IBM

IBM TotalStorage Virtual Tape Server:

Planning, Implementing, and Monitoring

Improved performance and capacity

Enhanced attachment capabilities

Advanced system management

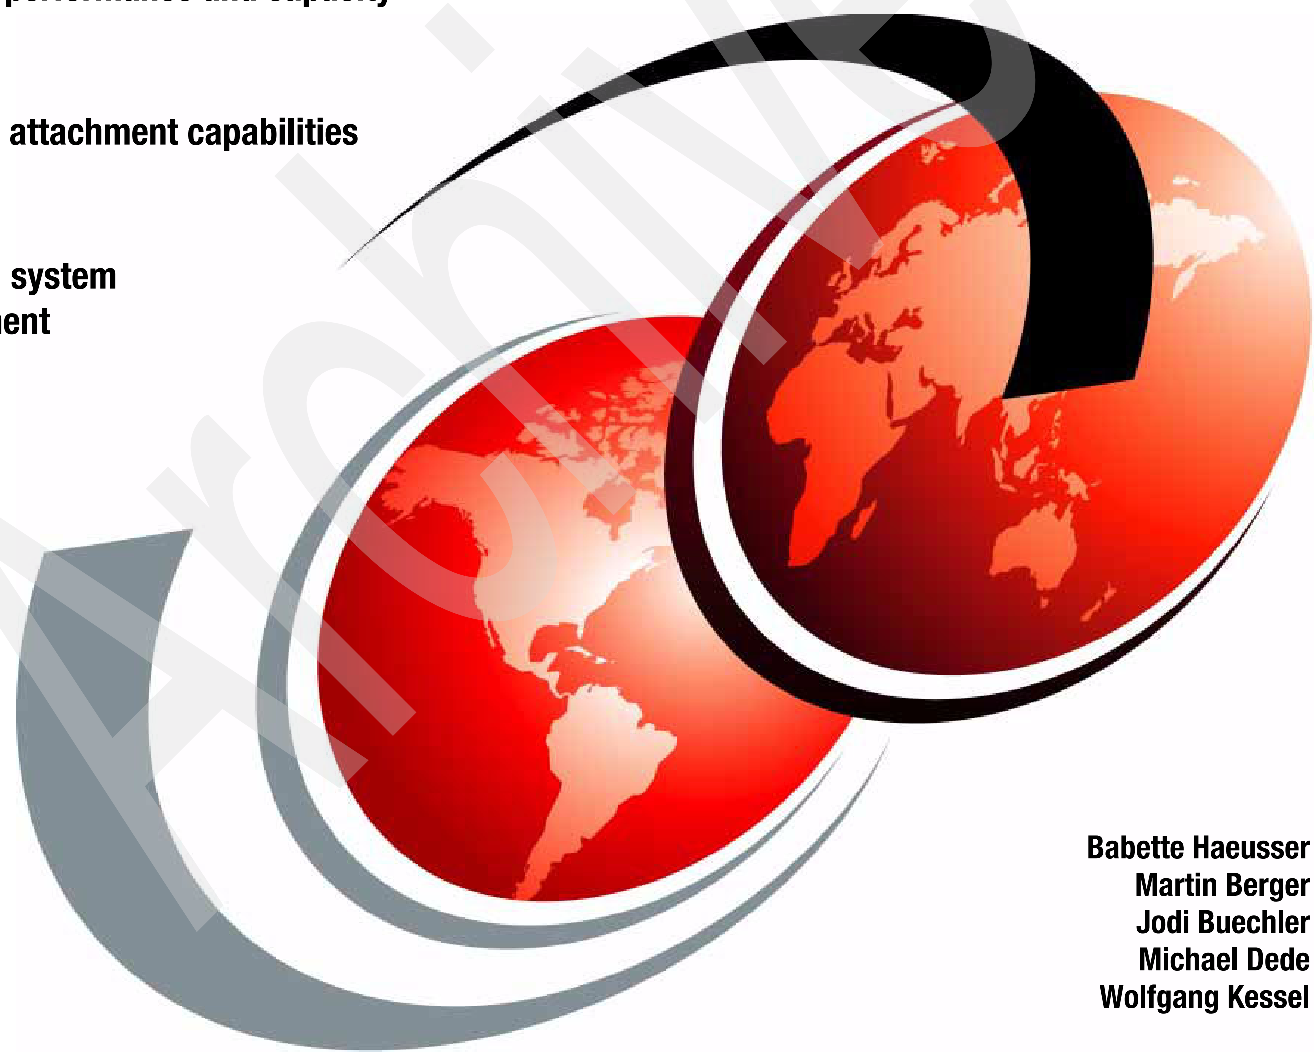

Babette Haeusser
Martin Berger
Jodi Buechler
Michael Dede
Wolfgang Kessel

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring

November 2005

SG24-2229-07

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

Eighth Edition (November 2005)

This edition applies to IBM TotalStorage Enterprise Virtual Tape Server models current at the time of publication.

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:

- @server®
- @server®
- Redbooks (logo) ™
- ^®
- pSeries®
- z/OS®
- z/VM®
- z/VSE™
- zSeries®
- AIX®
- AS/400®
- CICS®
- CUA®
- DB2®
- DFSMS/MVS®
- DFSMS/VM®
- DFSMSdfp™
- DFSMSdss™
- DFSMShsm™
- DFSMSrmm™
- DFSORT™
- Enterprise Storage Server®
- ES/3090™
- ES/9000®
- ESCON®
- FlashCopy®
- FICON®
- GDPS®
- IBM®
- IMS™
- Lotus®
- Magstar®
- MVS™
- OS/2®
- OS/390®
- OS/400®
- Parallel Sysplex®
- Redbooks™
- RETAIN®
- RMF™
- RS/6000®
- S/390®
- Seascape®
- System Storage™
- Tivoli®
- TotalStorage®
- VM/ESA®
- VSE/ESA™
- WebSphere®

The following terms are trademarks of other companies:

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

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

Intel, Pentium, 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.

BrightStor®, CA-1®, and BrightStor CA-Dynam®/TLMS are trademarks of Computer Associates International, Inc.

BrightStor®, CA-1® Tape Management Copycat Utility, and BrightStor CA-Dynam®/TLMS Tape Management Copycat Utility are trademarks of Computer Associates International, Inc.

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

Preface

This IBM Redbook is the seventh update (eighth edition) to the best-selling IBM® Magstar® VTS Implementation Guide, which was first published in 1997. Since the first VTS installations in May 1997, much has changed in the Virtual Tape Server area. The performance, capacity, and functionality of the VTS has been phenomenally enhanced. The focus of this book is to provide implementation and usage guidance for the latest changes to the IBM TotalStorage® Enterprise Automated Tape Library VTS models B10 and B20. Minor additions and corrections have also been applied throughout the book. If you discuss, sell, order, or plan for a Virtual Tape Server, please consult this redbook first!

The IBM TotalStorage Virtual Tape Server (VTS) is integrated with the 3494 Tape Library and with the IBM 3584 Tape Library using the IBM 3953 Library Manager. The host system perceives it as an IBM Tape Library with up to 256 virtual 3490E tape drives and up to 500,000 virtual 3490E cartridges, which can have an uncompressed capacity of 400, 800, 1000, 2000, or 4000 MB. The VTS fully exploits the capability of the IBM TotalStorage 3590 and 3592 tape drives and the capacity of their media.

The VTS introduced the virtual concept into tape subsystems, similar to virtual storage and virtual disk concepts that already have been implemented. The latest changes allow IBM 3592 tape drives to attach to the VTS installed with an IBM 3584 Tape Library, thus enhancing the VTS attachment options and performance once again.

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 in 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, herself focusing 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, and Michael

The authors of the previous four 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:

Yvonne Lyon
Deanna Polm
Emma Jacobs
International Technical Support Organization, San Jose Center

Joseph Bacco, Jennifer Bappert, Thomas Bish, Wayne Carlson, Erika Dawson, Jim Fisher, Dennis Haight, J.D. Metzger, Collette Morton, Christina Orosco, Jon Peake, Jeffrey Pilch, Craig Schultz, David Shields, Jesse Thrall
IBM Tucson

Immo Compas, Antoinette Kaschner, 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 SG24-2229-07
for IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring
as created or updated on December 1, 2005.

November 2005, Eighth 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 3584 / 3953 support
  - APM enhancements

1

Introducing the IBM TotalStorage Virtual Tape Server

The role and characteristics of tape storage have changed dramatically over the last 50 years. In the early days of computing, tape was the first electronic media not only to store data on, but it also supported the first Operating System. With the fast evolution of disk, tape became the media of choice for large data sets, exchange data, and backup. Cheaper disks with huge capacities became an economical alternative for tape; it was not glamorous at all. Despite all efforts to constrain the use of tape, its use is still growing exponentially. Moreover, high-end tape technology has exploded from 200 megabytes of raw capacity per cartridge up to today's 500 gigabytes of uncompressed data and even greater capacities are projected during the next couple of years.

When IBM first announced the IBM Magstar 3590 tape drive in April 1995, it was already clear that some means of using high-capacity tapes would also be required; this was expressed at that time as a statement of direction for Magstar "capacity enablement". Some applications and storage management products, such as Data Facility Storage Management Subsystem Hierarchical Storage Manager (DFSMShsm), are designed to optimize their own tape usage. However, our traditional uses for tape (archive, backup, recovery, and even batch processing) have not changed dramatically over the last 15 years and therefore have not kept pace with the technology. A survey of IBM customers shows that most of the cartridges are under-utilized, because it is not an easy job to stack multiple data sets or even volumes.

In the years since this high capacity tape has been available, tape data set stacking products and software solutions to increase tape utilization have evolved in response to customer requirements for more efficient ways of managing the information stored on tape. Currently, more than 80% of z/OS sites have a tape management system and some type of hierarchical storage management product in place.

However, market research companies estimate that only around 15% of these z/OS installations have any type of data set stacking. To date, tape data set stacking products do not seem to have achieved mass acceptance in the marketplace, perhaps because they usually involve post-processing activity using additional host resources and increasing tape mount activity overall, or perhaps because they are perceived as too expensive, or too complex to implement and to maintain.

For general tape environments, the introduction of a high capacity tape cartridge means that measured utilization will drop to less than 1% if current tape practices remain unchanged.

Consequently, to manage this continuing explosion in capacity effectively, it has become imperative to insulate the user from both the physical tape drive and the volume characteristics. IBM recognized this critical requirement and was the first to place its Virtual Tape Server (VTS model B16) in general availability in July 1997.

1.1 The “IBM TotalStorage 3494 Virtual Tape Server”

The VTS (shown in Figure 1-1) is a hardware-based solution that addresses not only tape cartridge utilization but also tape device utilization and hence overall tape subsystem costs. Because of its transparency to the host software, the VTS is readily supported and easy to use.

Figure 1-1 The 3494 and Virtual Tape Server

The VTS design is based on the IBM Storage Enterprise Architecture (SEASCAPE) concept. The principle to include built-in blocks of other product lines into the VTS exploits the rapid evolution of processors, disk, adapters and tape. The result is an outstanding increase in performance, capacity, and functionality since its first availability.

The VTS with the attached IBM Total Storage 3494 Tape Library (shown in Figure 1-1) for example can have a capacity of 163TB of uncompressed data. With the old 3490 technology you would need more than 200.000 cartridges.

1.1.1 VTS models

Since the introduction of the first IBM model B16 Virtual Tape Server, an additional three models have been made available. These are:

- Model B18 Virtual Tape Server (withdrawn from marketing November 30, 2001)
- Model B10 Virtual Tape Server
- Model B20 Virtual Tape Server

As an additional configuration, the Peer-to-Peer VTS is available with the above models. By enhancing data backup and recovery capabilities, the Peer-to-Peer VTS configuration provides an outstanding business continuity solution for tape customers as well as being designed to enhance data availability and to eliminate single points of failure. For an overview of the PtP VTS, refer to Appendix A, “Peer-to-Peer VTS solution” on page 363.

In addition, there is a High Capacity Cache Option available for VTS model B20. It is dedicated for content management applications and replaces the IBM 3995 optical archive.

The VTS models B10 and B20 are based on the IBM eserver pSeries® 660 architecture, which features significantly higher performance and expanded I/O capability. The Model B10 VTS is configured as a cost effective solution for modest throughput requirements, while the Model B20 VTS establishes higher standards for throughput performance as well as for the number of virtual devices managed.

1.1.2 VTS development

To illustrate the continuous product enhancement, the FICON® attached VTS Model B20 has a host data peak throughput up to forty times that of the original Model B16 VTS in 1997. Additional connectivity options are available with the introduction of FICON interfaces for the VTS models B10 and B20. Up to 256 virtual drives are available and up to 500,000 virtual volumes. The IBM TotalStorage Enterprise 3592 media used has an uncompressed capacity of 300 GB.

Other features have included these:

- SCSI Host Attachment feature
- Import/Export feature
- Extended host connectivity with up to eight FICON channels or sixteen ESCON® channels
- FICON Performance Accelerator Feature (FPAF)
- Advanced Policy Management (APM)
- Enhancements in Tape Volume Cache (TVC) capacity and performance
- Support of the IBM TotalStorage 3494 Tape Drive Model J1A
- Monitoring tools
- IBM TotalStorage 3584 Tape Library attachment

1.1.3 VTS benefits

VTS gives you the following benefits:

- **It fills the tape cartridge up to 100%**. Putting multiple virtual volumes into a stacked volume, VTS uses all of the available space on the cartridge. VTS uses IBM 3590 and 3592 cartridges when stacking volumes.
- **DFSMShsm, TMM, or Tape Management System is not required**. VTS is a hardware stacking function; no software prerequisites are required.
- **JCL modifications are not required**. The JCL references a virtual volume and does not need modification.
- **More tape drives are available**. Data stacking is not the only benefit; VTS emulates up to 256 virtual drives, allowing more tape volumes to be used concurrently.
- **There is little management overhead**. The stacking of data is performed by the VTS management software, without host knowledge.

- **No new software support is required**. VTS works with all current releases of z/OS and SMS. No additional software product is required for its implementation in your current environment.

The IBM VTS solution helps to reduce total cost of ownership, it needs no separate host management software, and it attaches to all platforms. It is the right choice for infrastructure simplification by the reduction of cartridges, drives, and tape libraries. It fits into information lifecycle management and is a need for disaster recovery requirements.

The VTS solution supports IBM zSeries® hosts. HP, AIX®, Sun™, NT, and Linux® hosts can also connect directly to the VTS and take advantage of the virtualization features.

1.2 VTS compared to traditional tape stacking methods

We compare traditional and software based tape stacking with the VTS. The comparison is made only for the most common techniques used in an IBM environment as shown in Table 1-1.

