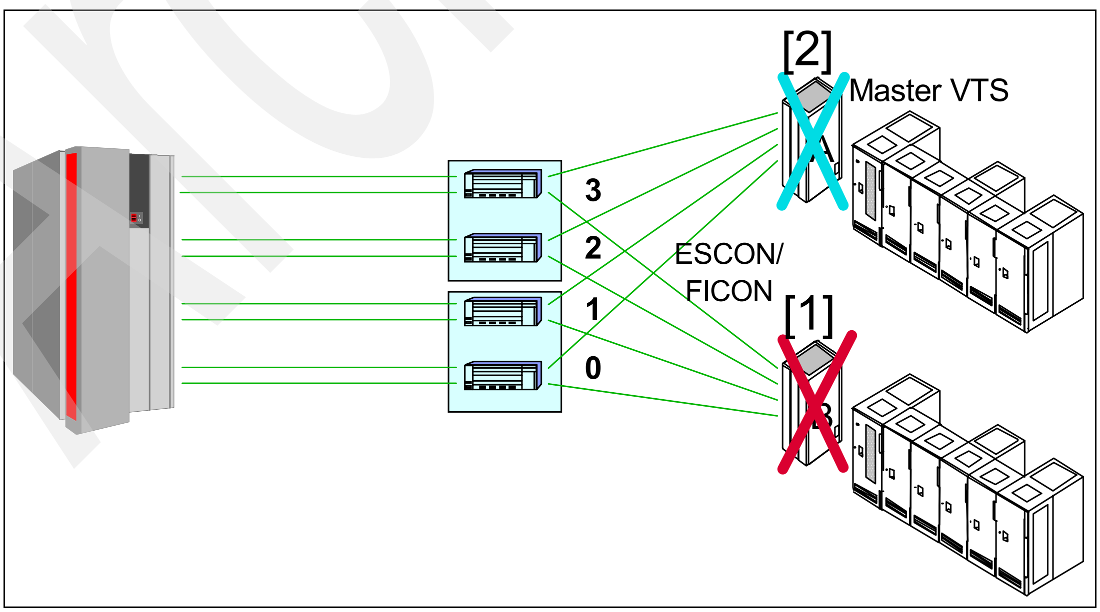

Figure 9-10 VTS failure

What fails?

One VTS, either VTS A or VTS B, fails. VTS A is the Master VTS.

What happens?

If VTS B [1] fails:

- All remaining Peer-to-Peer VTS components continue to operate.
- Copied data remains available.
- Logical volumes in the failed VTS that have not been copied will not be available.
- DDR is performed for tape drive addresses that were reading logical volumes with the failed VTS as the I/O VTS. When DDR is performed, the job will continue if a copy of the logical volume is available in the connected VTS that has not failed.
- Any jobs writing logical volumes with the failed VTS as the I/O VTS will fail and must be rerun.
- Logical volumes that had a failure of I/O operations in process due to the failed VTS are made available for subsequent use by other controllers.
- Performance degradation is possible because only one VTS must handle the total workload.
- Insert or eject of logical volumes is not possible.
- Preferencing is bypassed for VTCs that were preferencing the failed VTS.
- Call home support is invoked.

If the Master VTS, VTS A [2] fails:

- The same happens as described above when VTS B [1] fails.
- In addition, Master VTS switchover occurs, and VTS B becomes the Master VTS.

How to recover?

- Contact your SSR for the activation of the failing VTS.
- Rerun the failed workload.

> **Note:** After the failed VTS is repaired, immediate copy mode operation is automatically resumed by all VTCs they were in copy mode before the failure. This starts the copy operations for the deferred copies and will impact the bandwidth for data between the host and the Peer-to-Peer VTS if there is a large backlog of copies to be made.

If a Master VTS switchover has occurred, and if VTS A is the preferred Master, 30 minutes after restoring the failed link, the Master will be automatically switched back to VTS A.

9.2.8 Tape Library failure

When a Library Manager experiences a critical failure, communication with the VTS is lost. It is also possible that the link between the VTS and the Library Manager has failed. In this situation, the Library Manager is not able to inform the Peer-to-Peer VTS of the failure.

If the Library Manager is part of an HA1 configuration, or if the 3953-F05 contains two 3953-L05 Library Managers, the failure of the Library Manager causes a switchover to the Standby Library Manager to be executed. Once the switchover completes, the new Library Manager informs the Peer-to-Peer VTS that is online. It may take some minutes before the Peer-to-Peer VTS is informed of a switchover, and further minutes before the switchover is complete.

Dual control path attachment (FC5245) and dual control path facility (FC5246) provide a second Ethernet switch and necessary cables for a second independent LAN for the control path between the Library Manager and the tape subsystems installed in the library, thus enabling connectivity to the LM without interrupting operations.

The following scenario depicts 3494 tape libraries, but they apply equally well in a TS3500/3953 configuration. The only difference is the failure of the Library Managers attached to the VTS are located in the 3953 frame and not the TS3500 library.

Figure 9-11 shows the failing components.

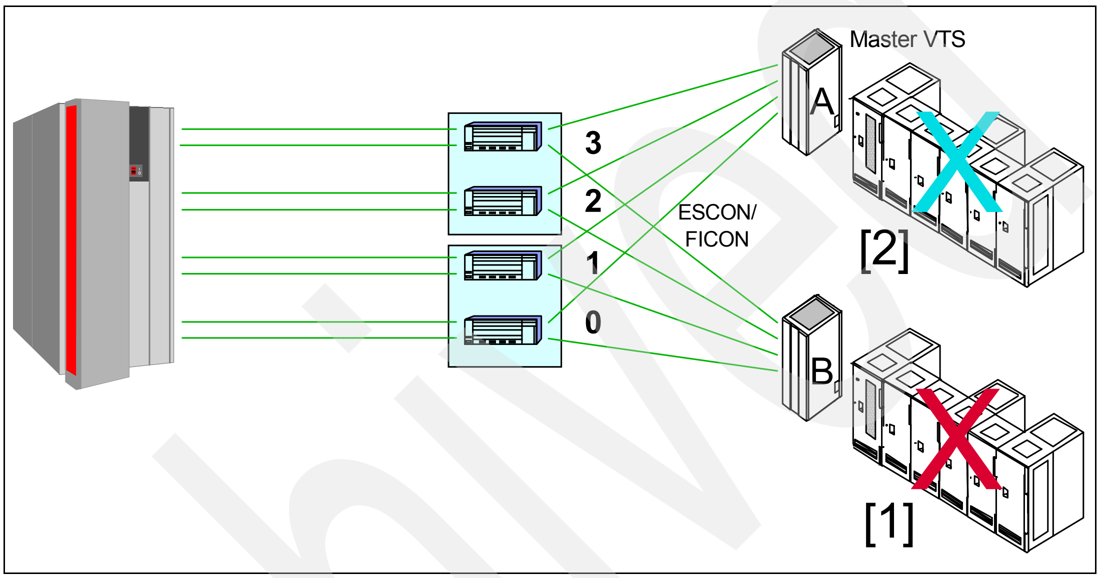

Figure 9-11 Tape Library failure

What fails?

A 3494 Tape Library or a TS3500/3953 Tape Library system housing one of the VTSs of the Peer-to-Peer VTS fails. The Master VTS, VTS A, is associated with library [2].

What happens?

If tape library [1] with VTS B fails:

- All remaining Peer-to-Peer VTS components continue to operate.
- Copied data remains available.
- VTS I/O operations in process will continue. Upon completion of operations in process, use of the VTS associated with the failed tape library will be discontinued by all VTCs.
- The mount of a logical volume requiring recall from a stacked volume in the failed tape library will fail.
- If in immediate copy mode, the dual copy of logical volumes will be deferred.
- Performance degradation is possible because only one VTS must handle the total workload.
- Insert or eject of logical volumes is not possible.
- Preferencing is bypassed for VTCs that were preferencing the VTS associated with the failed tape library.
- Call home support is invoked.

If the 3494 Tape Library or 3953 Library Manager [2] with Master VTS fails:

- The same as for first example, failing 3494 library [1] with VTS.
- VTS B [1] becomes the Master VTS.

How to recover?

- Contact your SSR for the activation of the failing 3494 Tape Library or 3953 Library Manager.
- Rerun the failed workload.

Note: After the failed tape library is repaired, immediate copy mode is automatically resumed by all VTCs they were in copy mode before the failure. This starts the copy operations for the deferred copies and will impact the bandwidth for data between the host and the Peer-to-Peer VTS if there is a large backlog of copies to be made.

If a Master VTS switchover has occurred, and if VTS A is the preferred Master, 30 minutes after restoring the failed link, the Master will be automatically switched back to VTS A.

9.2.9 Site failure

Figure 9-12 shows the failing components.

If you receive notice of an impending disaster such as a flood or tornado, you may wish to perform a Master VTS switchover if the Master VTS site is the site that could be affected. IBM support specialists can dial in to the Peer-to-Peer VTS to do this.

Note also that you could also configure and use Standby VTCs for additional resilience in this case, as described in 9.3, “Disaster Recovery testing with Standby VTCs” on page 341.

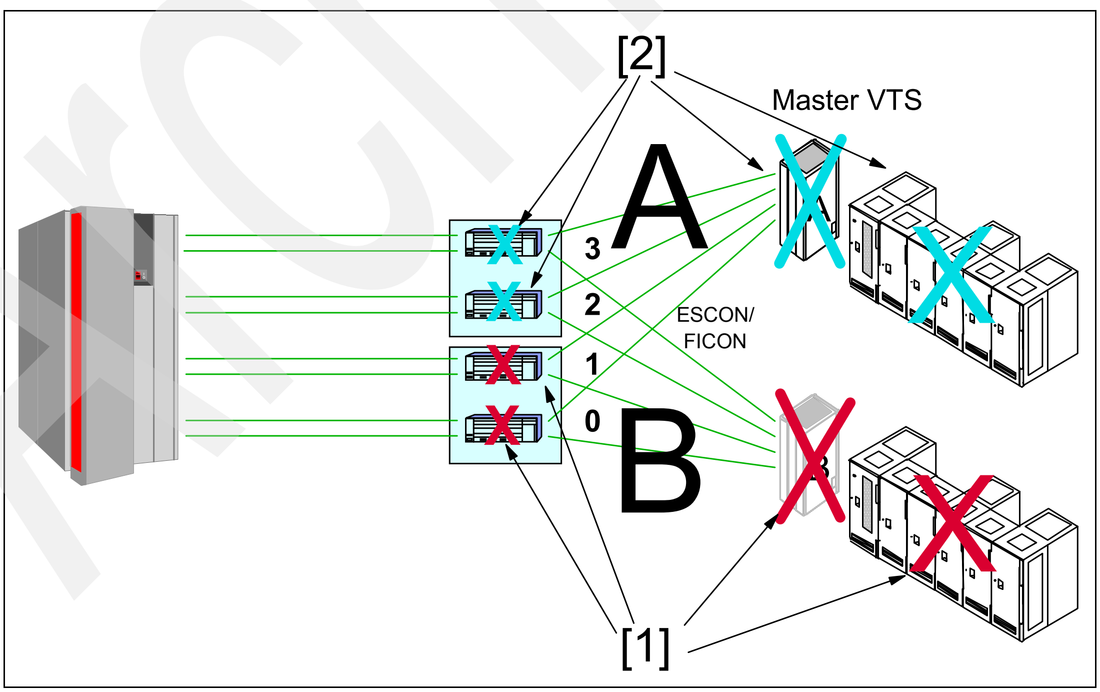

Figure 9-12 Site failure

What fails?

All the Peer-to-Peer VTS components in one site fail. Site A has the Master VTS. This scenario is applicable for site disasters.

What happens?

If the all Site B components [1] fail:

- All remaining Peer-to-Peer VTS components continue to operate.
- Copied data remains available.
- All channel activity on the failing host links are stopped.
- Host channel errors are reported or error information becomes available from the intermediate equipment.
- The virtual tape drive addresses for VTC 0 and VTC 1 will become unavailable.
- Logical volumes in VTS B which have not been copied will not be available.
- Jobs using the virtual tape drives for VTC 0 and VTC 1 will fail.
- Logical volumes which had failure of I/O operations in process due to the failed VTC are made available for subsequent use by other VTCs.
- Performance degradation is possible because the remaining site must handle the total workload.
- Insert or eject of logical volumes is not possible.
- Call home support is invoked.

If all Site A components [2] fail:

- Master VTS switchover fails.
- The remaining VTCs do not have a Master VTS connection.
- I/O operations in process on the remaining VTCs with VTS B as the I/O VTS will complete.
- All virtual tape drive addresses for the remaining VTCs will become unavailable ö
- All further processing is stopped.
- Virtual tape drive addresses may be boxed.
- Call home support is invoked.

How to recover?

After all Site B components [1] have failed:

- If a permanent disaster has happened, the disaster recovery procedure should be followed.
- Contact your SSR for repair of the failed components at site B.
- Rerun the failed workload.

After all Site A components [2] have failed:

- If a permanent disaster has happened, the disaster recovery procedure should be followed.
- Contact your SSR for repair of the failed components at site A.
- Consider operation of the remaining site B components in the special VTC operational modes described in 2.6.7, “Special VTC operational modes” on page 54. You can activate Read/Write Disconnected or Read-only mode through the PtP VTS Specialist panel described in 9.1.4, “Setting VTC mode” on page 325.

> **Note:** After the failed tape library is repaired, immediate copy mode is automatically resumed by all VTCs they were in copy mode before the failure. This starts the copy operations for the deferred copies and will impact the bandwidth for data between the host and the Peer-to-Peer VTS if there is a large backlog of copies to be made.

9.2.10 Other failures

These are failures in other parts of the subsystem.

Out of Empty Stacked Volumes

When a VTS runs out of scratch physical volumes, many of the operations of the distributed library cannot continue. Copying the virtual volumes to the stacked volumes is no longer possible. To prevent overflow of the cache, all logical mounts are suspended to this VTS. An intervention required message is generated, indicating that the distributed library is out of empty stacked volumes.

If this occurs, the Peer-to-Peer VTS takes the following actions.

- I/O processing to active virtual addresses continues.
- If the library out of stacked volumes is the Master VTS, a Master VTS switchover is initiated.
- Dual copy operations *to* the library that is out of stacked volumes are not performed.
- Dual copy operations *from* the library that is out of stacked volumes are not performed for logical volumes that are not in the Tape Volume Cache.
- If the Peer-to-Peer VTS is in immediate copy mode, the dual copy of logical volumes will be deferred.
- When the only valid version of a logical volume is in the Tape Volume Cache of the VTS that is out of empty stacked volumes, the logical volume is copied into the other VTS to allow a mount to be performed. The mount will fail if the logical volume is not in cache, since a recall is not possible.
- New mounts are accepted and processed on the other library (after a master switchover). The mount will fail if there is not yet a copy of the data in this library.
- Once your PtP environment has Advanced Policy Management installed, additional processing is done as part of Out-of-Scratch (OOS) handling. If an OOS condition in a physical pool occurs:
  - The VTCs will perform a master switchover to the VTS with no OOS pools unless any other conditions (such as link failures) prevent it.
  - With the slave VTS, logical volumes associated with that pool will not be balanced to the slave. This will result in mount skewing to the VTS with no OOS condition.
  - With the Master VTS (due to both VTSs having OOS pools, or the master switchover could not be done), then mounts to logical volumes associated with the OOS pools are held (just like a base VTS).
  - Any copies associated with pools in an OOS state are failed. These copies are then requeued and attempted at a later time.

An additional panel is available from the Real Time Statistics selection on the 3494 Specialist, which shows the media counts by stacked volume pool number. As can be seen in Figure 9-13, this panel shows, by pool:

- The pool number of the stacked volumes
- The media type of the stacked volumes
- The number of stacked volumes filling
- The number of stacked volumes marked as full
- The number of stacked volumes that are empty

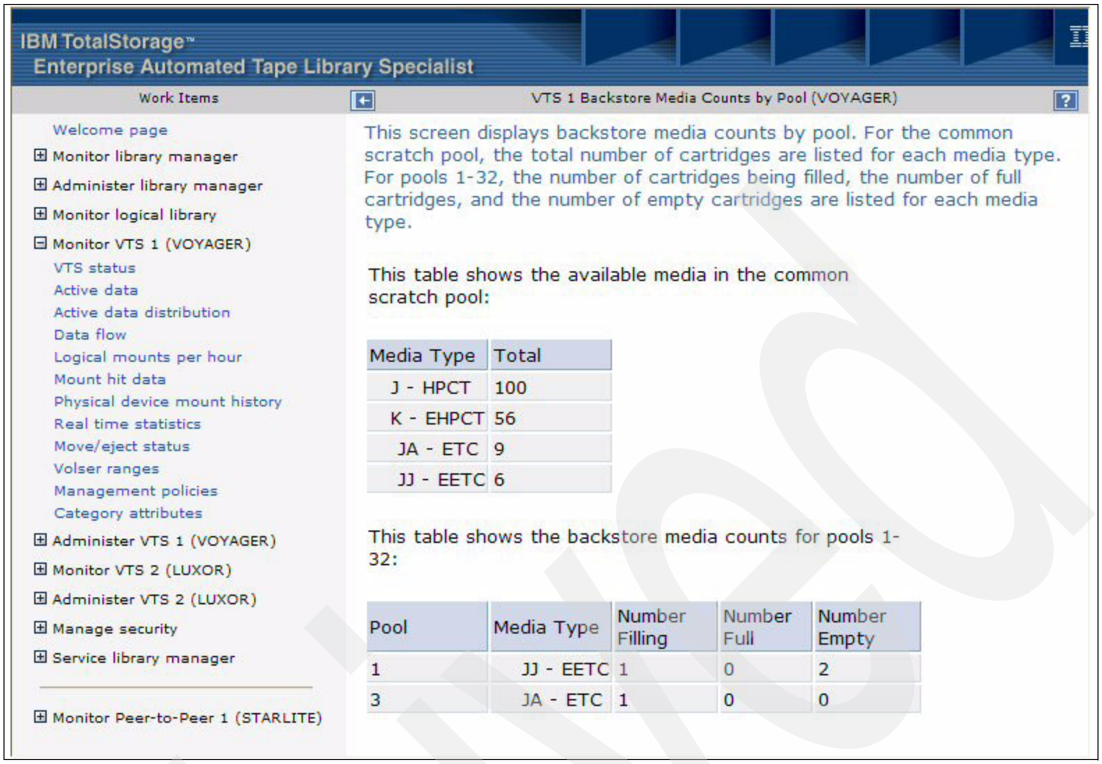

Figure 9-13 3494 Specialist media counts by pool

Loss of drives in a VTS

When the number of physical drives of a distributed VTS gets down to only one usable drive, the VTS will enter the *Insufficient Resources for Mount* state. Copying the virtual volumes to the stacked volumes is no longer possible. To prevent an overflow of the cache, all logical mounts are suspended to that VTS. An intervention required message is generated, indicating that the distributed library has insufficient resources.

The actions taken by the Peer-to-Peer VTS are the same as described in “Out of Empty Stacked Volumes” on page 340.

9.3 Disaster Recovery testing with Standby VTCs

If a Peer-to-Peer VTS is configured with one VTS remote from the other, an additional CX1 frame with up to four Standby VTCs can be placed at the remote site.

Standby VTCs are shown in Figure 9-14.

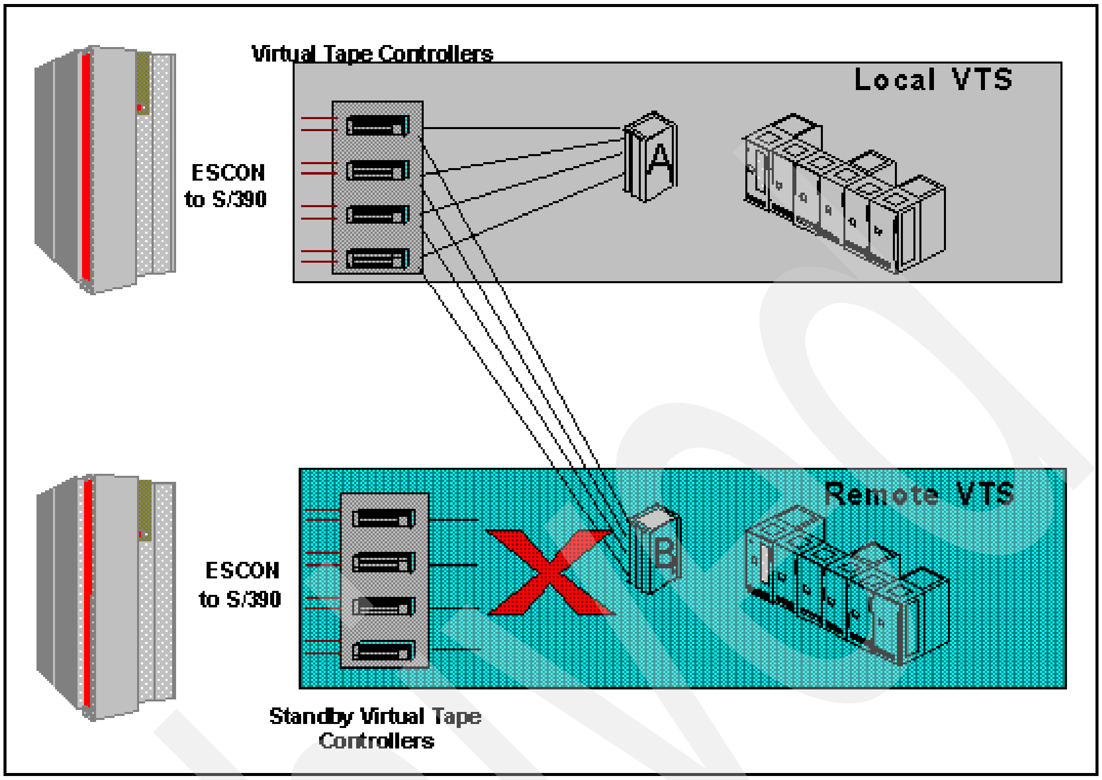

Figure 9-14 Standby VTCs

The purpose of Standby VTCs is to provide a way to access the remote VTS from a host system at the remote site in the event that all the primary VTCs are unavailable.

The Standby VTCs are attached to the remote host channels but not to the distributed VTS, and they have been set up with TCP/IP information to allow the PtP Specialist access. If you also have one or more host systems at the remote site and you have copies of the tape inventory data sets, you can make use of Standby VTCs either in the event of a disaster or to rehearse your disaster recovery plans. You might choose this configuration at a GDPS Standby site.

Four Standby VTCs will support the data bandwidth of the original VTC configuration if it had four VTCs.

> **Note:** The SSR *must* maintain the microcode levels on the Standby VTCs to ensure their availability.

9.3.1 Planning to use Standby VTCs

Before you can use Standby VTCs, you need to:

1. Install the CXn frame and four VTCs at the remote site.
2. Power up the VTCs and customize them with the TCP/IP information necessary to replace the local site VTCs.
3. Lay the additional channel cables that you need.

In each case, we assume that there is at least one z/OS system at the disaster recovery site that you will use to access the data in the disaster recovery site VTS. The systems must meet these requirements:

- At least one z/OS system must have the full support for Peer-to-Peer VTS installed. The others should have full or Toleration support installed.
- You must be able to recover the control information needed to locate and access your tape data. This includes:
  - Tape Management System catalog or control data set, describing where a volume is, its attributes, and how it should be managed.
  - ICF catalogs in which tape data sets are cataloged, describing the volume on which the data resides so the JCL that does not request a specific volume will continue to work.
  - The Tape Configuration Database, comprising the general VOLCAT and, optionally, specific VOLCATs. These contain information about tape volumes and tape libraries.
- Ensure that your storage management and security policies match those at the local site.
  - The recovery site SMS system should include the same tape library definitions and tape Storage Groups that were defined at the local site. The Storage Group ACS routine should assign tape Storage Groups correctly.

    However, if you are planning to be able to run in Read-Only mode, you can ensure that there is no attempt to write data, either by using tape management system facilities (such as DFSMSrmm CHANGEVOLUME volser ACCESS(READ)) or by placing the tape Storage Group in DISNEW status.
  - The RACF® profiles at the recovery site should match those at the local site so that jobs can access the tape data that they need.

You must also ensure that disk data is available for applications to run.

9.3.2 Using Standby VTCs in a disaster

In the event of a disaster at the local site, your IBM System Service Representative (SSR) and an IBM product specialist can bring up the Standby VTCs and allow the remote distributed VTS to be accessed using the Standby VTCs in Read/Write Disconnected mode. The Read/Write Disconnected mode is described in 2.6.7, “Special VTC operational modes” on page 54, and can only be used if the local site is not operational.

You must never use the Read/Write Disconnected mode for the two distributed VTSs separately. Otherwise, it could happen that the same logical volume gets modified on each site's VTS differently (either by getting selected as a scratch mount or modified by an application). Each VTS sets a flag in the volume's token indicating if the volume has been modified in R/W disconnected mode. If the flag is on for a volume in both tokens, the VTS has identified a data conflict and will block access to the volume. Development action is required to recover. If only one of the volume's token has the flag set, that VTS is assumed to have the most valid version of the volume.

The Standby VTCs must be cabled as follows:

1. Disconnect all the channel cables from the four or eight local site VTCs at the adapters of the disaster recovery site VTS. The cables should be labeled, because they must be replaced in exactly the same way when the local site returns.
2. Connect the four Standby VTCs to the adapters of the disaster recovery site VTS, using the upper port on each adapter.

3. Power up and IML the four VTCs and, when prompted, select the **Read/Write Disconnected mode**.
4. Check that all the Standby VTCs are operational and that the subsystem configuration is correct. Each Standby VTC should only be connected to the disaster recovery site VTS.

When this has been done, the tape devices at the disaster recovery site can be varied online to the z/OS host systems at that site, and tape jobs can be run on those systems, reading and writing data.

Eventually, you will want to recover to a full Peer-to-Peer VTS configuration. If the two VTSs have become considerably out of synchronization because of a high volume of write activity, or if a new empty VTS has been installed at the local site, the procedure shown below must be followed carefully, involving IBM product specialists as well as your SSR. The procedure assumes that the VTCs and the local VTS became unavailable at the same time. The most important consideration for the procedure is that the two VTSs must not be allowed to write data at the same time until they are merged back into a Peer-to-Peer configuration.

The procedure to return to an operational Peer-to-Peer VTS is as follows:

1. Make a VTS and four or eight VTCs available at the local site. This may be new equipment or the original equipment after repair or restoration of site power.
2. Vary offline all the devices at the disaster recovery site that are connected to the Standby VTCs.
3. Disconnect the Standby VTCs at the adapters of the disaster recovery site VTS.
4. Reconnect all the cables from the local VTCs to the adapters of the disaster recovery site VTS exactly as they were originally.
5. Disconnect all the cables from the local VTCs at the adapters of the local VTS. This leaves the local site VTCs connected only to the disaster recovery site VTS, which has been operational and will have new data.
6. Power up the local VTS. It has no cables attached at this stage.
7. Power up and IML the local VTCs in Read/Write Disconnected mode. Make sure that only the disaster recovery site VTS is online to the local VTCs.
8. After the IML of the local VTCs completes and the disaster recovery site VTS is online to all local site VTCs, you may vary virtual drives online and use the Peer-to-Peer VTS in Read/Write Disconnected mode.
9. Reconnect all the cables from the local site VTCs to the adapters of the local site VTS.
10. At any one of the local site VTCs, cancel the service preparation state for the local site VTS.
11. The VTCs at the local site now have data available to the host systems while the VTCs begin the task of resynchronizing the local VTS. The time needed for full synchronization of the dual copies to be achieved depends on how many logical volumes were written during the outage.

9.3.3 Rehearsing disaster recovery with Standby VTCs

In the following scenarios, we assume that the PtP VTS is installed at two different sites with the VTS at the local site being defined as the Master VTS.

There are two ways to rehearse the potential of loosing access to all the logical volumes in one side of the PtP VTS:

- By only reading the logical volumes in one PtP VTS
- By reading and writing logical volumes on one side of the PtP VTS

In either case, the environment to you use to rehearse for a potential disaster must be an isolated one. This would include a host system with its own DASD to be IPL'ed in a simulated production mode. It should have a copy of all the production data required to IPL the system.

Read-only method

When you want to test your disaster recovery capabilities at the remote disaster recovery site, you can do so without interrupting your production work. You will have full read and write access to your active PtP VTS at the local site while testing is done with a test system with the other half of the PtP. The Standby VTCs would be started in Read-Only mode.

> **Attention:** During any testing you must ensure that all the production VTCs are disconnected from the test environment or the testing could impact production work.

If your PtP VTS is part of a GDPS complex, testing by separating the PtP is not allowed.

DR Test procedure

Follow these steps to run your disaster recovery test:

1. Start VTS service preparation for the VTS at the remote disaster recovery site. This may take some time depending on the mode of preparation selected — normal, expedite, immediate, or forced. See 2.7.1, “Service preparation” on page 55, for details. The Peer-to-Peer VTS continues to run production read/write tape work using the local site VTS and VTCs.
2. After the remote site VTS has completed service preparation, connect Standby VTCs to VTS and test system. Depending on the configuration of you system you would follow one of these two methods:
   – Disconnect the four or eight cables from the local site VTCs at the adapters of the remote disaster recovery site VTS. The cables should be labeled, because they must be replaced in exactly the same way when the disaster test is finished. Connect the four Standby VTCs to the adapters of the remote disaster recovery site VTS, using the upper port on each adapter.
   – If the remote VTS connections are running through an ESCON Director or FICON Switch, you would block all the ports to the remote VTS and redirect the links to the Standby VTS.
3. Power up and IML the four Standby VTCs. The VTCs will not complete the IML because they cannot communicate with two VTSs or another VTC that has completed its IPL. There are two ways to select the mode of the VTC:
   a. The SSR can log on to one VTC and access the service panels, the Read-Only mode can then be selected. Once one VTC is set in this mode all the other Standby VTCs will complete there IML and enter Read-Only mode.
   b. With VTS microcode Release 7.4 the selection of the VTC mode of operation can be selected to from the VTS Peer-to-Peer Specialist panels. By connecting to one of the powered up Standby VTCs the **Read-Only** mode can be selected as in Figure 9-3 on page 327. Do not select Read/Write Disconnected mode.

4. Check that all the Standby VTCs are operational and that the subsystem configuration is correct. Each Standby VTC should only be connected to the remote disaster recovery site VTS.

When this has been done, the distributed library and the tape devices at the remote disaster recovery site can be varied online to the z/OS host systems at that site and read-only tape jobs can be run on those systems.

If you attempt to change a volume's status or contents, you will see the following error messages:

- Scratch Mount

  CBR4126I Library *Libraryname* drive is in read only mode.

- Write commands to a logical volume which had been mounted before Write Protect mode was entered:

  IOS000I devn,chp,err,cmd,stat,,dcbctfd,ser,mbe,eod, jobname,sens,WRITE PROTECTED

- LM LP command

  CBR1086I LIBRARY LMPOLICY results for volume *volser.*
  FUNCTION NOT COMPATIBLE WITH THE LIBRARY.

- Insert logical volume

  CBR3606I Entry of volume *volser* into library *libraryname* failed. Unable to set the volume category.
  CBR3726I Function incompatible error code 10 from library *libraryname* for volume *volser*.

Returning to normal after test

When the disaster recovery test has been completed, you will want to return to normal operation. The following procedure will make sure that deferred copies from the local site will be resynchronized:

1. Vary offline from z/OS all the devices at the remote disaster recovery site that are connected to the Standby VTCs.
2. Power off the Standby VTCs.
3. Disconnect the Standby VTCs at the adapters of the remote disaster recovery site VTS.
4. Reconnect the cables or re-establish links from the local VTCs exactly as they were before the disaster recovery test.
5. At any one of the local site VTCs, cancel the service preparation state for the remote site VTS.
6. The VTCs at the local site will bring the remote VTS online and beginning to copy all the deferred copies of the logical volumes created during the DRP test.

Using Read/Write Disconnected Mode

As described in 9.1.3, "Disaster recovery test enhancements" on page 324, is it not permitted to run the two VTSs of a PtP VTS independently of each other, writing to both VTSs independently. the same logical volume gets modified on each site's VTS differently (either by getting selected as a scratch mount or modified by an application). Each VTS sets a flag in the volume's token indicating if the volume has been modified in R/W disconnected mode. If the flag is on for a volume in both tokens, the VTS has identified a data conflict and will block access to the volume (era 91 would be indicated to a host that attempts to use the volume) Development action is required to recover. If only one of the volume's token has the flag set, that VTS is assumed to have the most valid version of the volume.

The only way to perform DR testing in Read/Write Disconnected mode would be to disconnect the remote VTS from the production VTCs, connect it to a test system through standby VTCs, and define it to run in Read/Write Disconnected mode. Even when only writing to a separate range of logical volumes, this way of testing is an unsupported scenario that can result in data loss and inconsistencies requiring engineering intervention to be resolved.

If you have the requirement to write into the PtP VTS during DR testing, consider using Read-only mode and divert tape writes to disk using modified ACS routines. If using disk is not an option, we recommend that you discuss with your SSR and Product Field Engineering (PFE) alternatives to achieve this goal.

Peer-to-Peer Specialist

The PtP Specialist can be used on the Standby VTCs as well as the original VTCs. However, the TCP/IP information for the Standby VTCs must be different from the original VTCs because duplicate IP addresses are not allowed. Therefore the Standby VTCs must be set up with unique IP addresses and host names.

The Library Manager(s) in the 3494 or 3953 run a Web server for the ETL Specialist which links to the PtP Specialist on Web servers in each VTC. The links are held in a Web server information table as VTC Controller Identifications. If unchanged, these links will point to the original VTCs. Recommendations for handling these links are as follows:

- If the Standby VTCs are being used to rehearse disaster recovery, it is recommended that the VTC Controller Identification values be unchanged and use browser bookmarks to ensure that the Standby VTCs are accessed. If nothing is done, the Library Manager links will point to the VTCs at the original, local site.
- In the event of a real disaster, the SSR should be asked to change the VTC Controller Identification values so that the links work correctly. If nothing is done, the links point to addresses at the site suffering the disaster, which is presumed unavailable. This means that the change must be reversed when returning to the original site after the disaster.

The TCP/IP definitions for the Standby VTCs should be set up as follows:

- The Standby VTCs must be assigned host names and IP addresses. The addresses should be documented in recovery documentation so that, in the event of needing to use the Standby VTCs, users know the correct addresses to use.
- The Peer-to-Peer Specialist worksheet from the *IBM 3494 Tape Library Introduction and Planning Guide*, GA32-0448, should be completed and given to the SSR.
- If a name server is being used, the new names and addresses must be entered.

9.4 Rehearsing for a disaster without Standby VTCs

With microcode Release 7.4 there is an option to allow you to test the PtP VTS resources without interfering with any production workload and without separating the PtP VTS. This is a more limited testing scenario but is very easy to configure and very easy to enter and exit this mode.

9.4.1 Write Protect mode

When you want to test your disaster recovery capabilities at the remote disaster recovery site, you can do so without interrupting your production work. You will have full read and write access to your PtP VTS with the production systems and any copy volumes will be made as usual while testing is done with a test system. There is no need to put one side of the PtP VTS in Service Prep mode. This method involves taking devices of one or more VTCs offline to the production environment and connect the selected VTCs to the test system or image. Then put the VTC into Write Protect mode and vary the devices online to the test system.

With this environment created, your test system will have access to both sides of the Peer-to-Peer, limited only by the number of devices defined to the VTCs you put into Write Protect mode. With the ease of creating this environment, this would be useful in testing the host software definitions and application setup in preparation for a D/R test where the PtP is separated.

> **Note:** The test system must be defined and configured to support the PtP VTS as is currently is in your production host. During the test you will only be using one or two of the VTCs for your testing.

9.4.2 Test steps

You would follow these steps to run your disaster recovery test:

1. Select a VTC that is connected to you test host.
2. Vary this VTC's devices offline to the production hosts.
3. Vary this VTC online to your test system.
4. Connect to the VTS PtP Specialist of the selected VTC and change this VTC's operations mode to enable Write Protect mode. See Figure 9-1 on page 326.

At this point the devices defined to this selected VTC are available for testing. You will have access to all logical volumes within the PtP VTS for read only jobs. Jobs will fail if they attempt to do any of the following:

- Any scratch mounts
- A host attempt to modify a logical volume's contents that was mounted while the mode was active
- A host attempt to change the attributes for the volume (category, constructs)

Depending on the attempt, you may see one of the following error messages:

- Scratch Mount

  ```
  CBR4126I LIBRARY Libraryname DRIVE IS IN READ ONLY MODE.
  ```

- Write commands to a logical volume which had been mounted before Write Protect mode was entered:

  ```
  IOS000I devn,chp,err,cmd,stat,,dcbctfd,ser,mbe,eod, jobname,sens,WRITE PROTECTED
  ```

- LM LP command

  ```
  CBR1086I LIBRARY LMPOLICY RESULTS FOR VOLUME volser.
           FUNCTION NOT COMPATIBLE WITH THE LIBRARY.
  ```

- Insert logical volume

  ```
  CBR3606I ENTRY OF VOLUME volser INTO LIBRARY libraryname FAILED. UNABLE TO SET THE
  VOLUME CATEGORY.
  CBR3726I FUNCTION INCOMPATIBLE ERROR CODE 10 FROM LIBRARY libraryname FOR VOLUME volser.
  ```

During your testing using the above method, your test jobs could be accessing either side of the PtP. If you wish to force the jobs to use one side of the PtP, the remote side for instance, you could block the path this VTC uses to access the local VTS. The following steps will force the VTC to only communicate with the remote VTS when you use it in Write Protect mode. This would enable you to verify that the remote side of the PtP VTS is ready for a real DR test.

Perform the following steps:

1. Ensure the Master VTS is at the remote site.

   A less disruptive method, avoiding Master VTS Switchover, is to set the link speed values on the VTC to indicate that the link to the remote VTS is 1 and the local VTS is 5. This will force all accesses of the data to the remote VTS (assuming that it has a valid copy of the data).
2. Take the VTC devices offline.
3. Enable Write Protect mode on the selected VTC.
4. Block the port this VTC uses to communicate to the local VTS, local to the production environment. Or if this connection does not go through an ESCON Director or FICON Switch, you would disconnect the fibre at the back of the VTS or VTC for this specific link
5. Bring the selected VTC's devices online to the D/R test image or host.

Now all tape requests have to be fulfilled by the remote VTS. The same conditions apply as they did for the first environment. No writing or modification to logical volumes is allowed through this VTC.

Another way to achieve that all requests are fulfilled by the remote VTS is to set the link speed parameters so that the link between the VTC and the VTS, the test data needs to come from, is set to **1,** and the link from that VTC to the production site VTS is **5**. This will force all read I/O to the test VTS even if a recall is needed.

9.4.3 Returning to normal after test

After all testing is complete and you are ready to put the PtP VTS back to normal:

1. Vary the VTC's devices offline to the test system.
2. If you have elected to disconnect or block any ports to force all requests to go to the remote VTS, then you must return these connections back to normal.
3. Perform master switchover, if required.
4. Connect to the VTS PtP Specialist of the VTC and change this VTC's operations mode by selecting **disable Write Protect mode**. See Figure 9-2.
5. Vary this VTC's devices online to the production hosts.
6. Verify connectivity.

9.5 Geographically Dispersed Parallel Sysplex

With the continued decrease in Information Technology costs as well as greater emphasis on application availability and disaster recovery, more and more enterprises are adopting a two-site strategy. The IBM Series 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 subsystem(s), automates Parallel Sysplex operational tasks, and automates failure recovery from a single point of control, thereby improving application availability. GDPS supports all transaction managers (for example, CICS® and IMS) and database managers (for example, DB2®, IMS, and VSAM).

9.5.1 What a GDPS is

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.

The GDPS manages cross-site data consistency through a combination of storage subsystem, sysplex, and environmental triggers. GDPS provides the cross-site automation that works in combination with the DASD error recovery procedures exploiting new storage subsystem functions (freeze/run) to manage cross-site data consistency.

The physical topology of a GDPS consists of a base or Parallel Sysplex cluster spread across two sites (known as site A and site B) with one or more Series z systems at each site, separated by up to 40 kilometers (km) when using a PPRC form. With XRC, the distances between sites can be increased to over 100 km depending on the channel extenders in use.

GDPS consists of production systems, standby systems, and controlling systems (see Figure 9-15).

- The production systems (1A,1B, and 1C in our example) execute the mission-critical workload.
- The Standby systems normally do not run any production work, and instead provide processing resources when a production system or a site is unavailable. A Standby system may be an active member of the Parallel Sysplex cluster that will acquire processing resources from expendable systems executing in logical partitions on the same processor, or an inactive member of the Parallel Sysplex cluster.

  When an additional processing resource is required, an active Standby system will expand by deactivating one or more logical partitions that are executing expendable work and configure their storage resources online (processor MIPS will also be absorbed) and/or an inactive Standby system will be IPLed by the system resetting one or more systems executing expendable work and IPLing the inactive system in its place.

  There must be sufficient processing resource capacity (such as processor power, main and expanded storage, and channel paths) available that can quickly be brought online to restart a system's or site's critical workload (typically by terminating one or more systems executing expendable—non-critical—work and acquiring its processing resource).
- The Controlling system, also known as the K system, is a system that is in the same sysplex as the systems that are being managed by GDPS. However, its primary role is to be able to control the recovery following an outage. For this reason, the controlling system has all its DASD in site 2, and it shares only the control data sets with other systems in the sysplex; this ensures that the controlling system will not be affected by any events that may impact the availability of the managed systems. By convention, all GDPS functions are initiated and coordinated by one controlling system. The first active system that is specified in the controlling system list will be the primary controlling system with the others, if any, being secondary controlling systems.

  Note: All GDPS operations should be carried out from this system.

All critical data resides on storage subsystem(s) in site A (the primary copy of data) and is mirrored to site B (the secondary copy of data) via PPRC synchronous remote copy or XRC asynchronous remote copy. This includes system infrastructure (for example, JES2 checkpoint and spool) and software subsystem infrastructure (for example, program libraries, PROCLIBs, control files, and DBMS data).

GDPS uses Parallel Sysplex cluster facilities to communicate between systems. Each GDPS system joins the GDPS Parallel Sysplex cluster group. When a system joins a Parallel Sysplex cluster, GDPS automation will transfer GDPS policy and all GDPS status to the joining system.

All GDPS systems, except for any inactive Standby systems, are running GDPS automation based upon Tivoli® NetView® for z/OS and System Automation for z/OS. Each system will monitor the Parallel Sysplex cluster, coupling facilities, and storage subsystems, and maintain GDPS status. This GDPS automation can coexist with an enterprise's existing automation product.

The data center must have a survival strategy (the operations function must survive any disaster to maintain processing). Ideally, the data center is located at a third site, and the back-up operations center is located at the same site as the controlling system to help maximize survivability.

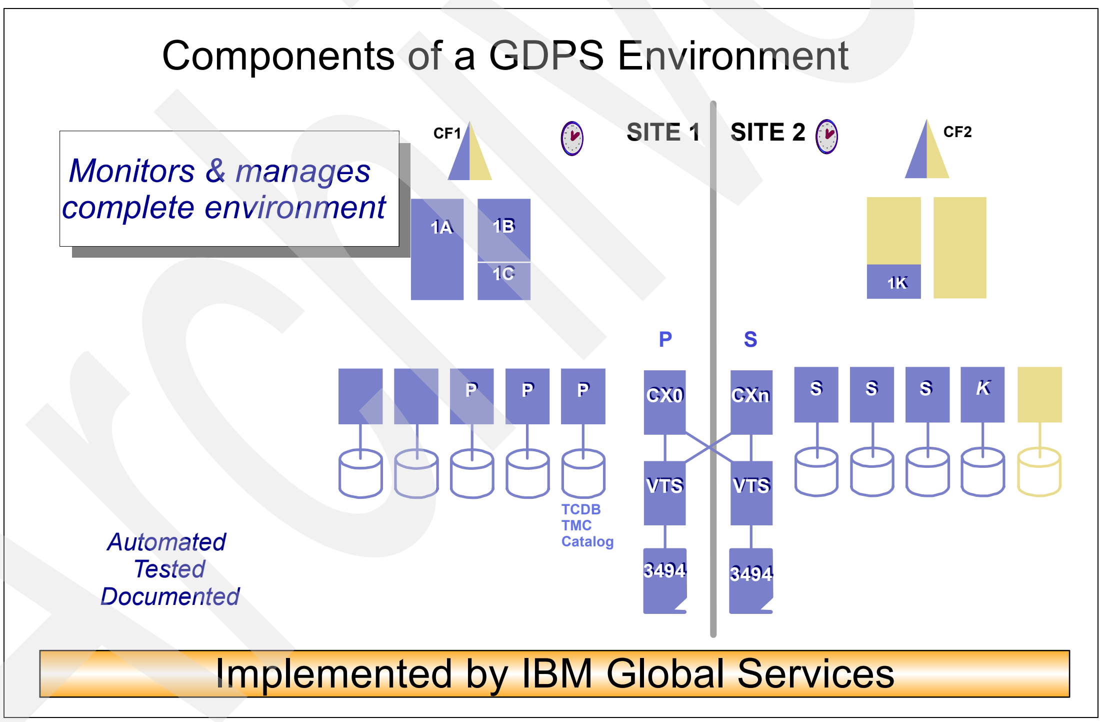

Figure 9-15 Components of a GDPS environment

9.5.2 GDPS functions

GDPS provides the following functions:

- Remote Copy Management Facility (RCMF): automates management of the remote copy infrastructure.
- Planned reconfiguration support: automates operational tasks from one single point of control.

- Unplanned reconfiguration support: 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, since 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 and/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):

- **Freeze and go:** GDPS will freeze the secondary copy of data when remote copy processing suspends and the critical workload will continue to execute, making updates to the primary copy of data; however, these updates will not be on the secondary DASD if there is a subsequent site A failure in which the primary copy of data is damaged or destroyed. This is the recommended option for those enterprises that can tolerate limited data loss or have established processes to recreate the data.

- **Freeze and stop unconditional:** GDPS will freeze the secondary copy of data when remote copy processing suspends and will quiesce the production systems resulting in the critical workload being stopped and thereby preventing any data loss. This option may cause the production systems executing the critical workload to be quiesced for transient events that interrupt PPRC processing, thereby adversely impacting application availability.
- **Freeze and stop conditional:** GDPS inspects the reason for the suspension of remote copy processing. If the suspension is caused by the storage subsystems that contain the secondary copy of data, processing is the same as for freeze and go; otherwise, processing is the same as for freeze and stop unconditional. This function is only valid in PPRC environments. It is the recommended option for those enterprises that cannot tolerate any data loss but require maximum availability.

If a site fails, the secondary copy of data will be frozen and the workload resumed or quiesced based upon policy, the failed system(s) will be removed from the Parallel Sysplex cluster, the primary copy of data will be switched, processing resources will be acquired at the surviving site (active Standby systems will expand and/or inactive Standby systems will be IPLed), the network will be switched if necessary, and the critical workload restarted.

9.5.3 Peer-to-Peer VTS within a GDPS configuration

A PtP VTS can participate in a GDPS configuration (Figure 9-16) and represents a highly available tape solution. It has to be operated in immediate copy mode with preference to the primary site. This mode within a GDPS configuration is called the primary operation mode (see "New GDPS terminology" on page 355). The primary VTS performs all host I/Os. The VTC-defined mode should be selected as the Peer-to-Peer Copy mode. This will prevent the Management Class defined copy mode from overriding the host or VTC-defined copy mode. All VTS-relevant control files like the TMC, TCDB, and tape volume catalogs have to be mirrored via PPRC.

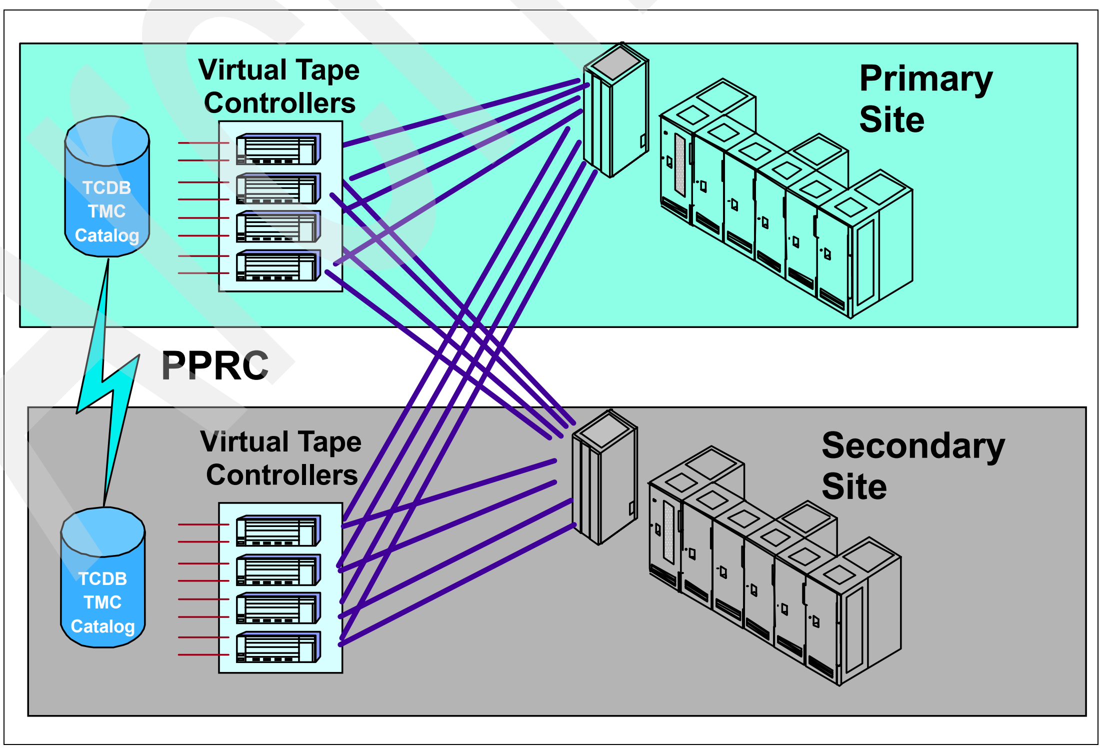

Figure 9-16 Peer-to-Peer VTS GDPS configuration

VTS terminology

The common Peer-to Peer VTS terminology as of today (see Figure 9-17) is:

- Composite library: Logical library image presented to the host
- Distributed library: Physical Enterprise Library IBM 3494 or TS3500/3953 with its associated Virtual Tape Server
- Master VTS: One VTS that serializes access to logical volumes
- User interface library: one of the distributed libraries, allows for insertion of logical volumes and is “master” of date and time
- I/O VTS: The VTS that processes the host I/O commands for a certain logical volume
- Copy mode
  - Immediate
  - Deferred
  - No copy
- Operating mode
  - Balanced
  - Preferred
  - Primary (this mode has been added to support GDPS)

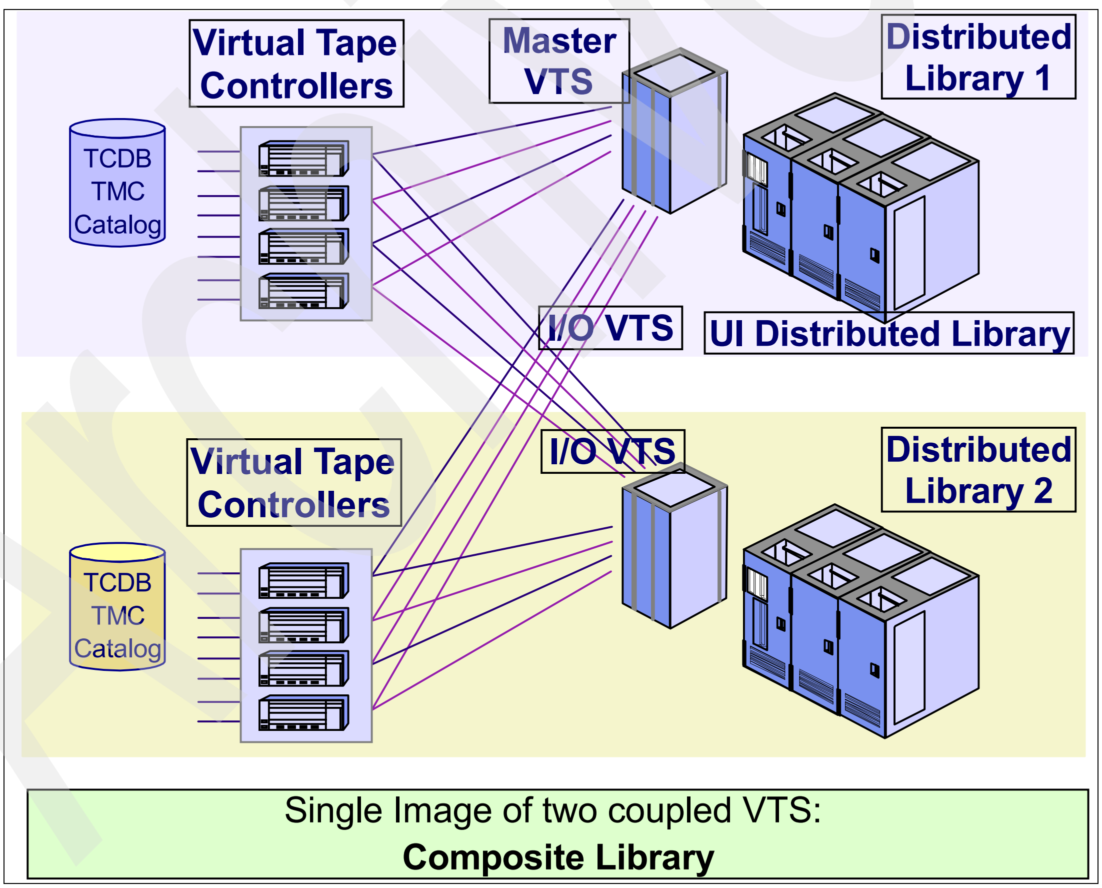

Figure 9-17 The VTS terminology

New GDPS terminology

The *primary VTS* is running in the VTS primary operation mode. This mode introduces some differences as compared to the “preferred” operation mode. For example, a primary VTS receives all host tape mounts and handles all read and writes while a VTS PtP in preferred mode allows host-VTS read I/Os to the secondary VTS. Let us assume a job with a modification request (`DISP=MOD`) onto a logical volume. The logical volume is stored in the cache of the secondary VTS but not in the primary. In a GDPS environment with a VTS PtP running in primary mode, the I/O accesses the primary VTS, which requests to copy the logical volume from the cache of the secondary VTS to the primary VTS before it is provided to the host for modification. In contrast, a PtP VTS in preferred mode would allow the host to access the secondary I/O VTS and provide the logical volume directly to the host for modification.

The *secondary VTS* is the other I/O VTS, which receives all logical tape copy volumes.

Duplexed/synchronized is the status of the composite library when all logical copies are completed.

Suspended is used to establish a data-consistent situation as soon as possible. Dual copy operations in progress will complete. New or queued operations will be held in the VTS copy queue. As soon as tape dual copy suspends, GDPS suspends all dual copy operations of the same freeze group. Suspend of tape copy might freeze disk PPRC volumes and vice versa, depending on the established policies.

The *freeze group* is a group of one or multiple composite libraries, which usually includes mirrored disk volumes. A VTS PtP cannot belong to multiple freeze groups.

PtP VTS enhancements for GDPS and DFSMS

A full description of these enhancements and new elements can be found in APARs OW49900 and OW49901.

The enhancements to the DISPLAY SMS command and to messages are as follows:

- Presents up to 512 devices
- Additional status lines
  - Volume is cache resident
  - Valid copy in each distributed library
  - Corrupted token volume category
- Added error messages
  - Out of Empty Stacked Volumes
  - Not enough physical drives

A new Application Program Interface (API) allows GDPS the change of operation modes:

- Enable/disable copy operations
- Query device and volume status

And it allows the caller to determine:

- Logical Volume in cache?
- Valid copy in each distributed library?
- List corrupted token volume category

GDPS enhancements for Peer-to-Peer VTS

The GDPS control code has been enhanced by the following items:

- New panels:
  - VTS PtP configuration
  - Line mode and global GDPS commands like query, suspend, switch, and resynch
- Monitor status and error conditions:
  - New installation policy options: VTS will/will not trigger freeze
  - Coordinated disk/tape freeze
- Planned and unplanned site switch:
  - Disk/tape failover in one, coordinated, single GDPS script
  - Query volumes for volume copies “in flight”
- New tape scripting keywords:
  - Swap the roles of the two distributed VTSs
  - Switch the primary VTS
  - Suspend dual copy operation

Single site Parallel Sysplex with PtP VTS

In this example (Figure 9-18) the workload is running within a single site Parallel Sysplex. Under normal conditions, all loads will run at site 1 and all primary copies will be at Site 1. Site 2 consists of a monitoring partition — the 1K partition, secondary copies of Disk and VTS logical volumes, and some spare capacity of CF, CPU, and disk.

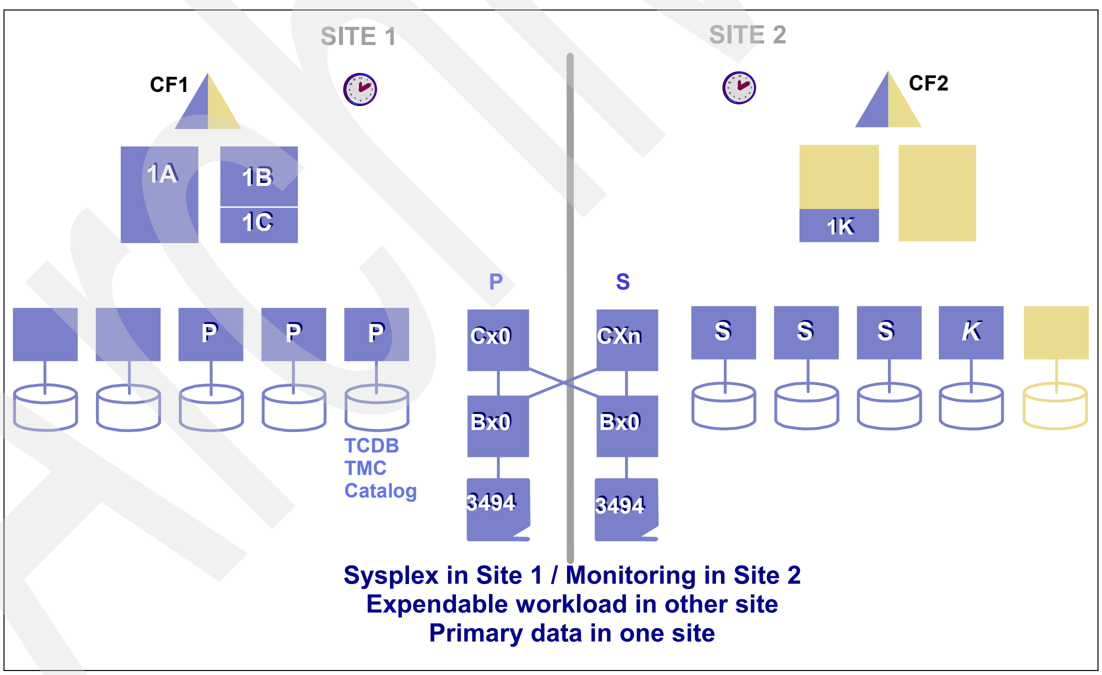

Figure 9-18 Single site Parallel Sysplex with PtP VTS

Multi-site Parallel Sysplex with PtP VTS

In this example (Figure 9-19) the workload is running within a multi-site Parallel Sysplex. The distance between these two sites is 20 km. Each site should hold some spare capacity to fully take over the load.

Even in a multi-site configuration, all primary data of one freeze group has to be in one site.

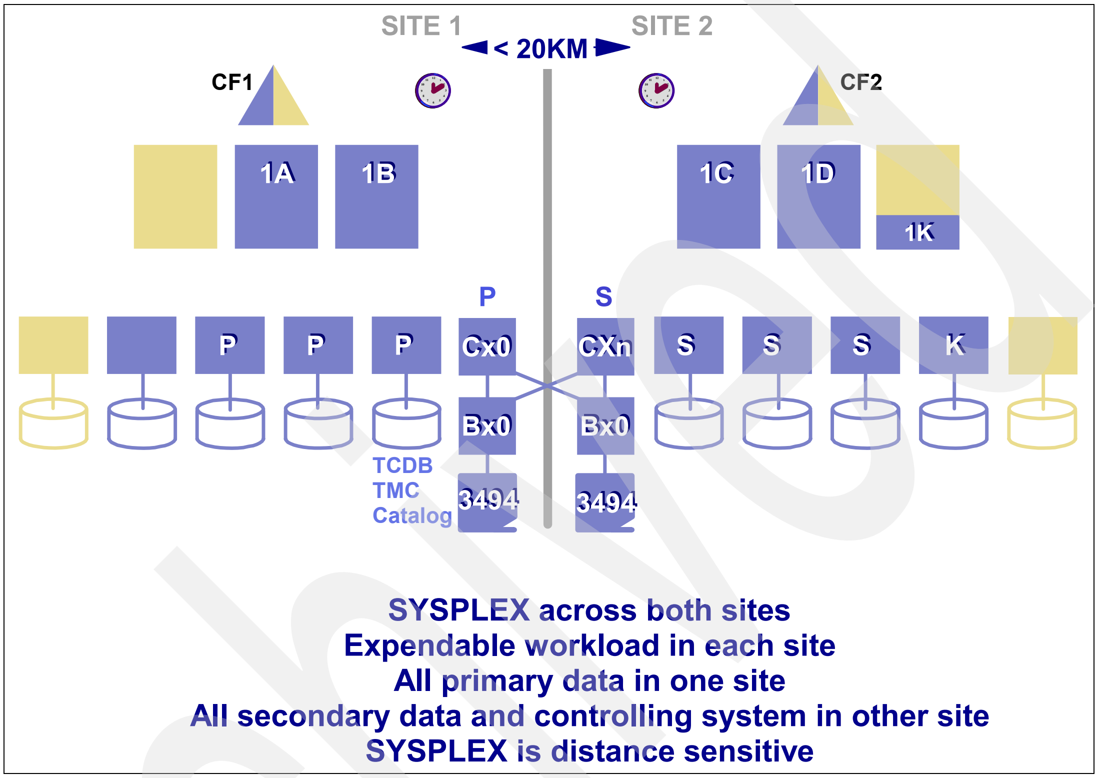

Figure 9-19 Multi site Parallel Sysplex with PtP VTS

Local configuration with Standby VTCs

There are two ways of implementing the VTCs within a Peer-to-Peer VTS configuration:

- All local with stand-by VTCs
- Fully activated (using VTCs in both sites)

We will discuss the advantages and disadvantages of both implementations. Most comments are valid for GDPS and non-GDPS managed VTS PtP configurations.

Under normal GDPS production conditions, all the data transfers are to the locally attached VTCs. The so-called Standby VTCs need to be installed, defined, active, and in concurrent access by the GDPS monitor partition to make sure that the monitor partition can initialize the switch in case of a failover. As this configuration requires less communication bandwidth between the two sites, it might be the preferred choice for a single workload Parallel Sysplex (see Figure 9-20).

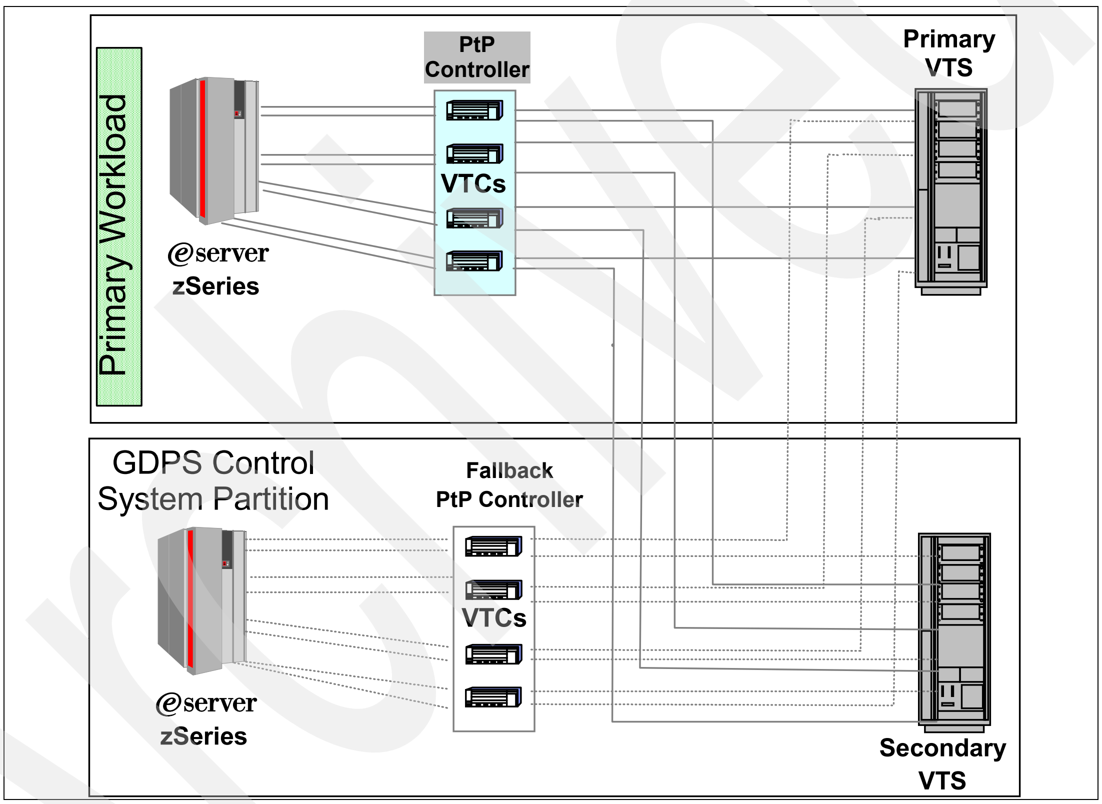

Figure 9-20 Local configuration with stand-by AX0s

Fully activated and used VTCs in both sites

This configuration (Figure 9-21) might be the preferred choice for a two-site parallel Sysplex. It can also be used for a single site workload where costs for communication links between sites are acceptable.

Even in this scenario all data from both sites needs to store its primary data in one site. In case of a failover, the customer has to expect a reduction in performance, available device addresses, and channels by approximately 50 percent.