Table 1-1 Benefits of VTS compared with traditional stacking methods

	VTS	TMM	User JCL stacking	Pre-alloc stacking SW	Post-alloc stacking SW
Cartridge 100% used	Yes	Yes	Yes	Yes	Yes
No JCL changes	Yes	Yes[a]	No	Yes	Yes
Media awareness	Yes	Yes	No	Yes	Yes
No extra management	Yes	No	No	No	No
No additional SW	Yes	Yes	Yes	No	No
No additional DASD	Yes	No	Yes	Yes	Yes
No additional mounts	Yes	Yes	Yes	Yes	No
No additional host processor cycles	Yes	No	No	No	No

a. Data sets must be under catalog control for TMM. You must change JCL that does not catalog their data sets.

TMM (Tape Mount Management) has the following characteristics:

- It fills the tape cartridge completely because TMM tape data is DFSMShsm managed.
- It is not dependent on media type. DFSMShsm is aware of MEDIA type used.
- In most cases, and if tape data sets are under catalog control, TMM methodology does not require JCL modifications.
- It needs SMS installed and active. TMM uses SMS to redirect selected allocations from tape to DASD.
- It needs DFSMShsm installed and active.
- TMM then uses DFSMShsm to move DASD data back to tape.
- It needs a DASD buffer to store data that later will be stacked on tape.
- TMM is not appropriate for all data. Some data cannot use TMM and some workloads do not benefit from using TMM. Refer to *Implementing System Managed Storage*, SC26-3123, for more information.

User JCL stacking has these characteristics:

- It requires many JCL modifications. To stack data onto tape, you have to modify the JCL considerably to use multiple labels on the same volume.
- It depends on media type. Every time the capacity of the media you are using changes, JCL modifications are required.
- It depends on the size of the data set. If the size of the data set being written changes, you must modify the JCL to fill the output tape correctly.
- It has a high management overhead. With the previous considerations, the management of tape stacking imposes a very high overhead.

Stacking software (pre-allocation or post-allocation) has these characteristics:

- It needs a specific tape management system installed and active. Some tape management systems have stacking functions. You have to install, activate, tailor the product, and possibly migrate from another tape management system to use these stacking functions.
- JCL modifications are required to read uncataloged data. Stacking is performed either before data is written on tape or tape data is consolidated after the job completes. The file number and the volser specified in the JCL are probably not the same as specified in the JCL. When you need to read from tape and the data set is not cataloged, you have to find the volume and label number that holds your data.
- It has a high management overhead. Human resources are required to tailor and manage the stacking process and the tape management system.

1.2.1 The z/OS software based virtual tape solution concept

The z/OS software based virtual tape solution claims to be a highly scalable, flexible and cost effective storage solution that allows customers to take advantage of virtual tape concepts while protecting their investment in existing z/Series hardware technology. The z/OS software based virtual tape solution redirects z/OS tape I/O to z/OS memory buffers. Virtual volume data is then staged from memory to standard zSeries Disk devices for virtual volume creation, interim storage and possible reuse.

Following virtual tape EOV processing, the z/OS software based virtual tape solution stacks the virtual tape onto a physical tape. This physical tape can be mounted on either stand-alone or automated tape drives. Once stacked on a tape cartridge the virtual tape volumes on the DASD 'buffer' remain available for recalls and the space is then released according to the disk management LRU management. After this space release takes place subsequent recalls for the virtual volume will require that the volume be staged from the tape to the buffer prior to system access.

1.2.2 Advantages and disadvantages of software based tape solutions

Perceived advantages of a z/OS software based virtual tape solution, in general, are:

- It minimizes new zSeries hardware investments and optimizes existing z/OS resources.
- It works with any zSeries disk subsystem.
- It works with any zSeries tape subsystem (no automation is needed).
- It is highly scalable.
- It has unlimited cache.
- It can emulate IBM 3480 or IBM 3490 Base models.
- It is easy to evaluate — there is no hardware to install.

Disadvantages of software based virtual tape solutions, compared to the VTS, are:

- z/OS software based virtual tape solutions are another TMM-like method with similar characteristics (it is zSeries processor intensive and requires substantial management resources). The VTS is an outboard tape management solution that doesn't require extra processor power and needs minimal resources to manage because it is a self managed system.
- There is no sharing of virtual volumes with other systems. The VTS can share virtual volumes across multiple zSeries with VM/VSE, VM, z/OS, TPF hosts and with Open System hosts. Special care must be taken to ensure that only the primary system controls scratch processing.
- Normal tape data flows three times over the channel infrastructure. With a VTS, tape data flows only once over the channel infrastructure.
- Incremental DASD is needed for the interim buffer; the tape data is not software compressed normally because it costs too much in processing cycles. A VTS with EHPO feature compresses Tape Volume Cache with a typical 3:1 compression ratio on the ESCON channel adapter cards.
- Incremental processor memory is needed for virtual device emulation. These resources are inherent in the VTS outboard design.
- Incremental processor cycles are needed for a z/OS software based virtual tape solution. VTS does not require extra processor cycles at all.
- Extra software costs are associated with a z/OS software based virtual tape solution. IBM's VTS needs no extra software.
- For a software based virtual tape solution, extra software needs to be installed and maintained by the customer on the host.
- Tape hardware changes need involvement of the hardware and software. Changes in a z/OS software based virtual tape solution and O system software are a frequent requirement, requiring skilled resources. Tape hardware changes in the IBM VTS are transparent to the operating system. The z/OS is only aware of virtual 3490E cartridges and tape drives and changes of the physical tape drive are transparent to the attached hosts.

Here are some questions one should ask considering a z/OS software based virtual tape solution:

- How much memory is needed to support up to 256 virtual tape drives?
- How many MIPS are required to support the desired workload?
- How much disk cache is required to provide the buffer for the virtual volumes?
- Can the extra channel load be accommodated without affecting the other batch processes?
- What will be the effect on resource contention and DASD subsystem throughput?
- What total throughput will be derived? Which software product will migrate the data to tape? What are the limitations of that software migration product? Can that software product fill the tapes fully and exploit the tape drive fully?
- How much of the operating budget will be spent on systems personnel to plan, implement, and manage the installation?
- What are the future functions of the z/OS software based virtual tape solution?

Conclusion: A z/OS software based virtual tape solution requires significant processor cycles and channel resources. The z/OS software based virtual tape solution supplies no savings with regard to disk or tape resources at all and will require extra manpower to manage. As shown above, a VTS solution relieves your resources.

2

Architecture, functional characteristics, and components

In this chapter we describe how the VTS works, which components it needs, and how data becomes virtualized. We review the first and second generation VTS and describe the most current VTS models 3494-B10 and B20:

- We describe the VTS Architecture.
- We show how the VTS works.
- We introduce the new terminology.
- We look at the logical and physical components.
- We introduce the VTS attachment to IBM 3494 and 3584 tape libraries.
- We explain the concept of Advanced Policy Management.

2.1 VTS concepts

The VTS was introduced available in 1996 as a new breed of storage solution that combines a high-speed disk cache with tape automation, tape drives, and intelligent storage management software running on a server. The disk cache associated with the VTS acts as a buffer to the tape drives, providing near-instantaneous performance for multiple, simultaneous *scratch-mount* requests and for specific mount requests for tape volumes that reside in the disk cache.

A VTS breaks the one-to-one connection between a logical tape drive and a physical tape drive, enabling logical access to significantly more tape drives than are physically installed. In addition, a VTS breaks the one-to-one connection between a tape cartridge and a tape volume. One key reason tapes are significantly under utilized is that a single application may *own* a particular drive and the associated tapes. If that application does not fully utilize the associated tape capacity, it is essentially wasted.

The VTS writes the volumes created by attached servers to the high-performance serial disk array, the Tape Volume Cache (TVC). If the volumes are re-referenced, the VTS reads them back at disk speeds. This cache-based process helps reduce the need to mount a cartridge to satisfy a mount request, eliminates many of the physical delays associated with tape I/O, and increases the performance of the data transfer operation.

Shortly after the creation of the volume in the TVC, where integrity and protection are provided by the RAID 5 array, the volume is written to the attached back-end tapes in the automated tape library as a *logical volume* in either a first-in or first-out order. There is neither wasted space in the cache nor on the physical tapes because only the actual amount of data is written. In addition, the data stored in the TVC and on the back-end tapes is compressing, greatly increasing the effective storage capacity of the VTS.

There are multiple benefits to this process. Multiple jobs can run simultaneously, because there is no longer a limiting number of real drives. The capacity of the physical cartridges is used up to their full capacity. Considering all jobs are written to cache, the backup performance is greatly increased. Migration to new technology is done by the system by copying all logical volumes in the background. Figure 2-1 describes the virtual tape concept.

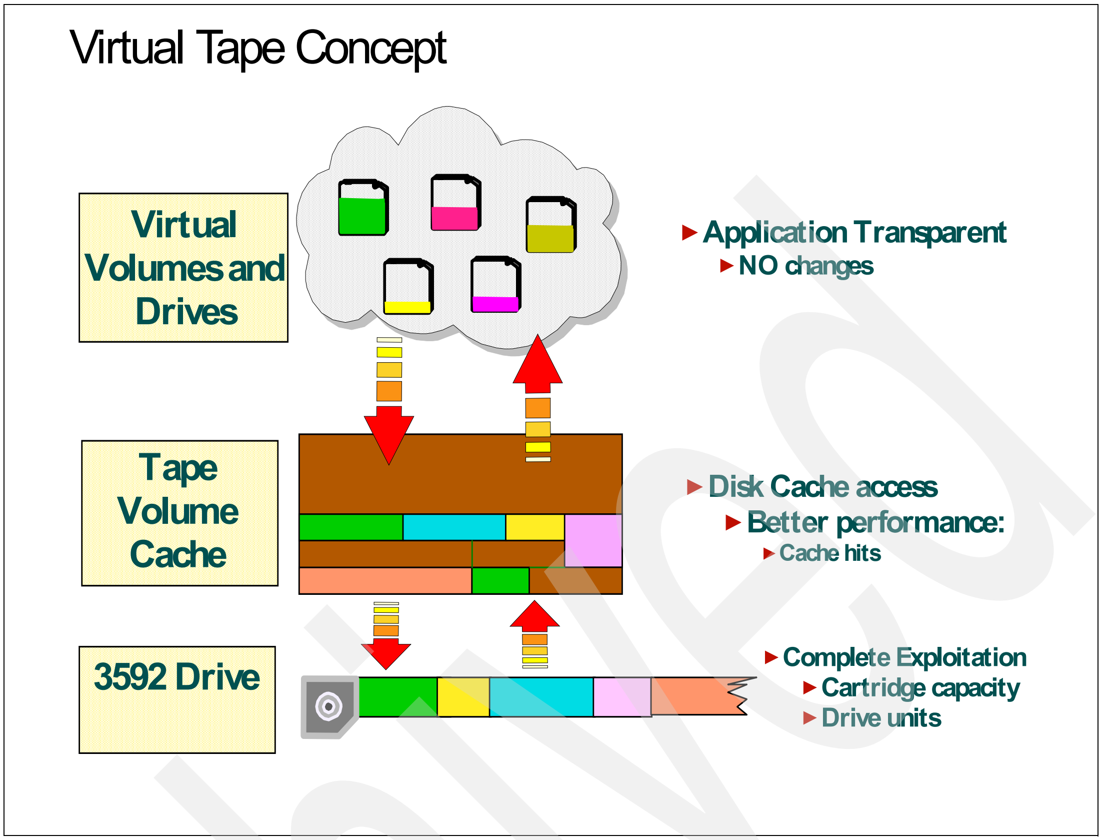

Figure 2-1 The virtual tape concept

With the Advanced Policy Management (APM) feature, you also have control over the handling of logical and physical volume by using DFSMS Automated Class Selection (ACS) routines. Grouping of logical volumes onto a specified range of physical volumes can be accomplished by the physical volume pooling support. Volumes remain in the Tape Volume Cache (TVC) for extended periods of time, depending on their re-reference pattern and the influence of the APM specifications.

Extended residency times help to ensure that volumes created in the batch window are “cache hits” when they are re-referenced later. By enabling longer residency times, larger cache sizes can provide faster access to critical volumes. The APM feature can also enable dual copies of data which will be written to separate physical volumes. If there is a need to explicitly erase expired data on a physical cartridge, the secure data erase function can be enabled on a physical pool basis. Though APM settings, support is provided to use larger logical volumes of up to 4,000 MB uncompressed capacity.

2.1.1 Logical components

All IBM virtual tape solutions are integrated within a 3494 Tape Library or, with R7.4 which is available since May 2005, within an IBM 3584/3953 Tape Library configuration. They appear to the host as a system-managed tape library with their own library name, their own set of tape volumes (each VTS can have up to 500,000 logical volumes) and their own set of tape drives. The host sees 64, 128, or 256 virtual 3490E drives, attached by up to:

- Sixteen Enterprise System Connection (ESCON) paths
- Eight SCSI paths that transfer data concurrently
- Eight FICON channels

You can combine ESCON with FICON or SCSI channels. However, FICON channel types cannot be intermixed with SCSI attachments.

The VTS can be shared between operating systems running on a zSeries platform and Open Systems by specifying up to four additional SCSI Host Attachment features. Each feature provides two SCSI buses and includes data compression. Up to 64 virtual drives can be assigned to Open Systems.

Currently two VTS subsystems can be integrated into one IBM 3494 library sharing the maximum number of 1,000,000 logical volumes equally and providing up to 256 virtual drives each. The VTSs can coexist effectively in the same IBM 3494 Tape Library with native IBM 3490E, native IBM 3590, and native 3592 tape cartridges and drives.

The IBM 3584 Tape Library supports up to four 3953 Library Managers per physical tape library. Two VTS subsystems can be included in one 3953 Library Manager partition sharing the maximum number of 1,000,000 logical volumes equally and providing up to 256 virtual drives each. Up to four IBM 3953 Library Managers and up to eight VTS systems can be attached to one IBM 3584 Tape Library, providing up to 4,000,000 virtual volumes in a single physical tape library. The VTSs can coexist effectively in the same IBM 3584 Tape Library with native IBM 3592 attached to J70 controller or to Open System hosts as well as LTO tape drives.

With access to a virtual library with 500,000 emulated cartridges and 256 emulated drives, the host can process much more tape workload in parallel without the expense of actually installing 256 real tape drives. With so many virtual tape drives, there is the opportunity to dedicate drives to different hosts and platforms rather than managing an often complex shared environment.

Figure 2-2 illustrates how the VTS appears to the ESCON or FICON-attached host, and to the SCSI-attached host. The VTS is managed by the Library Manager in the IBM library, so that the host sees a logical library, separate from any other drives within the IBM 3494 or IBM 3584, with its own library name (see 2.5, “Virtual Tape Server configuration in the library” on page 33).

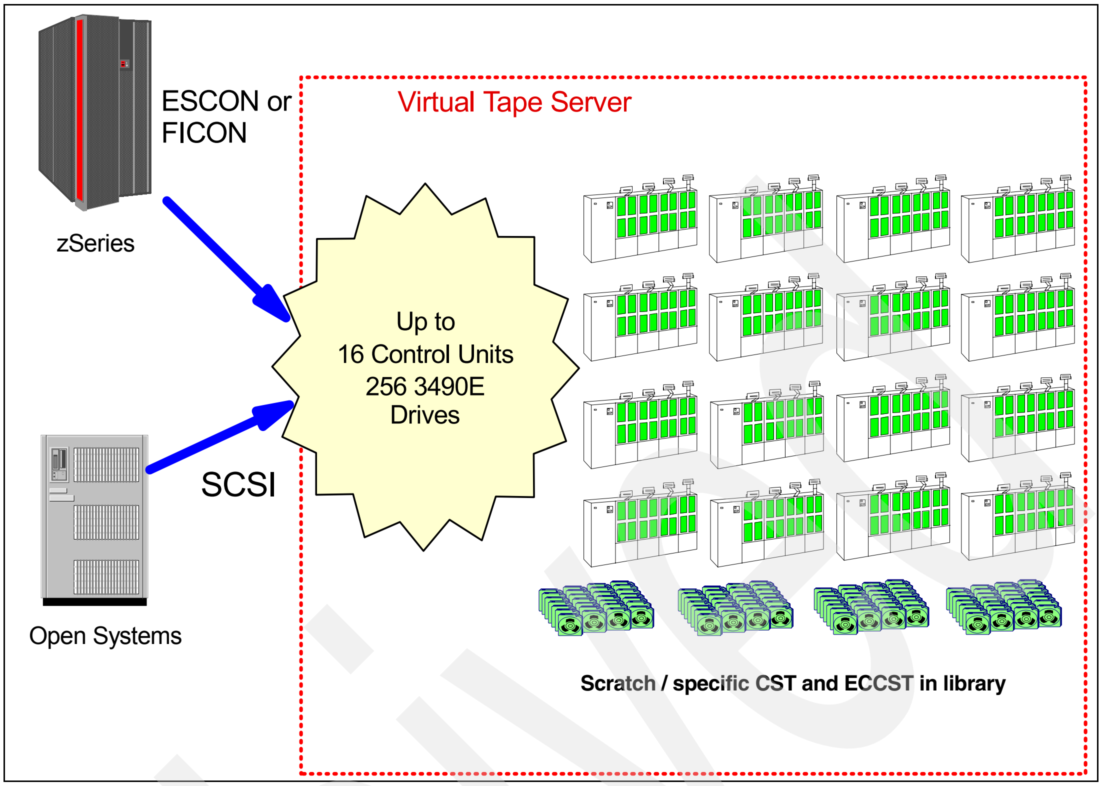

Figure 2-2 Logical components of the VTS as seen by the host system

The internal structure of the VTS is unknown to the host; the host sees the VTS as four to sixteen complete IBM 3490E strings inside an IBM tape library. The host is not aware of the physical cartridges in the VTS and does not need to provide support for the IBM 3590 drives or IBM 3592 drives nor is the host aware of the internal DASD used for the Tape Volume Cache (TVC). Therefore, to use the base functionality of the VTS, you do not have to make changes to the operating system software as long as it has the standard 3494 and 3490E support. Even a VTS attached to a IBM 3584 Tape Library using the IBM 3953 Library Manager does not require additional changes in the host software.

2.1.2 Mounting a scratch tape

The Library Manager uses categories to group volumes. After virtual volumes are defined through the Library Manager console, they are placed in the insert category and handled exactly as native cartridges. When the IBM VTS is varied online or during insert processing, the host operating system assigns scratch volumes to a particular category. When a request for a scratch is issued to the 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.

If the volume is mounted for the first time (that is, the volser has never been used before as a virtual volume), the VTS controller software generates a set of records for the volume in the TVC as if the volume had been initialized using the EDGINERS or IEHINITT program, with an IBM standard label (that is, the volume will contain a VOL1 record, an HDR1 record and a tape mark).

Once a virtual volume has been used for the first time (initialized), information about it is maintained in the TVC. After the volume has been copied to a physical cartridge, also called a *stacked volume* and the cache data space has been made available for new volumes, 4K size of data from the last use of the volume, called the *stub*, remains in the TVC. The stub contains the first several hundred bytes of data of the volume, including, at a minimum, the tape label information.

The z/OS operating system uses two of the Library Manager categories for 3490E cartridges: either Media type 1 (standard CST), or Media type 2 (ECCST). With the introduction of the VTS, a fast-ready attribute must be assigned to a category to tell the VTS storage management software that it does not need to recall any data from tape into the TVC before signaling mount complete to the host. In other words, the operating system can use the stub information in the TVC for volume verification and tape management system integrity checking before using a logical volume for a scratch mount and no physical tape mount is required.

Tip: You must assign the fast-ready attribute to all of your VTS scratch categories. This ensures optimal scratch mount performance for host scratch mount requests.

Having checked whether the selected volume has stub information in the TVC, or if not, having created the information, the VTS storage management software signals mount complete to the host, which can then start to write data to the virtual volume. Because a physical mount of a volume is not required, the nonspecific scratch mount request is handled much faster than in a conventional tape library.

With Advanced Policy management (APM), DFSMS construct names are passed to the Library Manager. Upon acceptance of the mount request, the Library Manager stores the constructs in its database.

When the host closes the volume, end-of-volume (EOV) processing is performed and the volume is rewound and unloaded, but without any physical tape movements. At this point, with Advanced Policy Management installed, all passed DFSMS construct names are checked and determinations made as to what action to take — for example, create a dual copy of the logical volume (see 2.6.2, “Selective dual copy” on page 47).

2.1.3 Buffering data in the Tape Volume Cache (TVC)

The TVC is the physical disk buffer used for transferring data to and from the host. It operates as an intelligent buffer or cache. Like a cache, it is designed to operate full of data and has unused space only when the VTS is first installed. The data in the TVC is not automatically freed when the data occupying that space has been copied to a stacked volume. Once the TVC has filled with tape volumes (virtual volumes), space taken up by older volumes is freed for new allocations, with the default selection policy being the least recently used (LRU) algorithm.

If a virtual volume is written and then accessed again within a short time, the data may still reside in the TVC even if it has been copied to a physical cartridge. The access time is less than that to mount a real tape volume because it avoids the necessity of a physical tape mount.

The time stamp is updated in this case and the logical volume is placed back on the copy queue if it had not been copied before.

If a logical volume has not been copied to a stacked volume and it is mounted again and modified or appended to, the data in the TVC is overwritten or additional data is added to the volume. After the application closes the volume or end of volume (EOV) occurs, the logical volume becomes eligible for copy again.

Note: The length of time that data remains in the TVC cache depends on how active the VTS is and how much data is being written. It can also be influenced by TVC Cache Management features such as FC4000 - 4004. More detailed information can be found in 2.6.3, “Tape volume cache management” on page 50.

Once the virtual volume has been initialized (after the first use of a logical volser), a stub file associated is always with it in the TVC; otherwise the full virtual volume information is present in the TVC. The stub uses space sufficient to include the VOL1, HDR1 and HDR2 records, any URL1 through URL8 records if present and a tape mark. Even if the maximum number (500,000) of logical volumes has been initialized, the stub files use only a fraction of the space available in the TVC.

2.1.4 Copying the data to tape

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

2.1.5 Mounting a specific tape

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

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

Note: Running an IEFBR14 with a virtual volume in the DD statement will not cause the logical volume to be recalled to the TVC: IEFBR14 issues a MOUNT request to the Library Manager, immediately followed by a REWIND, UNLOAD and DEMOUNT. The Library Manager cancels the virtual volume recall and the request is dropped from the VTS recall queue.

However, for a JES3 environment, depending on the implementation, this may be a way to pre-load logical volumes.

In a z/OS environment, in order to mount a specific volume in the IBM VTS, that volume must reside in a private category within the library. DFSMS OAM prevents a scratch volume from being mounted in response to a specific mount request. The library itself does not consider the category of a logical volume specified in a specific mount request.

2.1.6 Expired volumes

Expired volume status and recovery changed in conjunction with the announcement of the Model B10 VTS and Model B20 VTS. With Library Manager Licensed Internal Code (LIC) Level 526 and above, you can optionally have immediate or deferred release of expired volumes. This means that expired logical volumes will not be copied to new stacked volumes in the event of a reclamation. This may also have implications as to the number of logical volumes you wish to define to the VTS. For a full description of this facility, refer to 6.6, “Expired volume management” on page 245.

Library Manager without expired volume management

Prior to Library Manager LIC Level 526, when a logical volume expired, it was returned to scratch status. However, the data associated with that logical volume remained available on the stacked volume until the logical volser is reused as a scratch volume and new data was written on it. The space on the stacked volume used by a logical volume was not invalidated until the scratch volume had been rewritten with new data, which occurred only if unused virtual volumes were not available. It was therefore important to specify only as many virtual volumes as needed.

> **Note:** The best way to understand this is to consider what happens when a real physical 3490E volume expires and the physical cartridge is returned to scratch. In this case it is easy to see that the expired data remains available on the cartridge until the volume is actually reused for new data.

2.1.7 Secure data erase

As described above, expired data on a physical volume remains readable until the volume has been completely overwritten with new data, even though the VTS has a non-standard data format and that data would need IBM assistance to be recovered.

Some customers are concerned that a court order could expose them to liability and cost to be able to try to *find* an old version of a data volume. Another concern is security of old data.

VTS Release 7.4 adds physical volume erasure on a physical volume pool basis controlled by an additional reclamation policy. It requires the APM feature. 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 cartridge is not available as a scratch cartridge as long as its data is not erased.

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

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

The VTS keeps track of the amount of active data on a physical volume. It starts at 100% when a volume becomes full. Although the granularity of the percent of full the VTS tracks 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 by:

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

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

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

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

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

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

The 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 with care when using this new function. It is explained in more detail in 3.9.6, “Secure Data Erase” on page 100.

2.1.8 Reconciliation and reclamation

Any time a logical volume is modified, the data from the previous use of the logical volume which is on a stacked volume becomes obsolete. When you modify a logical volume, you have a new virtual volume in the TVC and the copy on the stacked volume is invalidated.

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

A reconciliation normally occurs sometime during a 24 hour period. It can occur more often if it is forced by a reclamation process. A reconciliation will occur before every reclamation process. This will make the reclamation process more effective as it will delete more invalidated logical volumes from the active volume list and therefore reduce the threshold percentage of active data on more stacked cartridges.

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

2.1.9 VTS internal space management

The VTS automatically manages two levels of storage: the TVC (disk) and the physical stacked volumes (tape cartridges). External intervention is not required for space management, although you have the option of suppressing tape reclamation at certain times of the day and you can add more physical cartridges as and when you need more tape space. In addition you have the option of changing the VTS threshold related to stacked volume occupancy.

2.2 VTS terminology

VTS brings a new concept of volumes and devices because there is a difference between the host view of the system and the reality of the installed hardware. When using a VTS, the host application writes tape data to virtual drives. The volumes written by the host are physically stored in a tape volume cache (a RAID disk buffer) and are called *Virtual Volumes*.

The storage management software in the VTS controller copies these virtual volumes in the TVC to the physical cartridges owned by the VTS subsystem. Once a virtual volume is copied from the TVC to tape, it is called a *Logical Volume*.

The host cannot distinguish between physical and virtual volumes, or physical and virtual drives, and treats them as if they were “real” cartridges and drives because the host's view of the hardware is virtual. Therefore, there is no need to make changes to the host operating system software. The management of the tape drives and physical cartridges is completely under control of the VTS storage management software.

Figure 2-3 describes the relationship between several virtual and real images of components that are shown to the system differently.

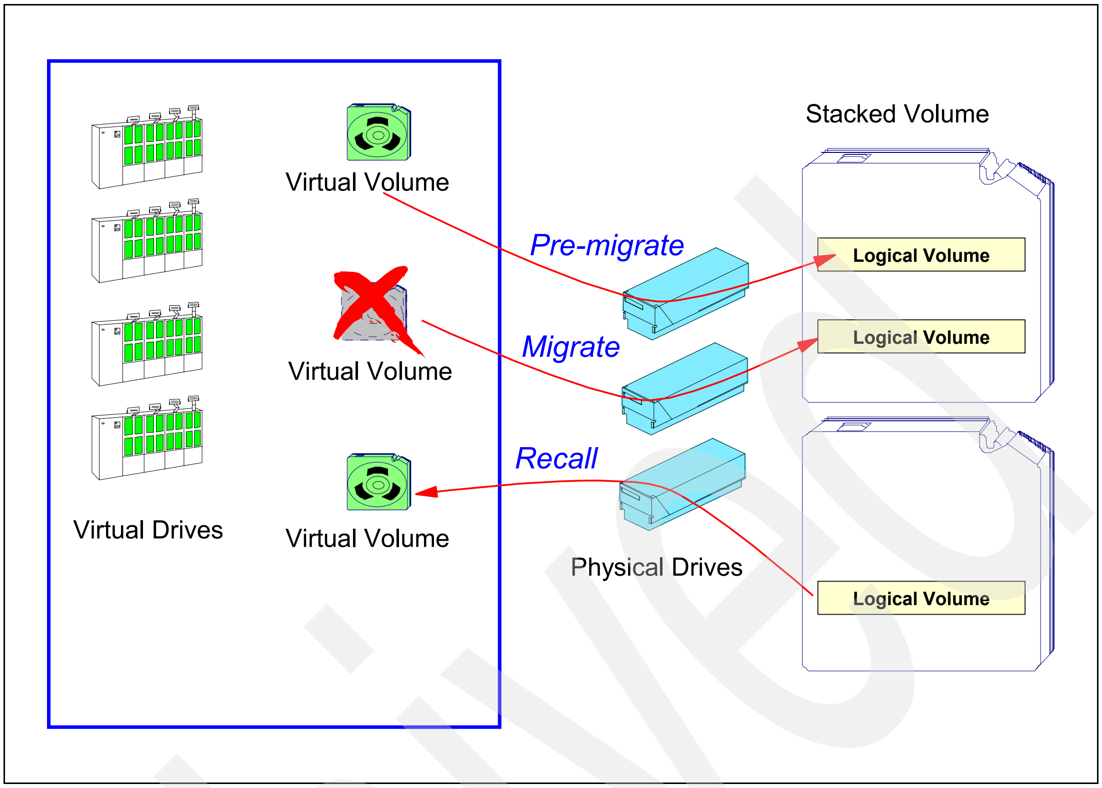

Figure 2-3 Relationship between virtual and real components without APM

When a virtual volume is written to a physical cartridge, it is *pre-migrated*. When it is removed from the TVC and only exists on the cartridge, it has been *migrated*. When a host system wants to access the logical volume, it has to be *recalled* into the TVC prior to completion of the mount request.

2.2.1 Tape Volume Cache (TVC)

The TVC is under complete and exclusive control of the VTS. When a virtual volume in the TVC is closed and demounted, it is scheduled to be copied to a stacked volume. Space for new virtual volume data or recalled existing logical volumes is achieved by retaining in the disk space only the stub for the LRU volumes. The only external interface with which to influence this is Tape Volume Cache Management, which comes with either the Advanced Function feature (FC4000) or Advanced Policy Management (FC4001-FC4004). For further details on this feature, see 2.6.3, “Tape volume cache management” on page 50.