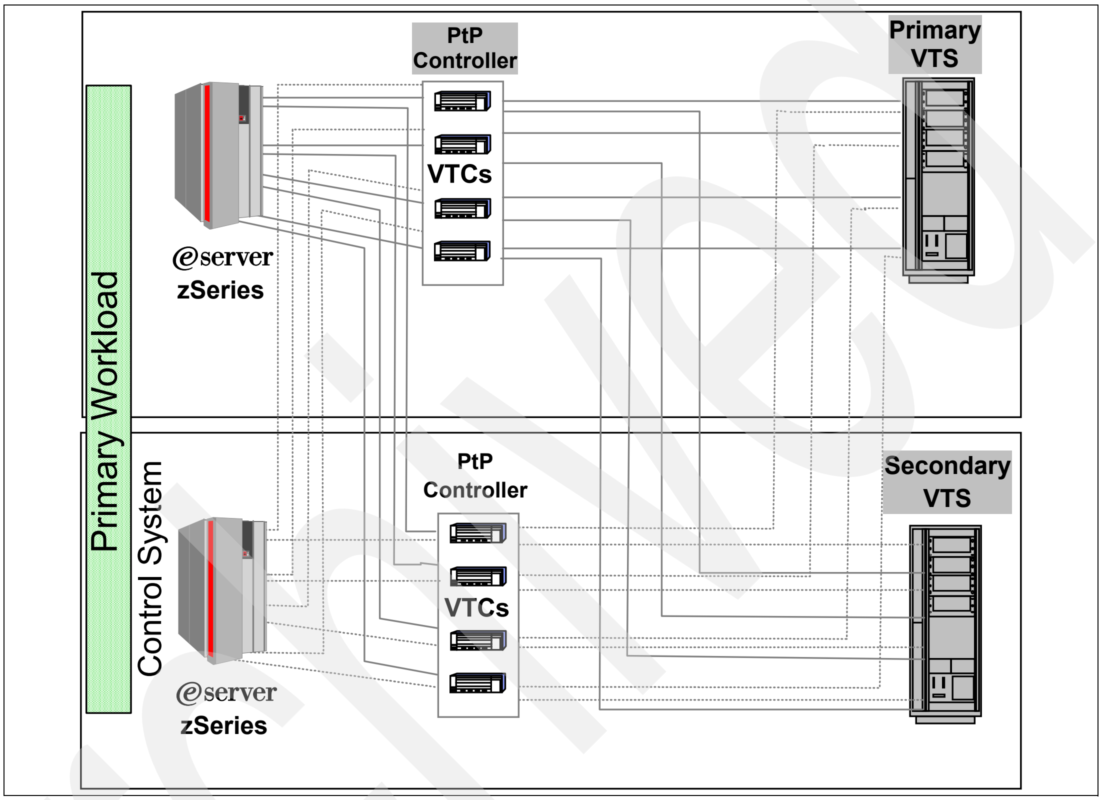

Figure 9-21 Fully activated/used AX0s in both sites

Differences between DASD and tape

There exists a tape data consistency exposure between PPRC-mirrored tape control data sets because tape dual copy creation is asynchronous at rewind/unload time, and for HSM ML2 workload and the like applications that are taking its own synchpoints.

As shown in Figure 9-22, there are major differences between the mirroring of a DASD I/O and those of a VTS PtP. The consistency discussion becomes even more complex by the fact that some of the VTS PtP specific control information (like the TCDB and the TMS catalog) have to be stored on mirrored disk and managed as DASD I/O's.

Data consistency exposures:

- In case of unplanned failures where the completion of the logical volume copy might not have been verified.
- For HSM migration level 2 workload, because HSM takes its own synchpoints and stores them on its control data sets. We do not recommend storage of HSM migration level 2 data, with critical service levels onto a VTS PtP under GDPS.

Characteristic	Enterprise Storage Server	Virtual Tape Server
Volume to device relationship?	Permanently assigned to device address	Assigned to a device at mount time
Where are write I/Os directed?	Primary disk subsystem	Primary VTS
When is copy made?	Synchronous to I/O	At rewind/unload time in immediate copy mode
Dependent writes sequence?	Dependent writes maintained	Not applicable to tape CDS resident on disk
"Freeze" implications?	Active I/O failed & re-driven transparently	Active copy operation completes and subsequent operations suspended

Figure 9-22 DASD and TAPE differences

Recovery considerations

In the event of an unplanned outage where suspend could not successfully complete, action has to be taken to clean up those logical volume copies and primary copies that were "on the fly." These kinds of listed actions are not new to GDPS and are valid for a VTS PtP with or without GDPS. However, GDPS helps to improve the process by taking advantage of the new DISPLAY SMS commands. GDPS provides a list of inconsistent tape volumes as a result of its queries to the VTS PtP.

- Inconsistent Tape CDSs, resident on PPRC disk, must be cleaned up before re-execution of batch jobs.
  - TCDB: Set volume back to scratch.
  - TMS: Delete all references using TMS utilities.
  - Tape application CDS.
  - Recovery actions depend on application, disposition mode (new, mod), and type of failure.
- GDPS provides a list of inconsistent tape volumes.
- HSM ML2 data is exposed to data loss in the case of a VTS primary failure.
  - HSM takes its own synchpoint.
  - Deletes disk resident data prior to rewind/unload.
  - GDPS or PtP do not make any difference.

General GDPS requirements

A GDPS environment requires different multiple components in place before a VTS PtP can be implemented. Moving from a non-GDPS managed VTS PtP to a GDPS-managed VTS PtP does not require additional components the customers already operates a complete GDPS PPRC or XRC environment.

- Parallel SYSPLEX across two sites for zero data loss requirements (PPRC)
- Replicate hardware across sites for redundancy
  - Processor(s), coupling facility, and sysplex timer
  - HMC automation infrastructure and associated processor support
  - Disk subsystems supporting PPRC level 2 (freeze command)
  - Peer-to-Peer VTS or tape
- z/OS Version 1 Release 0 or OS/390 Version 2 Release 6 or higher
- System Automation for OS/390 Version 1.3 or higher

- Tivoli NetView for OS/390 Version 1.2 or higher, or NetView 3.1
- Data required for restart must be disk resident and mirrored

> **Note:** A VTS PtP with or without GDPS requires IBM Systems Assurance.

GDPS installation support

The GDPS solution is delivered as a service offering by IBM Global Services. The services provided for GDPS may vary from country to country. Consult your local IBM business contact for information on GDPS services available.

9.5.4 GDPS DR testing consideration

A PtP VTS under GDPS control is required to run in Primary mode and in Immediate copy mode. Primary mode is more strict than preferred mode:

- If the PtP VTS is running in preferred mode and the preferred VTS fails, then a mount will be executed on the other VTS. A specific mount is also processed on the remaining VTS, if it has a current copy of the logical volume.
- If the PtP VTS is running in primary mode, all mounts will only be executed on the primary VTS. If the primary VTS becomes unavailable, these mounts will fail. Specific mount requests will also be executed on the Primary VTS, even if the other VTS has a copy of the logical volume in cache.

The primary requirement for GDPS was to insure that the data is created and accessed on the local VTS, therefore Primary mode was created so that this would happen. Also, with Primary mode, if the specified VTS becomes unavailable, then the mount request is failed.

During normal operation, you should not change the copy mode and the I/O VTS selection through Management Class specifications, when running under GDPS. If you use a Management Class that defines the No Copy Option and select the VTS, which is not the Primary VTS, as the I/O VTS on a scratch mount request, you can write a logical volume to the non-primary VTS without creating a secondary copy on the Primary VTS.

For a specific mount request, Management Class definitions are not honored. For specific mounts the VTC must select the VTS that is specified by the Primary mode. Since the data is not on the Primary VTS, it will be copied from VTS1 even though the Management Class specified No Copy. A host informational message is generated in this case:

```
VT0103 WARNING: PRIMARY VTS VOLUME xxxxxx COPIED FROM SECONDARY VTS TO PRIMARY VTS
DESPITE STG MGMNT CLASS SPECIFICATION TO INHIBIT COPY
```

> **Restriction:** A DR test scenario as shown in Figure 9-23 is therefore ***not*** possible, when running under GDPS.

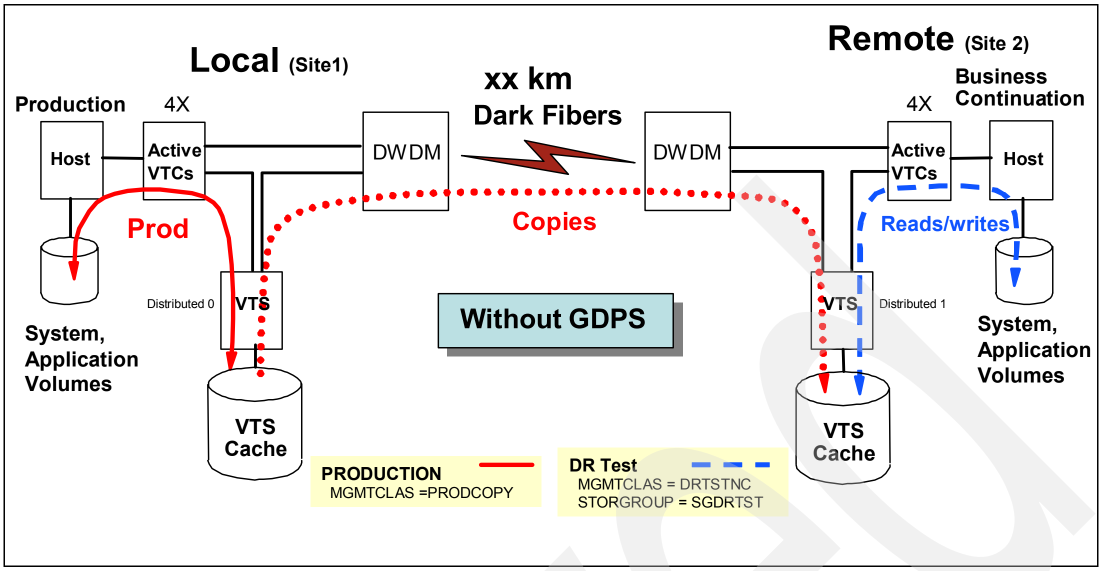

Figure 9-23 Possible DR test scenario without GDPS

In the scenario shown in Figure 9-23, the local site (Site 1) is the production site, and the Remote site is for DR purposes only. All data is written to VTS0 on the left side during normal operation, and all copies are written the VTS1on the right side. This is achieved through Management Class definitions.

During the DR test, production should continue as before. The DR data is written to VTS1 only and read from there as well. The Tape Management System is used to protect overwriting of existing production data. In addition, for the PtP VTS the following actions are taken:

- Define new Storage Group SGDRTST on Distributed library 1.
- Define new Management Class DRTSTNC which selects VTS1 as I/O VTS and specifies No Copy. This will ensure that scratch mounts are failed, if VTS 1 becomes unavailable and are not directed to VTS0.
- To direct specific mounts to VTS1:
  - Set the VTCs at the remote site to prefer VTS1.
  - Set the link speed of the VTCs at the remote site so that the links to Site 1 are the slowest and to Site 2 are is the fastest. The link settings must be set so that the link from the VTC to the desired VTS indicates that it is much faster than the link to the non-desired VTS. Link settings are from 1 to 5, where 1 is the fastest and 5 is the slowest. By setting the link to the desired VTS to 1 and the link to the non-desired VTS to 5, even if the data is in cache on the non-desired VTS, the VTC will pick the desired VTS and recall the data.

  With these two settings active, specific mounts will only go to VTS0, if VTS1 is unavailable, or if VTS1 does not have a valid copy of the logical volume.
- Set up the TMS to not allow modifications of existing logical volumes.
- Define the new constructs on the host and change the ACS routines to select those constructs for DR testing. Activate the new configuration.
- Perform the DR test.
- Delete logical volumes created during the test (Storage Group SGDRTST).

A

Availability configurations

In this appendix, we describe how you can use a Peer-to-Peer VTS 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 Peer-to-Peer VTS, 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 based on the level of data loss and service loss. The relationship between disaster recovery and continuous availability is also discussed. For an in-depth discussion on these topics, see the *Continuous Availability Systems Design Guide*, SG24-2085, and the *Continuous Availability S/390 Technology Guide*, SG24-2086. We hope that these discussions will help you to decide what level of availability you wish to seek and how you might begin to achieve your availability objectives for tape by implementing a Peer-to-Peer VTS.

The Peer-to-Peer VTS uses either ESCON or FICON connections between its components. It is important that the technical details of ESCON and FICON and its consequences are well understood before you plan a configuration. We use three theoretical examples to show the role of ESCON/FICON in a Peer-to-Peer environment. It should be understood that pictorially ESCON or FICON connectivity is depicted in the same manner, the differences between these protocols are not represent in the charts. More information on ESCON/FICON can be found starting at "FICON technology" on page 384.

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 may assist you in determining the preferred configuration for your environment.

A Peer-to-Peer VTS 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 describes the Enterprise System CONnection (ESCON) and the FIbre CONnection channel architecture. It provides information on distance issues and channel configuration alternatives using fiber optics technology that can be used for a Peer-to-Peer VTS configuration.

More information on using ESCON and FICON is available in the following publications:

- *Introduction to IBM S/390 FICON*, SG24-5176
- *Enterprise Systems Connection (ESCON) Implementation Guide*, SG24-4662
- *Continuous Availability S/390 Technology Guide*, SG24-2086
- *IBM zSeries Connectivity Handbook*, SG24-5444
- *ESCON Introduction*, GA23-0383
- *Introducing ESCON Director,* GA23-0363
- *IBM Magstar Tape Products Family: A Practical Guide*, SG24-4632
- *Guide to Sharing and Partitioning IBM Tape Library Dataserver*, SG24-4409
- *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
- *IBM Fiber Saver (2029) Implementation Guide*, SG24-5608
- *Introduction to IBM S/390 FICON*, SG24-5176
- *Magstar and IBM 3590 High Performance Tape Subsystem Technical Guide*, SG24-2506
- *IBM Magstar Virtual Tape Server and Enhancements to Magstar: New Era in Tape*, SG24-4917

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 from 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 explained in the following sections.

- **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 may not be as high as it should be since it may 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 hours per day 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 users experience neither scheduled nor 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 may 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 may also have a single point of failure that, despite the fault tolerance, can cause them as 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 make a fault-tolerant system only if it is implemented and managed properly.

Figure A-1 illustrates the relationship between the components of continuous availability.

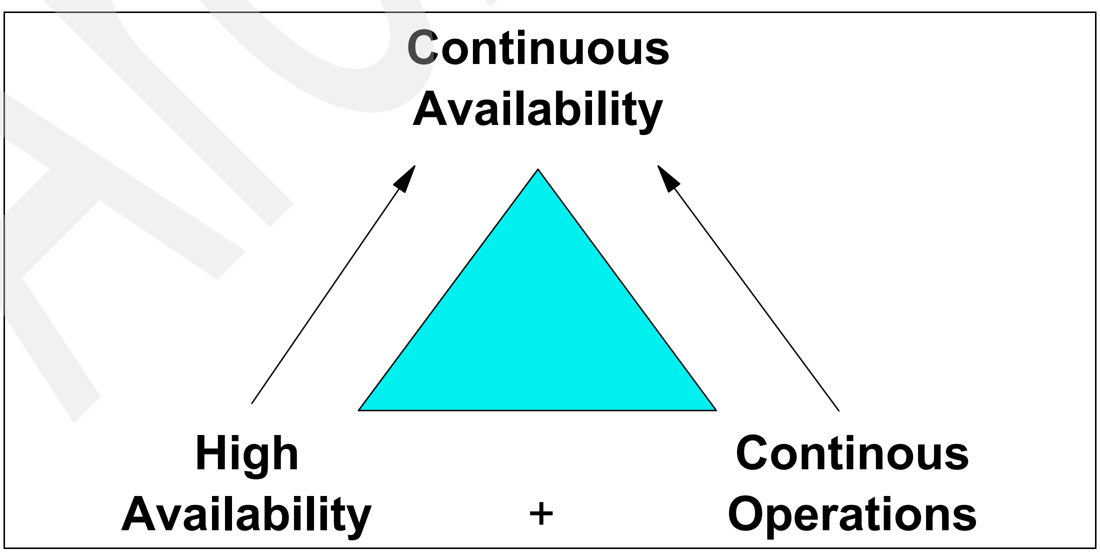

Figure A-1 Components of continuous availability

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 once 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).

Typical recovery time: The length of time for recovery is unpredictable. Indeed, you cannot be sure of being able to recover at all.

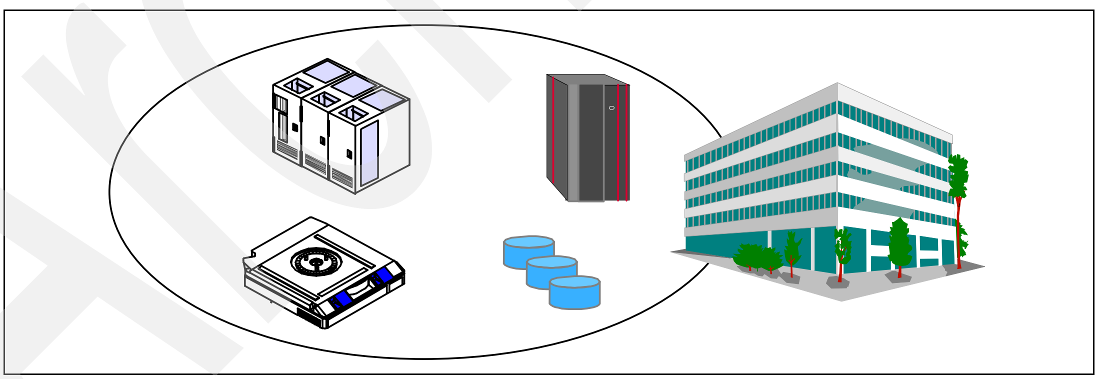

Figure A-2 Tier 0: No recovery solution

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).

Typical recovery time: the length of time till recovery is usually more than a week.

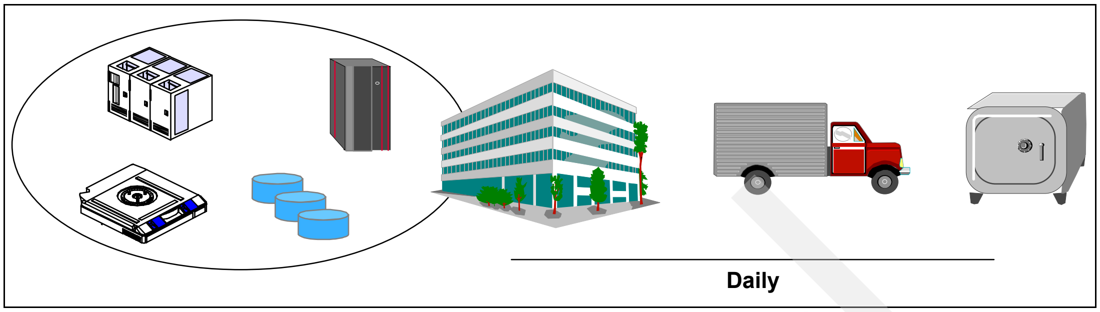

Figure A-3 Tier 1 recovery solution

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).

Typical recovery time: the length of time for recovery is usually more than one day.

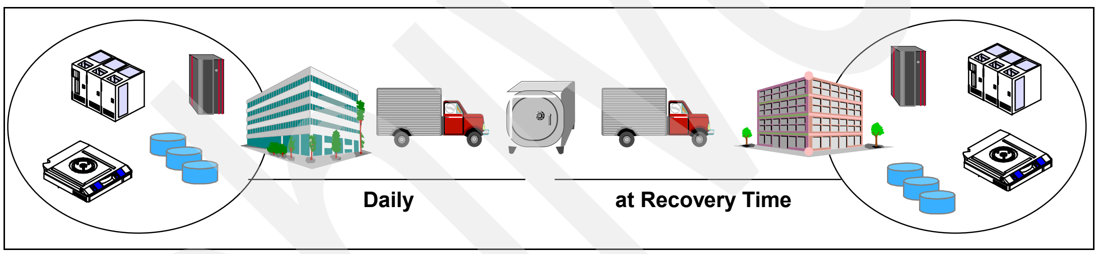

Figure A-4 Tier 2 recovery solution

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).

Typical recovery time: the length of time until recovery is usually about one day.

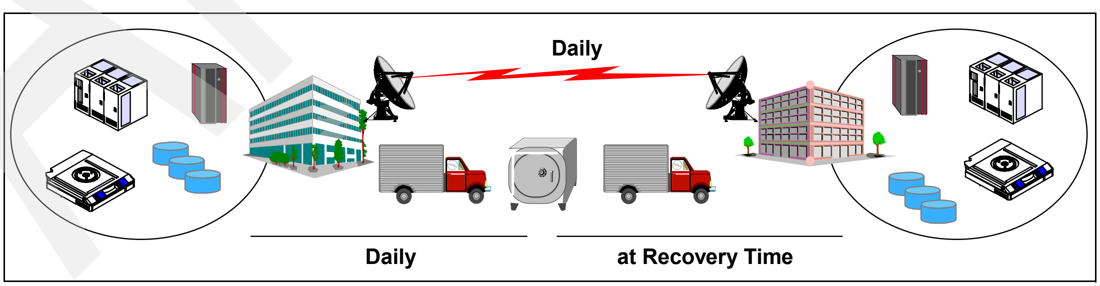

Figure A-5 Tier 3 recovery solution

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).

Typical recovery time: the length of time for recovery is usually up to one day.

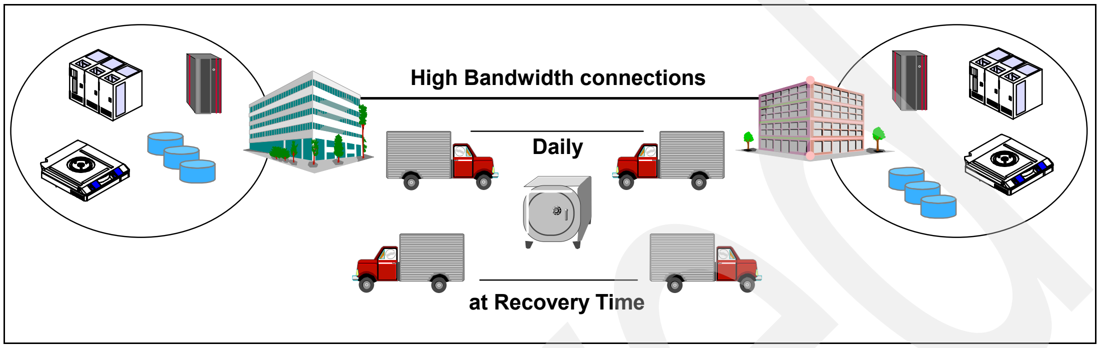

Figure A-6 Tier 4 recovery solution

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).

Typical recovery time: the length of time for recovery is usually less than 12 hours.

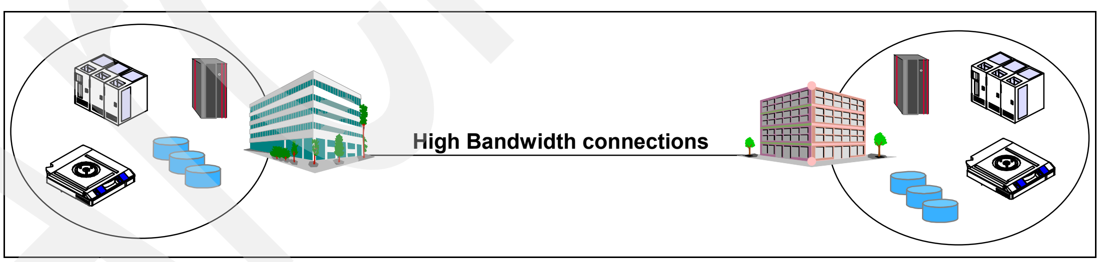

Figure A-7 Tier 5 recovery solution

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).

Typical recovery time: the length of time for recovery is usually a few minutes.

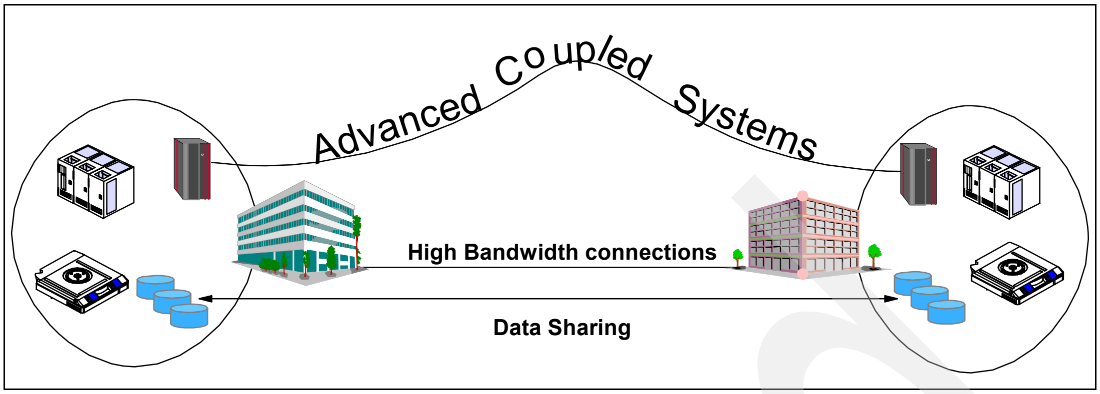

Figure A-8 Tier 6 recovery solution

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 the currency of their data. Still others need both a short recovery time and maximum data currency (Figure A-9).

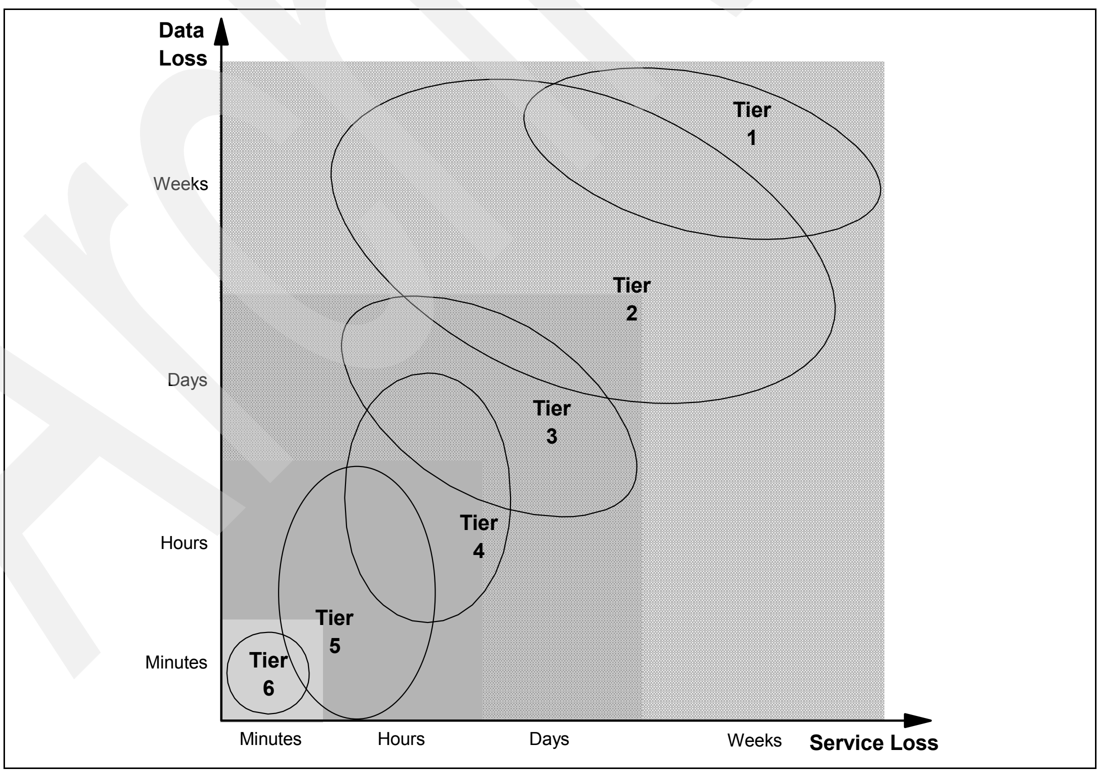

Figure A-9 Data loss and service loss

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 may 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 S/390 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.

A combination of S/390 technologies such as Parallel Sysplex, ESCON or FICON channel technology, IBM RAMAC® Virtual Array (RVA) with Extended Remote Copy (XRC), Peer-to-Peer Remote Copy (PPRC), and Snapshot functions, IBM Enterprise Storage Server (ESS) with XRC, PPRC and FlashCopy functions, and IBM Peer-to-Peer VTS, together with a good set of system automation products that include site and environmental control, can allow you to approach a tier 6 disaster recovery solution.

System attachment

The PTP VTS attaches to host systems through the channel attachment adapters of each VTC. Each VTC provides two host channel attachments. The channel types supported are:

- ESCON
- FICON

All channel attachments support a common compression algorithm and data written on one interface can be read through any other. Although ESCON and FICON cannot be intermixed in the same VTC, a PTP VTS configuration may have ESCON and FICON configured VTCs.

Distances supported

Table A-1 outlines the distances support with either ESCON or FICON as a direct connection and also with Dense Wave Division Multiplexer.

Table A-1 Distance Support

	ESCON	ESCON with DWDM	FICON with DWDM
Host Server to VTC	43 Km	75 Km	100 Km
VTC to VTS	26 Km	50 Km	100 Km

Extended distances can be provided by channel extenders certified for use by providers with a PtP VTS ESCON configuration.

ESCON attachment support

Each F/C 1001 VTC comes configured with two Extended Performance ESCON channel attachments. This attachment interface conforms to the definition provided in The Serial-I/O Architecture, PKD-0811-03 for LED drivers and receivers. The allowable maximum ESCON cable length between the VTC and either an ESCON director or an ESCON host channel is 2 km with 50.0 micron cable or 3 km with 62.5 micron cable. The maximum distance between a VTC and a host processor is 43 km. The 43 km distance assumes a 62.5 micron cable between the VTC and an ESCON director (3 km), an ESCON laser link (20 km) to another ESCON director and a final ESCON laser link (20 km) between the second ESCON director and the host processor.

FICON attachment support

A FICON (F/C 1011:1019) VTC can be configured with two short or two long wave FICON Channel attachments. Short and long wave channel attachments can be intermixed in the same VTC. Short wave attachments provide for a direct link of up to 500 m distance between the VTC and the host on 50 µ fiber and up to 175 m distance on 62.5 µ fiber. Long wave attachments provide for a direct link of up to 10 km distance between the VTC and the host on 9 µ fiber. See Table A-2 below for feature code support and distances. Short and long wave attachments provide for up to 100 km distance between the VTC and the host or between the VTC and a VTS using appropriate switches and Dense Wave Length Division Multiplexors (DWDM). One level of cascade support is provided (two dynamic switches)

Table A-2 FICON Adaptor distances

FICON Adapter	SingleMode Fibre 9 µ	MultiMode Fibre 50 µ	Multimode Fibre 62.5 µ
FC3415 LW	10 km (RPQ fro 20Km at 1Gb/sec)	550m (requires mode conditioning cable for 1Gb/sec	550m (requires mode conditioning cable for 1 Gb/sec)
FC3416 SW	N/A	500m for 1 Gb/sec 300 m for 2 Gb/sec	175 m for 1 Gb/sec 125 m for 2 Gb/sec

IBM 2029 Fiber Saver attachment

The IBM 2029 Fiber Saver can be used to provide connectivity over longer distances between an ESCON host system and a VTC and between a VTC and the Model B10, B18 or B20. The maximum end-to-end distance of the host system to a VTC interface, including the Fiber Saver pair length, is 74 km. The maximum end-to-end distance of the VTC to VTS interface, including the Fiber Saver pair length is 50 km. The maximum length may be comprised of any combination of laser or LED directors and the 2029 fiber pair. Further information on the 2029 Fiber Saver, consult the IBM Redbook *IBM Fiber Saver (2029) Implementation Guide*, SG24-5608

Peer-to-Peer VTS configurations

This section deals with some sample configurations of a Peer-to-Peer VTS. Diagrams depicted here can represent either FICON or ESCON connectivity. They are only showing possible configurations that are available with PtP VTS. You should review your requirements on connectivity with your IBM representative to further analyze your site specifically.

You must design a configuration that suits your requirements before you can plan and implement a Peer-to-Peer VTS. There is no single solution for all possible situations. Each organization has different requirements. And even with the same set of requirements, you can design various configurations depending on physical factors like the location of your computer sites and the possibility of using the currently available infrastructure.

Tape data is a part of your total data infrastructure. Any solution for your tape data must be in line with the measures you have taken or plan to take for your other data. Do you plan to use PPRC or XRC in your DASD environment? Did you take any measures for your optical data? As said before, tape data availability is a part of your data availability strategy and should be evaluated accordingly.

The basic elements for design depend on the level of continuous availability you want to achieve for your daily operations and the tier level of the disaster recovery solution you want to implement for your tape data.

The configuration samples in this section are designed to help you distinguish the differences between them. For each scenario, we describe the configuration, the options advised, and the considerations related to that scenario. In most of the configurations, the ESCON/FICON connections are presented as lines only, which includes the possibility of using directors. Remote connections are explicitly indicated, because the use of remote channel extenders impacts the use and the performance of the configuration considerably.

Configurations can be grouped as follows:

- Single site configurations
- Dual site configurations

Single site configurations

Here we describe the following scenarios:

- Single Peer-to-Peer VTS
- Multiple Peer-to-Peer VTSs
- Mixed configuration

Single Peer-to-Peer VTS

Figure A-10 represents a single Peer-to-Peer VTS in a single location with four VTCs.

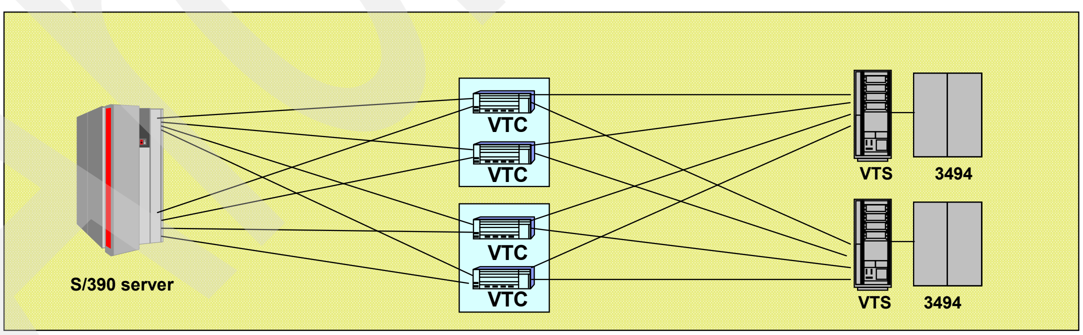

Figure A-10 Single Peer-to-Peer VTS configuration

Availability targets

The availability targets are:

- This configuration addresses the need for high data availability in the day-to-day operation. Tape data availability in both distributed libraries must be guaranteed. It can be compared with a request for a RAID 1 solution in the DASD environment. This configuration may be required if applications rely heavily on tape data.

- This configuration is not intended for disaster recovery—you need to address disaster measures separately. That can be done by making additional copies of your tape data on real volumes that are transported to another site (a tier 1 or tier 2 solution) or sent via telecommunication lines to the disaster location (a tier 3 solution).