2.2.2 Virtual volumes

A virtual volume is created in the TVC when the host writes data to the VTS subsystem. All host interaction with tape data in a VTS subsystem is through virtual volumes and virtual tape drives.

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

- It has a unique volume serial (volser) number known to the host.
- It is loaded and unloaded on a virtual device.
- It supports all tape write modes, including Tape Write Immediate.
- It contains all standard tape marks and data blocks.

- It supports an IBM standard label (and non-labeled tape as well).
- It can be appended to after it is initially written from the beginning of tape.
- The application is notified that the write operation is complete when the data has been written in the TVC.
- Each host written record has a logical block ID.
- End of volume is signaled when the total number of bytes written into the TVC (after compression by the EHPO), has reached 400 MB for an emulated cartridge system tape (CST), 800 MB for an emulated enhanced capacity cartridge system tape (ECCST) volume or 1000, 2000 or 4000 MB using a new option of the Advanced Functions with R7.4.

With data compression based on the IBMLZ1 algorithm by the ESCON, FICON and SCSI channel card in a VTS, the actual host data stored on a virtual CST or ECCST volume can vary from 1.2 GB up to 12 GB, respectively (assuming a 3:1 compression ratio).

The Advanced functions are enhanced by the extension of the default logical volume sizes of 400 and 800 MB 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.

Virtual volumes can exist only in a VTS. You can direct data to a virtual tape drive by directing it to a system managed storage (SMS) tape storage group inside the VTS, using the automatic class selection (ACS) routines in a system-managed tape environment or using Basic Tape Library Support (BTLS) in any other z/OS environment. With Advanced Policy Management you also can pass data class name, management class name, and storage class name as well as the storage group name. The VTS then uses these class names to determine Outboard Policy Management (OPM) on the virtual volume. The details on Advanced and Outboard Policy Management are covered within 2.6, “Advanced Policy Management (APM)” on page 38.

You can define up to 500,000 virtual volumes per VTS, making even a small library look very large. The volsers for the logical volumes are defined through the Library Manager Console. You should also associate a “Fast Ready” attribute with a scratch category of VTS virtual volumes, to speed the scratch mount process, as explained in more detail in 4.3.5, “Define Fast Ready categories” on page 133. Although you may want to use logical volume sizes larger than 400 or 800 MB you still define CST and/or ECCST emulated cartridges to VTS, simulating MEDIA1 with 400 MB capacity, or MEDIA2 with 800 MB capacity. Virtual volumes go through the same cartridge entry processing as native cartridges inserted in a library.

2.2.3 Logical volumes

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

Figure 2-4 illustrates the relationship between virtual and logical volumes.

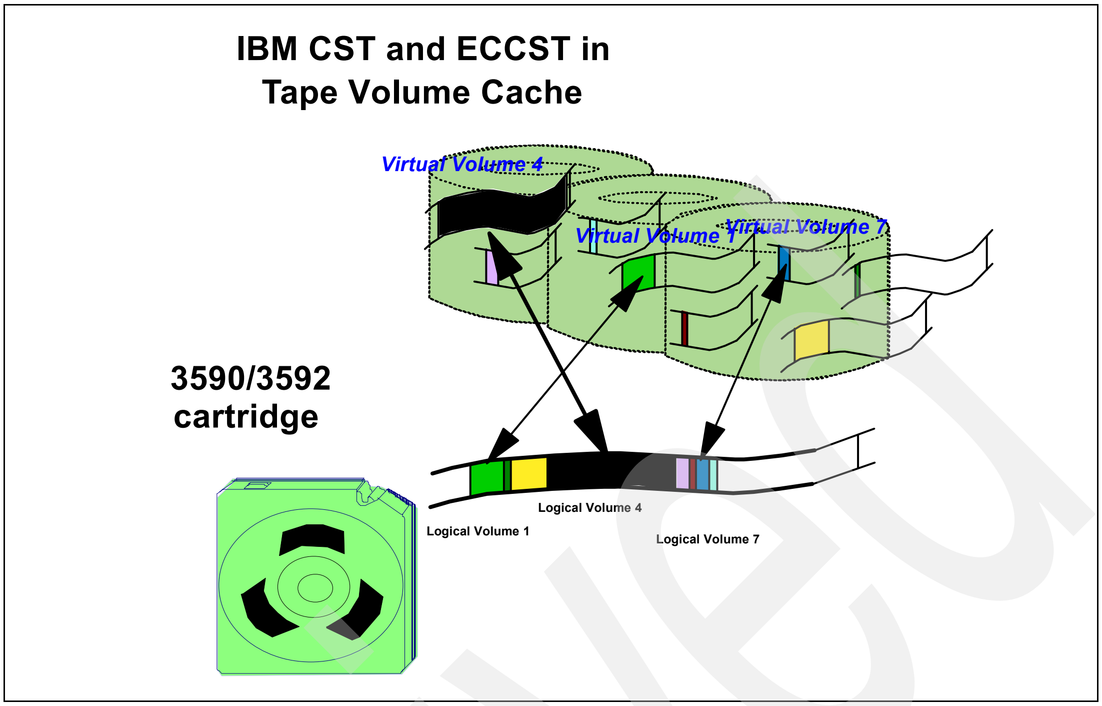

Figure 2-4 Data in the TVC and on physical tape

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

As virtual volumes are copied from the TVC to a physical cartridge, they are stacked on the cartridge end to end, taking up only the space written by the host application. This arrangement maximizes utilization of a cartridge's storage capacity. The storage management software within the VTS manages the location of the logical volumes on the physical cartridges. The only control the user has over the location of the data is where volume pooling is used as part of Advanced Policy Management. Without Advanced Policy Management there is no control.

2.2.4 Increased logical volume sizes

With R7.4, you have the choice of additional logical volume sizes through 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 according to the size specified in the data class on the Library Manager. All definitions are outboard and may be defined at the Library Manager console or via the Web specialist.

A logical size of 1000 MB is available automatically after the new code Release 7.4 is installed. Volume sizes of 2000 and 4000 MB have to be enabled by an IBM System Service Representative (SSR). A volume size of 4000 MB requires a VTS model B20 with 3592 drives and the maximum TVC capacity of 1.7 TB uncompressed. The total logical volume storage capacity is increased by this function up to factor ten (from 400 to 4000 MB multiplied with the number of logical volumes defined). However, you should consider that the physical cartridge capacity of the VTS needs to match the increased local volume sizes.

These new large logical volume sizes still use the only two media types (MEDIA1 and MEDIA2), which are inserted with the “volume Insertion” function during the installation and implementation process for the VTS. There is no change in the processing of a scratch request:

- When a request for a scratch volume 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 this category.
- The LM tells the VTS the logical volume size to use. The default is based on the media type MEDIA1 or MEDIA2 as specified at insert time and is overridden by the data class specification.
- 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 Library Manager had provided for the volume, the VTS indicates that logical end of volume (LEOV) has been reached, and the host closes the volume and opens another volume to continue writing.

Enabling 2000/4000 MB volumes reduces the pre-migration throttling threshold by 300 to 400 GB, which in turn may shorten the peak period of the B20 VTS. The B10 peak period is also reduced. Recall of larger datasets may take more time, due to increased amount of data to stage in. But less interrupts may occur for subsequent mounts if a large dataset spans small logical volumes with 800 MB because all data resides in cache.

2.2.5 Stacked volumes

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

Through the Library Manager Console, you define which physical cartridges are to be used by the VTS. Logical volumes stored on those cartridges are mapped by the internal storage management software. With Advanced Policy Management installed, your stacked volumes can be assigned to individual pools. Logical volumes can then be assigned to the stacked volume pools. In the case of out-of-scratch scenarios, pools can be set up to enable “borrowing” from other pools. The methodology and configuration of volume pooling is covered in 4.3.3, “Define stacked volume ranges” on page 129.

> **Note:** Stacked cartridges containing active data cannot be removed from the VTS. Therefore, you cannot eject a stacked volume containing valid data. You can, however, eject a stacked cartridge for maintenance or operational purposes, but only from the Library Manager panels. The data is moved to another cartridge and the empty volume is then ejected. Refer to 6.4.5, “Ejecting VTS stacked volumes” on page 232 for more information about ejecting a stacked volume.

Exported stacked volumes

The physical cartridges used by the VTS to hold exported logical volumes are called *exported stacked volumes*. Once exported, these logical volumes no longer exist within the VTS.

2.2.6 Virtual drives

From a host perspective, the VTS looks like two, four, eight, or sixteen IBM 3490E tape control units, each with sixteen drives attached through ESCON / FICON or SCSI channels. Virtual tape drives are defined just like physical IBM 3490 controller addresses, through hardware configuration definition (HCD), or through the I/O configuration program (IOCP) and the MVSCP or HCPRIO macro. Defining a preferred path for the virtual drives gives you no benefit because the IBM 3490 control unit functions inside the VTS are emulated to the host.

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

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

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

2.2.7 Physical drives

The physical IBM 3590 or 3592 tape drives used by the VTS are installed in the IBM 3494-D12 frame within a IBM 3494 Tape Library. Physical IBM 3592 tape drives used by the VTS are installed in the IBM 3494-D22 frame within a IBM 3494 Tape Library or in an IBM 3584-L22/D22 frame. The physical tape drives are not addressable by any attached host system. They are completely under the control of the VTS subsystem.

The Advanced Library Management System (ALMS) of the 3584 Tape Library allows for dynamic reconfiguration of tape drives. This includes the drives originally dedicated and configured for the attached VTS.

Attention: Do not use ALMS to re-assign physical tape drives attached to a VTS. This function is used for Open Systems hosts and can cause unpredictable complications when used to re-configure VTS-attached tape drives without the assistance of trained IBM service personnel.

2.3 VTS models

Currently, the VTS models B10 and B20 are available. The first and second generation VTS models B16 and B18 are no longer available, but existing models can still be upgraded to the latest models.

2.3.1 Model B16 VTS and Model B18 VTS

The Model B16 VTS was part of the first generation of Seascape® solutions for storage consolidation. It provided a hierarchical storage management system to attached ESCON hosts by integrating IBM 3590 tape drives, fault tolerant RAID disk storage, and a RISC based controller. It appeared to host systems as two fully configured 3490E tape subsystems. It supported 72 GB or 144 GB of TVC, three to six 3590 drives, up to two ESCON attachments, and was located within the IBM 3494 Tape Library beside its IBM TotalStorage Enterprise Tape Drive Expansion Frame (D12 Frame). It was incorporated into the robotic drive train and contained cartridge storage cells as well as the VTS. In this book, we named all VTS Models as VTS except for references to the original Model B16 VTS. Due to its limitation in functions, Model B16 VTS will be described separately. An existing Model B16 VTS can be converted to a Model D12, and any data migrated to a new VTS.

The Model B18 VTS was one of the second generation of Seascape solutions. Virtual configurations provide a maximum of 128 virtual devices and up to 250,000 logical volumes per Model B18 VTS. Each virtual volume can have a maximum capacity of 2.4 GB (assuming 3:1 compression). It supports up to 1.7 TB of TVC (5.2 TB when compressed), three to twelve 3590 drives, up to eight ESCON channels or four SCSI ports, and is located in a stand-alone frame; within 14 m of its associated D12 Frame(s). The optional Extended High Performance Option (EHPO) and Advanced Function features enable efficient use of tape volume cache and improves average mount response time.

The VTS Model B18 is not supported on R7.2 or later levels of code; no new functions in R7.2 or beyond are supported. Fix levels built on R7.1 will continue to be provided as needed. Standalone B18s can no longer be upgraded to participate in a Peer-to-Peer VTS. Feature Codes 4010 to 4013 (Peer-to-Peer Copy Enablement) and 5264 (Additional 64 virtual drives) have been withdrawn from marketing effective May 2004. A Model B18 VTS can be upgraded to the Model B20 VTS.

2.3.2 VTS models B10 and B20

This new generation of VTS products provides an entry model, the Model B10 VTS, and an enhanced model, the Model B20 VTS. These models incorporate data compression and Performance Accelerator features as a standard, as well as an improved controller.

The VTS is attachable to a IBM 3494 Tape Library with IBM 3590 or 3592 tape technology or to an IBM 3584 Tape Library with the IBM 3953 Library Manager and with IBM 3592 tapes. Supported 3590-B1A, E1A, or H1A drives reside in a D12 Frame, 3592 drive models J1A and E05 reside in an L22 or D22 frame.

The Model B10 VTS supports up to 432 GB of TVC (1.3 TB assuming 3:1 compression), four to six 3590 drives, or four to 12 3592 drives, up to 64 virtual devices and up to four ESCON channels or eight SCSI ports or four FICON channels.

The Model B20 VTS supports up 1.7 TB of TVC (5.2 TB assuming 3:1 compression), six to twelve 3590 drives, or four to 12 3592 drives, or a combination of up to six 3590 and up to twelve 3592 tape drives, and up to 16 ESCON channels or eight SCSI ports or eight FICON channels, as well as up to 256 virtual devices.

2.3.3 VTS internal elements

The VTS is available in different configurations in terms of the number and type of channels, host system attachments, sizes of TVC and the number of physical tape drives that are dedicated to the VTS. Figure 2-5 shows the maximum configuration for a B20 VTS when only 3590 tape drives are installed. Maximum configuration is highly dependent on VTS models, for example, up to 432GB cache, four ESCON/FICON channels, and six 3590 or twelve 3592 tape drives are supported on Model B10 VTS.

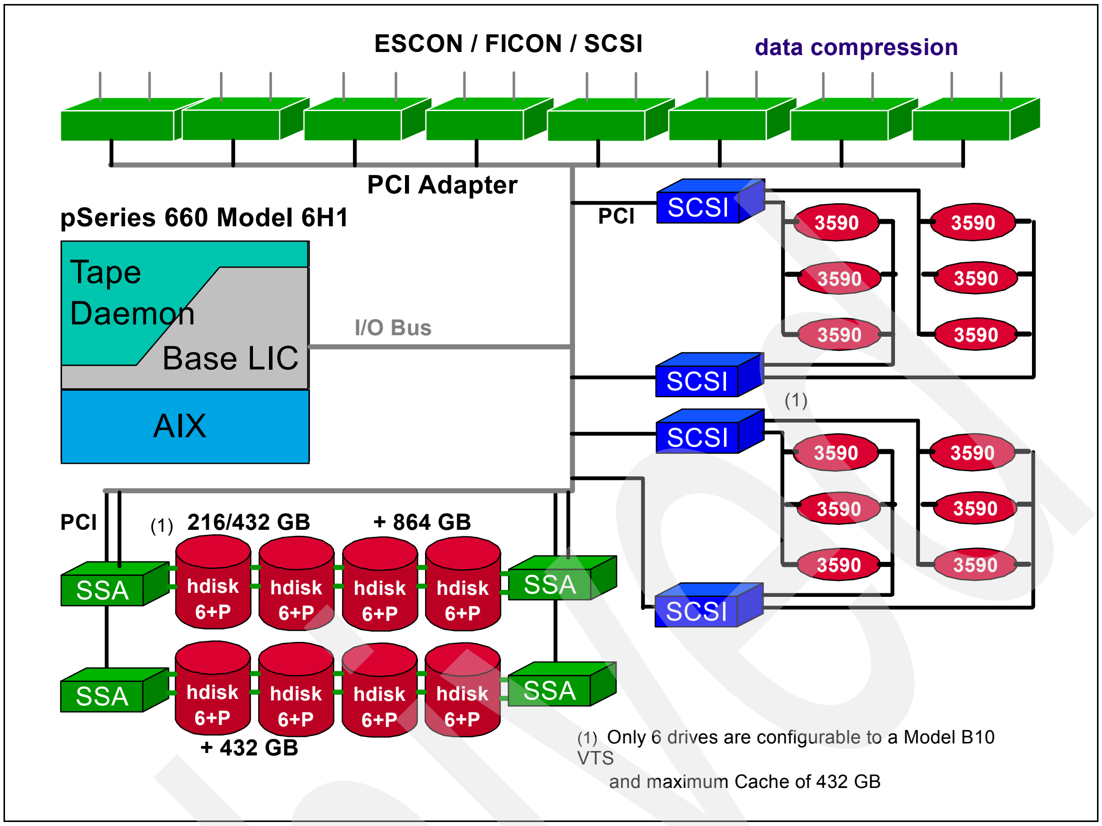

Figure 2-5 VTS configuration with 3590 only

With the 3592 drive support and heterogeneous drive support for B20 VTS, IBM introduced Fibre connection for the VTS as described in Figure 2-6. B10 may have twelve 3592 drives or up to six 3590 drives. The Model B20 supports either twelve 3590 or 3592 drives, or twelve 3592 drives and six 3590 drives.

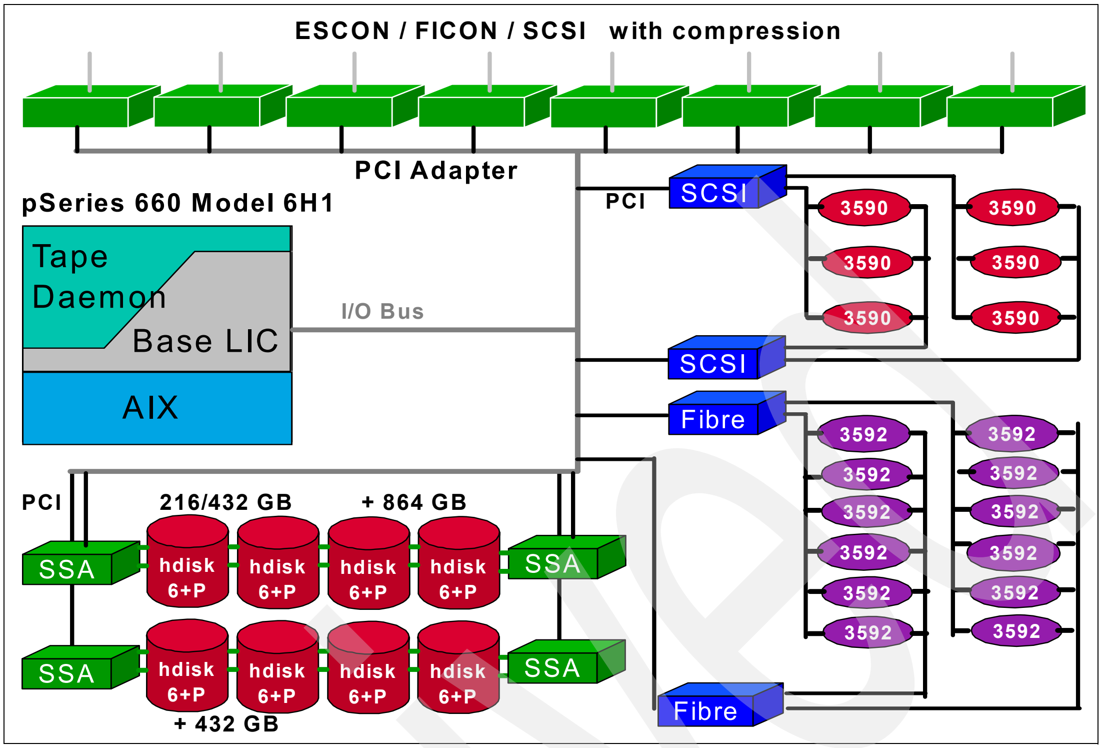

Figure 2-6 Configuration with heterogeneous drive support

2.3.4 Peer-to-Peer VTS

As an extension of the VTS, the Peer-To-Peer 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 Peer-To-Peer Virtual Tape Server is specifically designed for business continuity and high data availability requirements. 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 Peer-to-Peer Virtual Tape Server improves system reliability and availability, as well as data access. To help protect current hardware investments, existing VTSs can be upgraded for use in this new configuration.

In this new configuration, two VTSs are connected in a single complex, at the same location or geographically separate locations, by utilizing Virtual Tape Controllers to provide ESCON or FICON attachment to one or more z/OS hosts. The Peer-To-Peer VTS automatically generates a copy of all virtual volumes, resulting in identical volume images in each VTS. The Virtual Tape Controllers help maintain dual copies in complete synchronization. In addition, each VTS in the configuration can continue to operate if the other VTS is unavailable because of maintenance, routine service, or a system upgrade. In this way, the VTS solution can significantly improve the availability of critical tape data. For more information about the IBM TotalStorage Peer-to-Peer VTS, refer to the *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115.

2.3.5 VTS High Capacity Cache Option

This solution is dedicated for content management applications and aims at replacing zSeries attached IBM 3995 optical libraries. It is transparent to existing OAM applications to allow easy migration and can replace up to five 3995 model C38s.

The High Capacity Cache Option feature (FC4020) can be ordered on a new, standalone Model B20 VTS only and provides an initial cache size of 1.68 TB uncompressed and 2000 logical volumes. With FC4021, you can an additional increment of 1.68 TB of cache and 2000 logical volumes to the VTS. Up to three FC4021 can be installed per VTS to allow for a total capacity of 6.7 TB of cache and 8000 logical volumes.

If you have the requirement to place objects onto write once read many (WORM) media, you can initiate a copy out of this VTS onto native 3592 tape drives which support media.

For more information about this VTS solution, contact you IBM representative.

2.4 Hardware components

The two VTS models themselves come in two different frames: *Model B10 VTS or Model B20 VTS*. They contain the VTS controller and its associated storage management software and the RAID disk arrays that make up the TVC. The Model B10 VTS or Model B20 VTS frame is a stand-alone unit that can be located a maximum of 14 m from the associated 3590 drive frame and up to 25 m from a 3592 drive frame, when installed with an IBM 3494 Tape Library. If the VTS is connected to an IBM 3584/3953 configuration, it can be located up to 62 m (200 feet) away from the D22 frames of the IBM 3584 Tape Library as shown in Figure 2-7.

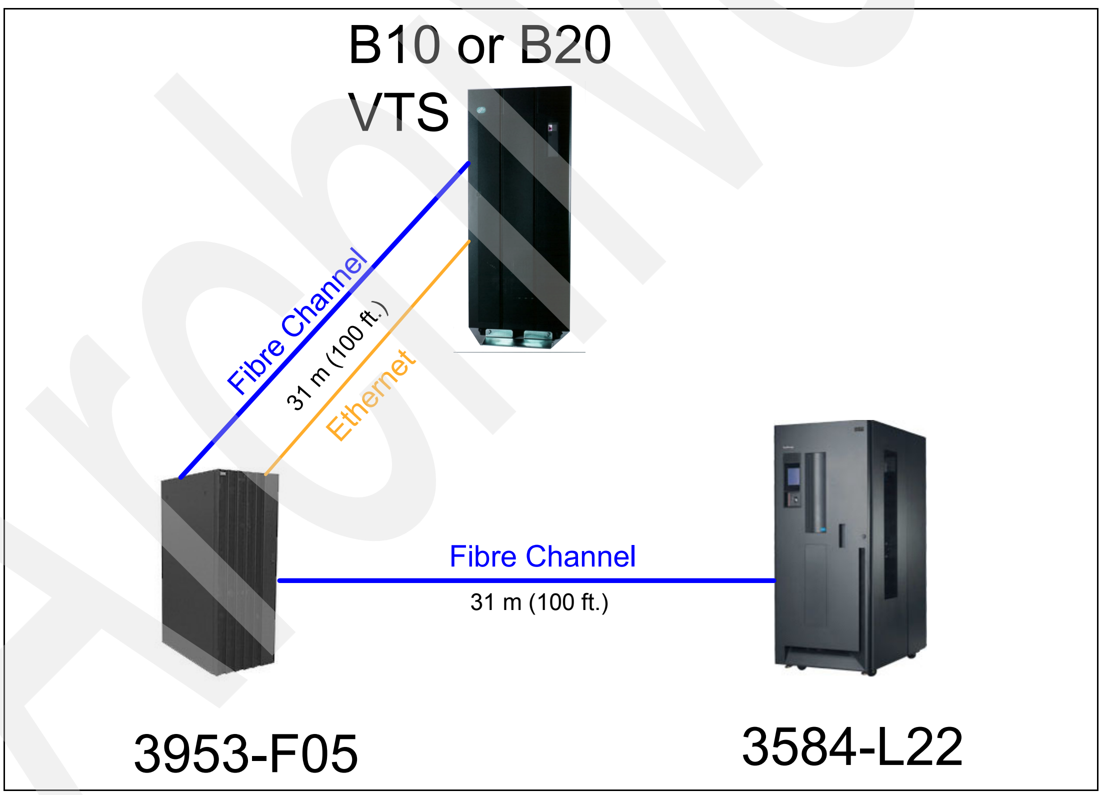

Figure 2-7 VTS distance to IBM 3584 and 3953

The VTS is available in multiple configurations, options include: TVC size, physical drives, host attachment and performance features. Those configurations are described in detail in 3.1.2, “Model B10 VTS and Model B20 VTS configurations” on page 59.

Several frames are needed in addition to the Bxx frame for the VTS: library frames for its drives and the IBM 3494 Library Manager hardware, which is installed in the IBM 3953 Library Manager frame or if attached to a IBM 3494 library in the Library Base Frame (Lxx Frame).

A number of interface types are provided for to enable connectivity to a wide variety of hosts. For z/OS, you can connect up to sixteen ESCON or eight FICON channels. All channel types provide Ziv-Lempel IBMLZ1 compaction at the channel card, resulting in a significant performance improvement in terms of data throughput and Tape Volume Cache (TVC) use. The VTS also can be attached to the SCSI host by up to eight SCSI bus. This VTS can be dedicated to SCSI host only or shared with z/OS host, but sharing of SCSI and FICON interfaces on the VTS is not supported. VTS attached to a 3584 Tape Library does not support SCSI hosts. Figure 2-8 shows a diagram of the physical hardware used for the VTS.

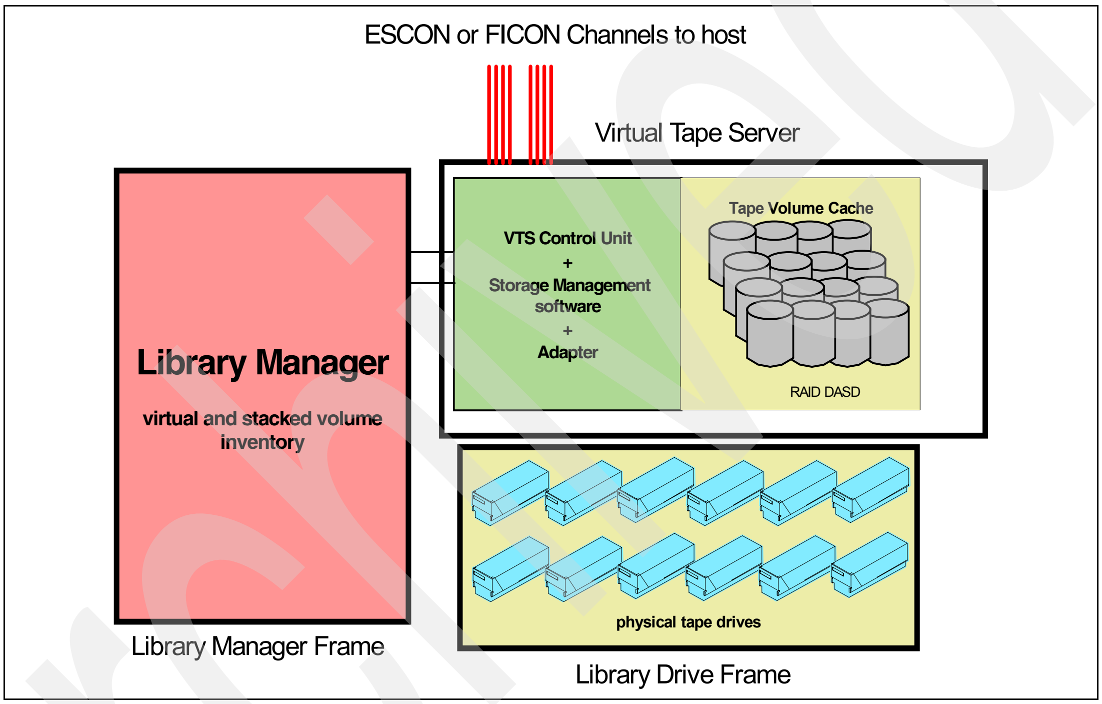

Figure 2-8 Hardware components of the VTS transparent to the host system

Each of the major components shown are described in the subsections that follow. More detailed information about the VTS configuration is provided in 3.1, "Hardware configurations" on page 58.

2.4.1 VTS controller

The VTS controller with its associated Licensed Internal Code (LIC) contains the key components of the VTS. It consists of:

- A RISC-based processor, which provides ESCON / FICON as well as SCSI attachments, 3490E emulation to the host system and attachment to the TVC disks. The processor in the Model B10 and B20 VTS uses an IBM @server pSeries 660 Model 6H1.
- Storage management software, which manages the TVC contents and stacked volumes, controls the movement of data between disk (the TVC) and tape (3590 or 3592 cartridges), and automatically fills the physical tapes. The storage management techniques, built on more than 20 years of IBM experience, include many of the features and functions found in IBM's industry-leading DFSMShsm and Tivoli® Storage Manager (TSM) software products.

2.4.2 Tape Volume Cache (TVC)

The Tape Volume Cache (TVC) of the VTS subsystem is the key element that enables the utilization of the high capacity tape technology. Buffering host-created volumes and later stacking them makes it possible for the cartridge capacity of the technology to be fully utilized. The TVC is a disk buffer where the host 3490E emulated tape volumes are written before they are copied to the physical tape cartridges. The host operating system sees IBM 3490E tape drives, but actually, the space is represented by storage space in RAID fault-tolerant Serial Storage Architecture (SSA) disk subsystems. All host interaction is through the virtual control unit; the host never writes directly to the physical tape drives inside the VTS.