Configuration description

The configuration is as follows:

- The S/390 servers represent one or more OS/390 images. These images can share in an SMSplex or partition the composite library as is presently available on stand-alone VTSs.
- Figure A-10 on page 373 shows two channel connections per VTC to the host. The controller allows you to use only one channel connection rather than two, but you should be aware that you have configured a single point of failure if you do so.
- With four VTCs, the controllers can be housed in one Model CXn frame or split over two Model CXn frames. Two frames offer a slight advantage over one frame because the VTCs can be located in different places.
- With eight VTCs, the controllers must be split across at least two CXn frames. If you have multiple Peer-to-Peer VTSs, perhaps each CXn might contain VTCs from two separate Peer-to-Peer VTSs.
- Each Model CXn frame has two power sources.

Configuration considerations

The configuration considerations are as follows:

- The use of immediate mode of operation for the composite library ensures that the copy is available at completion of rewind/unload of the tape.
- The use of one channel connection per VTC limits the requirement for the number of channels available in the host. However, it also limits availability of the VTC, if the connection fails, you lose 16 virtual addresses.
- Neither of the VTSs is defined as the preferred VTS. Non-preferred I/O provides the best performing option.
- One of the distributed 3494 libraries must be assigned as the UI distributed library. It is likely that it will be the library closest to the operator console bridge.

Dual Peer-to-Peer VTS

Figure A-11 shows a dual Peer-to-Peer VTS in a single location. Compared to the previous scenario, another Peer-to-Peer VTS has been added in the same location.

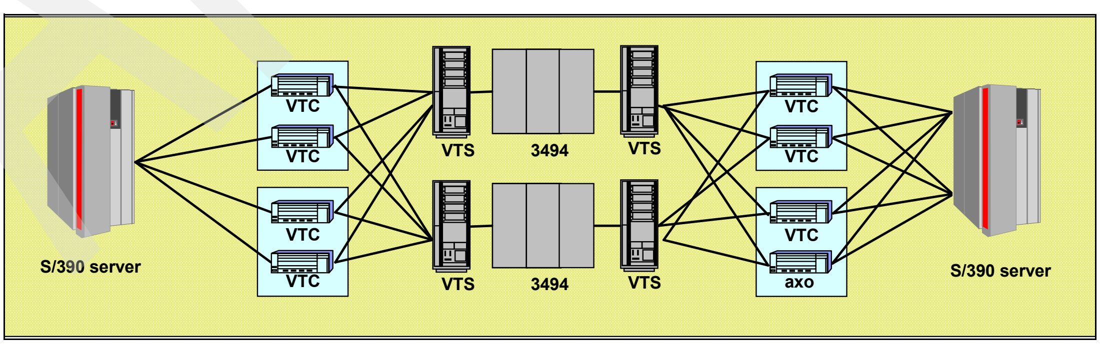

Figure A-11 Dual Peer-to-Peer VTS configuration

Availability targets

The availability targets are:

- The data availability targets for the composite libraries can be different. One library can have a requirement for high data availability in day-to-day operations, while the other composite library has a requirement for high batch throughput.
- From a disaster recovery point of view, both Peer-to-Peer VTSs are the same. Other measures should be put in place for this.

Configuration description

The configuration is:

- Both S/390 servers have OS/390 images running. The workload of these images can be different. The images in a single S/390 may share their connected library in a SMSplex or they may partition the library. Images in the other S/390 server are not connected to the same Peer-to-Peer VTS. They can be part of the same SMSplex but they do not share the Peer-to-Peer VTS libraries.
- The S/390 server on the left in Figure A-11 on page 374 has four channels connected to the VTC, thereby limiting the number of channels required at the server that may be necessary if there is a shortage of channels. Compare this to the S/390 server on the right, which has eight channel attachments.
- The VTC can be placed in one Model CXn frame per composite library. It is possible to split the VTCs over two Model CXn frames. Each Model CXn frame has two VTCs from one composite library and two VTCs from the other. If there were eight VTCs per Peer-to-Peer VTS, we would need four CXn frames and might choose to distribute VTCs across CXn frames in a similar way although each Peer-to-Peer VTS would already have VTCs in multiple CXn frames.
- There are two 3494 tape libraries. Each has two Model Bxx VTSs. Each VTS of a Peer-to-Peer VTS must be attached to a different library.

Configuration considerations

The configuration considerations are:

- The use of one channel connection per VTC, as shown for the left S/390 server, limits the requirement for the number of channels available in the host. However, it also limits availability of the VTC, if the ESCON connection fails, you lose 16 virtual addresses.
- One of the composite libraries can use the immediate mode of operation while the other can use deferred mode. The 3494 Library Manager is able to communicate with distributed VTSs running in different modes of operation.
- The use of immediate mode of operation for the composite library ensures that the copy is available at the completion of rewind/unload for the volume, but it increases the duration of a job by the time needed to perform the copy function. Deferred copy mode schedules the copy function asynchronously.
- The best choice is to have all VTCs installed with no preference. The load is then balanced over the two distributed libraries. Although it is strongly advised not to do so, it may happen that the distributed libraries within a Peer-to-Peer VTS have unequal cache sizes. Configuring VTCs to prefer the VTS with the bigger cache size reduces the risk of throttling caused by cache overload.
- One of the distributed libraries from each composite library must be assigned as the UI distributed library. It is likely that it will be the 3494 tape library nearest to the operator console bridge.

- A 3494 Tape Library has some single points of failure. Many of them, but not all, can be removed by installing availability options like the high availability option (HA1), a second accessor and dual grippers. A failure of a 3494 Tape Library impacts both Peer-to-Peer VTSs.

Mixed configuration

In Figure A-12, a Peer-to-Peer VTS, a stand-alone Model Bxx VTS, and native drives share two 3494 tape libraries at a single location.

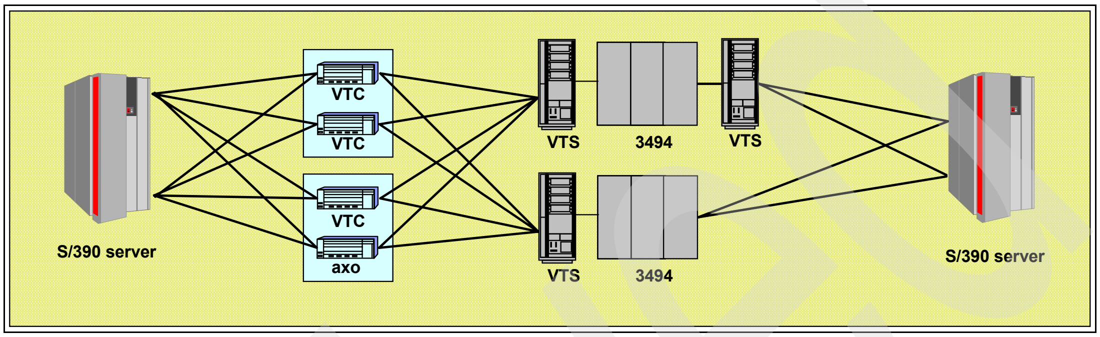

Figure A-12 Mixed configuration at a single location

Availability targets

The availability targets are as follows:

- The data availability targets for Peer-to-Peer VTS are the same as in the previous configuration.
- The data availability targets for the stand-alone Model Bxx VTS and for the native drives are not further discussed here.

Configuration description

The configuration is as follows:

- Both S/390 servers have OS/390 images running. The images may be part of the same SMSplex but they do not share the logical libraries if they are not running on the same S/390 server.
- For the VTCs, the same applies as discussed in “Single Peer-to-Peer VTS” on page 373.
- There are three logical libraries: the Peer-to-Peer VTS, the stand-alone VTS, and the native drives. One 3494 Tape Library includes a distributed library and the stand-alone library. The other 3494 Tape Library includes a distributed library and the native drives.

Configuration considerations

The configuration considerations are as follows:

- For the Peer-to-Peer VTS, the same considerations apply as described in “Single Peer-to-Peer VTS” on page 373.
- The 3494 tape libraries are not dedicated solely to the Peer-to-Peer VTS. It is possible that the availability of the Peer-to-Peer VTS may be impacted by problems with the other logical libraries.

Dual site configurations

Here we describe the following scenarios:

- External distributed library
- External distributed library with standby VTCs
- Remote distributed library
- External distributed library and VTCs
- Mirrored site

External distributed library

In Figure A-13, one of the distributed libraries is placed in an external location. This library is connected to the VTCs with remote links. This configuration has four VTCs; we could also have a configuration with eight VTCs.

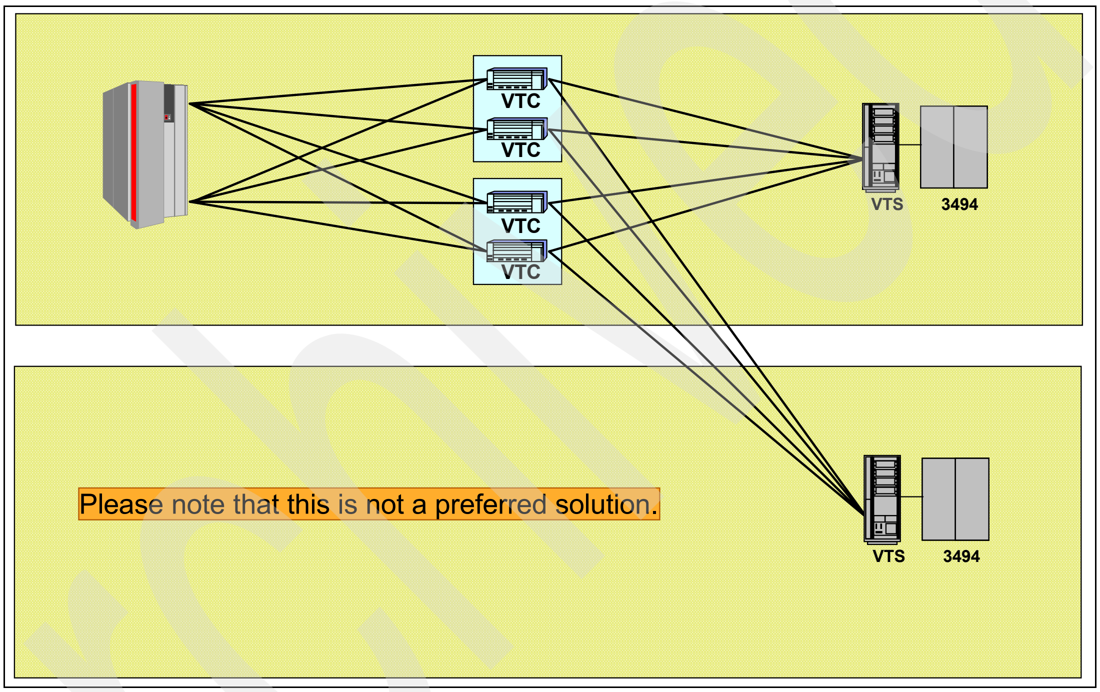

Figure A-13 External distributed library

Availability targets

The availability targets are as follows:

- This configuration addresses the need for having a copy of the tape data available outside the main computer center for disaster recovery reasons. This can be considered as a tier 3 solution for tape data.
- An outage of one of the components of the distributed library in the main location must be handled by the configuration within hours.
- Recover from a full disaster may take from days to one week. There must be arrangements with a recovery service provider like IBM Continuity Services for providing IT facilities in case of a disaster.

Configuration description

The configuration is as follows:

- The S/390 servers represent one or more OS/390 images. The images can share in a SMSplex or partition the composite library as is today available on stand-alone B10, B18, B20 VTSs.
- All VTCs are placed in the main location. There may be several reasons for choosing this configuration.
  - The distance between the two sites requires remote channel connections. When you place VTCs in the remote location, there is a need for additional infrastructure between the two sites. The costs involved for this may prohibit this solution.
  - If you place VTCs in the remote location and the distance between the locations is near its limit, it introduces throughput delays between the host and the remote VTCs, this is especially relevant for ESCON connections. It can also causes delays between the remote controllers and the distributed library in the main location. This is in addition to the delay already experienced between the VTCS in the main location and the distributed library in the remote location.
  - With a configuration containing four VTCs, you may not want to be exposed to the temporary loss of the remote location, given that you would risk being unable to continue writing to your Peer-to-Peer VTS because you would not have three AX0s available to perform a master switchover. See 9.2.9, “Site failure” on page 338, for details. This could happen if you did not control the remote location. The use of eight VTCs alleviates this concern.
  - The external location may not be owned by you. Placing VTCs at this site adds additional risks to your daily operation.

Configuration considerations

The configuration considerations are as follows:

- The availability targets allow the use of either mode of operation for the composite library: immediate copy mode or deferred copy mode. Deferred mode is the option of choice if you are able to continue your workload without tape data that has not been copied to the remote distributed library.
- The choice between balanced I/O and preferred is dependent on the expected configuration and the response times on the channels to the remote location. Excessive delays dictate the use of preferred I/O in all VTCs for the distributed library in the main location.
- However, if you choose deferred copy mode, it may be advantageous to give the distributed library in the remote location preference. This results in a copy of the tape data in the remote location as early as possible.
- Additional equipment is required in case of a disaster. A distributed library is not able to function without the presence of a pair of VTCs. These controllers must be installed before you can access the tape data in the remaining VTS in the remote location.

External distributed library with Standby VTCs

This is a variation on the previous case, very useful if the remote site has OS/390 processors. Here we add four Standby VTCs in a CXn frame at the remote location. The Standby VTCs are not cabled to the distributed VTS in normal operation. This configuration and the way it is used are described in 9.3, “Disaster Recovery testing with Standby VTCs” on page 341.

Availability targets

The availability targets for this configuration are better than for the previous configuration as access to data can be restored in a short while.

Configuration description

The configuration is:

- The difference between this configuration and the previous one is the use of the Standby VTCs. If there are one or more host systems at the remote site, Standby VTCs permit the use of the remote distributed library alone in the event of a disaster at the primary site or for scheduled disaster recovery testing at the remote site.
- A spare set of VTCs is installed at the remote site but not cabled to the VTS unless there is a disaster at the primary site.

Configuration considerations

The configuration considerations are as follows:

- The availability targets allow the use of either mode of operation for the composite library: immediate copy mode and deferred copy mode. Deferred mode is the option of choice here if you are able to continue your workload without tape data that has been produced in the period just before the outage of one of the distributed libraries.
- Recabling is required in case of a disaster. A distributed library is not able to function without the presence of VTCs. In this case, the controllers are already installed and need only be cabled and IMLed before you can access the tape data in the remaining VTS in the remote location.

Remote distributed library

In the configuration presented in Figure A-14, the external VTS is placed in a remote location requiring the use of switching units, DWDMs or channel extenders.

> **Note:** Please remember that you will need to discuss the use of any of these options in a Peer-to-Peer VTS environment with the vendors to ensure that the channel extenders have been tested and certified for use in this environment.

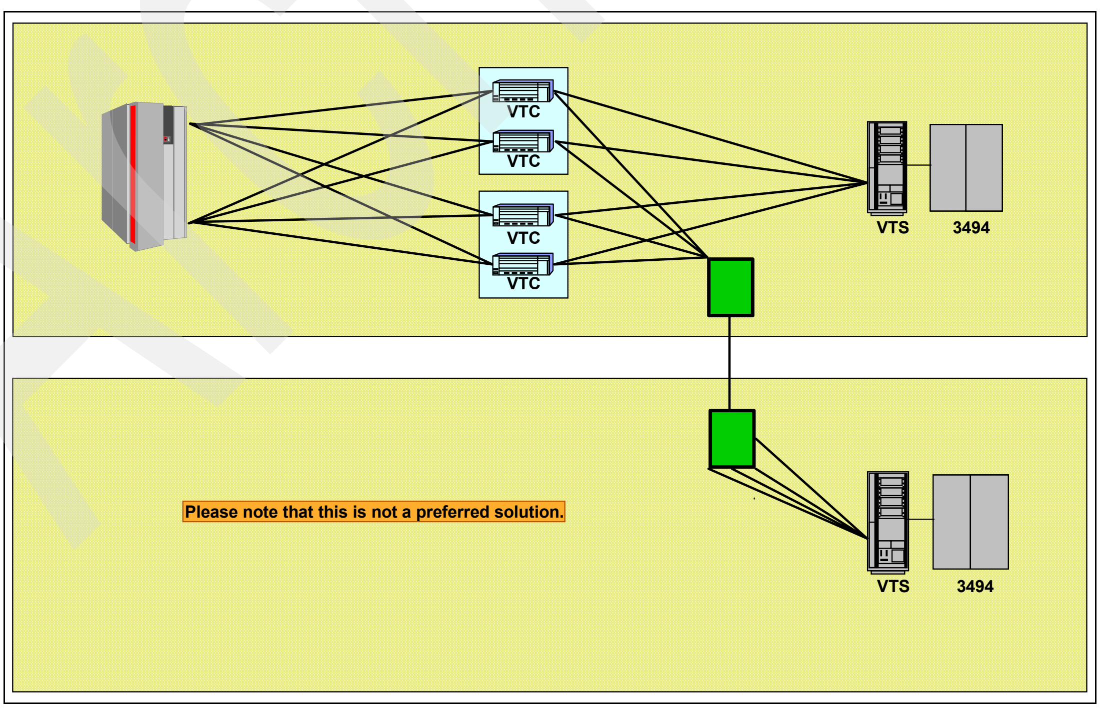

Figure A-14 Remote distributed library

Availability targets

The availability targets for this configuration are the same as in the previous configuration described on “Availability targets” on page 377.

Configuration description

The configuration is as follows:

- The main difference between this configuration and the previous one is the use of the remote channel extenders. It was calculated that the remote connection was able to handle the load even in peak periods.
- The “Configuration description” on page 378 fully applies to this configuration.

Configuration considerations

The configuration considerations are as follows:

- The availability targets allow the use of either mode of operation for the composite library: immediate copy mode and deferred copy mode. Deferred mode is the option of choice here if you are able to continue your workload without tape data that has been produced in the period just before the outage of one of the distributed libraries.
- Current performance measurements for distances of up to 25 km show that the choice between non-preferred I/O and preferred is in favor of non-preferred I/O for the distributed library, despite the use of extenders.
- Also in this configuration, additional equipment is required in the event of a disaster. A distributed library is not able to function without the presence of two VTCs. These controllers must be installed before you can access the data in the remaining VTS in the remote location.

External distributed library and AX0 controllers

In the configuration presented in Figure A-15, two VTCs are placed in the location of the external distributed VTS. Basically, this is an improved scenario compared to the one shown in “External distributed library” on page 377.

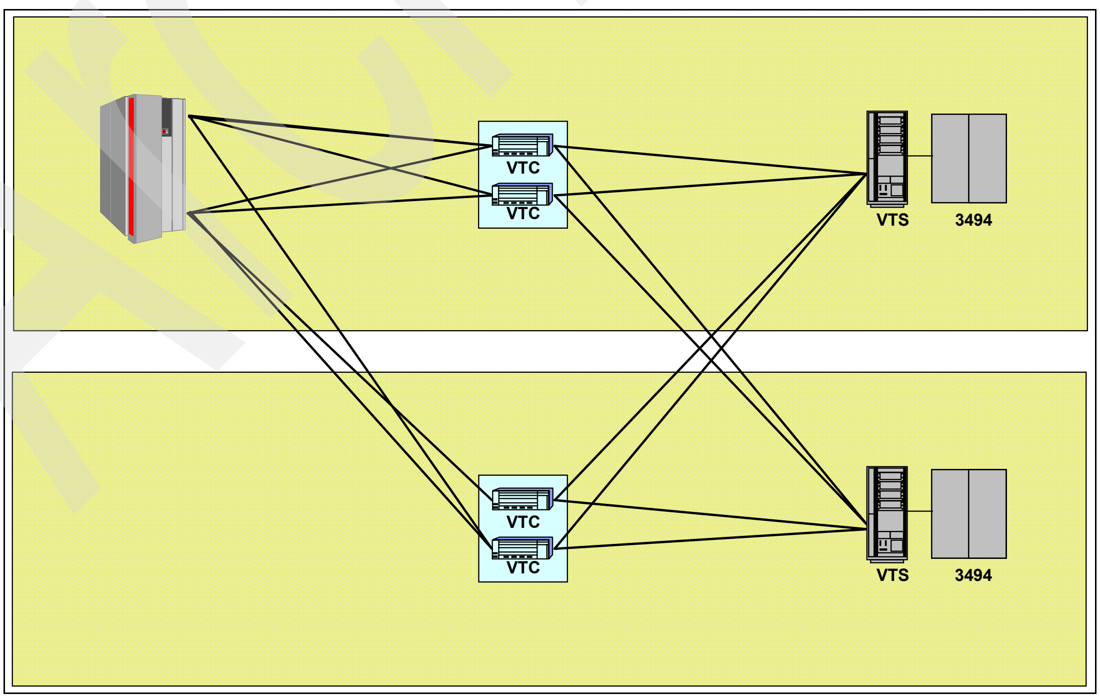

Figure A-15 External distributed library and four AX0 controllers

We could improve still further on this configuration by having eight VTCs (four at the local site and four at the remote site).

Availability targets

The availability targets as follows:

- This configuration addresses the need for having a copy of the tape data available outside the main computer center for disaster recovery reasons. This can be considered a tier 3 solution for tape data.
- A partial disaster with one of the components of the distributed library in the main location must be handled by the configuration within hours.
- Recovery from a full disaster may take from hours to days. There are arrangements made with a recovery service provider like IBM Continuity Services for providing ICT facilities (processing power) in case of a disaster.

Configuration description

The configuration is as follows:

- The main differentiator between this configuration and the one described in "External distributed library" on page 377, is the placement of two or four VTCs in Model CXn frames in the external site.

Configuration considerations

The configuration considerations are as follows:

- The availability targets allow the use of either mode of operation for the composite library: immediate copy mode or deferred copy mode. Deferred mode is the option of choice here if you are able to continue your workload without tape data that has been produced in the period just before the outage of one of the distributed libraries.
- Current performance measurements for distances up to 25 km show that the choice between non-preferred I/O and preferred is in favor of non-preferred I/O for the distributed library, despite the use of extenders.

Mirrored site

Figure A-16 presents a dual site configuration where all major components are duplicated in both sites. There is no main site because they are mirrors of each other.

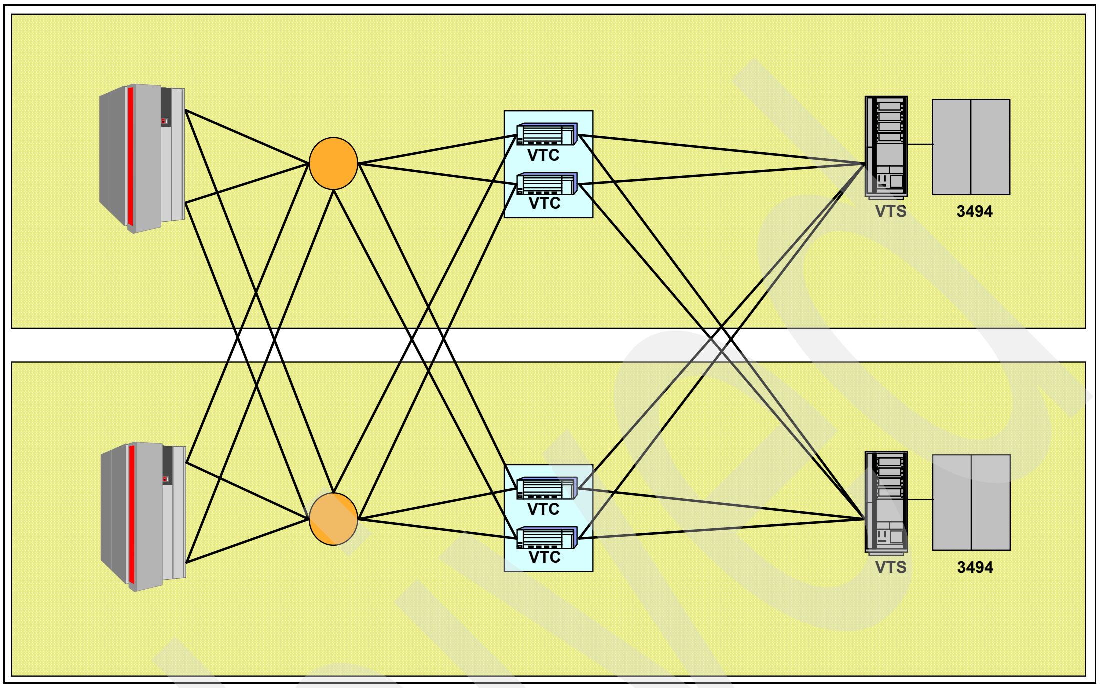

Figure A-16 Mirrored external site

The figure shows directors on the left-hand (host) side of the configuration, as we expect that they will already be present in many installations to allow easy reconfiguration. They could equally well appear on the right of the diagram as well. DWDM connectivity could also be used in this scenario if required.

This configuration might be included in a Geographically Dispersed Parallel Sysplex (GDPS) configuration. GDPS provides switching capability from one site to another site, for planned and for unplanned outages. It provides the capability to manage the configuration, automates parallel sysplex operational tasks, and automates the recovery from a failure through a single point of control. GDPS supports all transaction managers (for example, CICS and IMS) and database managers (for example DB2, IMS, and VSAM). For more information, see 9.5, "Geographically Dispersed Parallel Sysplex" on page 349.

Availability targets

The availability targets are as follows:

- This configuration addresses the need for having a high availability data center on a tier 5 or tier 6 level. All data, both DASD and tape data is mirrored in both sites.
- A partial disaster with one of the components of the distributed library in the main location must be recovered immediately.
- Recovery from a full disaster may take less than a couple of hours, and additional arrangements are probably made for the critical applications to continue without a service interruption.

Configuration description

For the configuration:

- The S/390 servers represents one or more OS/390 images. The images can share in a SMSplex or partition the composite library as is today available on stand-alone B10, B18, B20 VTSs.
- The production OS/390 has at least an active backup system available in the other site. It is likely that the images are part of a parallel sysplex environment.
- Complex data centers usually have directors in use for connecting the I/O equipment to the hosts.
- There are many connections between the two sites. They are not only required for DASD and tape mirroring but also needed for channel-to-channel connections, printers and network controllers.
- Eight VTCs housed in Model CXn frames are distributed between the two sites, either four and four or six and two. The CXn frames use the dual power facility.
- The installation of additional director functionality between the Model AX0 controllers and the VTSs improves the configuration flexibility.

Configuration considerations

The configuration considerations are as follows:

- The availability targets allow the use of either mode of operation for the composite library: immediate copy mode or deferred copy mode. For your critical workload with a requirement for a tier 6 solution, the only option is immediate copy mode. Remember that this mode is set the same in all VTCs. If less critical partitions share the Peer-to-Peer VTS, they will also use immediate copy mode. The influence of the load of the less critical partitions, particularly on the batch window of your production partition, may be undesirable. Only a dedicated Peer-to-Peer VTS for your critical workloads prevents this workload from being impacted by another load.
- Most likely, you will choose no preference for the I/O VTS if the sites are within acceptable connectivity distances. This assures that you exploit the available resources to an optimum level.
- The location of your operations and support personnel will govern which of the distributed libraries you choose to be the UI library. It influences the initial selection of the Master VTS, but as explained in 2.6.2, “Master VTS switchover” on page 50, the Master function can switch over to the other distributed library after a failure of one of the components of the composite library.

ESCON technology

ESCON is an IBM architecture and fiber optic technology. It is an integral part of the Enterprise System Architecture/390 (ESA/390). It allows a data center installation to replace electrical-copper cables with fiber-optic cables. Beyond this physical aspect, ESCON provides new capabilities to connect I/O equipment to a host processor at distances beyond 100 km. With these increased distances you can establish new disaster recovery methods by remote copying your DASD and tape data.

Although ESCON provides capabilities to connect over long distances, the following factors affect the distance that can be achieved:

- Physical size of fiber
- Fiber transmission mode
- Connection mode

FICON technology

The increased performance, flexibility, and distance capability of FIbre CONnection (FICON) channel attachment is now available to the IBM TotalStorage Peer-to-Peer Virtual Tape Server 3494 (PtP VTS). The IBM TotalStorage Virtual Tape Frame 3494 Model CX1 integrates new Virtual Tape Controller (VTC) features with FICON attachments. The VTCs can attach the IBM TotalStorage Virtual Tape Server (VTS) 3494 Models B10 or B20 to host system FICON channels.

The new PtP VTS configurations provide:

- FICON attachment for increased performance and distance
- Cascading and 2Gb/s switch support for greater attachment flexibility
- Dual Path Concentrator to improve availability

When compared with ESCON, FICON reduces protocol overhead. This capability, combined with technology advances, allows a fiber channel link to perform at up to 100 MB/sec full duplex. This approximately equates to up to five concurrent ESCON bulk-data I/O operations, or up to eight concurrent ESCON transaction I/O operations (depending on the channel utilization of the ESCON channels). Enhanced-distance connectivity is enabled by the fact that the FICON channel can span up to 100 kilometers before significant data rate droop takes effect. However, such an implementation would require DWDM.

FICON channels can use existing 9 micron single-mode fiber, or 50 or 62.5 micron multi-mode fibers. Any existing fibers in use by ESCON could be used for FICON. Note that 50 and 62.5 micron fiber is only supported up to 550 meters; 9 micron is supported up to 10 kilometers (or 20 kilometers with an RPQ). FICON will be the foundation for high-performance channels. The architecture and implementation allows for:

- Initial implementation to ESCON devices via the IBM 9032 model 5 ESCON director acting as a bridge
- Point-to-point to I/O subsystems with S/390 FICON interfaces
- Switched point-to-point via a FICON Director to I/O subsystems with S/390 FICON interfaces
- DWDM connection

Table 3-10 on page 112 details FICON adaptor distances.

Physical size of a fiber

An optical fiber functions as a kind of waveguide for light. It is usually made of silica glass. The fiber itself has a central core and a surrounding cladding of slightly different glass material. The physical size of an optical fiber is determined by the diameter of the core and cladding, expressed in microns (μ). A fiber optic cable having a core diameter of 62.5 μ and a cladding of 125 μ is designated as an 62.5/125 μ optical fiber. Other fibers commonly used are 50/125 μ and 9/125 μ. Figure A-17 shows the fiber optic technology.

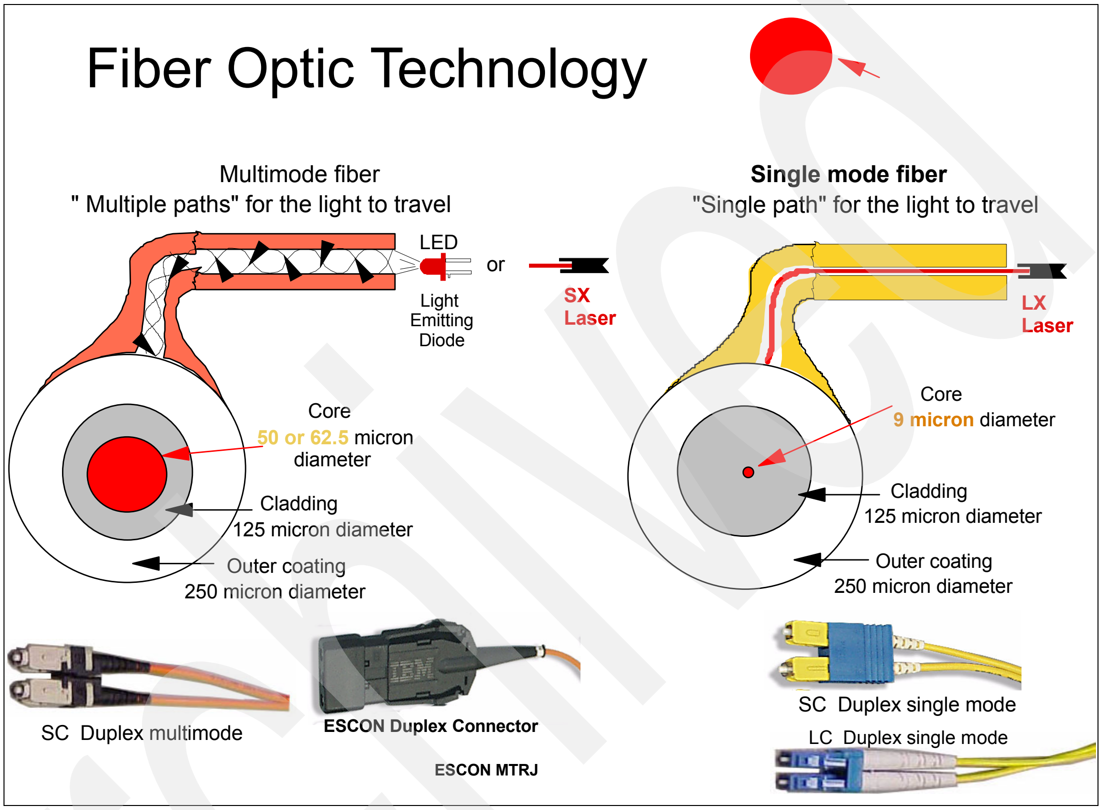

Figure A-17 Fiber Optic Technology

The jumper cable has the same function as the pair of cables that make up a parallel channel. A jumper cable consists of two unidirectional fibers, one that carries data in one direction (transmit), and the other that carries data in the opposite direction (receive). The unidirection is due to the way the fibers are connected to transmitters and receivers, not to any property of the fiber itself.

Jumper cables are typically used to connect:

- A channel to a control unit
- A channel to a director
- A director to a director
- A director to a control unit
- A channel to a VTC
- A VTC to a Model B10 or B20 VTS

Fiber transmission mode

Two modes can be used to send light signals through an optical fiber: Single mode (mono mode) or multimode. The optical fibers used are accordingly called single-mode or multimode fibers. These fibers have different physical dimensions and light transmission characteristics.

The intensity of the light decreases when the light passes through the fiber. The longer the fiber, the lower the light intensity. The light intensity attenuation results in length limitations of the physical link because the receiver needs a minimum level of light intensity to correctly detect the signals. Figure A-18 shows how the two modes are affected over distance.