While resident in the TVC, the user data is protected by RAID-5. The TVC is configured into two RAID groups. This RAID configuration gives continuous data availability to the users. If one data disk in a RAID group becomes unavailable, the user data can be re-created dynamically from the remaining disk data contents and the parity disk. In addition, the RAID group contains already mounted (hot) spare disks to replace any damaged disk, so the RAID controller can rebuild the data from the damaged disk on the hot spare. In this way, the VTS can continue working while the IBM service representative replaces the damaged hard disk assembly in the VTS.

With this approach, in addition to fulfilling the objective of making full use of high-capacity tape cartridges, there are these benefits:

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

Cache Management allows management of the contents of the tape volume cache to improve average mount response time. DFSMS policy management can be used to assign virtual volumes to preference groups. These groups can control the priority of when volumes are deleted from cache after they are copied to tape when space is needed in cache. For more information on this topic, see 2.6.3, “Tape volume cache management” on page 50.

2.4.3 IBM TotalStorage Enterprise Tape Drives

The IBM VTS is configured with four to twelve IBM 3590 tape drives or IBM 3592 drives where logical volumes get recalled from and copied to out of the TVC.

These drives are attached to the VTS controller only and are dedicated to the function of the VTS. There is SCSI attachment for IBM 3590 tape drives and Fibre Channel attachment for IBM 3592 drives. They are not visible to, and therefore cannot be used by, any external host system. Special considerations apply for IBM 3592 drives in a 3494 Tape Library. See 3.2.1, “VTS with 3590 tape drives” on page 61.

The VTS can exploit either IBM 3590 Base Models and/or IBM 3590 High Capacity Models. The drives for a Model B18 or B10 VTS ***must*** be the same drive type (3590 or 3592), and, in case of 3590 drives, of the same model. In case of 3592 drives, models J1A and E05 can be intermixed. The IBM System Storage™ TS1120 Tape Drive, also called 3592-E05, when attached to a VTS, will automatically come up in J1A Emulation mode and only operate in this mode.

> **Important:** IBM 3590 tape drives are ***not*** supported inside an IBM 3584 Tape Library.

IBM 3590 Model B1A

The IBM 3590 Base Models read and write data in 128-track mode on physical tape cartridges. The IBM 3590 Base Models are characterized by a drive data rate of 9 MB/s and can store 30 GB of data on a standard HPCT cartridge and can store 60 GB of data on a Extended Length Cartridge (EHPCT), assuming a 3:1 compression ratio. The 3590-B1A drive is housed in a D12 Frame.

Note: The 3590B models have been withdrawn from Marketing.

IBM 3590 Model E1A

The IBM 3590 High Capacity Models read and write data in 256-track mode on physical tape cartridges. The IBM 3590 High Capacity Models are characterized by a drive data rate of 14 MB/s and can store 60 GB of data on a standard HPCT cartridge and can store 120 GB of data on a Extended Length Cartridge (EHPCT), assuming a 3:1 compression ratio. Currently installed IBM 3590 Base Models can be upgraded to IBM 3590 High Capacity Models. The IBM 3590 High Capacity Models read in 128-track mode from physical cartridges which were written by IBM 3590 Base Models. The 3590-E1A drive is housed in a D12 Frame.

The main reason to upgrade to IBM 3590 High Capacity Models is doubling the data storage capacity of the VTS owned tape cartridges.

IBM 3590 Model H1A

The IBM 3590 High Capacity Models read and write data in 384-track mode on physical tape cartridges. The IBM 3590 High Capacity Models are characterized by a drive data rate of 14 MB/s and can store 90 GB of data on a standard HPCT cartridge and can store 180 GB of data on an Extended Length Cartridge (EHPCT), assuming a 3:1 compression ratio. Currently installed IBM 3590 Base Models can be upgraded to IBM 3590 High Capacity Models. The IBM 3590 High Capacity Models read in 128 or 256-track mode from physical cartridges which were written by IBM 3590 Base Models. Various tape capacities are compared in Table 2-1 on page 32. The 3590-H1A drive is housed in a D12 Frame.

IBM 3592 Model J1A

The IBM TotalStorage 3592 Tape Drive Model J1A is the first model of a new generation of enterprise tapes with characteristics which outperform the IBM 3590 tape drives. The 3592 tape drive and media were introduced in September of 2003 for both direct zSeries as well as Open Systems attachment environments. 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. It reads and writes data in 512-track mode on the new physical tape cartridges. The fibre channel 3592 drives must all reside in a single D22 frame if 3494 Tape Library attached or they can reside in a L22 and/or multiple D22 frames if 3584 attached. In May of 2004 an additional cartridge, the Enterprise Economy Tape Cartridge, was made available which provides a non-compressed storage capacity of 60GB. Support for the 3592 drive and the two cartridge capacities are provided under the VTS system at the VTS R7.3 and associated Library Manager 530 code levels since October 2004.

IBM 3592 Model E05

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

- The E05 models can not only read cartridges written by J1A drives but can also write in J1A Emulation mode. This means that the drive can dynamically change its mode of operation per physical cartridge.
- The 3592-E05 tape drives can be intermixed with 3592-J1A tape drives:
  - In the same tape library frame
  - Behind the same Model J70 controller
  - Behind the same VTS.

Note: When attached to a VTS, the E05 drives always operate in J1A Emulation mode.

The Model E05 tape drives use the same cartridges as the Model J1A drives and provide, when operating in E05 mode, an uncompressed cartridge capacity of 500 GB (JA media) or 100 GB (JJ media). The native drive data rate is 100 MB/s when running in E05 mode and up 50 MB/s when running in J1A Emulation mode. Support for the 3592-E05 tape drive is provided with VTS R7.4+5.

Coexistence and migration

As the capability of the underlying tape technology has progressed from the 3590 Model B1A tape drive and the High Performance Cartridge Tape through the E1A and H1A drive models and the Extended High Performance Cartridge Tape, the base VTS and its high availability Peer-to-Peer VTS (PTP VTS) configuration has been able to transparently utilize the increased capacity. One of the aspects of the migration from the 3590 B1A models through the H1A models, was the ability of the newer models to read the data cartridges written by the older models.

With the introduction of the new 3592 drive technology, that is no longer the case. Cartridges written by 3590 drives are not compatible with 3592 drives and vice versa. To assist in the utilization and migration to the new 3592 technology, the model B20 VTS supports 3590 drives and media and 3592 drives and media both under the control of the same VTS and provide for the migration of data from 3590 written cartridges to 3592 written cartridges. The next generations of 3592 will again extend the capacity and bandwidth but will rely on the same media as used today. Consequently the migration onto next generations of 3592 will be supported in the same way as with 3590 family.

"Stacked volume pool properties" on page 45 describes the additional configurations and control functions the VTS provides for the integration of and migration to the 3592 drives and media. It also contains suggestions on the use of the Advanced Policy Management functions to manage the usage and migration aspects of a VTS with both technologies.

Physical tape cartridges

The physical cartridges belonging to the VTS are stored anywhere in the storage cells of the IBM 3494 library where the VTS is installed. When the VTS is installed inside an IBM 3584 Tape Library, the cartridges can reside in any storage cell inside a Model L22 or D22 frame. They are not separated from any other cartridges in the library, but they are dedicated for use by the VTS. Their position and status are recorded and maintained by the Library Manager. These stacked cartridges are not visible to attached hosts. IBM 3590 HPCT and EHPCT cartridges can coexist with IBM 3592 ETC and EETC.

Table 2-1 compares various cartridge types and their capacities which are available.

Table 2-1 Cartridge media types and native capacities

Cartridge type	Media ID	Length	Capacity	Comments	DFSMS Media type
CST	1 or blank	168 m	400 MB	Cartridge System Tape	Media1
ECCST	E	335 m	800 MB	Enhanced Capacity CST	Media2
HPCT	J	300 m	10 GB	High Performance Cartridge Tape	Media3
EHPCT	K	600 m	20 GB	Extended Length HPCT	Media4
ETC	JA	609 m	300 GB	Enterprise Tape Cartridge	Media5
EWTC	JW	609 m	300 GB	Enterprise WORM Tape Cartridge	Media6
EETC	JJ	246 m	60 GB	Enterprise Economic Tape Cartridge	Media7
EEWTC	JR	246 m	60 GB	Enterprise Economic WORM Tape Cartridge	Media8

Restriction: The VTS currently does not support Write Once Read Many (WORM) cartridges. These cartridge types can only be used with native IBM 3592 tape drives, which are not attached to a VTS.

Media types

The IBM TotalStorage Enterprise Tape Drive 3592 models J1A and E05 uses four media cartridge types: JA, JJ, JW, and JR. All four cartridge types contain the same dual coat advanced partile media. The media is housed in a 3592 cartridge shell, which is close but not identical to current 3590 cartridges in size and shape. The only externally visible difference between the four 3592 cartridge types is the color of the certain parts of the cartridge.

Table 2-1 lists all the different media types which are available. For the VTS, the following rules apply:

- Logical volumes can be MEDIA1 or MEDIA2. The actual length of a logical volume is determined by the data class specification and can be 400, 800, 1000, 2000, or 4000 MB depending on VTS configuration and LIC level.
- With 3590 tape drives attached to a VTS, MEDIA3 and MEDIA4 cartridges can be used. Cartridge selection is based on the APM definitions for a specific pool or random, if APM is not implemented.
- With 3592 tape drives attached to a VTS, MEDIA5 and MEDIA7 cartridges can be used. Cartridge selection is based on the APM definitions for a specific pool or random, if APM is not implemented.

- With both IBM 3590 and 3592 tape drives attached to a Model B20 VTS, MEDIA3, MEDIA4, MEDIA5, and MEDIA7 cartridges can be used. APM is a required prerequisite for this configuration, and cartridge selection is based on the APM definitions for a specific pool.
- MEDIA6 and MEDIA8 are not supported in a VTS configuration.

2.5 Virtual Tape Server configuration in the library

So far we have discussed the logical and physical attributes of the VTS. Now we describe the different VTS attachment options to the IBM 3494 library and to the IBM 3584 Tape Library.

The VTS is always installed in a 3494 Tape Library or a 3584/3953 Tape Library system. It may share the library with a second VTS or with 3590 or 3592 tape drives (within a 3494 Tape Library) that are not part of the VTS configuration. Or it can be part of 3584/3953 Tape Library together with native 3592 tape drives connected to a zSeries host and LTO or 3592 tape drives connected to Open Systems hosts. To operate a VTS, you need to understand how the Library Manager within a 3494, or the Library Manager within a 3953 frame connected to a 3584 Tape Library interacts with the host operating system.

2.5.1 Library Manager

You perform the set of unique operational tasks associated with a VTS through the 3494 or 3953 Library Manager console, the primary operational interface to the VTS. The storage administrator and operator use the Library Manager Console to define the VTS environment, to initiate VTS functions and to view VTS statistics. The VTS in turn reports its error and status conditions through the Library Manager and the operator uses the Library Manager Console to determine required actions.

In the case of the 3494, the Library Manager console is located at the rear of the L1x frame. If the 3494 is also equipped with an HA1 frame, a second Library Manager and console is located at the rear of the HA1 frame at the other end of the library. Depending on which state the Library Managers are in (Active or Standby) determines which functions are available. Primarily, Operator and Administrative functions are performed from the Active Library Manager. Service functions can be performed on either.

In the case of the 3584/3953 environment, the Library Manager console is located at the rear of the 3953-F05 frame. As in the case with the 3494, a second Library Manager can be ordered (FC5065, 5066 and 5067) to allow for high availability. All Operator, Administrative and Service function can be done with this console. Since in the Library Manager in the 3953 has no direct control over the accessor in the 3584, further Operator, Administrative and Service functions will be accessed through the 3584 Specialist. Accessing the Specialist is through a separate PC which is LAN attached to the 3584. A Master Console (FC2714 or FC2718) can be installed inside the 3953 along with a KVM Switch to facilitate accessing all the required areas from the one console in the rear of the 3953.

If you have ordered the Remote Library Manager Console feature you can control the Library Manager remotely. For more information about ordering the Remote Library Manager Console feature, see the *IBM TotalStorage 3494 Tape Library Introduction and Planning Guide*, GA32-0279.

For further information about setting up the Remote Library Manager Console, talk to your IBM service representative. For information about operating and starting the Remote Library Manager Console, see the *IBM TotalStorage 3494 Tape Library Dataserver Operator's Guide*, GA32-0280.

The VTS and Library system may also be monitored from the World Wide Web using the TotalStorage Specialist family of products. For details using the tape library specialist to monitor the VTS performance, see 7.7, “Using the Tape Library Specialist to monitor the VTS” on page 311.

An additional possibility to monitor key Library Manager events is to use an SNMP monitor station and the SNMP trap facility. Refer to 6.3.3, “Simple network management protocol (SNMP) traps” on page 219 for details.

2.5.2 VTS attached to an IBM 3494 Tape Library

The VTS can be installed by adding frames to a currently installed 3494, or it can be installed with a new library. It is designed to work as part of the IBM library infrastructure so that emulated drives and logical volumes can coexist effectively in the same library with native 3490E, 3590, and 3592 cartridges plus the native drives (SCSI, Fiber Channel, ESCON or FICON attached); the mixed environment is fully managed by the Library Manager. In fact, the combination of native 3590 and/or 3592 and VTS emulated drives can be particularly powerful, being able to cope with both high-throughput work that can fill cartridges and lower-throughput work that does not fill cartridges. Table 5-5 on page 208 summarizes our recommendations on where to store different types of data.

In a mixed environment, the Library Manager manages the VTS and native drives within the library and presents the image of up to three logical libraries to the host systems, each with its own library name. One library contains the native 3490E or 3590 drives and cartridges and the other two libraries are VTS libraries each with 64, 128 or 256 emulated 3490E drives and associated logical volumes. Figure 2-9 shows an example of a seven-frame mixed library.