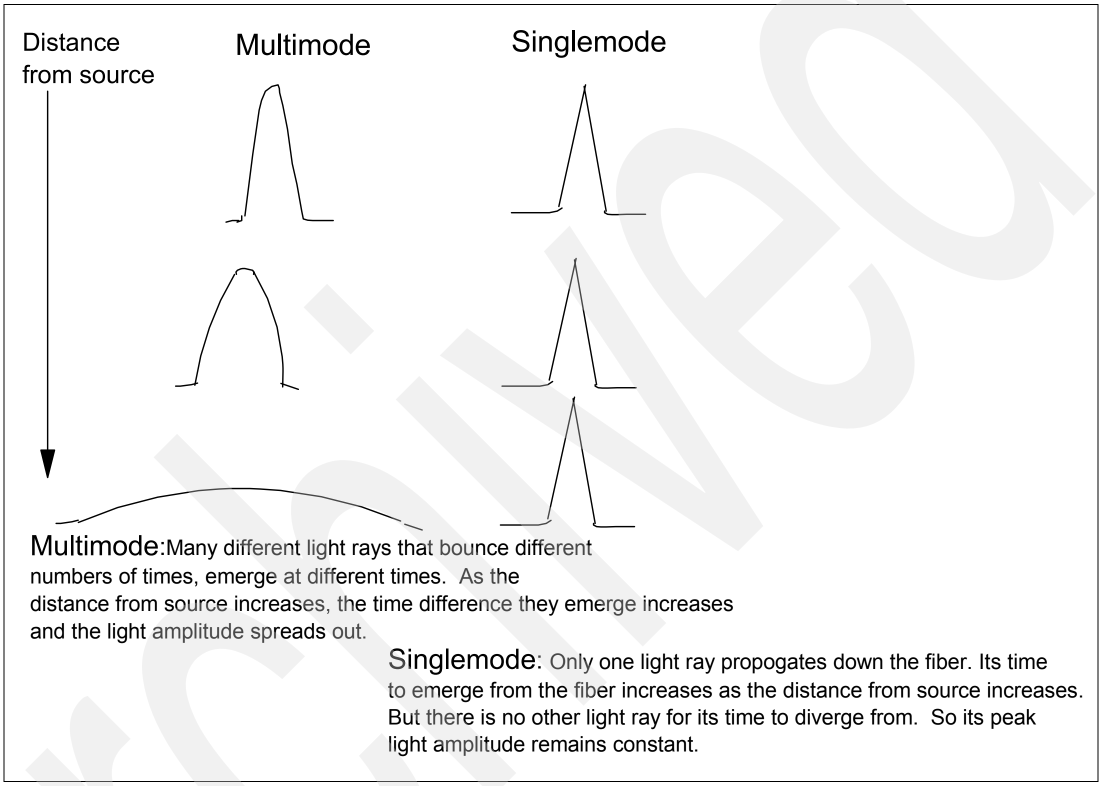

Figure A-18 Multimode Fiber Light Amplitude loss

The multimode fibers supported by IBM are either 62.5/125 µ or 50/125 µ. The light source used for multimode fiber is usually a light emitting diode (LED). LED jumpers are orange, and their duplex connectors are generally black.

The maximum distance for a multimode fiber link is 3 km if 62.5/125 µ fibers are used and 2 km if 50/125 µ fibers are used. The maximum distance for a multimode fiber link is 2 km if both 62.5/125 µ and 50/125 µ fibers are used.

Single-mode fiber usually has a core diameter of 8 to 10 µ and a cladding diameter of 125 µ. The light source used for single-mode fiber is a laser. IBM supports single-mode 9/125 µ fibers for use in an ESCON/FICON environment. Using the grating properties of light at this Wavelength and opening size one ray enters. Since one ray enter, it comes out at the same time, no matter the distance, relative to itself since there is only one. LC connector on each end (one centimeter wide) used with new Hosts and new switches and newer wiring panels and also to connect directly from VTC to VTS. The maximum distance for a single-mode fiber link is 20 km.

ESCON/FICON connection modes

You can connect an I/O control unit to a host channel through:

- A point-to-point connection
- A point-to-point connection with channel sharing
- A director using a static connection
- A director using a dynamic connection
- A FICON bridge connection

A point-to-point connection

A point-to-point connection is a fiber connection between a host channel and an I/O control unit (Figure A-19). This is called a physical link and depicts both ESCON and FICON. Through this physical link a host can communicate with only the devices attached to that I/O control unit. This diagram shows a point-to-point connection.

Figure A-19 Point-to-point connection

Point-to-point connection with channel sharing

Channel sharing can be achieved by using the Enhanced Multiple Image Facility (EMIF or MIF). With EMIF, the processor channel subsystem provides physical path sharing by extending the logical addressing capability of the channel architecture to logical partitions in host images. Each partition has its own logical channel subsystem and its own view of each shared channel and each control unit connected to the shared channel subsystem (Figure A-20). Nevertheless, it is still possible to define non-shared channels.

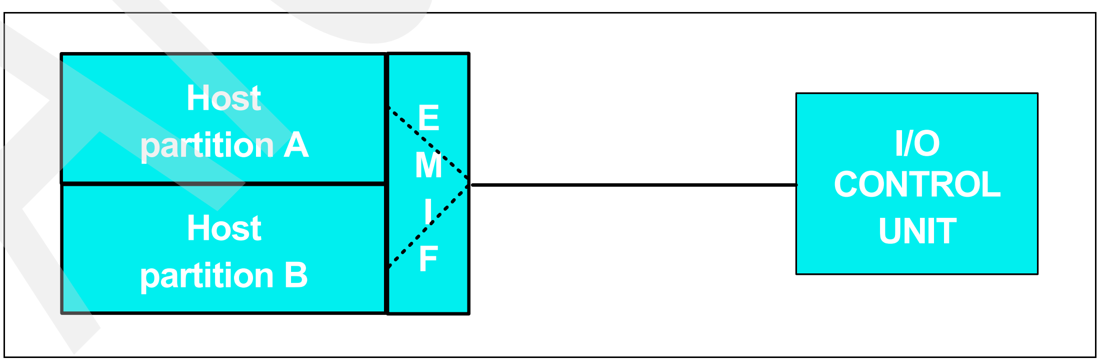

Figure A-20 Point-to-point connection with channel sharing

ESCON director using a static connection

A static connection through an ESCON director involves having one fiber connected from a host channel to a ESCON director and a second fiber from an ESCON director to an I/O control unit. In Figure A-21, the host channel is connected through ESCON director port D7 and ESCON director port C3 to an I/O control unit. You must set up an ESCON director for a static connection (dedicated connection) for these two ports, by using the ESCON director console or a managing program running on a OS/390 server.

A static connection through an IBM 9032 ESCON director is more flexible than a point-to-point connection. You can achieve a different connection by statically switching to another port, without reconnecting the fiber.

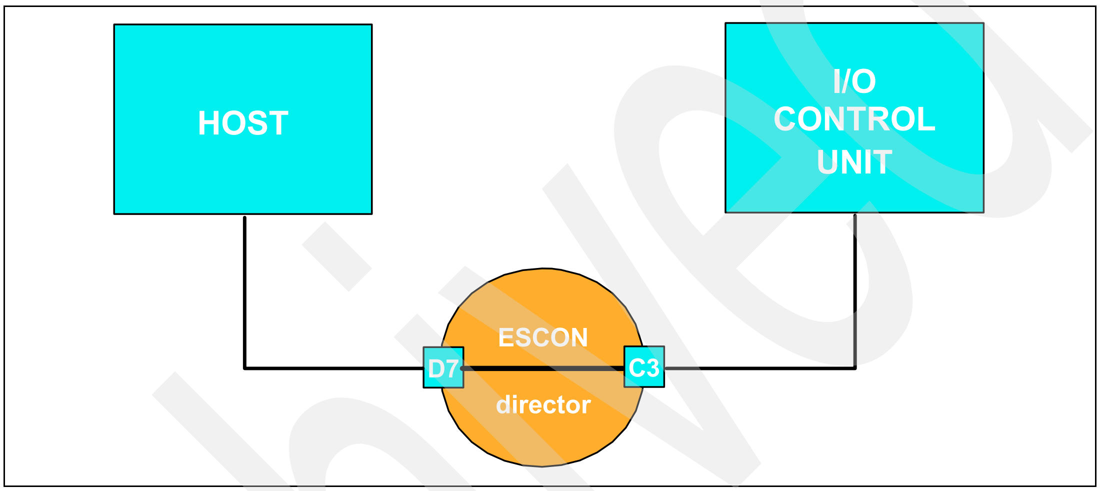

Figure A-21 Static connection through ESCON director

FICON director using a static connection

A static connection through a FICON director can now be used between the host and VTC and between the VTC and VTS. Currently only McData ED6064 and InRange FC9000 Model 64, 128 and 256 switches (in any combination) are supported in a PtP VTS subsystem. These are defined/altered using GUI switch setup, the fabric is bi-directional and can be used for DWDM connections. Pictorially this configuration is the same as Figure A-21 with the only difference being ESCON director is replaced with the FICON director used and as such another diagram for FICON is not shown.

ESCON/FICON director using a dynamic connection

A dynamic connection through an ESCON/FICON director, as seen in Figure A-22, involves having the same fiber arrangement as for a static connection. However, instead of setting up an ESCON/FICON director for a static connection, you have to code the connections to ESCON/FICON director destination ports C3 and E5 in the link statement of the control unit macro in the IOCDS or equivalent HCD panel. In this mode, a host can communicate with multiple devices behind different I/O control units by using the same physical link between the host channel and the director. The ESCON/FICON director dynamically connects to the destination port, as required.

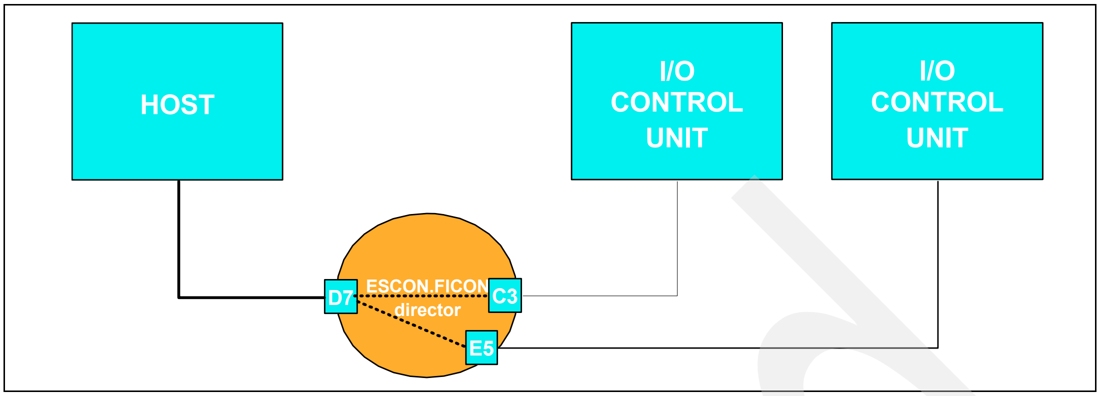

Figure A-22 Dynamic connection through an ESCON/FICON director

ESCON/FICON distance alternatives

Because the light intensity is lower for increased distances, some kind of amplifier or repeater is needed to achieve distances longer than:

- 2 km for a 50/125 µ multimode fiber
- 3 km for a 62.5/125 µ multimode fiber
- 20 km for a 9/125 µ single-mode fiber

There are several options available for extending this distance:

- Using an ESCON/FICON director
- Using an ESCON channel extender
- Using a remote channel extender

Using an ESCON director

The main purpose of an ESCON director is to provide switching functions within an ESCON network. In addition to its switching function it acts as an amplifier or repeater, because the light signals are received on one port and transmitted by another, thus allowing distances of 2 km, 3 km, or 20 km to be reached.

A connection between the host and the I/O control unit can contain two ESCON directors (Figure A-23). The maximum distance between these directors is 3 km for a 62.5/125 µ multimode fiber, 2 km for a 50/125 µ multimode fiber, and 20 km for a 9/125 µ single mode fiber, using XDF ports on the ESCON director.

Remember that there are no XDF ports available on the host and the I/O control units. The maximum distance with two directors between the host and the control unit is 26 km (62.5/125 µ) or 24 km (50/125 µ). Depending on the transmission mode, configuration, adapters, and the quality of the fibers used, the distance that can be reached may be less than that maximum.

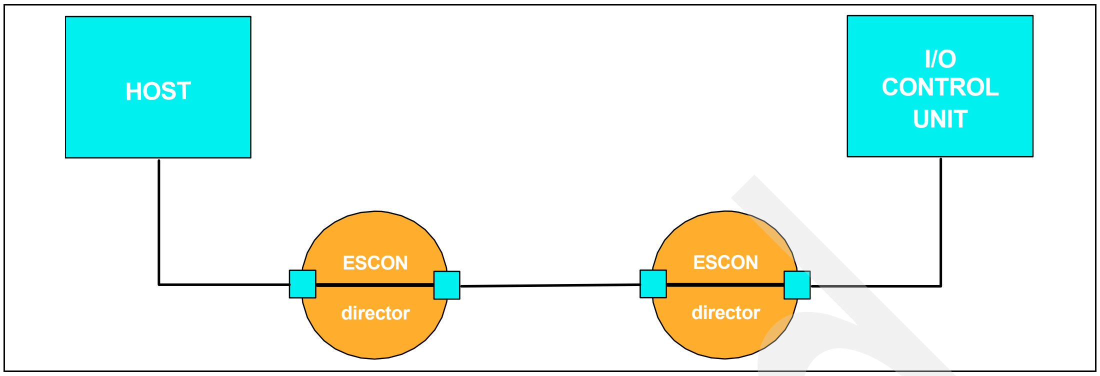

Figure A-23 Connection through two ESCON directors

On the S/390 Generation 5 and Generation 6 servers, FIber CONnection (FICON) channels are available. FICON increases the link rate from 20 MB all the way up to 100 MB/sec. The effective data transfer rates are expected to be between 60 MB/sec and 80 MB/sec. The FICON implementation enables full duplex data transfer, so data travels in both directions simultaneously, rather than as with the ESCON half-duplex data transfer.

The S/390 server with FICON channels allows the accessing of control units with an ESCON interface via an ESCON director FICON bridge adapter card. Using a 9 µ single mode fiber FICON supports an unrepeated channel distance of 10 km. This can be increased to 20 KM via RPQ.

Figure A-24 shows that with FICON, fewer connections are needed. A single FICON channel is able to handle more than one ESCON control unit without performance impact on the throughput of the I/O between the host and the control unit.

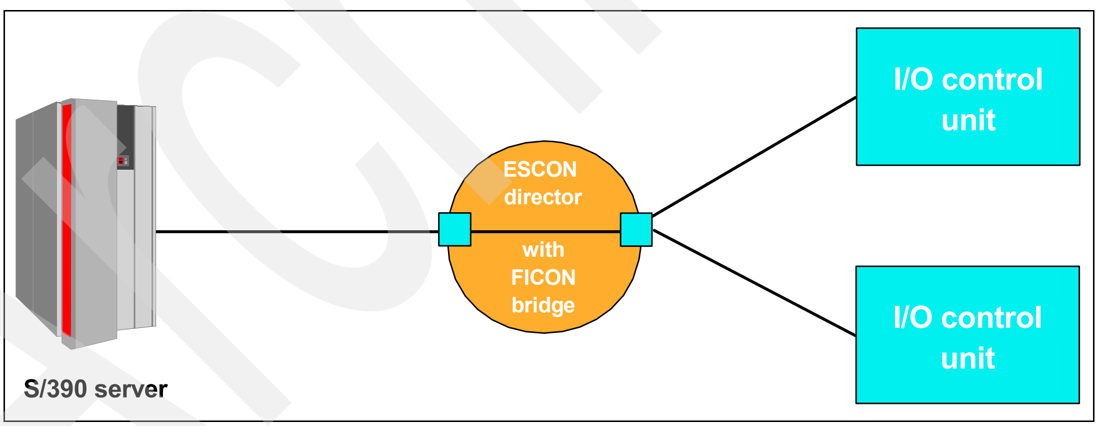

Figure A-24 Connection using ESCON director with FICON bridge

Using an ESCON channel extender

An ESCON channel extender is a signal repeater and fiber type converter that can be used to support ESCON-XDF configurations (Figure A-25). You can buy these extenders from various vendors active in this market segment.

The channel speed achieved is determined by the length of the channel. Measurement shows that a decrease of about 2 percent per km is experienced for data transfers of 32 KB blocks. For tape, the maximum distance between the host and the control unit should not exceed a distance of 43 km.

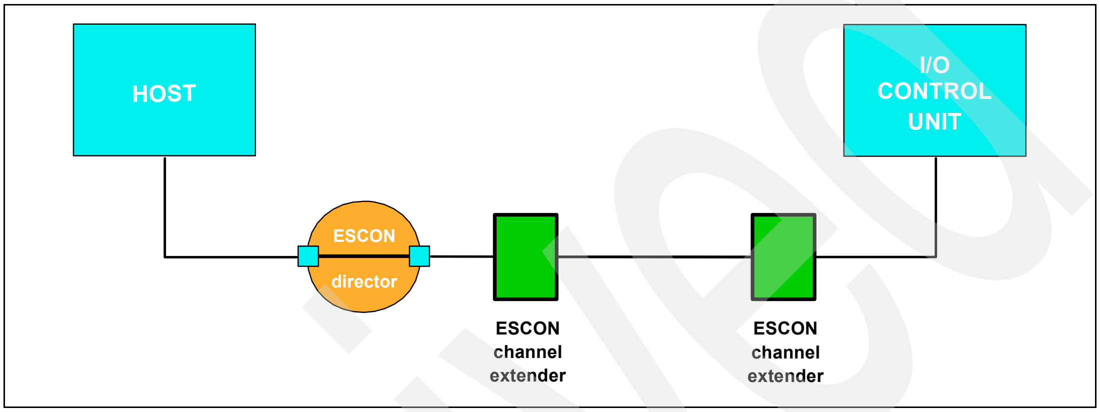

Figure A-25 Connection through ESCON channel extenders

The IBM 2029 Fiber Saver is supported for use with Peer-to-Peer VTS configurations. The maximum supported distance between a host and an VTC is 75 km. The maximum supported distance between a VTC and a VTS is 50 km.

Note: The discussion of channel extenders in this book is general in nature, and you should not assume that all channel extenders will work in a Peer-to-Peer VTS environment. New commands and protocols are used, and their correct use must be verified.

You should also be aware that some extenders may be certified for use between a host and VTCs but not between the VTCs and distributed VTSs.

Your channel extender vendor will be able to tell you whether particular models of extender have been tested and certified to work with a Peer-to-Peer VTS.

FICON channel connection with Dense Wave Division Multiplexer (DWDM)

The B10 and B20 PtP VTSs can attach to host systems through FICON channels. Short and long wavelength attachments provide for up to 100km distance between the VTS and the host or between the VTC and VTS using appropriate switches and DWDMs. Figure A-26 shows the numerous options for the physical distance between the VTS, VTC and the Host. FICON protocol buffer credits insure that performance is still very good at 100Km on multiple job runs. PtP VTSs are supported with al DWDMs.

Important: Coarse Wave Division Multiplexer (CWDM) and Time Division Multiplexer (TDM) are not supported with PtP VTS.

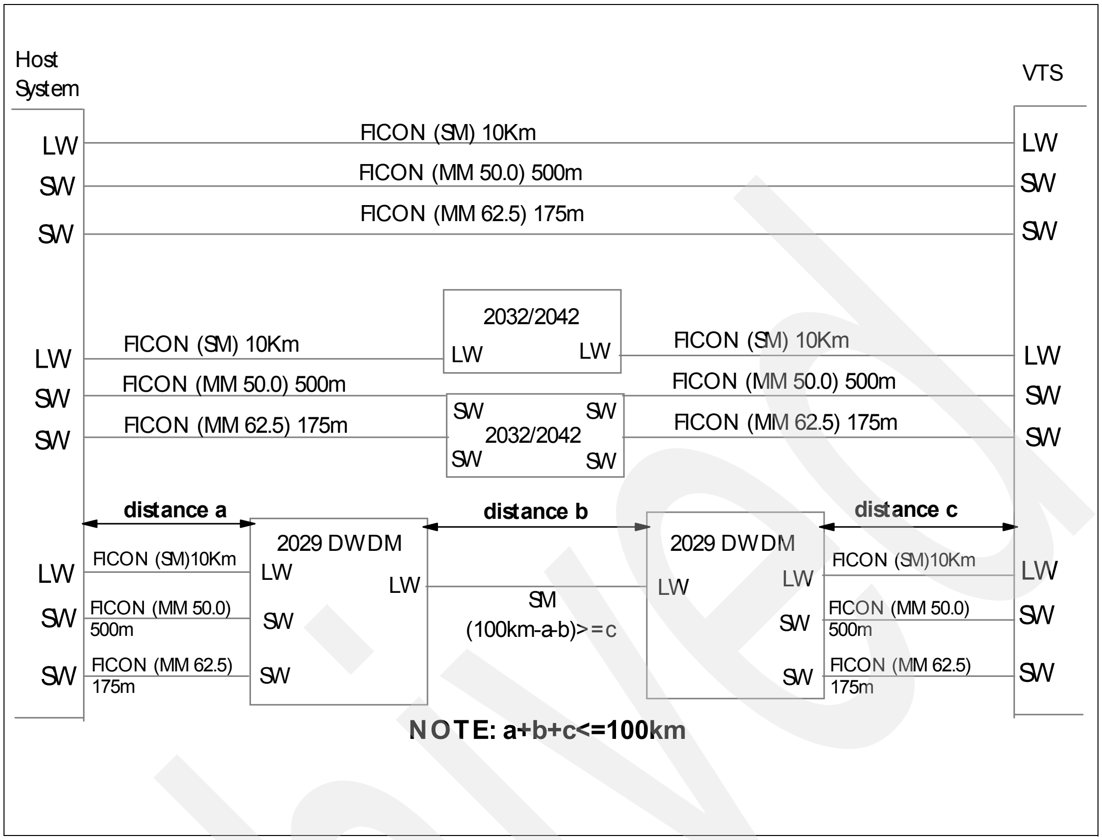

Figure A-26 FICON connectivity with distance

Using a remote channel extender

We use the term *remote channel extender* to group together equipment that provides switching and concentrator functions over telecommunication lines. The equipment sends the data from one or more ESCON links over one or more high speed data communication lines.

The data rates achieved are generally lower than with ESCON links. The technology used for the communication links is the limiting factor. It is, however, able to reach longer distances compared to ESCON channel extenders. The remote extenders are also a solution when ESCON links in a non-campus environment are not available or cannot be cost justified.

Although IBM does not provide remote channel extenders, we work together with various vendors implementing remote solutions at customer sites.

The justification for using remote channel extenders, the presence of the vendors, the costs and the availability of fiber technology in a non-campus environment, and the availability of high speed data communication lines varies from country to country. You may consult IBM certified specialists for designing the best solution in your environment, meeting your availability requirements within your budget limits (Figure A-27).

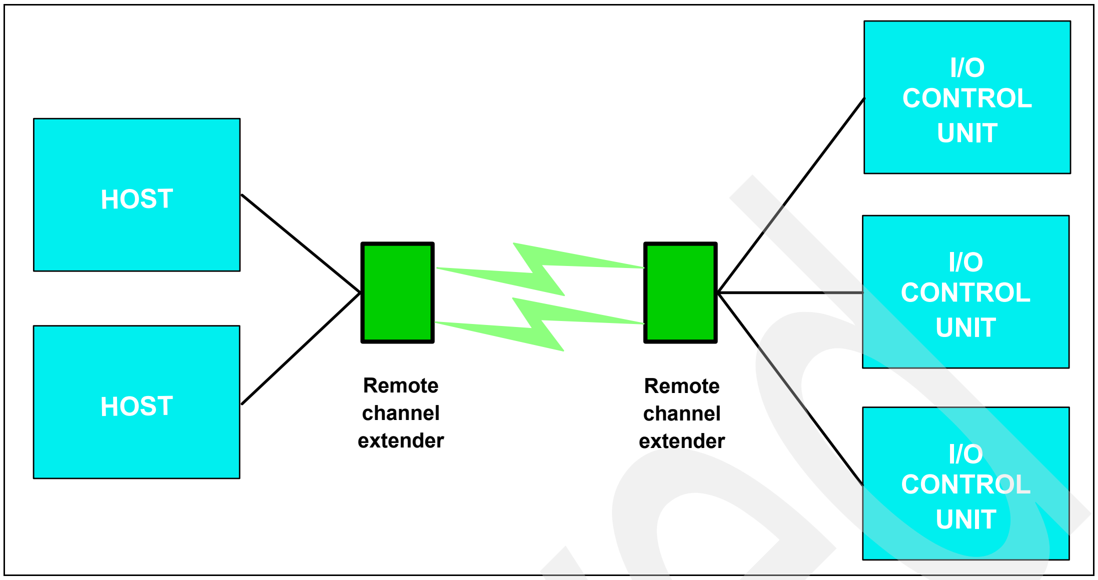

Figure A-27 Connection through a remote channel extender

Basic extender configurations

There are two cases when an extender may be needed (Figure A-28):

- Extending the distance between the host and the VTS (either base VTS or PtP VTS)
- Extending the PtP with an extender between the VTC and VTS

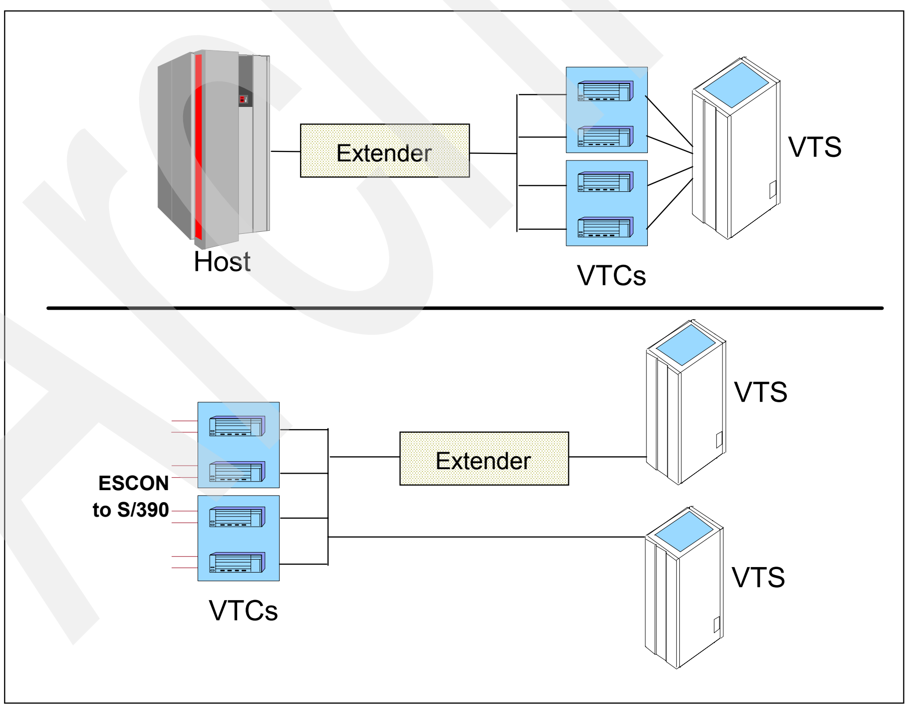

Figure A-28 Basic extender configurations

Before the VTS LIC 2.23 and LM 526 availability, only new ESCON commands existed between VTC and VTS in a Peer-to-Peer subsystem. The channel extension between host and VTC did not require any special extender code. But the VTC to VTS extension required special extender code and configurations.

The new LIC of VTS 2.23 and LM 526 introduces new ESCON commands from host to VTS/VTC and from VTC to VTS. Thus new extender code will be required to operate such an environment. The sequence of upgrades is dependent on the VTS to be upgraded, base VTS or Peer-to-Peer VTS, and the configuration of the extender.

Host extended configurations

The extender code always needs to be upgraded before the VTS or VTC code (Figure A-29).

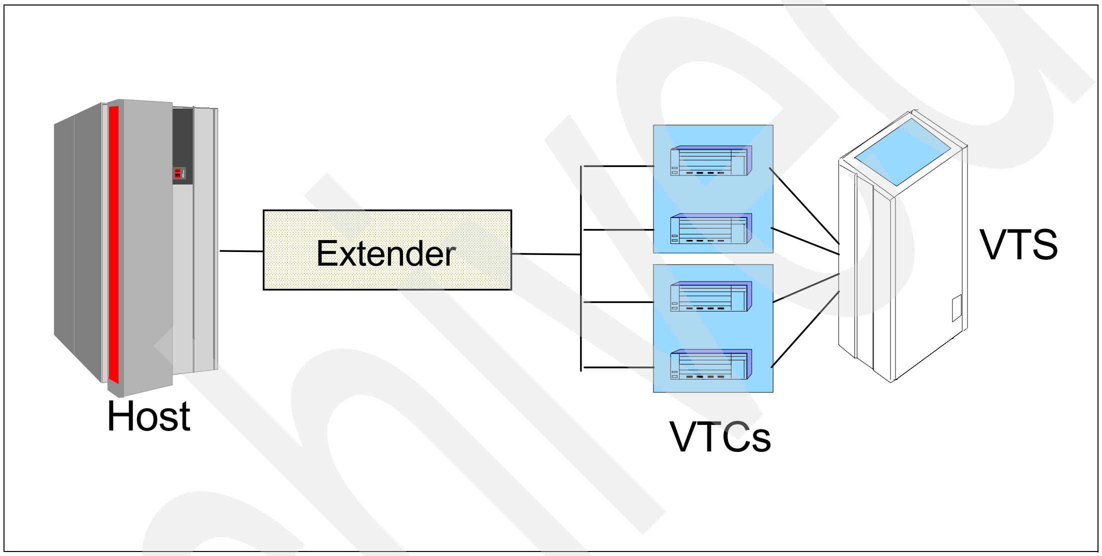

Figure A-29 Host extended configurations

PtP VTC extended configurations

In this environment (Figure A-30), the extender code should be upgraded before VTS and VTC. The VTC at VTS LIC level 2.26 and LM 527 are APM aware and will probe the VTS for APM support. They will receive a positive APM support to the probe, but the new commands will fail if the extender is not upgraded first.

The VTCs will use new APM commands independently of the host support level.

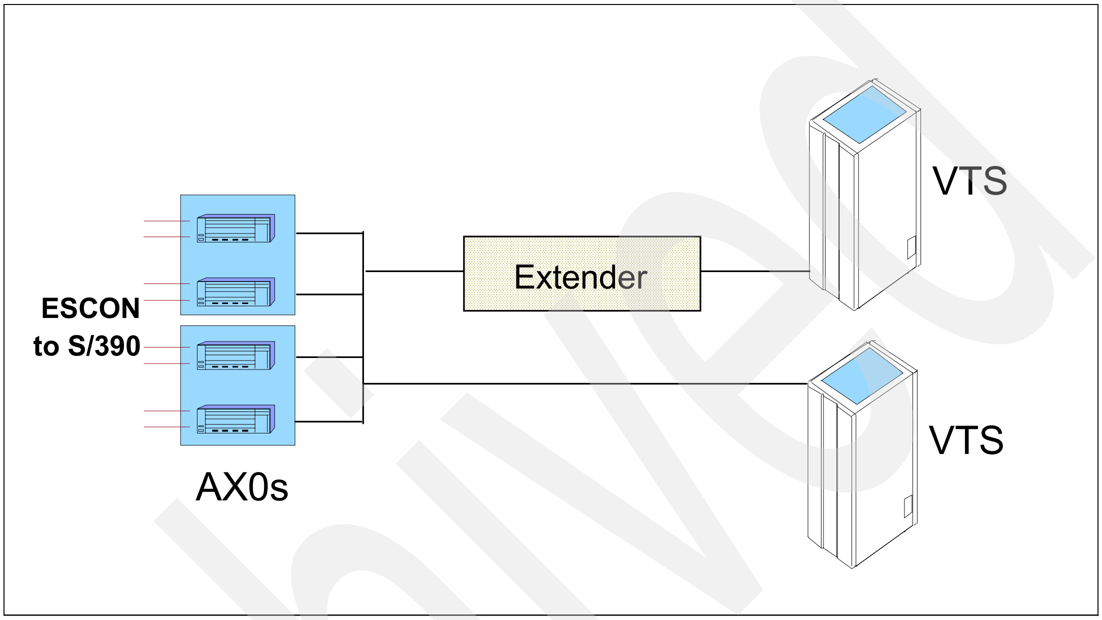

Figure A-30 VTC extended configurations

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 or OS/390 users may 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 which 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 Type E.
0003	DFSMS/MVS	Indicates scratch MEDIA3. MEDIA3 is the IBM 3590 High Performance Tape Cartridge Type J.
0004	DFSMS/MVS	Indicates scratch MEDIA4. MEDIA4 is the IBM 3590 Extended High Performance Tape Cartridge Type K.
0005	DFSMS/MVS	Indicates scratch MEDIA5. MEDIA5 is the IBM 3592 Enterprise Tape Cartridge Type JA.

Category (in hex)	Used by:	Definition
0006	DFSMS/MVS	Indicates scratch MEDIA6. MEDIA6 is the IBM 3592 Enterprise WORM Tape Cartridge Type JW.
0007	DFSMS/MVS	Indicates scratch MEDIA7. MEDIA7 is the IBM 3592 Enterprise Economy Cartridge Type JJ.
0008	DFSMS/MVS	Indicates scratch MEDIA8. MEDIA8 is the IBM 3592 Enterprise Economy WORM Tape Cartridge Type JR.
0009 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

Category (in hex)	Used by:	Definition
008D	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHD
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 VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH00
00A1	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH01
00A2	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH02
00A3	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH03
00A4	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH04
00A5	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH05
00A6	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH06
00A7	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH07
00A8	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH08
00A9	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH09
00AA	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH10
00AB	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH11
00AC	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH12
00AD	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH13
00AE	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH14
00AF	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH15
00B0	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH16

Category (in hex)	Used by:	Definition
00B1	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH17
00B2	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH18
00B3	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH19
00B4	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH20
00B5	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH21
00B6	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH22
00B7	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH23
00B8	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH24
00B9	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH25
00BA	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH26
00BB	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH27
00BC	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH28
00BD	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH29
00BE	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH30
00BF	Native VSE/ESA	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 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	ADSM for AIX	Indicates a private volume. Volumes in this category are managed by ADSM.