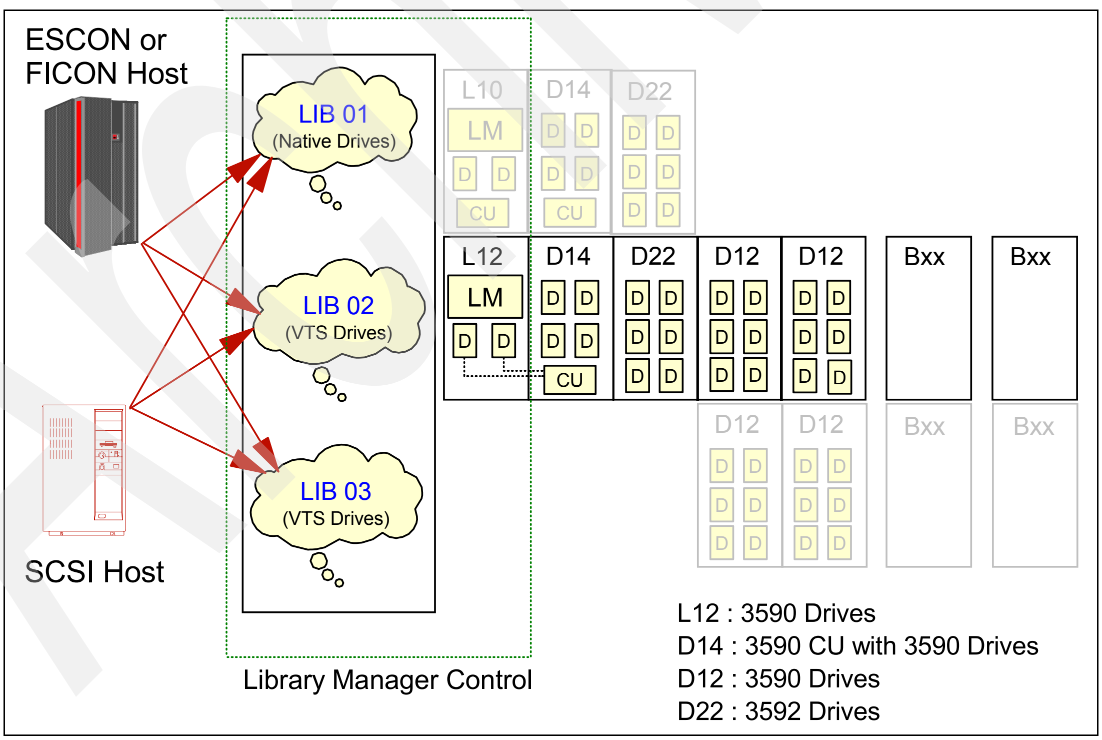

Figure 2-9 Virtual Tape Server configuration in the 3494 Tape Library

In this diagram, the ESCON/FICON host sees three libraries: two VTS libraries (LIB02 and LIB03) consisting of seven frames: Model Bxx VTSs and D12 Frames shown as the four “grayed” out frames at the bottom of Figure 2-9 and a native library (LIB01) with four native 3590 drives in the D14 frame, and two 3590 drives in the L12 frame connected to the CU in the D14 frame. The SCSI host also sees three libraries: a VTS libraries (LIB02 and LIB03) consisting of two Model Bxx VTSs and two D12 Frames shown as the four “grayed” out frames at the bottom of Figure 2-9 and a native library (LIB01) with six native 3592 drives in third frame. The native drives belonging to LIB01 are shown as the four “grayed” out frames in the top of Figure 2-9. The Library Manager is installed in the head frame of the library (L10).

IBM TotalStorage Enterprise Tape Library Expansion Frame D12: This is the library expansion frame; D12 Frame, which can house up to six IBM 3590-B1A, IBM 3590-E1A, or 3590-H1A tape drives. In this case, the drives it contains are dedicated to the use of the VTS. Up to two D12 Frames with up to twelve drives can be installed for VTS. All 3590 tape drives attached for use by the VTS must be of the same track density capability. IBM 3590 Model Bxx tape drives and IBM 3590 Model Exx tape drives cannot coexist in the same VTS.

IBM TotalStorage Enterprise Tape Library Expansion Frame D22: This frame can house four to twelve IBM 3592-J1A or E05 drives connected via two SAN switches to the VTS. The drives it contains are dedicated to the use of the VTS. For the migration from already installed IBM 3590 tape drives a second D12 Frames with up to six drives can be installed for a B20 VTS. This heterogeneous drive support with model B20 may have a maximum of 18 drives connected. A second D12 frame with 3590 drives has to be disconnected before the D22 and 3592 are connected to the B20 VTS. Mixed drive support is not available for B10 VTS.

A D12 and D22 frame belonging to a VTS can be installed anywhere along the length of the library. When two D12 Frames are used for VTS drives, the second D12 Frame must be located immediately adjacent to the first D12 Frame in the 3494 Tape Library. When a D12 and D22 frame are used for VTS drives, they do not have to be located immediately adjacent in the 3494 Tape Library. The stand-alone VTS frame does not count toward the maximum number of 16 frames for a 3494 Tape Library. The D12 and D22 frames contributes to the cartridge storage space for the whole library; the cartridges belonging to the VTS can be stored in any of the frames that make up the total library capacity. The physical position and status of the cartridges are controlled by the Library Manager.

Note: The VTS Models B10 and B20 as well as the associated drives do not occupy control unit function ports (RS-422 or RS-232) because they communicate with the Library Manager through a local area network (LAN) connection.

If you just want to install a single VTS system without any other library or native drives, the minimum configuration is three frames: an L1x Frame, one tape drive D12 Frame or D22 Frame and the stand-alone VTS frame. The L1x Frame can be installed without any drives inside; if so, the Library Manager presents only one library image to the ESCON host system with a single library name. The host sees a single library with up to sixteen ESCON paths or up to eight SCSI paths or up to eight FICON paths and up to 256 3490E drives. See 3.1.1, “VTS configuration options” on page 58 for more details on configuring hardware and features when planning to install a VTS.

We can now summarize the rules discussed in this section:

- The IBM L1x Frame can be left without drives.
- The D12 or D22 frame can be installed anywhere along the length of the library. If two D12 frames are installed in the same library, they must be adjacent. D12 and D22 frames do not have to be adjacent.

- The Model Bxx VTS frame must reside within a distance of 14 m from the associated D12 frame and up to 25 m from the D22 frame.
- A maximum of two VTS systems can be installed in any one 3494.
- For a Peer-to-Peer VTS, both VTS systems must reside in different physical tape libraries, however one can be attached to an IBM 3494 Tape Library and one to an IBM 3584 Tape Library.

2.5.3 VTS attached to an IBM 3584 Tape Library

With Release 7.4 IBM introduce the new IBM 3953 tape system which provides zSeries attachment capability for the 3584. The IBM 3584 Tape Library has been enhanced over many years since its initial launch in September of 2000. Initially a two-frame LTO-based library, it has grown into a true enterprise class library throughout time with IBM 3592 enterprise class tape drives, up to sixteen frame capacity, and, most recently, dual robotic accessors. This announcement takes the 3584 the last 2 steps to being a true enterprise class architecture by adding zSeries host attachment and VTS attachment capability. Figure 2-10 shows a 3494, a 3584 and a 3953 frame.

Figure 2-10 IBM 3494-Lxx, IBM 3584-Lxx, and IBM 3953-L05 frame

The IBM 3953 tape system consists of two models:

- Model F05 is the frame which holds all of the relevant hardware and firmware (including the Library Manager) to provide everything necessary for zSeries or VTS attachment.
- Model L05 is the Library Manager, both hardware and associated firmware that provides similar function as the Library Manager of the IBM 3494 Tape Library. It is NOT orderable and shareable (interchangeable) with its 3494 counterpart.

Also associated with these models are numerous feature codes which add additional, required or optional functionality.

IBM 3953 Model F05 Frame

The IBM 3953 Tape Frame comes in two different versions.

- **Base frame:**

  IBM 3953 Model F05 base frame is the first 3953 frame in a configuration. It houses one Library Manager or, optionally, two Library Managers for redundancy, multiple Ethernet switches, up to six Fibre Switches when an IBM 3592 Model J70 controller and VTSs are installed, one KVM Switch, and monitor and keyboard. It can house one Model J70 Controller and the TotalStorage Master Console, if installed. The switches required for VTS attachment are also housed in the base frame.

- **Expansion frame:**

  IBM 3953 Model F05 expansion frame can be used to install additional IBM 3592 Model J70 controllers. A single expansion frame can house up to three 3592-J70 controllers, up to two Ethernet switches, and up to six Fibre Switches.

Both the base frame and the expansion frame have the same model type F05 but are differentiated by having a descriptive feature code.

There may be up to six frames in total associated with one IBM 3953 and one IBM 3584 logical library. These frames do not need to be installed next to each other, nor do they need to be alongside the IBM 3584 Tape Library. The first frame must be a base frame, and houses the Library Manager.

IBM 3953 Model L05 Library Manager

The Library Manager controls all or part of the IBM 3584 Tape Library according to how the logical libraries are set up; for example the physical library may be divided between Open Systems hosts as well as zSeries. One Library Manager controls one logical library, but there may be up to four Library Manager controlled logical libraries in one IBM 3584 Tape Library, as well as a number of Open Systems controlled logical libraries.

It should also be pointed out that the 3584 has some unique capability when attached to a single or multiple 3953 tape systems. Note that a single 3584 Tape Library can support up to four 3953 tape systems concurrently. A sixteen- frame 3584 Tape Library with IBM 3592 models J1A or E05 tape drives installed can support a combination of up to 6256 cartridges or up to 192 tape drives, with a maximum of 12 drives per 3584 frame. For 3592 drives storing a full native capacity of 300GBs per cartridge, that is over 1.8 PB native capacity stored in a single IBM 3584 Tape Library.

As previously stated, since each 3953 Library Manager enables connectivity of up to 2 VTS systems, a single 3584 Tape Library can support up to 8 VTSs.

Add to the zSeries and VTS enablement that the 3584 has a faster performing robotic mechanical mechanism and improved cartridge per square foot density over the 3494 and you have the world’s leading enterprise class zSeries tape system. The dual robotic accessor feature of the IBM 3584 significantly reduces the time necessary for the 3584 Tape Library to be offline in order to add additional expansion or D-frames. The improvement in upgrade time over 3494 is quite dramatic.

2.5.4 Logical libraries in the 3584/3953

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

- Splitting a single IBM 3584 Tape Library by putting logical walls between drives, and slots within a single physical library enclosure. This is managed within the IBM 3584 hardware itself, either with or without ALMS.

- The library name or ID seen by a zSeries host when it is attached to a tape library. For each library, up to three library IDs, called logical libraries are supported. This is managed through the Library Manager and the zSeries host configuration software.

In the IBM 3584 you must first identify a hardware-based logical library for the IBM 3953 to manage. Note that there may be only one logical library defined which occupies the entire physical library and is dedicated to zSeries.

The IBM 3953 Library Manager, and the zSeries host are aware only of one library. When one or two VTSs, as well as native IBM 3592 drives are attached, the Library Manager defines its own logical libraries or partitions: one for the native drives, and one for each attached VTS subsystem.

In both cases the terms logical library and partition are used interchangeably. It is very important when referring to logical libraries or partitions to understand in what context the term is being used.From a host operating system point of view the VTSs and J70 controller which it is connected to via the Library Manager in the new 3953 frame there is no change versus the three different logical libraries within a 3494 Tape Library.

Each 3953 can support up to 3 logical library partitions consisting of 2 VTSs (B10 or B20 only) and a native logical library with 3592-J70 control units. These three logical library partitions themselves make up one logical library within a 3584 Tape Library. This is in addition to any Open Systems logical libraries already configured in the 3584. A total of four 3953 subsystems can be connected to one 3584 for a maximum of 12 different 3953 managed Logical libraries. More in depth description of the new connectivity ca be found in *IBM TotalStorage 3584 and 3953 for zSeries Guide, SG24-6789*.

2.6 Advanced Policy Management (APM)

We mentioned APM within several sections in this chapter. Now we give you an introduction of APM and its functionality.

August 2002 saw the General Availability of the Advanced Policy Management (APM) enhancements for the IBM TotalStorage Virtual Tape Server (VTS). Subsequent microcode releases enhanced and added more functionality. As of today, the following management functions are available:

- Physical volume pooling with
  - Various reclamation policies
  - Secure data erase
- Selective Dual Copy
- Peer-to-Peer Copy mode control
- Tape Volume Cache Management
- Extended logical volume sizes
- Bulk Volume Information Retrieval (BVIR)
- Export/Import

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. However, some customers want to better utilize the resources of a VTS or have a need to control the management of logical volumes. APM is available with feature codes FC4001 through FC4004. This includes all the features mentioned above. Existing users of FC4000 (Advanced Functions) can upgrade to this feature.

The Implementation is based on following principles:

- Do not introduce new concepts or customer interfaces regarding storage management. Expand the use of the existing System Managed Storage (SMS) constructs to control the new functions.
- Keep the control of functions in the Automatic Class Selection routines.
- Do not make the use of new functions dependent on third party software or solutions if at all possible.
- Enable the capability for a tape management system to influence the selection of the storage management constructs through the ACS routines.
- For the functions that are analogous to cached DASD functions, use the same construct. For example, volume grouping should be controlled by a storage group construct because it deals with the grouping together of data volumes.
- Do not require that the controls be used. The VTS subsystem and library have sensible defaults for the controls.
- Customer controls override default management policies.
- Provide new functions without requiring host software changes to be able to exploit the functions.
- Provide a way for platforms other than MVS to exploit the advanced functions without having to depend on massive software support.

With APM enabled in the z/OS environment for SMS-managed tape, DFSMS construct names are passed to the LM. Figure 2-11 details the file allocation, the data flow and how the DFSMS constructs are passed to the Library Manager.

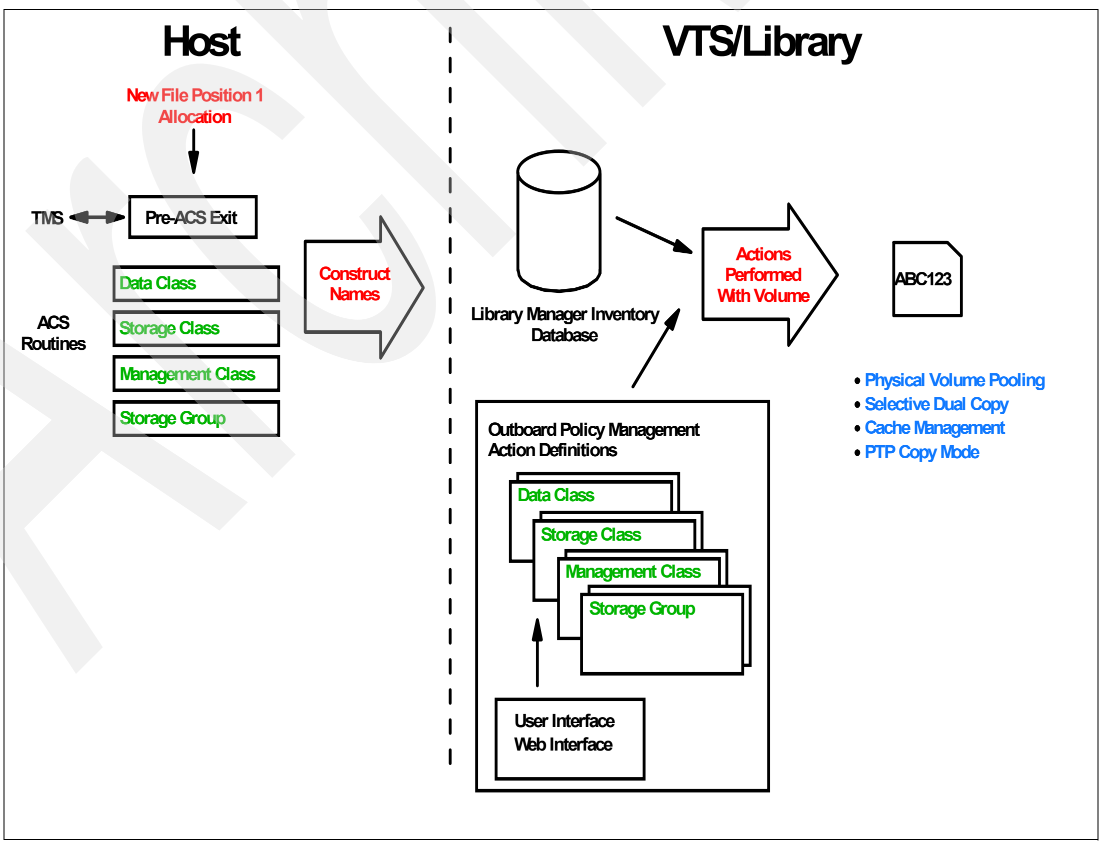

Figure 2-11 APM control flow

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 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

Once Advanced Policy Management is installed and configured, the VTS can perform different functions based on the constructs assigned (see Figure 2-12). In this example, virtual volumes 1 and 2 are assigned to the same pool (SG assignment is Pool 1). However, since virtual volume 1 has a been assigned a different management class (MC = Dualcopy), a second copy of the volume is made to the assigned dual copy pool (Pool 2).