Category (in hex)	Used by:	Definition
012D	ADSM for AIX	Indicates an IBM 3490 scratch volume. Volumes in this category are managed by ADSM.
012E	ADSM for AIX	Indicates an IBM 3590 scratch volume. Volumes in this category are managed by ADSM.
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.
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.

Category (in hex)	Used by:	Definition
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 may be updated by operator interaction via Librarian Workstation Support.
FF01	Virtual Tape Server	Stacked Volume Insert category for a Virtual Tape Server. 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. 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. This category is used by the VTS for its scratch stacked volumes. This category is not used if LIC is 527 or higher.
FF04	Virtual Tape Server	Stacked Volume Private category for a Virtual Tape Server. This category is used by the VTS for its private stacked volumes. If LIC level is 527 or higher, this category includes both scratch and private volumes.
FF05	Virtual Tape Server	Stacked Volume Disaster Recovery category for a Virtual Tape Server. 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	Used by any VTS for its backup volumes.
FF07	Virtual Tape Server	Used by any VTS for its transaction log volumes.
FF08	Virtual Tape Server	This category is used by the VTS when it has determined that a stacked volume has an unreadable or invalid internal label.
FF09 to FF0F	-	Reserved for future hardware functions

Category (in hex)	Used by:	Definition
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. **Engineering Note:** If the Library Export operation is cancelled or fails, any volumes assigned to this category are re-assigned to the category they were in prior to the export operation.
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 were entered into the library via the convenience I/O station and placed in the Unassigned category.

Category (in hex)	Used by:	Definition
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. **Engineering Note:** If the Library Import operation is cancelled or fails, any volumes assigned to this category are deleted from the library inventory.
FF16	Virtual Tape Server	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 would be 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	Corrupted-Token Volume Category In a Peer to Peer VTS, volumes are assigned to this category by an AX0 controller when it has 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.
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.

Category (in hex)	Used by:	Definition
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 have been 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

DEVSERV QLIB command

This appendix lists the syntax and the parameter explanation 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.

> **Important:** Do not use this state save command for just testing purposes. It will impact the performance of your VTS/ATL, since it consumes time to take the dump in the hardware.

Example: C-1 APAR OA09599 text

```
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
          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

       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
```

<u>Parameters:</u>

```
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.
```

<u>Sub-Parameters:</u>

```
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

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).
```

D

SMF type 94 record layout

The composite PtP tape library, as well as each of the two distributed VTS tape libraries that interconnect to form the IBM PtP Tape Server, accumulate statistics over a period of one hour. These statistics represent the activity of the IBM Tape Library Dataserver resulting from all hosts attached to the PtP. This results in three different SMF type 94 records being calculated and recorded on the hour:

1. One record containing the statistical data for distributed VTS library 1
2. One record containing the statistical data for distributed VTS library 2
3. One record containing the statistical data for the composite PtP configuration

When FC4001 is installed and Outboard Policy Management is enabled, a new SMF94 record will be recorded. The new record is an SMF94 subtype 2, and will be recorded in addition to the existing SMF94 subtype 1. The SMF94 subtype 2 record is used to record Volume Pool Statistics.

The existing subtype 1 record is also changed. When FC4001 is installed and Outboard Policy Management is enabled, some existing statistics will not be reported. The unreported fields will contain binary zero. Existing field SMF94HSF in subtype 1 records will contain the value 2 when FC4001 is installed and Outboard Policy Management is enabled.

Record mapping

This section describes the mapping of the type 94 record.

Header and self-defining section

This section contains the common SMF record headers fields and the triplet fields (offset/length/number), if applicable, that locate the other sections on the record. The mapping macro IECSMF94 resides in SYS1.MACLIB.

The Table is updated with the latest changes of APAR OA10565, which is the support for VTS release 7.4.

The output shown below in Example D-1 is built by a macro exec which produces this kind of output. The advantage is that when the macro changes, the new output can be created much faster.

Example: D-1 Mapping of the type 94 record

```
Offsets     Type      Length Name            Description

==================================================================================
     MACRO NAME = IECSMF94
  PROPRIETARY V3 STATEMENT
  LICENSED MATERIALS - PROPERTY OF IBM
  "RESTRICTED MATERIALS OF IBM"
  5695-DF1
  (C) COPYRIGHT 1993 IBM CORP.
  END PROPRIETARY V3 STATEMENT
     STATUS = JDZ1110
     SMF record type 94(5E), SUBTYPE(01), is written approximately
     every hour when the 34xx tape library data server reports
     statistics. The 34xx reports statistics via a statistical
     data message which is read and recorded by the MVS tape error
     recovery program (module IGE0001E).
   CHANGE ACTIVITY =
   $01=OY60348,JDZ1110,SJPLLRU: Prevent mult SMF T94 writes
                                Related APAR: OY60349
   $02=OY62249 ,JDZ1110,SJPLLRU: User gets IEV025, IEV041, and
                                 IEV158 msgs during assembly
   $03=OY62842,JDZ1110,SJPLLRU: Refresh of OY62249, no changes
                                were made for this APAR
   $04=OW27369,HDZ11B0,TUCKEE: FCR0252 - added Virtual Tape
                                Server Statistics
   $IX=OW36360,HDZ11D0,SJPLRDB: VTS IMPORT/EXPORT SPE
                                FCR 258-03
   $XX=OW42071,HDZ11E0,SJPLRDB: enhanced VTS statistics
                                FRC267-02 gemini changes
   $05=OW52989,HDZ11F0,SJPLRDB: 8 AX0 support.
   $30=OW54059,HDZ11G0,SJPLRDB: Orca2(F/C 4001) statistics.
   $50=OA08777,HDZ11F0,SJPLRDB: FCR279 support
   $51=OA08777,HDZ11F0,SJPLRDB: FCR282 support
   $52=OA08777,HDZ11H0,SJPLRDB: FCR0401 VTS release 7.3
   $53=OA10565,HDZ11G0,SJPLRDB: FCR0406 VTS release 7.4
    SMF TYPE 94 RECORD, PLS DEFINITION
==================================================================================
```

```
   0    (0) STRUCTURE   124  SMFRCD94
   0    (0) CHARACTER    24  SMF94RPS       Type 94 Record Header
   0    (0) UNSIGNED      2  SMF94LEN       Record Length
   2    (2) UNSIGNED      2  SMF94SEG       Segment Descriptor
   4    (4) CHARACTER     1  SMF94FLG       Header flag byte
        1...  ....           SMF94FNM       ..Subsys name follows std hdr
        .1..  ....           SMF94FST       ..Subtypes utilized
        ..11  1111           *              ..reserved
   5    (5) CHARACTER     1  SMF94RTY       record type = 94 (hex 5E)
   6    (6) UNSIGNED      4  SMF94TME       Time since midnight,
                                            in hundredths of a second,
                                            when the record was moved into
                                            the SMF buffer.
  10    (A) CHARACTER     4  SMF94DTE       Date when the record was moved
                                            into the SMF buffer, in the
                                            form cyydddF (where F is the
                                            sign).

Offsets     Type      Length Name           Description

  14    (E) CHARACTER     4  SMF94SID       system id (from SMFPRMxx parm)
  18   (12) CHARACTER     4  SMF94WID       subsystem id, worktype indic.
  22   (16) CHARACTER     2  SMF94STP       record subtype 1 = 34xx
                                            Library Statistics 2 = volume
                                            pool statistics@30a see
                                            SMF94S2
==================================================================================
     SELF DEFINING SECTION (for format=1 statistics)
==================================================================================

  24   (18) CHARACTER   100  SMF94SDS       Self Defining Section
  24   (18) UNSIGNED      4  SMF94SDL       Self Defining Section length
  28   (1C) UNSIGNED      4  SMF94POF       Offset to product Section
  32   (20) UNSIGNED      2  SMF94PLN       Length of product Section
  34   (22) UNSIGNED      2  SMF94PON       Number of product Sections
  36   (24) UNSIGNED      4  SMF94HOF       Offset to header section
  40   (28) UNSIGNED      2  SMF94HLN       Length of header section
  42   (2A) UNSIGNED      2  SMF94HON       Number of header sections
  44   (2C) UNSIGNED      4  SMF94SOF       Offset to self descr info
  48   (30) UNSIGNED      2  SMF94SLN       Length of self descr info
  50   (32) UNSIGNED      2  SMF94SON       Number of self descr info
  52   (34) UNSIGNED      4  SMF94LOF       Offset to library stats
  56   (38) UNSIGNED      2  SMF94LLN       Length of library stats
  58   (3A) UNSIGNED      2  SMF94LON       Number of library stats
  60   (3C) UNSIGNED      4  SMF94MOF       Offset to mount stats
  64   (40) UNSIGNED      2  SMF94MLN       Length of mount stats
  66   (42) UNSIGNED      2  SMF94MON       Number of mount stats
  68   (44) UNSIGNED      4  SMF94DOF       Offset to demount stats
  72   (48) UNSIGNED      2  SMF94DLN       Length of demount stats
  74   (4A) UNSIGNED      2  SMF94DON       Number of demount stats
  76   (4C) UNSIGNED      4  SMF94EOF       Offset to eject stats
  80   (50) UNSIGNED      2  SMF94ELN       Length of eject stats
  82   (52) UNSIGNED      2  SMF94EON       Number of eject stats
  84   (54) UNSIGNED      4  SMF94AOF       Offset to audit stats
  88   (58) UNSIGNED      2  SMF94ALN       Length of audit stats
  90   (5A) UNSIGNED      2  SMF94AON       Number of audit stats
  92   (5C) UNSIGNED      4  SMF94IOF       Offset to insert stats
  96   (60) UNSIGNED      2  SMF94ILN       Length of insert stats
  98   (62) UNSIGNED      2  SMF94ION       Number of insert stats
 100   (64) UNSIGNED      4  SMF94VOF       Offset to VTS stats
```

```
  104   (68) UNSIGNED      2  SMF94VLN       Length of VTS stats
  106   (6A) UNSIGNED      2  SMF94VON       Number of VTS stats
  108   (6C) UNSIGNED      4  SMF94XOF       Offset to import/export
  112   (70) UNSIGNED      2  SMF94XLN       Length of import/export
  114   (72) UNSIGNED      2  SMF94XON       Number of import/export
  116   (74) UNSIGNED      4  SMF942OF       Offset to enhanced stats
  120   (78) UNSIGNED      2  SMF942LN       Length of enhanced stats
  122   (7A) UNSIGNED      2  SMF942ON       Number of enhanced stats
==================================================================================
     PRODUCT SECTION
==================================================================================


Offsets     Type      Length Name            Description

    0    (0) STRUCTURE    20  SMF94PSS
    0    (0) UNSIGNED      2  SMF94TYP       record subtype
    2    (2) CHARACTER     2  SMF94RVN       record version number = 01
    4    (4) CHARACTER     8  SMF94PNM       product name 'JDZ1110 '
   12    (C) CHARACTER     8  SMF94MVS       MVS operating system name
==================================================================================
    HEADER SECTION (bytes 9-11 in source record)
==================================================================================


Offsets     Type      Length Name            Description

    0    (0) STRUCTURE     3  SMF94HDR
    0    (0) UNSIGNED      1  SMF94HSF       Statistics format. Set to
                                             2 when F/C4001 is enabled.
    1    (1) CHARACTER     2  SMF94HHI       Hour index in binary,
                                             incremented every hour. Value
                                             ranges from 0 to 23 (17 hex).
==================================================================================
    SELF DESCRIPTION INFORMATION (bytes 12-37 in source record)
==================================================================================


Offsets     Type      Length Name            Description

    0    (0) STRUCTURE    27  SMF94SLF
    0    (0) CHARACTER     6  SMF94SLT       System-managed Tape Library
                                             type number. For example,
                                             003495 represents the 3495
                                             Tape Library Dataserver.
    6    (6) CHARACTER     3  SMF94SLM       System-managed Tape Library
                                             model number. For example, L30
                                             represents model L30.
    9    (9) CHARACTER     3  SMF94SMA       System-managed Tape Library
                                             manufacturer. Always equals
                                             IBM.
   12    (C) CHARACTER     2  SMF94SPL       System-managed Tape Library
                                             plant of manufacture. For
                                             example, 13 represents San
                                             Jose, California and 77
                                             represents Valencia, Spain.
   14    (E) CHARACTER    12  SMF94SNO       System-managed Tape Library
                                             sequence number. Uniquely
                                             identifies a system-managed
                                             tape library.
```

```
   26   (1A) CHARACTER     1  *              Reserved
==================================================================================
    LIBRARY STATISTICS (bytes 38-53 in source record)
==================================================================================


Offsets      Type      Length Name           Description

    0    (0) STRUCTURE    16  SMF94LIB
    0    (0) UNSIGNED      2  SMF94LID       Number of drives currently
                                             installed in a system-managed
                                             tape library.
    2    (2) UNSIGNED      2  SMF94LMD       Number of drives currently
                                             mounted in a system-managed
                                             tape library.
    4    (4) UNSIGNED      2  SMF94LM1       Maximum number of drives
                                             mounted during the last hour.
    6    (6) UNSIGNED      2  SMF94LM2       Minimum number of drives
                                             mounted during the last hour.
    8    (8) UNSIGNED      2  SMF94LM3       Average number of drives
                                             mounted during the last hour.
   10    (A) UNSIGNED      2  SMF94LT1       Maximum amount of time, in
                                             seconds, that a tape volume
                                             was mounted on a drive during
                                             the last hour. The mount time
                                             of a volume is the time when
                                             the system completed mounting
                                             a volume on a drive until the
                                             the time when the system
                                             managed tape library receives
                                             an order from the host to
                                             demount the volume.
   12    (C) UNSIGNED      2  SMF94LT2       Minimum amount of time, in
                                             seconds, that a tape volume
                                             was mounted on a drive during
                                             the last hour. The mount time
                                             of a volume is the time when
                                             the system completed mounting
                                             a volume on a drive until the
                                             time when the system managed
                                             tape library receives an order
                                             from the host to demount the
                                             volume.

Offsets      Type      Length Name           Description

   14    (E) UNSIGNED      2  SMF94LT3       Average amount of time, in
                                             seconds, that all tape volumes
                                             were mounted on drives during
                                             the last hour. The mount time
                                             of a volume is the time when
                                             the system completed mounting
                                             a volume on a drive until the
                                             time when the system managed
                                             tape library receives an order
                                             from the host to demount the
                                             volume.
==================================================================================
    MOUNT STATISTICS (bytes 54-73 in source record)
==================================================================================
```

```
Offsets     Type      Length Name           Description

    0    (0) STRUCTURE    20  SMF94MNT
    0    (0) UNSIGNED      2  SMF94MPR       The total number of mount
                                             requests currently pending.
    2    (2) UNSIGNED      2  SMF94MP1       Maximum number of mount
                                             requests pending during the
                                             last hour.
    4    (4) UNSIGNED      2  SMF94MP2       Minimum number of mount
                                             requests pending during the
                                             last hour.
    6    (6) UNSIGNED      2  SMF94MP3       Average number of mount
                                             requests pending during the
                                             last hour.
    8    (8) UNSIGNED      2  SMF94MTO       Total number of mounts during
                                             the last hour.
   10    (A) UNSIGNED      2  SMF94MIN       Index mounts during the
                                             last hour. An index
                                             mount is a mount accomplished
                                             using the automatic cartridge
                                             loader of a 3490 tape drive
   12    (C) UNSIGNED      2  SMF94MPM       Pre-mounts during last hour. A
                                             single pre-mount operation
                                             causes a volume to be added to
                                             the automatic cartridge loader
                                             of a 3490 tape drive.
   14    (E) UNSIGNED      2  SMF94MT1       Maximum amount of time, in
                                             seconds, required to perform
                                             any single mount operation
                                             during the last hour.
   16   (10) UNSIGNED      2  SMF94MT2       Minimum amount of time, in
                                             seconds, required to perform
                                             any single mount operation
                                             during the last hour.
   18   (12) UNSIGNED      2  SMF94MT3       Average amount of time, in
                                             seconds, required to perform
                                             any single mount operation
                                             during the last hour.
=================================================================================
    DEMOUNT STATISTICS (bytes 74-93 in source record)
=================================================================================


Offsets     Type      Length Name           Description

    0    (0) STRUCTURE    20  SMF94DMT
    0    (0) UNSIGNED      2  SMF94DPR       The total number of demount
                                             requests currently pending.
    2    (2) UNSIGNED      2  SMF94DP1       Maximum number of demount
                                             requests pending during the
                                             last hour.
    4    (4) UNSIGNED      2  SMF94DP2       Minimum number of demount
                                             requests pending during the
                                             last hour.
    6    (6) UNSIGNED      2  SMF94DP3       Average number of demount
                                             requests pending during the
                                             last hour.
    8    (8) UNSIGNED      2  SMF94DTO       Total number of demounts
```

```
                                               during the last hour.
   10     (A) UNSIGNED       2   SMF94DIN      Index demounts during the last
                                               hour. An index demount moves a
                                               volume from the feed station
                                               to the output stack of the
                                               automatic cartridge loader of
                                               a 3490 tape drive.
   12     (C) UNSIGNED       2   SMF94DPM      Post-demounts during the last
                                               hour. A post-demount operation
                                               moves a volume from the output
                                               stack of the automatic
                                               cartridge loader of a 3490
                                               tape drive.
   14     (E) UNSIGNED       2   SMF94DT1      Maximum amount of time, in
                                               seconds, required to perform
                                               any single demount operation
                                               during the last hour.
   16    (10) UNSIGNED       2   SMF94DT2      Minimum amount of time, in
                                               seconds, required to perform
                                               any single demount operation
                                               during the last hour.
   18    (12) UNSIGNED       2   SMF94DT3      Average amount of time, in
                                               seconds, required to perform
                                               any single demount operation
                                               during the last hour.

Offsets     Type      Length Name              Description

==================================================================================
    EJECT STATISTICS (bytes 94-109 in source record)
==================================================================================


Offsets     Type      Length Name              Description

    0     (0) STRUCTURE    16   SMF94EJT
    0     (0) UNSIGNED      2   SMF94EPR      The total number of eject
                                              requests currently pending. An
                                              eject operation moves one
                                              volume from the system-managed
                                              tape library to an output
                                              station for removal by an
                                              operator.
    2     (2) UNSIGNED      2   SMF94EP1      Maximum number of eject
                                              requests pending during the
                                              last hour.
    4     (4) UNSIGNED      2   SMF94EP2      Minimum number of eject
                                              requests pending during the
                                              last hour.
    6     (6) UNSIGNED      2   SMF94EP3      Average number of eject
                                              requests pending during the
                                              last hour.
    8     (8) UNSIGNED      2   SMF94ETO      Total number of ejects during
                                              the last hour.
   10     (A) UNSIGNED      2   SMF94ET1      Maximum amount of time, in
                                              seconds, required to perform
                                              any single eject operation
                                              during the last hour.
   12     (C) UNSIGNED      2   SMF94ET2      Minimum amount of time, in
                                              seconds, required to perform
```

```
                                             any single eject operation
                                             during the last hour.
   14    (E) UNSIGNED      2  SMF94ET3       Average amount of time, in
                                             seconds, required to perform a
                                             single eject operation during
                                             the last hour.
=================================================================================
    AUDIT STATISTICS (bytes 110-125 in source record)
=================================================================================


Offsets     Type      Length Name            Description

    0    (0) STRUCTURE    16  SMF94AUD
    0    (0) UNSIGNED      2  SMF94APR       The total number of audit
                                             requests currently pending.
                                             When the host requests an
                                             audit operation , the accessor
                                             moves to a shelf location and
                                             ensures that a volume is
                                             present.
    2    (2) UNSIGNED      2  SMF94AP1       Maximum number of audit
                                             requests pending during the
                                             last hour.
    4    (4) UNSIGNED      2  SMF94AP2       Minimum number of audit
                                             requests pending during the
                                             last hour.
    6    (6) UNSIGNED      2  SMF94AP3       Average number of audit
                                             requests pending during the
                                             last hour.
    8    (8) UNSIGNED      2  SMF94AT0       Total number of audits during
                                             the last hour.
   10    (A) UNSIGNED      2  SMF94AT1       Maximum amount of time, in
                                             seconds, required to perform
                                             any single audit operation
                                             during the last hour.
   12    (C) UNSIGNED      2  SMF94AT2       Maximum amount of time, in
                                             seconds, required to perform
                                             any single audit operation
                                             during the last hour.
   14    (E) UNSIGNED      2  SMF94AT3       Average amount of time, in
                                             seconds, required to perform a
                                             single audit operation during
                                             the last hour.
=================================================================================
    INPUT STATISTICS (bytes 126-127 in source record)
=================================================================================


Offsets     Type      Length Name            Description

    0    (0) STRUCTURE     2  SMF94INP
    0    (0) UNSIGNED      2  SMF94INS       Number of insert stores during
                                             last hour. This number is the
                                             number of volumes moved from
                                             an input station to a location
                                             inside the system managed tape
                                             library.
=================================================================================
   VIRTUAL TAPE SERVER STATS (bytes 128-211 in source record )
```

```
   The VTS statistical record contains information about physical
   drive usage and physical mount performance of the library.
   The current record has a length of 256 bytes, of which only
   the first 128 bytes contain defined values. When one or more
   VTS subsystems are installed in the library, the existing
   statistics continue to reflect all physical movement of
   cartridges within the physical library including that of the
   VTS subsystems. The virtual mount statistics and operational
   information about the VTS are reported in bytes 128:211.
   Although up to two VTS subsystems can be configured in a
   library, the statistics reported in bytes 128:211 are only
   for the specific VTS subsystem the request was issued to.
   For statistical data requests issued to a non-VTS device,
   bytes 128:211 contain X'00'.
==================================================================================


Offsets     Type      Length Name           Description

    0    (0) STRUCTURE    84  SMF94VTS
    0    (0) CHARACTER     1  SMF94VNO       cluster lib number (128)
    1    (1) CHARACTER     5  SMF94VLS       Library sequence number for
==================================================================================
        library segment for which VTS statistics are being reported.
==================================================================================

    6    (6) UNSIGNED      1  SMF94VTI       Number of underlying physical
==================================================================================
        tape devices currently installed in the VTS subsystem.
        Contains zero with F/C4001
==================================================================================

    7    (7) UNSIGNED      1  SMF94VTA       Number of underlying physical
==================================================================================
        tape devices currently available for use by VTS. Contains
        zero with F/C4001
==================================================================================

    8    (8) UNSIGNED      1  SMF94VTX       Maximum number of underlying
==================================================================================
        physical tape devices mounted concurrently in this VTS
        during last hour. Contains zero with F/C4001
==================================================================================

    9    (9) UNSIGNED      1  SMF94VTN       Minimum number of underlying
==================================================================================
        physical tape devices mounted concurrently in this VTS
        during last hour. Contains zero with F/C4001
==================================================================================

   10    (A) UNSIGNED      1  SMF94VTV       Average number of underlying
==================================================================================
        physical tape devices mounted concurrently in the VTS during
        last hour. Value is determined by summing number of concur-
        rently mounted physical drives every 10 seconds, and
        dividing resultant sum by 360 during hourly statistics
        generation. Contains zero with F/C4001
==================================================================================

   11    (B) CHARACTER     1  SMF94VR2       Reserved, set to X'00'
```

```
   12    (C) UNSIGNED      2  SMF94VMX       Maximum time in seconds used
==================================================================================
        to perform a mount on a physical drive in VTS in the last
        hour. Time is accrued from time mount request is accepted
        until it is complete. Mount time is accredited to hour that
        mount completes. Contains zero with F/C4001
==================================================================================

   14    (E) UNSIGNED      2  SMF94VMN       Minimum time in seconds used
==================================================================================
        to perform a mount on a physical drive in VTS in the last
        hour. Time is accrued from time mount request is accepted
        until it is complete. Mount time is accredited to hour that
        mount completes. Contains zero with F/C4001
==================================================================================

   16   (10) UNSIGNED      2  SMF94VMV       Average time in seconds used
==================================================================================
        to perform a mount on a physical drive in VTS in the last
        hour. Time is accrued from time mount request is accepted
        until it is complete. Mount time is accredited to hour that
        mount completes. Contains zero with F/C4001
==================================================================================


Offsets     Type      Length Name            Description

   18   (12) UNSIGNED      2  SMF94VPS       The number of physical mount
==================================================================================
        requests completed in last hour to satisfy stage mounts.
        Contains zero with F/C4001
==================================================================================

   20   (14) UNSIGNED      2  SMF94VPM       The number of physical mounts
==================================================================================
        completed in last hour to satisfy migrate requests. Contains
        zero with F/C4001
==================================================================================

   22   (16) UNSIGNED      2  SMF94VPR       The number of physical mounts
==================================================================================
        completed in last hour to satisfy reclamation mounts.
        Contains zero with F/C4001
==================================================================================

   24   (18) UNSIGNED      1  SMF94VDC       The number of virtual devices
==================================================================================
        configured in this VTS at the time request for statistics
        was received (current.) Contains zero with F/C4001
==================================================================================

   25   (19) UNSIGNED      1  SMF94VDX       The maximum number of virtual
==================================================================================
        drives that were concurrently mounted in this VTS during the
        last hour. Contains zero with F/C4001
==================================================================================

   26   (1A) UNSIGNED      1  SMF94VDN       The minimum number of virtual
==================================================================================
        drives that were concurrently mounted in this VTS during the
```

```
       last hour. Contains zero with F/C4001
===================================================================================

   27   (1B) UNSIGNED      1  SMF94VDA       The average number of virtual
===================================================================================
       drives that were concurrently mounted in the VTS during last
       hour. Value is determined by summing number of concurrently
       mounted virtual devices every 10 seconds, and dividing
       resultant sum by 360 during hourly statistics generation.
       Contains zero with F/C4001
===================================================================================

   28   (1C) UNSIGNED      2  SMF94VVX       Maximum time in seconds that a
===================================================================================
       virtual drive was mounted in VTS during the last hour. Time
       is accrued from completion of mount until demount is issued.
       Mount time is accredited to hour that demount is issued.
===================================================================================

   30   (1E) UNSIGNED      2  SMF94VVN       Minimum time in seconds that a
===================================================================================
       virtual drive was mounted in VTS during the last hour. Time
       is accrued from completion of mount until demount is issued.
       Mount time is accredited to hour that demount is issued.
===================================================================================

   32   (20) UNSIGNED      2  SMF94VVA       Average time in seconds that a
===================================================================================
       virtual drive was mounted in VTS during the last hour. Time
       is accrued from completion of mount until demount is issued.
       Mount time is accredited to hour that demount is issued.
===================================================================================

   34   (22) UNSIGNED      2  SMF94VRX       Maximum time in seconds for a

Offsets     Type      Length Name            Description

===================================================================================
       mount to complete on a virtual drive in the last hour. Time
       is accrued from time mount request is accepted until it is
       completed. Mount time is accredited to the hour that mount
       is completed.
===================================================================================

   36   (24) UNSIGNED      2  SMF94VRN       Minimum time in seconds for a
===================================================================================
       mount to complete on a virtual drive in the last hour. Time
       is accrued from time mount request is accepted until it is
       completed. Mount time is accredited to the hour that mount
       is completed.
===================================================================================

   38   (26) UNSIGNED      2  SMF94VRA       Average time in seconds for a
===================================================================================
       mount to complete on a virtual drive in the last hour. Time
       is accrued from time mount request is accepted until it is
       completed. Mount time is accredited to the hour that mount
       is completed.
===================================================================================
```

```
   40   (28) UNSIGNED      2  SMF94VFR       The number of virtual mounts
==================================================================================
        in last hour using VTS Fast- Ready facility. Fast-Ready is
        used for mount-from-category requests for which specified
        category has the Fast-Ready attribute set, or for specific
        volume requests for which the specified volume is, at time
        mount request is received, assigned to a category that has
        Fast-Ready attribute set.
==================================================================================

   42   (2A) UNSIGNED      2  SMF94VMH       The number of virtual mounts
==================================================================================
        last hour that completed for specific requested volumes
        found resident in Tape Volume Cache (specific mount hits.)
==================================================================================

   44   (2C) UNSIGNED      2  SMF94VMS       The number of virtual mounts
==================================================================================
        last hour completed with specific requested logical volumes
        staged from a physical tape back to Tape Volume Cache
==================================================================================

   46   (2E) UNSIGNED      2  SMF94VMP       The number of virtual volumes
==================================================================================
        for which premigrate was completed in the last hour.
==================================================================================

   48   (30) UNSIGNED      4  SMF94VBW       Total number of bytes written
==================================================================================
        successfully through host channels to virtual volumes in an
        integral multiple of 4096 bytes during the last hour. If
        number of bytes written is not an integer multiple of 4096,
        the number is rounded up.
==================================================================================

   52   (34) UNSIGNED      4  SMF94VBR       Total number of bytes read
==================================================================================
        successfully through host channels from virtual volumes in
        integral multiple of 4096 bytes during the last hour. If
        number of bytes read is not an integer multiple of 4096, the
        number is rounded up.
==================================================================================

   56   (38) UNSIGNED      4  SMF94VTW       Total number of bytes written

Offsets     Type      Length Name           Description

==================================================================================
        successfully by VTS to its attached physical drives in
        integral mulitple of 4096 bytes during last hour. If number
        of bytes written is not an integer multiple of 4096, the
        number is rounded up. Bytes are accredited to hour in which
        the underlying premigrates of virtual volumes complete.
        Contains zero with F/C4001
==================================================================================

   60   (3C) UNSIGNED      4  SMF94VTR       Total number of bytes read
==================================================================================
        successfully by VTS from its attached physical drives in an
        integral mulitple of 4096 bytes during last hour. If number
```

```
       of bytes read is not an integer multiple of 4096, the number
       is rounded up. Bytes are accredited to the hour in which the
       underlying premigrates of virtual volumes complete. Contains
       zero with F/C4001
=================================================================================

   64   (40) UNSIGNED      2  SMF94VCA       The average age, in minutes,
                                             of
=================================================================================
       the times of last reference of the vitual volumes in the
       tape Volume Cache at the end of the reported hour. Contains
       zero with F/C4001
=================================================================================

   66   (42) UNSIGNED      2  SMF94VCZ       The average size (mb) of the
=================================================================================
       virtual volumes in the Tape Volume Cache at the end of the
       reported hour. Contains zero with F/C4001
=================================================================================

   68   (44) UNSIGNED      2  SMF94VNM       The number of virtual volumes
=================================================================================
       in Tape Volume Cache at the end of reported hour. Contains
       zero with F/C4001
=================================================================================

   70   (46) CHARACTER     2  SMF94VR3       Reserved, set to X'00'
   72   (48) UNSIGNED      4  SMF94VBA       The number of bytes of data
=================================================================================
       on the active logical volumes which are on VTS stacked
       volumes at the end of the reported hour. (200:203)
=================================================================================

   76   (4C) UNSIGNED      4  SMF94VLA       The number of active logical
=================================================================================
       volumes which are on VTS stacked volumes at the end of the
       reported hour.(204:207)
=================================================================================

   80   (50) UNSIGNED      4  SMF94VEC       The total estimated amount
=================================================================================
       storage capacity of the empty 3590 cartridges managed by the
       VTS in an integral multiple of 1,048,576 bytes (1 MBytes) at
       the end of the reported hour
=================================================================================

=================================================================================
   VTS Import/Export STATS (bytes 212-239 in source record )
   The VTS Import/Export record contains informationa about 1 hours
   import/export activity in a VTS library.
=================================================================================


Offsets     Type      Length Name            Description

    0    (0) STRUCTURE    44  SMF94X00
    0    (0) UNSIGNED      2  SMF94IM1       count of physical volumes
                                             processed during import oper-
                                             ations that completed in the
                                             last hour. (212:213)
```

```
     2    (2) UNSIGNED      2  SMF94EX1       count of the number of
                                               physical volumes that contain
                                               the successfully exported
                                               logical volumes exported
                                               during the last hour.
                                               (214:215)
     4    (4) UNSIGNED      4  SMF94IM2       count of the number of logical
                                               volumes successfully imported
                                               during import operations that
                                               completed during the last
                                               hour. (216:219)
     8    (8) UNSIGNED      4  SMF94EX2       count of the number of logical
                                               volumes that were successfully
                                               exported for export operations
                                               completed in the last hour.
                                               (220:223)
    12    (C) UNSIGNED      4  SMF94IM3       megabytes of data imported for
                                               import operations that
                                               completed in the last hour.
                                               (224:227)
    16   (10) UNSIGNED      4  SMF94EX3       megabytes of data exported
                                               during export operation that
                                               completed in the last hour.
                                               (228:231)
    20   (14) UNSIGNED      4  SMF94IM4       megabytes of data that was
                                               moved from one physical
                                               stacked volume to another as
                                               part of the import operations
                                               completed in the last hour.
                                               (232:235)

 Offsets     Type      Length Name           Description

    24   (18) UNSIGNED      4  SMF94EX4       megabytes moved from one
                                               physical stacked volume to
                                               another as part of the export
                                               operations completed in the
                                               last hour. (236:239)
    28   (1C) CHARACTER     8  *              Reserved, set to x'00'
                                               (240:247)
    36   (24) UNSIGNED      2  SMF94ACA       Accessor A mounts. The count
                                               of the number of mount
                                               operations accessor A
                                               completed during the last
                                               hour. (248:249)
    38   (26) UNSIGNED      2  SMF94ACB       Accessor B mounts. The count
                                               of the number of mount
                                               operations accessor A
                                               completed during the last
                                               hour. (250:251)
    40   (28) CHARACTER     4  *              Reserved, set to x'00'
                                               (252:255)
=================================================================================
   VTS Enhanced statistics (bytes 256-xxx in source record )
   The VTS Enhanced statistics contain information about 1 hours
   activity in a VTS library.
=================================================================================


 Offsets     Type      Length Name           Description
```

```
  0     (0) STRUCTURE  1132  S94STATS
  0     (0) UNSIGNED      2  S94BSRAT       256 backstore comp. ratio
  2     (2) UNSIGNED      2  S94HARAT       258 adapter comp. ratio
  4     (4) UNSIGNED      2  S94TVCS        260 tape volume cache size
  6     (6) UNSIGNED      1  S94ESCON       262 number of escon chans
  7     (7) UNSIGNED      1  S94SCSI        263 number of scsi chans
  8     (8) UNSIGNED      4  S94NUMBS       264 channel blocks written
 12     (C) UNSIGNED      1  S940KB         268 0 to 2K
 13     (D) UNSIGNED      1  S942KB         269 2k to 4k
 14     (E) UNSIGNED      1  S944KB         270 4k to 8k
 15     (F) UNSIGNED      1  S948KB         271 8k to 16k
 16    (10) UNSIGNED      1  S9416KB        272 16k to 32k
 17    (11) UNSIGNED      1  S9432KB        273 32k to 64k
 18    (12) UNSIGNED      1  S9464KB        274 greater than 64k
 19    (13) UNSIGNED      1  S94RCPRT       275 recall predominate
 20    (14) UNSIGNED      1  S94WROVT       276 write overrun
 21    (15) UNSIGNED      1  *              277 reserved
 22    (16) CHARACTER     2  *              278 reserved (for allignment)
 24    (18) UNSIGNED      4  S94AVRCT       280 average recall
 28    (1C) UNSIGNED      4  S94AVWOT       284 average write overrun
 32    (20) UNSIGNED      4  *              288 reserved
 36    (24) UNSIGNED      4  S94TOTAT       292 overall
 40    (28) UNSIGNED      2  S94MAXFR       296 max fast ready mount time
 42    (2A) UNSIGNED      2  S94MINFR       298 min fast ready mount time
 44    (2C) UNSIGNED      2  S94AVGFR       300 avg fast ready mount time
 46    (2E) UNSIGNED      2  S94MAXCH       302 max cache hit mount time
 48    (30) UNSIGNED      2  S94MINCH       304 min cache hit mount time
 50    (32) UNSIGNED      2  S94AVGCH       306 avg cache hit mount time
 52    (34) UNSIGNED      2  S94MAXRM       308 max recall mnt mount time
 54    (36) UNSIGNED      2  S94MINRM       310 min recall mnt mount time
 56    (38) UNSIGNED      2  S94AVGRM       312 avg recall mnt mount time
 58    (3A) UNSIGNED      2  S94ADV05       314 05% active data volume
 60    (3C) UNSIGNED      2  S94ADV10       316 10% active data volume
 62    (3E) UNSIGNED      2  S94ADV15       318 15% active data volume
 64    (40) UNSIGNED      2  S94ADV20       320 20% active data volume
 66    (42) UNSIGNED      2  S94ADV25       322 25% active data volume
 68    (44) UNSIGNED      2  S94ADV30       324 30% active data volume
 70    (46) UNSIGNED      2  S94ADV35       326 35% active data volume
 72    (48) UNSIGNED      2  S94ADV40       328 40% active data volume
 74    (4A) UNSIGNED      2  S94ADV45       330 45% active data volume
 76    (4C) UNSIGNED      2  S94ADV50       332 50% active data volume
 78    (4E) UNSIGNED      2  S94ADV55       334 55% active data volume
 80    (50) UNSIGNED      2  S94ADV60       336 60% active data volume
 82    (52) UNSIGNED      2  S94ADV65       338 65% active data volume
 84    (54) UNSIGNED      2  S94ADV70       340 70% active data volume
 86    (56) UNSIGNED      2  S94ADV75       342 75% active data volume
 88    (58) UNSIGNED      2  S94ADV80       344 80% active data volume
 90    (5A) UNSIGNED      2  S94ADV85       346 85% active data volume
 92    (5C) UNSIGNED      2  S94ADV90       348 90% active data volume
 94    (5E) UNSIGNED      2  S94ADV95       350 95% active data volume
 96    (60) UNSIGNED      2  S94ADV00       352 100% active data volume
 98    (62) UNSIGNED      1  S94THRES       354 reclaim threshold
 99    (63) UNSIGNED      1  *              355 reserved (set to zero)
100    (64) UNSIGNED      2  S94SRTCT       356 scratch stacked volume ct
102    (66) UNSIGNED      2  S94PRICT       358 private stacked volume ct
104    (68) UNSIGNED      4  S94MTVCA       360 max tape vol cache age

Offsets     Type      Length Name           Description
```

```
  108   (6C) BITSTRING     1  S94CMGTS       364 cache mang. settings
         1...  ....           S94VALPL       valid bit - indicates that the
                                             following sub-fields contain
                                             valid data.
         .111  ....           S94CVMPL       copy volume management
                                             preference level(0-7)
         ....  1...           *              reserved(set to zero)
         ....  .111           S94RVMPL       recall volume management
                                             preference level(0-7)
  109   (6D) CHARACTER     3  *              365:367 reserved
=================================================================================
        The following fields report the code levels of the VTS
        and Library Manager that generated the statistical
        record.
=================================================================================

  112   (70) UNSIGNED      2  S94LVVCM       368:369 VTS code modification
                                             value.
  114   (72) UNSIGNED      2  S94LVVCF       370:371 VTS code fix value
  116   (74) UNSIGNED      2  S94LVLMV       372:373 Library Manager code
                                             version value.
  118   (76) UNSIGNED      2  S94LVLMR       374:375 Library Manager code
                                             release value.
  120   (78) CHARACTER     8  *              376:383 reserved
=================================================================================
        Composite library Statistics
=================================================================================

  128   (80) UNSIGNED      4  S94CLLVC       384 logical vols to be copied
  132   (84) UNSIGNED      4  S94CLDTC       388 data to copy
  136   (88) UNSIGNED      2  S94CLMT0       392 mounts for VTS-0
  138   (8A) UNSIGNED      2  S94CLMT1       394 mounts for VTS-1
  140   (8C) CHARACTER     4  *              396 reserved
=================================================================================
        400:655 contains report for up to 8 AX0's in composite lib.
        Changes to the following fields for FCR279 also include
        field definitions at the end of this macro(flagged @50).
        Statistics for AX0(0)
=================================================================================

  144   (90) UNSIGNED      4  S94CLDC0       0:3 data copied by AX0
  148   (94) UNSIGNED      4  S94CLVC0       4:7 volumes copied by AX0
  152   (98) UNSIGNED      4  S94CLRD0       8:11 read data transferred
  156   (9C) UNSIGNED      4  S94CLWD0       12:15 write data transferred
  160   (A0) UNSIGNED      2  S94CLCM0       16:17 category mounts for AX0
  162   (A2) UNSIGNED      2  S94CLSM0       18:19 specific cache mounts
  164   (A4) UNSIGNED      2  S94CLRM0       20:21 specific recall mounts
  166   (A6) UNSIGNED      2  S94CLCR0       22:23 Compression ratio
  168   (A8) BITSTRING     1  *              24:24 Distributed library
                                             Preferences
         1111  ....           S94CLPI0       Preferred I/O VTS
         ....  1111           S94CLPM0       Preferred Master VTS

Offsets     Type      Length Name            Description

  169   (A9) BITSTRING     2  *              25:26 Configured Settings.
                                             The reported settings@50a in
                                             this 2 byte field
                                             are only valid if
                                             bit S94CLVA0 is set.
```

```
        11.. ....          S94CLDM0       default copy mode
        ..1. ....          S94CLFS0       force scratch to the
                                          preferred VTS.
        ...1 ....          S94CLPV0       VTS I/O selection
                                          criteria is PRIMARY
        .... 11..          S94CLCO0       Controller
                                          Operational Mode
        .... ..1.          S94CLWP0       Write Protected Mode
  169   (A9) BITSTRING     1  *           reserved(set to zero)@53c
        .... ...1          S94CLVA0       field valid indicator@50a
  171   (AB) UNSIGNED      1  *           27:27 relative link speeds
        1111 ....          S94CLLS0_AX0VTS0
                                          AX0 to VTS0 linkspeed
        .... 1111          S94CLLS0_AX0VTS1
                                          AX0 to VTS1 linkspeed
  172   (AC) CHARACTER     4  *           28:31 reserved(set to
                                          zero)@50a
=================================================================================
        Statistics for AX0(1)
=================================================================================

  176   (B0) UNSIGNED      4  S94CLDC1    0:3 data copied by AX0
  180   (B4) UNSIGNED      4  S94CLVC1    4:7 volumes copied by AX0
  184   (B8) UNSIGNED      4  S94CLRD1    8:11 read data transferred
  188   (BC) UNSIGNED      4  S94CLWD1    12:15 write data transferred
  192   (C0) UNSIGNED      2  S94CLCM1    16:17 category mounts for AX0
  194   (C2) UNSIGNED      2  S94CLSM1    18:19 specific cache mounts
  196   (C4) UNSIGNED      2  S94CLRM1    20:21 specific recall mounts
  198   (C6) UNSIGNED      2  S94CLCR1    22:23 Compression ratio
  200   (C8) BITSTRING     1  *           24:24 Distributed library
                                          Preferences
        1111 ....          S94CLPI1       Preferred I/O VTS
        .... 1111          S94CLPM1       Preferred Master VTS
  201   (C9) BITSTRING     2  *           25:26 Configured Settings.
                                          The reported settings@50a in
                                          this 2 byte field
                                          are only valid if
                                          bit S94CLVA1 is set.
        11.. ....          S94CLDM1       default copy mode
        ..1. ....          S94CLFS1       force scratch to the
                                          preferred VTS.
        ...1 ....          S94CLPV1       VTS I/O selection
                                          criteria is PRIMARY
        .... 11..          S94CLCO1       Controller
                                          Operational Mode
        .... ..1.          S94CLWP1       Write Protected Mode
  201   (C9) BITSTRING     1  *           reserved(set to zero)@53c
        .... ...1          S94CLVA1       field valid indicator@50a
  203   (CB) UNSIGNED      1  *           27:27 relative link speeds

Offsets     Type      Length Name         Description

        1111 ....          S94CLLS1_AX0VTS0
                                          AX0-VTS0 link speed
        .... 1111          S94CLLS1_AX0VTS1
                                          AX0-VTS1 link speed
  204   (CC) CHARACTER     4  *           28:31 reserved(set to
                                          zero)@50a
=================================================================================
        Statistics for AX0(2)
```

```
=================================================================================

  208   (D0) UNSIGNED      4  S94CLDC2       0:3 data copied by AX0
  212   (D4) UNSIGNED      4  S94CLVC2       4:7 volumes copied by AX0
  216   (D8) UNSIGNED      4  S94CLRD2       8:11 read data transferred
  220   (DC) UNSIGNED      4  S94CLWD2       12:15 write data transferred
  224   (E0) UNSIGNED      2  S94CLCM2       16:17 category mounts for AX0
  226   (E2) UNSIGNED      2  S94CLSM2       18:19 specific cache mounts
  228   (E4) UNSIGNED      2  S94CLRM2       20:21 specific recall mounts
  230   (E6) UNSIGNED      2  S94CLCR2       22:23 Compression ratio
  232   (E8) BITSTRING     1  *              24:24 Distributed library
                                             Preferences
         1111  ....           S94CLPI2       Preferred I/O VTS
         ....  1111           S94CLPM2       Preferred Master VTS
  233   (E9) BITSTRING     2  *              25:26 Configured Settings.
                                             The reported settings@50a in
                                             this 2 byte field
                                             are only valid if
                                             bit S94CLVA2 is set.
         11..  ....           S94CLDM2       default copy mode
         ..1.  ....           S94CLFS2       force scratch to the
                                             preferred VTS.
         ...1  ....           S94CLPV2       VTS I/O selection
                                             criteria is PRIMARY
         ....  11..           S94CLCO2       Controller
                                             Operational Mode
         ....  ..1.           S94CLWP2       Write Protected Mode
  233   (E9) BITSTRING     1  *              reserved(set to zero)@53c
         ....  ...1           S94CLVA2       field valid indicator@50a
  235   (EB) UNSIGNED      1  *              27:27 relative link speeds
         1111  ....           S94CLLS2_AX0VTS0
                                             AX0-VTS0 link speed
         ....  1111           S94CLLS2_AX0VTS1
                                             AX0-VTS1 link speed
  236   (EC) CHARACTER     4  *              28:31 reserved(set to
                                             zero)@50a
=================================================================================
        Statistics for AX0(3)
=================================================================================

  240   (F0) UNSIGNED      4  S94CLDC3       0:3 data copied by AX0
  244   (F4) UNSIGNED      4  S94CLVC3       4:7 volumes copied by AX0
  248   (F8) UNSIGNED      4  S94CLRD3       8:11 read data transferred
  252   (FC) UNSIGNED      4  S94CLWD3       12:15 write data transferred
  256  (100) UNSIGNED      2  S94CLCM3       16:17 category mounts for AX0
  258  (102) UNSIGNED      2  S94CLSM3       18:19 specific cache mounts
  260  (104) UNSIGNED      2  S94CLRM3       20:21 specific recall mounts

Offsets     Type      Length Name            Description

  262  (106) UNSIGNED      2  S94CLCR3       22:23 Compression ratio
  264  (108) BITSTRING     1  *              24:24 Distributed library
                                             Preferences
         1111  ....           S94CLPI3       Preferred I/O VTS
         ....  1111           S94CLPM3       Preferred Master VTS
  265  (109) BITSTRING     2  *              25:26 Configured Settings.
                                             The reported settings@50a in
                                             this 2 byte field
                                             are only valid if
                                             bit S94CLVA3 is set.
```

```
        11..  ....          S94CLDM3       default copy mode
        ..1.  ....          S94CLFS3       force scratch to the
                                           preferred VTS.
        ...1  ....          S94CLPV3       VTS I/O selection
                                           criteria is PRIMARY
        ....  11..          S94CLCO3       Controller
                                           Operational Mode
        ....  ..1.          S94CLWP3       Write Protected Mode
  265  (109) BITSTRING     1  *            reserved(set to zero)@53c
        ....  ...1          S94CLVA3       field valid indicator@50a
  267  (10B) UNSIGNED      1  *            27:27 relative link speeds
        1111  ....          S94CLLS3_AXOVTS0
                                           AXO-VTS0 link speed
        ....  1111          S94CLLS3_AXOVTS1
                                           AXO-VTS1 link speed
  268  (10C) CHARACTER     4  *            28:31 reserved(set to
                                           zero)@50a
=================================================================================
       Statistics for AXO(4)
=================================================================================

  272  (110) UNSIGNED      4  S94CLDC4     0:3 data copied by AXO
  276  (114) UNSIGNED      4  S94CLVC4     4:7 volumes copied by AXO
  280  (118) UNSIGNED      4  S94CLRD4     8:11 read data transferred
  284  (11C) UNSIGNED      4  S94CLWD4     12:15 write data transferred
  288  (120) UNSIGNED      2  S94CLCM4     16:17 category mounts for AXO
  290  (122) UNSIGNED      2  S94CLSM4     18:19 specific cache mounts
  292  (124) UNSIGNED      2  S94CLRM4     20:21 specific recall mounts
  294  (126) UNSIGNED      2  S94CLCR4     22:23 Compression ratio
  296  (128) BITSTRING     1  *            24:24 Distributed library
                                           Preferences
        1111  ....          S94CLPI4       Preferred I/O VTS
        ....  1111          S94CLPM4       Preferred Master VTS
  297  (129) BITSTRING     2  *            25:26 Configured Settings.
                                           The reported settings@50a in
                                           this 2 byte field
                                           are only valid if
                                           bit S94CLVA4 is set.
        11..  ....          S94CLDM4       default copy mode

Offsets     Type      Length Name          Description

        ..1.  ....          S94CLFS4       force scratch to the
                                           preferred VTS.
        ...1  ....          S94CLPV4       VTS I/O selection
                                           criteria is PRIMARY
        ....  11..          S94CLCO4       Controller
                                           Operational Mode
        ....  ..1.          S94CLWP4       Write Protected Mode
  297  (129) BITSTRING     1  *            reserved(set to zero)@53c
        ....  ...1          S94CLVA4       field valid indicator@50a
  299  (12B) UNSIGNED      1  *            27:27 relative link speeds
        1111  ....          S94CLLS4_AXOVTS0
                                           AXO-VTS0 link speed
        ....  1111          S94CLLS4_AXOVTS1
                                           AXO-VTS1 link speed
  300  (12C) CHARACTER     4  *            28:31 reserved(set to
                                           zero)@50a
=================================================================================
       Statistics for AXO(5)
```

```
==================================================================================

  304  (130) UNSIGNED      4  S94CLDC5       0:3 data copied by AX0
  308  (134) UNSIGNED      4  S94CLVC5       4:7 volumes copied by AX0
  312  (138) UNSIGNED      4  S94CLRD5       8:11 read data transferred
  316  (13C) UNSIGNED      4  S94CLWD5       12:15 write data transferred
  320  (140) UNSIGNED      2  S94CLCM5       16:17 category mounts for AX0
  322  (142) UNSIGNED      2  S94CLSM5       18:19 specific cache mounts
  324  (144) UNSIGNED      2  S94CLRM5       20:21 specific recall mounts
  326  (146) UNSIGNED      2  S94CLCR5       22:23 Compression ratio
  328  (148) BITSTRING     1  *              24:24 Distributed library
                                             Preferences
         1111  ....           S94CLPI5       Preferred I/O VTS
         ....  1111           S94CLPM5       Preferred Master VTS
  329  (149) BITSTRING     2  *              25:26 Configured Settings.
                                             The reported settings@50a in
                                             this 2 byte field
                                             are only valid if
                                             bit S94CLVA5 is set.
         11..  ....           S94CLDM5       default copy mode
         ..1.  ....           S94CLFS5       force scratch to the
                                             preferred VTS.
         ...1  ....           S94CLPV5       VTS I/O selection
                                             criteria is PRIMARY
         ....  11..           S94CLCO5       Controller
                                             Operational Mode
         ....  ..1.           S94CLWP5       Write Protected Mode
  329  (149) BITSTRING     1  *              reserved(set to zero)@53c
         ....  ...1           S94CLVA5       field valid indicator@50a
  331  (14B) UNSIGNED      1  *              27:27 relative link speeds
         1111  ....           S94CLLS5_AX0VTS0
                                             AX0-VTS0 link speed
         ....  1111           S94CLLS5_AX0VTS1
                                             AX0-VTS1 link speed

Offsets     Type      Length Name            Description

  332  (14C) CHARACTER     4  *              28:31 reserved(set to
                                             zero)@50a
==================================================================================
        Statistics for AX0(6)
==================================================================================

  336  (150) UNSIGNED      4  S94CLDC6       0:3 data copied by AX0
  340  (154) UNSIGNED      4  S94CLVC6       4:7 volumes copied by AX0
  344  (158) UNSIGNED      4  S94CLRD6       8:11 read data transferred
  348  (15C) UNSIGNED      4  S94CLWD6       12:15 write data transferred
  352  (160) UNSIGNED      2  S94CLCM6       16:17 category mounts for AX0
  354  (162) UNSIGNED      2  S94CLSM6       18:19 specific cache mounts
  356  (164) UNSIGNED      2  S94CLRM6       20:21 specific recall mounts
  358  (166) UNSIGNED      2  S94CLCR6       22:23 Compression ratio
  360  (168) BITSTRING     1  *              24:24 Distributed library
                                             Preferences
         1111  ....           S94CLPI6       Preferred I/O VTS
         ....  1111           S94CLPM6       Preferred Master VTS
  361  (169) BITSTRING     2  *              25:26 Configured Settings.
                                             The reported settings@50a in
                                             this 2 byte field
                                             are only valid if
                                             bit S94CLVA6 is set.
```

```
        11..  ....          S94CLDM6       default copy mode
        ..1.  ....          S94CLFS6       force scratch to the
                                           preferred VTS.
        ...1  ....          S94CLPV6       VTS I/O selection
                                           criteria is PRIMARY
        ....  11..          S94CLCO6       Controller
                                           Operational Mode
        ....  ..1.          S94CLWP6       Write Protected Mode
  361  (169) BITSTRING     1  *              reserved(set to zero)@53c
        ....  ...1          S94CLVA6       field valid indicator@50a
  363  (16B) UNSIGNED      1  *              27:27 relative link speeds
        1111  ....          S94CLLS6_AX0VTS0
                                           AX0-VTS0 link speed
        ....  1111          S94CLLS6_AX0VTS1
                                           AX0-VTS1 link speed
  364  (16C) CHARACTER     4  *              28:31 reserved(set to
                                           zero)@50a
==================================================================================
       Statistics for AX0(7)
==================================================================================

  368  (170) UNSIGNED      4  S94CLDC7       0:3 data copied by AX0
  372  (174) UNSIGNED      4  S94CLVC7       4:7 volumes copied by AX0
  376  (178) UNSIGNED      4  S94CLRD7       8:11 read data transferred
  380  (17C) UNSIGNED      4  S94CLWD7       12:15 write data transferred
  384  (180) UNSIGNED      2  S94CLCM7       16:17 category mounts for AX0
  386  (182) UNSIGNED      2  S94CLSM7       18:19 specific cache mounts
  388  (184) UNSIGNED      2  S94CLRM7       20:21 specific recall mounts
  390  (186) UNSIGNED      2  S94CLCR7       22:23 Compression ratio
  392  (188) BITSTRING     1  *              24:24 Distributed library
                                           Preferences

Offsets     Type      Length Name            Description

        1111  ....          S94CLPI7       Preferred I/O VTS
        ....  1111          S94CLPM7       Preferred Master VTS
  393  (189) BITSTRING     2  *              25:26 Configured Settings.
                                           The reported settings@50a in
                                           this 2 byte field
                                           are only valid if
                                           bit S94CLVA7 is set.
        11..  ....          S94CLDM7       default copy mode
        ..1.  ....          S94CLFS7       force scratch to the
                                           preferred VTS.
        ...1  ....          S94CLPV7       VTS I/O selection
                                           criteria is PRIMARY
        ....  11..          S94CLCO7       Controller
                                           Operational Mode
        ....  ..1.          S94CLWP7       Write Protected Mode
  393  (189) BITSTRING     1  *              reserved(set to zero)@53c
        ....  ...1          S94CLVA7       field valid indicator@50a
  395  (18B) UNSIGNED      1  *              27:27 relative link speeds
        1111  ....          S94CLLS7_AX0VTS0
                                           AX0-VTS0 link speed
        ....  1111          S94CLLS7_AX0VTS1
                                           AX0-VTS1 link speed
  396  (18C) CHARACTER     4  *              28:31 reserved(set to
                                           zero)@50a
==================================================================================
       VTC code levels in a PTP VTS.
```

```
        The following set of fields report the code levels of
        the VTCs in a PTP VTS subsystem that generated the
        statistical record. For each of the 8 sets of fields,
        the field can return zeros for the following cases:
        1.The request was for a distributed library of a PTP.
        2.One or more VTC was not installed or was unavailable.
        3.The PTP VTS does not support providing code levels.
==================================================================================

  400  (190) UNSIGNED      2  S94CMV_VTC0    656:657 code mod value
  402  (192) UNSIGNED      2  S94CFV_VTC0    658:659 code fix value
  404  (194) UNSIGNED      2  S94CMV_VTC1    660:661 code mod value
  406  (196) UNSIGNED      2  S94CFV_VTC1    662:663 code fix value
  408  (198) UNSIGNED      2  S94CMV_VTC2    664:665 code mod value
  410  (19A) UNSIGNED      2  S94CFV_VTC2    666:667 code fix value
  412  (19C) UNSIGNED      2  S94CMV_VTC3    668:669 code mod value
  414  (19E) UNSIGNED      2  S94CFV_VTC3    670:671 code fix value
  416  (1A0) UNSIGNED      2  S94CMV_VTC4    672:673 code mod value
  418  (1A2) UNSIGNED      2  S94CFV_VTC4    674:675 code fix value
  420  (1A4) UNSIGNED      2  S94CMV_VTC5    676:677 code mod value
  422  (1A6) UNSIGNED      2  S94CFV_VTC5    678:679 code fix value
  424  (1A8) UNSIGNED      2  S94CMV_VTC6    680:681 code mod value
  426  (1AA) UNSIGNED      2  S94CFV_VTC6    682:683 code fix value
  428  (1AC) UNSIGNED      2  S94CMV_VTC7    684:685 code mod value
  430  (1AE) UNSIGNED      2  S94CFV_VTC7    686:687 code fix value
==================================================================================
        688:783 reserved (set to X'00')
==================================================================================


Offsets     Type      Length Name           Description

  432  (1B0) CHARACTER    96  *              reserved (set to X'00')
==================================================================================
        784:799 when F/C 4001 is installed & enabled, contains
        information about:
        Virtual Tape Device Configuration and Usage.
==================================================================================

  528  (210) UNSIGNED      2  S94OPM_VDC     Virtual Drives Configured
                                             Contains the number of virtual
                                             devices configured in the VTS
                                             subsystem.
  530  (212) UNSIGNED      2  S94OPM_MAXVDM  Max Virtual Drives Mounted
                                             con- currenty during the last
                                             hour.
  532  (214) UNSIGNED      2  S94OPM_MINVDM  Min Virtual Drives Mounted
                                             con- currenty during the last
                                             hour.
  534  (216) UNSIGNED      2  S94OPM_AVGVDM  Avg Virtual Drives Mounted
                                             con- currenty during the last
                                             hour.
  536  (218) CHARACTER     8  *              reserved (set to X'00')
==================================================================================
        800:863 when F/C 4001 is installed & enabled, contains
        information about:
        Physical tape devices usage characteristics.
        Each group of 32 bytes contains the statistics for the
        physical tape device usage of a specific device type in
        the VTS. Space is defined for 2 device types.
```

```
================================================================================

  544  (220) CHARACTER    32  *(2)           reserved (set to X'00')
  544  (220) UNSIGNED      1  S94OPM_DCI     00 device class identifier:
                                             X'00' no device class X'11'
                                             3590 Model B1A X'13' 3590
                                             Model E1A X'14' 3590 Model H1A
                                             X'20' 3592 Model J1A
  545  (221) UNSIGNED      1  S94OPM_PDI     01 physical devs installed
  546  (222) UNSIGNED      1  S94OPM_CAFU    02 currently avail for use
  547  (223) UNSIGNED      1  S94OPM_MAXCM   03 maximum concurrent mounts
                                             during the last hour.
  548  (224) UNSIGNED      1  S94OPM_MINCM   04 minimum concurrent mounts
                                             during the last hour.
  549  (225) UNSIGNED      1  S94OPM_AVGCM   05 average concurrent mounts
                                             during the last hour.
  550  (226) UNSIGNED      2  S94OPM_MAXTTM  06:07 max time(in secs) to
                                             mount during the last hour.

Offsets     Type      Length Name            Description

  552  (228) UNSIGNED      2  S94OPM_MINTTM  08:09 min time(in secs) to
                                             mount during the last hour.
  554  (22A) UNSIGNED      2  S94OPM_AVGTTM  10:11 avg time(in secs) to
                                             mount during the last hour.
  556  (22C) UNSIGNED      2  S94OPM_STGMNTS 12:13 number of stage
                                             mounts during the last hour.
  558  (22E) UNSIGNED      2  S94OPM_MIGMNTS 14:15 # of migration
                                             mounts during the last hour.
  560  (230) UNSIGNED      2  S94OPM_RECMNTS 16:17 # of reclamation
                                             mounts during the last hour.
  562  (232) UNSIGNED      2  S94OPM_SDEMNTS 18:19 # of physical mounts
                                             during the last hour to
                                             satisfy secure data erasure.
  564  (234) UNSIGNED      4  S94OPM_PPWRITEN20:23 megabytes
                                             premigrated from tape volume
                                             cache to a primary pool during
                                             the last hour.
  568  (238) UNSIGNED      4  S94OPM_SPWRITEN24:27 megabytes
                                             premigrated from tape volume
                                             cache to a secondary pool
                                             during the last hour.
  572  (23C) CHARACTER     4  *              reserved (set to X'00')
================================================================================
        864:1375 when F/C 4001 is installed & enabled, contains
        information about:
        VTS cache usage characteristics.
================================================================================

  608  (260) CHARACTER    64  S94OPM_PLCU(8) 00:64 Pref. Level Cache Usage
                                             w/o F/C 4001 set to X'00'
  608  (260) BITSTRING     4  S94OPM_PMC     00:03 preference mang.control
  612  (264) UNSIGNED      4  S94OPM_VVIC    04:07 virtual vols in cache
  616  (268) UNSIGNED      4  S94OPM_DRIC    08:11 data resident in cache
  620  (26C) UNSIGNED      4  S94OPM_TVCA4   12:15 tape volume cache age
                                             (4 hour rolling average)
  624  (270) UNSIGNED      4  S94OPM_VM4     16:19 volumes migrated from
                                             cache (last 4 hours)
  628  (274) UNSIGNED      4  S94OPM_TVCA48  20:23 tape volume cache age
                                             (48 hour rolling avg.)
```

```
  632  (278) UNSIGNED      4  S94OPM_VM48    24:27 volumes migrated over
                                             the last 48 hrs
  636  (27C) UNSIGNED      4  S94OPM_TVCA35  28:31 tape volume cache age
                                             (35 day rolling avg.)

Offsets     Type      Length Name           Description

  640  (280) UNSIGNED      4  S94OPM_VM35    32:35 volumes migrated from
                                             cache over last 35 days
  644  (284) UNSIGNED      2  S94OPM_FRMT    36:37 average fast ready
                                             mount time.
  646  (286) UNSIGNED      2  S94OPM_FRMNTS  38:39 # of fast-ready
                                             mounts
  648  (288) UNSIGNED      2  S94OPM_CHTIME  40:41 avg. cache hit mount
                                             time
  650  (28A) UNSIGNED      2  S94OPM_CHMNTS  42:43 cache hit mounts
  652  (28C) UNSIGNED      2  S94OPM_CMTIME  44:45 avg.cache miss mount
                                             time
  654  (28E) UNSIGNED      2  S94OPM_CMMNTS  46:47 cache miss mounts
  656  (290) CHARACTER    16  *              48:63 reserved (set to X'00')
==================================================================================
        1376:1387 IART Controlled Mount Statistics.
==================================================================================

 1120  (460) UNSIGNED      2  S94OPM_IARTAFRT1376:1377 Average
                                             fast-ready mount time.
 1122  (462) UNSIGNED      2  S94OPM_IARTFRM 1378:1379 Fast-ready
                                             mounts.
 1124  (464) UNSIGNED      2  S94OPM_IARTCHMT1380:1381 Average cache
                                             hit mount time.
 1126  (466) UNSIGNED      2  S94OPM_IARTCHM 1382:1383 Cache hit mounts.
 1128  (468) UNSIGNED      2  S94OPM_IARTCMMT1384:1385 Average cache
                                             miss mount time.
 1130  (46A) UNSIGNED      2  S94OPM_IARTCMM 1386:1387 cache miss
                                             mounts.
 1132  (46C) CHARACTER        S94END
==================================================================================
    SMF TYPE 94 (Subtype 2) RECORD, Volume Pool Statistics
    Note: this data structure is valid when SMF94 subtype
          field (SMF94STP) contains a 2.
==================================================================================


Offsets     Type      Length Name           Description

    0    (0) STRUCTURE    28  SMF94S2
==================================================================================
     SELF DEFINING SECTION
==================================================================================

    0    (0) CHARACTER    28  SMF94S2_SDS    Self Defining Section
    0    (0) UNSIGNED      4  SMF94S2_SDL    Self Defining Section length
    4    (4) UNSIGNED      4  SMF94S2_POF    Offset to product Section
    8    (8) UNSIGNED      2  SMF94S2_PLN    Length of product Section
   10    (A) UNSIGNED      2  SMF94S2_PON    Number of product Sections
   12    (C) UNSIGNED      4  SMF94S2_HOF    Offset to header section
   16   (10) UNSIGNED      2  SMF94S2_HLN    Length of header section
   18   (12) UNSIGNED      2  SMF94S2_HON    Number of header sections
   20   (14) UNSIGNED      4  SMF94S2_SOF    Offset to statistics
   24   (18) UNSIGNED      2  SMF94S2_SLN    Length of statistics
```

```
   26   (1A) UNSIGNED      2  SMF94S2_SON    Number of statistics
==================================================================================
    SMF TYPE 94 (Subtype 2) Product Section
==================================================================================


Offsets     Type      Length Name           Description

    0    (0) STRUCTURE    20  SMF94S2_PSS
    0    (0) UNSIGNED      2  SMF94S2_TYP    record subtype = 01
    2    (2) CHARACTER     2  SMF94S2_RVN    record version number = 01
    4    (4) CHARACTER     8  SMF94S2_PNM    product name 'fmid '
   12    (C) CHARACTER     8  SMF94S2_MVS    MVS operating system name
==================================================================================
    SMF TYPE 94 (Subtype 2) Header Section
    Reserved for information generated by software.
==================================================================================


Offsets     Type      Length Name           Description

    0    (0) STRUCTURE    32  SMF94S2_HDR
    0    (0) CHARACTER    32  *              reserved
==================================================================================
    SMF TYPE 94 (Subtype 2) Volume Pool Statistics (VPS)
    NOTE: there may be more than one of these sections recorded.
          The number of sections is contained in SMF94S2_SON.
          SMF94S2_SOF contains the offset to the first section, and
          SMF94S2_SLN contains the length of each section.
==================================================================================


Offsets     Type      Length Name           Description

    0    (0) STRUCTURE  2096  SMF94S2_DATA
    0    (0) CHARACTER    16  SMF94S2_MSGHDR 0:15 VPS message header data
    0    (0) UNSIGNED      2  SMF94S2_LENGTH 0:01 message length
    2    (2) CHARACTER     1  SMF94S2_MSGFMT 2:02 message format ('10'x)
    3    (3) CHARACTER     1  SMF94S2_MSGCODE3:03 message code ('29'x)
    4    (4) CHARACTER     4  SMF94S2_MSGID  4:07 message-id
    8    (8) CHARACTER     1  *              8:08 Reserved
    9    (9) CHARACTER     3  SMF94S2_LIBID  9:11 library sequence number
   12    (C) CHARACTER     1  SMF94S2_VPSFMT 12:12 volume pool stats
                                             format
   13    (D) CHARACTER     1  *              13:13 Reserved
   14    (E) UNSIGNED      2  SMF94S2_HHI    14:15 hour index
   16   (10) CHARACTER    32  *              16:47
   16   (10) CHARACTER    10  SMF94S2_LRTD   16:25 last reconcile
                                             time/data
   26   (1A) UNSIGNED      2  SMF94S2_MNVP   26:27 max pools in partition
   28   (1C) CHARACTER     1  *              28
   29   (1D) UNSIGNED      1  SMF94S2_VPSET  29 pool set (1of2,2of2,etc)
   30   (1E) CHARACTER     2  *              30:31
   32   (20) CHARACTER     1  *              32
   33   (21) UNSIGNED      1  SMF94S2_BPMI0  33 media type for BPSVC0
   34   (22) UNSIGNED      2  SMF94S2_BPSVC0 34:35 number of scratch
                                             stacked volumes of type BPMI0
   36   (24) CHARACTER     1  *              36 reserved
   37   (25) UNSIGNED      1  SMF94S2_BPMI1  37 media type for BPSVC1
   38   (26) UNSIGNED      2  SMF94S2_BPSVC1 38:39 number of scratch
```

```
                                            stacked volumes of type BPMI1
   40   (28) CHARACTER     1  *             40 reserved
   41   (29) UNSIGNED      1  SMF94S2_BPMI2  41 media type for BPSVC2
   42   (2A) UNSIGNED      2  SMF94S2_BPSVC2 42:43 number of scratch
                                            stacked volumes of type BPMI2
   44   (2C) CHARACTER     1  *             44 reserved
   45   (2D) UNSIGNED      1  SMF94S2_BPMI3  45 media type for BPSVC3
   46   (2E) UNSIGNED      2  SMF94S2_BPSVC3 46:47 number of scratch
                                            stacked volumes of type BPMI3
=================================================================================
    This section contains an array containing statistics for
    16 volume pools. Each entry is 112 bytes long.
=================================================================================

   48   (30) CHARACTER  1792  SMF94S2_ARRAY  48:2095 pool
                                            array
   48   (30) CHARACTER   112  SMF94S2_VPS(16)0:127 array
                                            entry
   48   (30) CHARACTER     1  *             0 reserved (set to '00'x)
   49   (31) UNSIGNED      1  SMF94S2_VPN    1 pool number (starts at 1)
   50   (32) CHARACTER     2  *             2:3 reserved (set to '00'x)

Offsets     Type      Length Name           Description

   52   (34) UNSIGNED      4  SMF94S2_ALVIP  4:7 # of logical volumes
                                            on physical vols assigned to
                                            the pool.
   56   (38) UNSIGNED      4  SMF94S2_ADIVP  8:11 Mbytes of pool data
   60   (3C) UNSIGNED      4  SMF94S2_DWTPLH 12:15 Mbytes written to
                                            pool in last hour.
   64   (40) UNSIGNED      1  SMF94S2_PDCI   16 pool device class.
   65   (41) UNSIGNED      1  SMF94S2_MI0    17 media 0 counts:
   66   (42) UNSIGNED      2  SMF94S2_PSSSVC018:19 static scratch
                                            stacked
   68   (44) UNSIGNED      2  SMF94S2_PSPSVC020:21 static private
                                            stacked
   70   (46) UNSIGNED      2  SMF94S2_PBSSVC022:23 borrowed scratch
                                            stacked
   72   (48) UNSIGNED      2  SMF94S2_PBPSVC024:25 borrowed private
                                            stacked
   74   (4A) CHARACTER     1  *             26 reserved (set to '00'x)
   75   (4B) UNSIGNED      1  SMF94S2_MI1    27 media 1 counts:
   76   (4C) UNSIGNED      2  SMF94S2_PSSSVC128:29 static scratch
                                            stacked
   78   (4E) UNSIGNED      2  SMF94S2_PSPSVC130:31 static private
                                            stacked
   80   (50) UNSIGNED      2  SMF94S2_PBSSVC132:33 borrowed scratch
                                            stacked
   82   (52) UNSIGNED      2  SMF94S2_PBPSVC134:35 borrowed private
                                            stacked
   84   (54) CHARACTER     4  *             36:39 reserved (set to '00'x)
   88   (58) UNSIGNED      2  SMF94S2_AAORD  40:41 avg.age of residual
                                            data
   90   (5A) UNSIGNED      2  SMF94S2_MAORD  42:43 max age of residual
                                            data
   92   (5C) UNSIGNED      2  SMF94S2_AAOFPSV44:45 avg.age of full
                                            private stacked volumes.
   94   (5E) UNSIGNED      2  SMF94S2_MAOFPSV46:47 max age of full
                                            private stacked volumes.
   96   (60) CHARACTER     1  *             48 reserved (set to '00'x)
```

```
  97   (61) UNSIGNED      1  SMF94S2_VPRTP  49 pool reclaim threshold %
  98   (62) CHARACTER    40  SMF94S2_ADD    50:89 Active Data Distribution
  98   (62) UNSIGNED      2  SMF94S2_ADD00  50:51 0- 5% active data
 100   (64) UNSIGNED      2  SMF94S2_ADD05  52:53 5-10% active data
 102   (66) UNSIGNED      2  SMF94S2_ADD10  54:55 10-15% active data
 104   (68) UNSIGNED      2  SMF94S2_ADD15  56:57 15-20% active data
 106   (6A) UNSIGNED      2  SMF94S2_ADD20  58:59 20-25% active data
 108   (6C) UNSIGNED      2  SMF94S2_ADD25  60:61 25-30% active data
 110   (6E) UNSIGNED      2  SMF94S2_ADD30  62:63 30-35% active data
 112   (70) UNSIGNED      2  SMF94S2_ADD35  64:65 35-40% active data

Offsets     Type      Length Name           Description

 114   (72) UNSIGNED      2  SMF94S2_ADD40  66:67 40-45% active data
 116   (74) UNSIGNED      2  SMF94S2_ADD45  68:69 45-50% active data
 118   (76) UNSIGNED      2  SMF94S2_ADD50  70:71 50-55% active data
 120   (78) UNSIGNED      2  SMF94S2_ADD55  72:73 55-60% active data
 122   (7A) UNSIGNED      2  SMF94S2_ADD60  74:75 60-65% active data
 124   (7C) UNSIGNED      2  SMF94S2_ADD65  76:77 65-70% active data
 126   (7E) UNSIGNED      2  SMF94S2_ADD70  78:79 70-75% active data
 128   (80) UNSIGNED      2  SMF94S2_ADD75  80:81 75-80% active data
 130   (82) UNSIGNED      2  SMF94S2_ADD80  82:83 80-85% active data
 132   (84) UNSIGNED      2  SMF94S2_ADD85  84:85 85-90% active data
 134   (86) UNSIGNED      2  SMF94S2_ADD90  86:87 90-95% active data
 136   (88) UNSIGNED      2  SMF94S2_ADD95  88:89 95-100%active data
 138   (8A) UNSIGNED      1  SMF94S2_PPP    90 Physical pool properties
 139   (8B) CHARACTER     1  *              91 reserved('00'x)
 140   (8C) UNSIGNED      1  SMF94S2_RPN    92 Reclamation pool number
 141   (8D) CHARACTER     1  *              93 reserved('00'X)
 142   (8E) UNSIGNED      2  SMF94S2_RPDSLA 94:95 Reclaim policy,
                                            days since last accessed
 144   (90) UNSIGNED      2  SMF94S2_RPDSLW 96:98 Reclaim policy,
                                            days since last written.
 146   (92) UNSIGNED      2  SMF94S2_RPDSLDI98:99 Reclaim policy,
                                            days since last data
                                            invalidate
 148   (94) UNSIGNED      1  SMF94S2_RPMADP 100 Reclaim policy,
                                            Minimum active data percent.
 149   (95) CHARACTER    11  *              101:111 reserved('00'x)
Constants
Length  Type      Value      Name           Description

==================================================================================
    CONSTANTS
==================================================================================

    1    DECIMAL   94         SMF94C94       Used in SMF94RTY field
==================================================================================
   Definitions for media types found in SMF94S2_BPMI0,
   SMF94S2_BPMI1, SMF94S2_BPMI2 and SMF94S2_BPMI3.
==================================================================================

    1    HEX       00         SMF94S2_BPM00  no media type
                                             configured
    1    HEX       10         SMF94S2_BPM10  3590'J' media
    1    HEX       11         SMF94S2_BPM11  3590'K' media
    1    HEX       20         SMF94S2_BPM20  3592'JA' media
    1    HEX       21         SMF94S2_BPM21  reserved.
    1    HEX       22         SMF94S2_BPM22  3592'JJ' media
    1    HEX       23         SMF94S2_BPM23  reserved.
```

```
================================================================================
   Definitions for the preferred I/O VTS which is reported in
   S94CLPI0-S94CLPI7 and in S94CLPM0-S94CLPM7
================================================================================

      BIT       1000       S94PREF_VTS0   VTS 0 is preferred
      BIT       1001       S94PREF_VTS1   VTS 1 is preferred
      BIT       1111       S94PREF_NONE   No Preference
================================================================================
   Definitions for the default copy mode which is reported in
   S94CLDM0-S94CLDM7
================================================================================

      BIT       00         S94COPY_DEFER  deferred copy mode
      BIT       01         S94COPY_IMMED  immediate copy mode
================================================================================
   Definitions for controller operation mode which is reported
   in S94CLCO0-S94CLCO7
================================================================================

      BIT       00         S94COM_NORM    normal mode
      BIT       01         S94COM_READ    read-only mode
      BIT       10         S94COM_RWD     read/write disconnected
================================================================================
   Definitions for relative link speeds which are reported in
   S94CLLS0_AX0VTS0 - S94CLLS7_AX0VTS0 and in
   S94CLLS0_AX0VTS1 - S94CLLS7_AX0VTS1
================================================================================

      BIT       0001       S94LS1         ESCON < 5km, FICON < 30km
      BIT       0010       S94LS2         ESCON <10km, FICON < 80km
      BIT       0011       S94LS3         channel extender <
                                          1000km,
                                          ESCON over DWDM < 10km,
                                          FICON < 100km
      BIT       0100       S94LS4         channel extender >
                                          1000km,
                                          ESCON over DWDM < 15km
      BIT       0101       S94LS5         ESCON over DWDM > 15km
```

Cross Reference

```
CROSS REFERENCE

                                HEX       HEX
NAME                         OFFSET     VALUE      LEVEL
____                         ______     _____      _____
SMFRCD94                          0                    1
SMF94ACA                         24                    2
SMF94ACB                         26                    2
SMF94ALN                         58                    3
SMF94AOF                         54                    3
SMF94AON                         5A                    3
SMF94APR                          0                    2
SMF94AP1                          2                    2
SMF94AP2                          4                    2
SMF94AP3                          6                    2
```

SMF94AT0	8		2
SMF94AT1	A		2
SMF94AT2	C		2
SMF94AT3	E		2
SMF94AUD	0		1
SMF94DIN	A		2
SMF94DLN	48		3
SMF94DMT	0		1
SMF94DOF	44		3
SMF94DON	4A		3
SMF94DPM	C		2
SMF94DPR	0		2
SMF94DP1	2		2
SMF94DP2	4		2
SMF94DP3	6		2
SMF94DTE	A		3
SMF94DT0	8		2
SMF94DT1	E		2
SMF94DT2	10		2
SMF94DT3	12		2
SMF94EJT	0		1
SMF94ELN	50		3
SMF94EOF	4C		3
SMF94EON	52		3
SMF94EPR	0		2
SMF94EP1	2		2
SMF94EP2	4		2
SMF94EP3	6		2
SMF94ET0	8		2
SMF94ET1	A		2
SMF94ET2	C		2
SMF94ET3	E		2
SMF94EX1	2		2
SMF94EX2	8		2
SMF94EX3	10		2
SMF94EX4	18		2
SMF94FLG	4		3
SMF94FNM	4	80	4
SMF94FST	4	40	4
SMF94HDR	0		1
SMF94HHI	1		2
SMF94HLN	28		3
SMF94HOF	24		3
SMF94HON	2A		3
SMF94HSF	0		2

CROSS REFERENCE

NAME	HEX OFFSET	HEX VALUE	LEVEL
SMF94ILN	60		3
SMF94IM1	0		2
SMF94IM2	4		2
SMF94IM3	C		2
SMF94IM4	14		2
SMF94INP	0		1
SMF94INS	0		2
SMF94IOF	5C		3
SMF94ION	62		3
SMF94LEN	0		3

SMF94LIB	0		1
SMF94LID	0		2
SMF94LLN	38		3
SMF94LMD	2		2
SMF94LM1	4		2
SMF94LM2	6		2
SMF94LM3	8		2
SMF94LOF	34		3
SMF94LON	3A		3
SMF94LT1	A		2
SMF94LT2	C		2
SMF94LT3	E		2
SMF94MIN	A		2
SMF94MLN	40		3
SMF94MNT	0		1
SMF94MOF	3C		3
SMF94MON	42		3
SMF94MPM	C		2
SMF94MPR	0		2
SMF94MP1	2		2
SMF94MP2	4		2
SMF94MP3	6		2
SMF94MT0	8		2
SMF94MT1	E		2
SMF94MT2	10		2
SMF94MT3	12		2
SMF94MVS	C		2
SMF94PLN	20		3
SMF94PNM	4		2
SMF94POF	1C		3
SMF94PON	22		3
SMF94PSS	0		1
SMF94RPS	0		2
SMF94RTY	5		3
SMF94RVN	2		2
SMF94SDL	18		3
SMF94SDS	18		2
SMF94SEG	2		3
SMF94SID	E		3
SMF94SLF	0		1
SMF94SLM	6		2
SMF94SLN	30		3
SMF94SLT	0		2
SMF94SMA	9		2
SMF94SNO	E		2

CROSS REFERENCE

NAME	HEX OFFSET	HEX VALUE	LEVEL
SMF94SOF	2C		3
SMF94SON	32		3
SMF94SPL	C		2
SMF94STP	16		3
SMF94S2	0		1
SMF94S2_AAOFPSV	5C		4
SMF94S2_AAORD	58		4
SMF94S2_ADD	62		4
SMF94S2_ADD00	62		5
SMF94S2_ADD05	64		5

SMF94S2_ADD10	66	5
SMF94S2_ADD15	68	5
SMF94S2_ADD20	6A	5
SMF94S2_ADD25	6C	5
SMF94S2_ADD30	6E	5
SMF94S2_ADD35	70	5
SMF94S2_ADD40	72	5
SMF94S2_ADD45	74	5
SMF94S2_ADD50	76	5
SMF94S2_ADD55	78	5
SMF94S2_ADD60	7A	5
SMF94S2_ADD65	7C	5
SMF94S2_ADD70	7E	5
SMF94S2_ADD75	80	5
SMF94S2_ADD80	82	5
SMF94S2_ADD85	84	5
SMF94S2_ADD90	86	5
SMF94S2_ADD95	88	5
SMF94S2_ADIVP	38	4
SMF94S2_ALVIP	34	4
SMF94S2_ARRAY	30	2
SMF94S2_BPMI0	21	3
SMF94S2_BPMI1	25	3
SMF94S2_BPMI2	29	3
SMF94S2_BPMI3	2D	3
SMF94S2_BPSVC0	22	3
SMF94S2_BPSVC1	26	3
SMF94S2_BPSVC2	2A	3
SMF94S2_BPSVC3	2E	3
SMF94S2_DATA	0	1
SMF94S2_DWTPLH	3C	4
SMF94S2_HDR	0	1
SMF94S2_HHI	E	3
SMF94S2_HLN	10	3
SMF94S2_HOF	C	3
SMF94S2_HON	12	3
SMF94S2_LENGTH	0	3
SMF94S2_LIBID	9	3
SMF94S2_LRTD	10	3
SMF94S2_MAOFPSV	5E	4
SMF94S2_MAORD	5A	4
SMF94S2_MI0	41	4
SMF94S2_MI1	4B	4
SMF94S2_MNVP	1A	3
SMF94S2_MSGCODE	3	3

CROSS REFERENCE

NAME	HEX OFFSET	HEX VALUE	LEVEL
SMF94S2_MSGFMT	2		3
SMF94S2_MSGHDR	0		2
SMF94S2_MSGID	4		3
SMF94S2_MVS	C		2
SMF94S2_PBPSVC0	48		4
SMF94S2_PBPSVC1	52		4
SMF94S2_PBSSVC0	46		4
SMF94S2_PBSSVC1	50		4
SMF94S2_PDCI	40		4
SMF94S2_PLN	8		3

SMF94S2_PNM	4		2
SMF94S2_POF	4		3
SMF94S2_PON	A		3
SMF94S2_PPP	8A		4
SMF94S2_PSPSVC0	44		4
SMF94S2_PSPSVC1	4E		4
SMF94S2_PSS	0		1
SMF94S2_PSSSVC0	42		4
SMF94S2_PSSSVC1	4C		4
SMF94S2_RPDSLA	8E		4
SMF94S2_RPDSLDI	92		4
SMF94S2_RPDSLW	90		4
SMF94S2_RPMADP	94		4
SMF94S2_RPN	8C		4
SMF94S2_RVN	2		2
SMF94S2_SDL	0		3
SMF94S2_SDS	0		2
SMF94S2_SLN	18		3
SMF94S2_SOF	14		3
SMF94S2_SON	1A		3
SMF94S2_TYP	0		2
SMF94S2_VPN	31		4
SMF94S2_VPRTP	61		4
SMF94S2_VPS	30		3
SMF94S2_VPSET	1D		3
SMF94S2_VPSFMT	C		3
SMF94TME	6		3
SMF94TYP	0		2
SMF94VBA	48		2
SMF94VBR	34		2
SMF94VBW	30		2
SMF94VCA	40		2
SMF94VCZ	42		2
SMF94VDA	1B		2
SMF94VDC	18		2
SMF94VDN	1A		2
SMF94VDX	19		2
SMF94VEC	50		2
SMF94VFR	28		2
SMF94VLA	4C		2
SMF94VLN	68		3
SMF94VLS	1		2
SMF94VMH	2A		2
SMF94VMN	E		2
SMF94VMP	2E		2

CROSS REFERENCE

NAME	HEX OFFSET	HEX VALUE	LEVEL
SMF94VMS	2C		2
SMF94VMV	10		2
SMF94VMX	C		2
SMF94VNM	44		2
SMF94VNO	0		2
SMF94VOF	64		3
SMF94VON	6A		3
SMF94VPM	14		2
SMF94VPR	16		2
SMF94VPS	12		2

SMF94VRA	26		2
SMF94VRN	24		2
SMF94VRX	22		2
SMF94VR2	B		2
SMF94VR3	46		2
SMF94VTA	7		2
SMF94VTI	6		2
SMF94VTN	9		2
SMF94VTR	3C		2
SMF94VTS	0		1
SMF94VTV	A		2
SMF94VTW	38		2
SMF94VTX	8		2
SMF94VVA	20		2
SMF94VVN	1E		2
SMF94VVX	1C		2
SMF94WID	12		3
SMF94XLN	70		3
SMF94XOF	6C		3
SMF94XON	72		3
SMF94X00	0		1
SMF942LN	78		3
SMF942OF	74		3
SMF942ON	7A		3
S94ADV00	60		2
S94ADV05	3A		2
S94ADV10	3C		2
S94ADV15	3E		2
S94ADV20	40		2
S94ADV25	42		2
S94ADV30	44		2
S94ADV35	46		2
S94ADV40	48		2
S94ADV45	4A		2
S94ADV50	4C		2
S94ADV55	4E		2
S94ADV60	50		2
S94ADV65	52		2
S94ADV70	54		2
S94ADV75	56		2
S94ADV80	58		2
S94ADV85	5A		2
S94ADV90	5C		2
S94ADV95	5E		2
S94AVGCH	32		2

CROSS REFERENCE

NAME	HEX OFFSET	HEX VALUE	LEVEL
S94AVGFR	2C		2
S94AVGRM	38		2
S94AVRCT	18		2
S94AVWOT	1C		2
S94BSRAT	0		2
S94CFV_VTC0	192		2
S94CFV_VTC1	196		2
S94CFV_VTC2	19A		2
S94CFV_VTC3	19E		2
S94CFV_VTC4	1A2		2

```
S94CFV_VTC5                          1A6                    2
S94CFV_VTC6                          1AA                    2
S94CFV_VTC7                          1AE                    2
S94CLCM0                              A0                    2
S94CLCM1                              C0                    2
S94CLCM2                              E0                    2
S94CLCM3                             100                    2
S94CLCM4                             120                    2
S94CLCM5                             140                    2
S94CLCM6                             160                    2
S94CLCM7                             180                    2
S94CLCO0                              A9        08          3
S94CLCO1                              C9        08          3
S94CLCO2                              E9        08          3
S94CLCO3                             109        08          3
S94CLCO4                             129        08          3
S94CLCO5                             149        08          3
S94CLCO6                             169        08          3
S94CLCO7                             189        08          3
S94CLCR0                              A6                    2
S94CLCR1                              C6                    2
S94CLCR2                              E6                    2
S94CLCR3                             106                    2
S94CLCR4                             126                    2
S94CLCR5                             146                    2
S94CLCR6                             166                    2
S94CLCR7                             186                    2
S94CLDC0                              90                    2
S94CLDC1                              B0                    2
S94CLDC2                              D0                    2
S94CLDC3                              F0                    2
S94CLDC4                             110                    2
S94CLDC5                             130                    2
S94CLDC6                             150                    2
S94CLDC7                             170                    2
S94CLDM0                              A9        80          3
S94CLDM1                              C9        80          3
S94CLDM2                              E9        80          3
S94CLDM3                             109        80          3
S94CLDM4                             129        80          3
S94CLDM5                             149        80          3
S94CLDM6                             169        80          3
S94CLDM7                             189        80          3
S94CLDTC                              84                    2
S94CLFS0                              A9        20          3
                         CROSS REFERENCE
                         _____ _________
                                     HEX       HEX
NAME                              OFFSET     VALUE      LEVEL
____                              ______     _____      _____
S94CLFS1                              C9        20          3
S94CLFS2                              E9        20          3
S94CLFS3                             109        20          3
S94CLFS4                             129        20          3
S94CLFS5                             149        20          3
S94CLFS6                             169        20          3
S94CLFS7                             189        20          3
S94CLLS0_AX0VTS0                      AB        80          3
S94CLLS0_AX0VTS1                      AB        08          3
S94CLLS1_AX0VTS0                      CB        80          3
```

S94CLLS1_AXOVTS1	CB	08	3
S94CLLS2_AXOVTS0	EB	80	3
S94CLLS2_AXOVTS1	EB	08	3
S94CLLS3_AXOVTS0	10B	80	3
S94CLLS3_AXOVTS1	10B	08	3
S94CLLS4_AXOVTS0	12B	80	3
S94CLLS4_AXOVTS1	12B	08	3
S94CLLS5_AXOVTS0	14B	80	3
S94CLLS5_AXOVTS1	14B	08	3
S94CLLS6_AXOVTS0	16B	80	3
S94CLLS6_AXOVTS1	16B	08	3
S94CLLS7_AXOVTS0	18B	80	3
S94CLLS7_AXOVTS1	18B	08	3
S94CLLVC	80		2
S94CLMT0	88		2
S94CLMT1	8A		2
S94CLPI0	A8	80	3
S94CLPI1	C8	80	3
S94CLPI2	E8	80	3
S94CLPI3	108	80	3
S94CLPI4	128	80	3
S94CLPI5	148	80	3
S94CLPI6	168	80	3
S94CLPI7	188	80	3
S94CLPM0	A8	08	3
S94CLPM1	C8	08	3
S94CLPM2	E8	08	3
S94CLPM3	108	08	3
S94CLPM4	128	08	3
S94CLPM5	148	08	3
S94CLPM6	168	08	3
S94CLPM7	188	08	3
S94CLPV0	A9	10	3
S94CLPV1	C9	10	3
S94CLPV2	E9	10	3
S94CLPV3	109	10	3
S94CLPV4	129	10	3
S94CLPV5	149	10	3
S94CLPV6	169	10	3
S94CLPV7	189	10	3
S94CLRD0	98		2
S94CLRD1	B8		2
S94CLRD2	D8		2
S94CLRD3	F8		2
S94CLRD4	118		2

CROSS REFERENCE

NAME	HEX OFFSET	HEX VALUE	LEVEL
S94CLRD5	138		2
S94CLRD6	158		2
S94CLRD7	178		2
S94CLRM0	A4		2
S94CLRM1	C4		2
S94CLRM2	E4		2
S94CLRM3	104		2
S94CLRM4	124		2
S94CLRM5	144		2
S94CLRM6	164		2

```
S94CLRM7                            184                   2
S94CLSM0                             A2                   2
S94CLSM1                             C2                   2
S94CLSM2                             E2                   2
S94CLSM3                            102                   2
S94CLSM4                            122                   2
S94CLSM5                            142                   2
S94CLSM6                            162                   2
S94CLSM7                            182                   2
S94CLVA0                             AA       01          3
S94CLVA1                             CA       01          3
S94CLVA2                             EA       01          3
S94CLVA3                            10A       01          3
S94CLVA4                            12A       01          3
S94CLVA5                            14A       01          3
S94CLVA6                            16A       01          3
S94CLVA7                            18A       01          3
S94CLVC0                             94                   2
S94CLVC1                             B4                   2
S94CLVC2                             D4                   2
S94CLVC3                             F4                   2
S94CLVC4                            114                   2
S94CLVC5                            134                   2
S94CLVC6                            154                   2
S94CLVC7                            174                   2
S94CLWD0                             9C                   2
S94CLWD1                             BC                   2
S94CLWD2                             DC                   2
S94CLWD3                             FC                   2
S94CLWD4                            11C                   2
S94CLWD5                            13C                   2
S94CLWD6                            15C                   2
S94CLWD7                            17C                   2
S94CLWP0                             A9       02          3
S94CLWP1                             C9       02          3
S94CLWP2                             E9       02          3
S94CLWP3                            109       02          3
S94CLWP4                            129       02          3
S94CLWP5                            149       02          3
S94CLWP6                            169       02          3
S94CLWP7                            189       02          3
S94CMGTS                             6C                   2
S94CMV_VTC0                         190                   2
S94CMV_VTC1                         194                   2
S94CMV_VTC2                         198                   2
                         CROSS REFERENCE
                         _____ _________
                                    HEX      HEX
NAME                             OFFSET    VALUE      LEVEL
____                             ______    _____      _____
S94CMV_VTC3                         19C                   2
S94CMV_VTC4                         1A0                   2
S94CMV_VTC5                         1A4                   2
S94CMV_VTC6                         1A8                   2
S94CMV_VTC7                         1AC                   2
S94CVMPL                             6C       40          3
S94END                              46C                   2
S94ESCON                              6                   2
S94HARAT                              2                   2
S94LVLMR                             76                   2
```

S94LVLMV	74		2
S94LVVCF	72		2
S94LVVCM	70		2
S94MAXCH	2E		2
S94MAXFR	28		2
S94MAXRM	34		2
S94MINCH	30		2
S94MINFR	2A		2
S94MINRM	36		2
S94MTVCA	68		2
S94NUMBS	8		2
S94OPM_AVGCM	225		3
S94OPM_AVGTTM	22A		3
S94OPM_AVGVDM	216		2
S94OPM_CAFU	222		3
S94OPM_CHMNTS	28A		3
S94OPM_CHTIME	288		3
S94OPM_CMMNTS	28E		3
S94OPM_CMTIME	28C		3
S94OPM_DCI	220		3
S94OPM_DRIC	268		3
S94OPM_FRMNTS	286		3
S94OPM_FRMT	284		3
S94OPM_IARTAFRT	460		2
S94OPM_IARTCHM	466		2
S94OPM_IARTCHMT	464		2
S94OPM_IARTCMM	46A		2
S94OPM_IARTCMMT	468		2
S94OPM_IARTFRM	462		2
S94OPM_MAXCM	223		3
S94OPM_MAXTTM	226		3
S94OPM_MAXVDM	212		2
S94OPM_MIGMNTS	22E		3
S94OPM_MINCM	224		3
S94OPM_MINTTM	228		3
S94OPM_MINVDM	214		2
S94OPM_PDI	221		3
S94OPM_PLCU	260		2
S94OPM_PMC	260		3
S94OPM_PPWRITEN	234		3
S94OPM_RECMNTS	230		3
S94OPM_SDEMNTS	232		3
S94OPM_SPWRITEN	238		3
S94OPM_STGMNTS	22C		3
S94OPM_TVCA35	27C		3

CROSS REFERENCE

NAME	HEX OFFSET	HEX VALUE	LEVEL
S94OPM_TVCA4	26C		3
S94OPM_TVCA48	274		3
S94OPM_VDC	210		2
S94OPM_VM35	280		3
S94OPM_VM4	270		3
S94OPM_VM48	278		3
S94OPM_VVIC	264		3
S94PRICT	66		2
S94RCPRT	13		2
S94RVMPL	6C	04	3

S94SCSI	7		2
S94SRTCT	64		2
S94STATS	0		1
S94THRES	62		2
S94TOTAT	24		2
S94TVCS	4		2
S94VALPL	6C	80	3
S94WROVT	14		2
S940KB	C		2
S9416KB	10		2
S942KB	D		2
S9432KB	11		2
S944KB	E		2
S9464KB	12		2
S948KB	F		2

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

- *Continuous Availability S/390 Technology Guide*, SG24-2086
- *Continuous Availability Systems Design Guide*, SG24-2085
- *DFSMS Release 10 Technical Update*, SG24-6120
- *Enterprise Systems Connection (ESCON) Implementation Guide*, SG24-4662
- *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409
- *IBM @server® zSeries Connectivity Handbook*, SG24-5444
- *IBM Fiber Saver (2029) Implementation Guide*, SG24-5608
- *IBM TotalStorage Tape Products Family: A Practical Guide*, SG24-4632
- *DFSMS/MVS V1R5 Technical Guide*, SG24-4892
- *IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring,* SG24-2229
- *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632
- *Introduction to IBM S/390 FICON*, SG24-5176
- *Introduction to SAN Distance Solutions*, SG24-6408
- *FICON Native Implementation and Reference Guide*, SG24-6266
- *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning and Implementation*, SG24-6789

Other resources

These publications are also relevant as further information sources:

- *z/OS DFSMSdfp Storage Administration Reference*, SC26-7402
- *z/OS Object Access Method Planning, Installation and Storage Administration Guide for Tape Libraries*, SC35-0427
- *Implementing System-Managed Storage*, SC26-7407
- *z/OS DFSMShsm Storage Administration Guide*, SC35-0421
- *z/OS DFSMShsm Implementation and Customization Guide*, SC35-0418
- *z/OS DFSMSrmm Guide and Reference*, SC26-7404

- *z/OS DFSMSrmm Implementation and Customization Guide*, SC26-7405
- *Recovery and Reconfiguration Guide*, SA22-7623
- *z/OS DFSMSdss Storage Administration Reference*, SC35-0424
- *z/OS DFSMSdss Storage Administration Guide*, LY35-0116
- *z/OS Object Access Method Planning, Installation and Storage Administration Guide for Object Support*, SC26-0426
- *IBM System Storage TS3500 Tape Library Introduction and Planning Guide,* GA32-0559
- *IBM System Storage TS3500 Tape Library Operator Guide*, GA32-0560
- *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Introduction and Planning Guide,* GA32-0472
- *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473
- *IBM TotalStorage 3494 Tape Library Introduction and Planning Guide*, GA32-0448
- *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-00449
- *3490E Installation Planning and Operator's Guide*, GA32-0378
- *S/390 Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference*, GC38-0401
- *z/OS Hardware Configuration Definition User's Guide*, SC33-7988
- z/OS *MVS Initialization and Tuning Reference*, SA22-7592
- *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141
- *IBM TotalStorage Virtual Tape Server Performance.* This white paper can be found under "White Papers" at:

  http://www.ibm.com/support/techdocs
- *IBM TotalStorage Enterprise Automated Tape Library (3494) Systems Assurance Product Review (SAPR) Guide*, SA01-005-05, which is available from your IBM representative or can be found at:

  http://w3-1.ibm.com/support/assure/assur30i.nsf/PubAllNum/SA185?OpenDocument

Referenced Web sites

These Web sites are also relevant as further information sources:

- IBM internal site for tape tools. To locate the tools, select:
  Geography → **Storage** → **Tape and Optical** → **Tape Tools**

  http://w3-1.ibm.com/sales/systems/ibmsm.nsf
- Product manuals, including OAM, SMF, and messages:

  http://www.s390.ibm.com.bkserv
- Product specifications and information:

  http://www.storage.ibm.com/hardsoft/tape/3494_vts/ptp.html
- Technical flashes concerning VTS:

  http://www.ibm.com/support/techdocs

- Tools available for monitoring and reporting:

  `ftp://ftp.software.ibm.com/storage/tapetool/`

  `http://w3-1.ibm.com/sales/systems/ibmsm.nsf/docnames/tapetools`
- VOLREUSE, a new analysis tool:

  `ftp://ftp.software.ibm.com/storage/tapetool`
- VTSSTATS program download:

  `ftp://ftp.software.ibm.com/storage/tapetool/`
- White Papers:

  `http://www-1.ibm.com/support/search/index.html`

How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

ibm.com/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the **CD-ROMs** button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide

Find out about the IBM TS3500 with the IBM 3953 Tape System

Optimize your VTS management with APM

Improve performance with IBM 3592 tape drives

This IBM Redbook discusses enhancements to the IBM TotalStorage Virtual Tape Server (VTS) in the area of automated tape copying for z/OS systems.

These enhancements broaden the usage of the VTS extensively and make the VTS a possible replacement product for any existing z/OS tape solution. However, they also add another level of complexity in terms of planning for and implementing the Peer-to-Peer VTS solution.

Starting with a detailed description of the architecture and the design of the Peer-to-Peer VTS, we give you a comprehensive overview of all the planning tasks, provide guidance on a wide variety of possible configuration choices, and discuss several possible recovery scenarios as well as considerations for day-to-day operations.

A good knowledge of IBM virtual tape technology is assumed.

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

SG24-6115-03 ISBN 0738497088