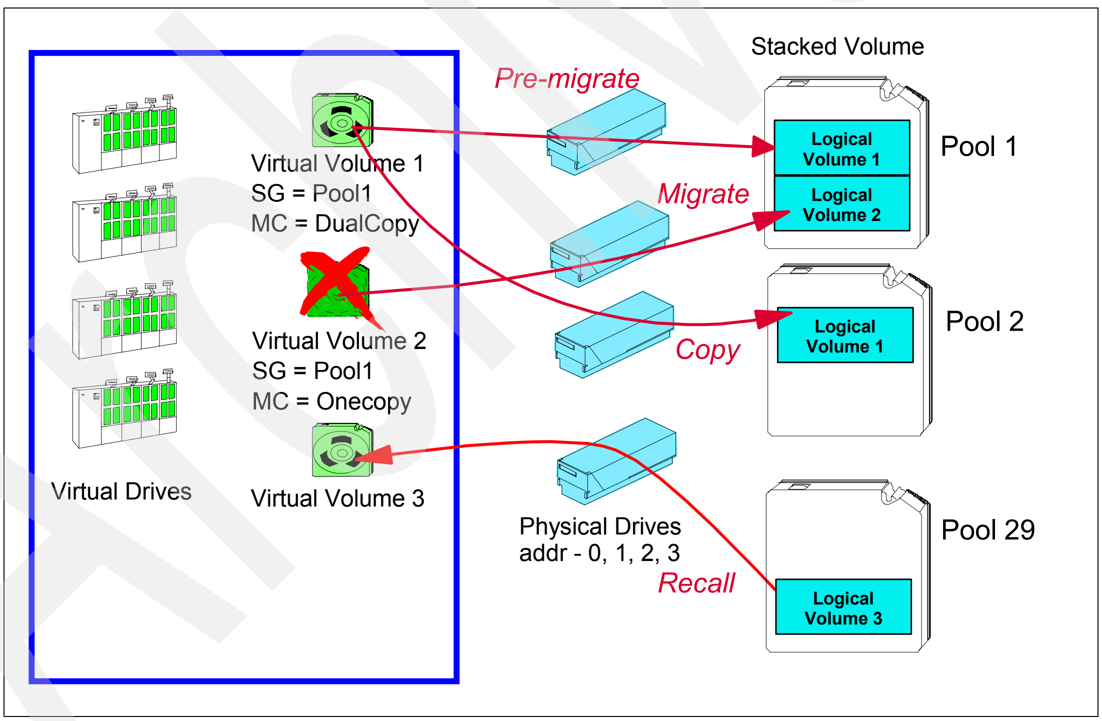

Figure 2-12 Relationship between virtual and real components with APM

This process and how to configure the constructs is covered in 4.4.2, "Defining storage groups" on page 145 and following.

Outboard policy management

The enhancement introduces the concept of Outboard Policy Management. This is defined as the functions and components on the 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 Enterprise Tape Library (ETL) 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
- Manage management classes:
  - Host - No effect. Pass construct name.
  - Outboard - Control of selective dual copy and Peer-to-Peer copy modes for logical volume
- 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

> **Note:** Construct names are limited to a maximum of 256 names per construct type (including the default of all blanks).

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

Table 2-2 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 moved outboard and preference levels added.
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 is enabled.

Enterprise Tape Library (ETL) Specialist

The IBM TotalStorage 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 APM has also introduced a new set of screens for the VTS component of the ETL Specialist. 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.

2.6.1 Physical volume pooling

With Outboard Policy Management enabled, you 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. This is shown in Figure 2-13 on page 44. Physical pooling of stacked volumes is identified by defining a pool number. This definition can be set through the LM or the tape library specialist. Currently, valid pools are 0-32.

Each VTS within a 3494 or 3584 has its own set of pools. There is a Common Scratch Pool (pool 00) which is a reserved pool containing only scratch stacked volumes. There are also 32 general purpose pools (pools 1-32).

Without volume pooling, the VTS stores virtual volumes on any stacked volume available to it. This creates an intermix of logical volumes from differing sources, for example, LPAR, customer, applications. The user cannot influence the physical location of the logical volume. Physical Volume Pooling addresses the following concerns:

- Different customer data on same physical volume, can contravene some outsourcing contracts
- Customers like to be able to “see, feel, touch” their data by having only their data on their media.
- Charging for tape is complicated. Traditionally users are charged by number of volumes they have in tape library; the customer can create and consolidate multiple volumes on a single VTS logical volume reducing their media charges.
- Export of logical volumes takes a considerable amount of time because the physical tape on which a selected logical volume resides has to be mounted; grouping logical volumes eligible for export in the same physical pool may decrease the number of physical tape mounts associated with the export.
- The VTS will use any Media type available to it and small logical volumes on Extended media (K cartridges) can take a considerable time to recall.

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, JA or JJ), 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.

Physical volume database

The function of the Physical Volume Database in the VTS is to:

- Track pooling information such as home and current pool definitions.
- Track media type.
- Track error events and error states.
- Provide a lock mechanism used to co-ordinate movement and error recovery activity on physical volumes.

The Physical Volume Database is created as part of the upgrade to VTS 2.26 and exists with this LIC level and above. The database is maintained as a logical database within the Token Database which was added for Peer-to-Peer VTS support. The Token Database is used by the VTS but token data is only used for Peer-to-Peer VTS. The backup of the Token Database is performed at the same time as for the other VTS Databases and is done to a physical volume.

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 tape library specialist when defining Stacked Volume Pool Properties.

At VTS code level 2.26 or later, scratch cartridges are no longer kept in category number FF03. Both scratch and private stacked volumes have category FF04 assigned, independent of their media type.

General purpose pools (Pools 1-32)

There are 32 General Purpose pools available per VTS. These pools can contain both physical scratch and private stacked volumes. All volumes, private and scratch, are assigned category FF04. When initially creating these pools, it is important to ensure that the correct borrowing properties are defined to the pool. See 4.4.1, “Defining stacked volume pool properties” on page 143.

This facility is also available to:

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

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

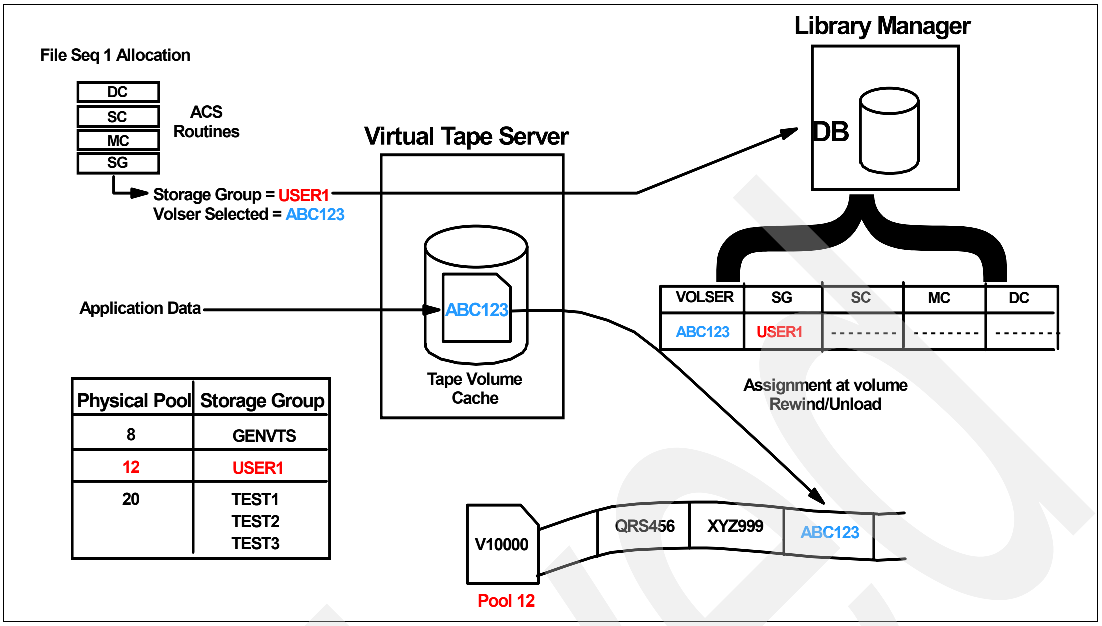

Figure 2-13 Logical volume allocation to specific physical volume pool flow

For those sites that are APM enabled, the following applies 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.

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).

> **Note:** Volume pooling is supported in both base VTS and Peer-to-Peer configurations. We recommend that you use the definitions for both VTSs in a Peer-to-Peer VTS configuration. For more information about IBM TotalStorage Peer-to-Peer VTS, refer to the *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115.

Borrowing and returning

With the Advanced Policy Management, new rules are introduced to manage out of scratch scenarios: the concept of *borrow and return*. 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.

Borrowing principles

The term borrow implies the eventual return of a volume. The following principles apply to borrowing:

- The primary pool must maintain two empty usable physical volumes in the pool. This pool must have been active in the last 48 hours:
  - Activity is measured as a time period since the last premigration activity has occurred.
  - After a period of inactivity (currently 48 hours), the pool quits borrowing.

- Service the pools with the most constraining borrowing rules first:
  a. Borrow for pools where only the “1st media” type has been specified.
  b. Borrow for pools where both “1st and 2nd media” types have been specified.
  c. Borrow for pools where “Any” has been specified.
- Always borrow physical volumes with the lowest usage count for a given media type.
- Borrowing only occurs from the Common Scratch pool (Pool 00).

Return principles

The following principles apply to the return of physical volumes:

- Physical volumes are returned to the pool from which they were borrowed from
- Physical volumes with the highest usage count are returned
- Only return physical borrowed volumes which are empty.
- There must be more than three empty usable physical volumes in an active pool before volumes are returned.
- When a pool has been inactive for more than 72 hours, all borrowed physical volumes are returned.

The borrow and return process flow is described in Figure 2-14.

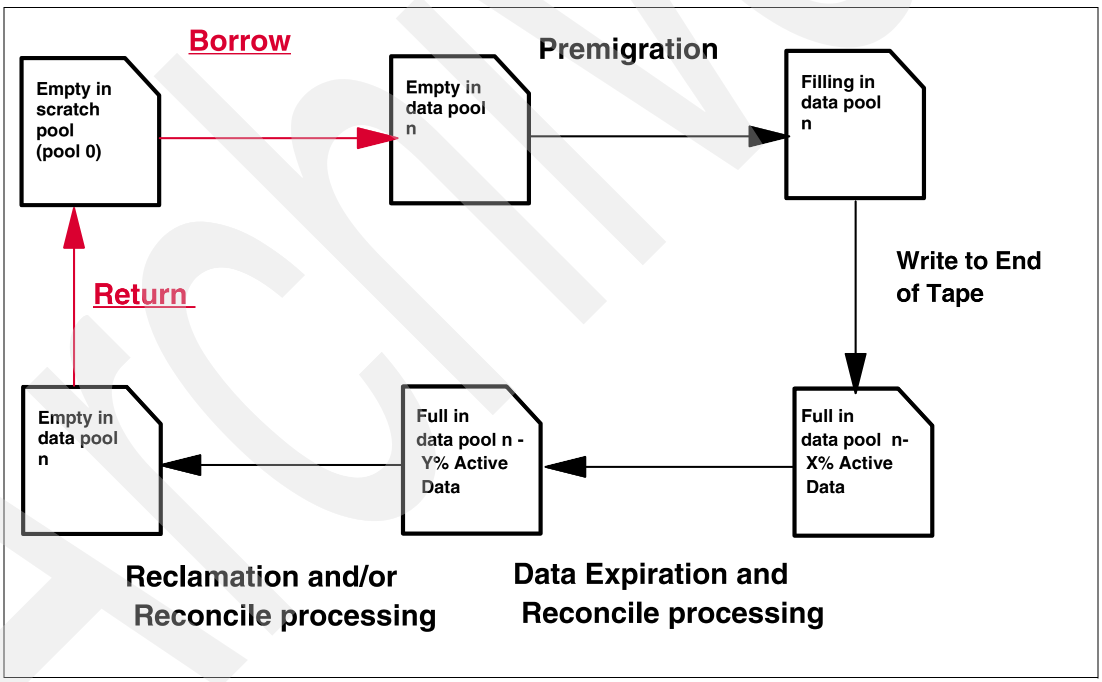

Figure 2-14 Borrowing and return flow

Stacked volume pool properties

Physical 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 HSM recycle target and 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.

Space reclamation is defined per pool as a percentage value. If the amount of active data per cartridge falls below this percentage, the cartridge is eligible for reclamation.

In addition, and to ease migration to new cartridge types or tape device tapes, several parameters can be used to control when a physical cartridge is to be reclaimed. These parameters apply independently of each other; if one of the criteria listed is met, the cartridge is eligible for reclamation. These parameters are:

- **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 minimum 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 minimum active data field.

- **Days before Secure Data Erase:**

  If you specify a non-zero value, reclamation will be triggered by this definition as well. After the number of days specified has passed since the first byte on a cartridge in this pool was invalidated, the cartridge will be reclaimed and subsequently physically erased. A value from 1 to 365 days can be specified. A value of zero disables the function.

Note: 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.

Figure 2-15 shows an example of the pool properties when displayed with the ETL Specialist.

Click on a pool (row) in the table below to modify the current properties for that pool.

Refresh   Last refresh: Freitag, 9. September 2005 19:54:56

Stacked Volume Pool Properties

	Pool	Media Class	First Media	Second Media	Borrow Indicator	Reclaim Pool	Maximum Devices	Days Before Secure Data Erase	Days Without Access	Age of Last Data Written	Days Without Data Inactivation	Maximum Active Data
Modify	1	3590	J - HPCT	None	Borrow, Return	1	All Devices	0	0	0	0	0
Modify	2	3590	K - EHPCT	None	Borrow, Return	2	All Devices	0	0	0	0	0
Modify	3	3592	JA	None	Borrow, Return	3	All Devices	0	0	0	0	0
Modify	4	3592	JJ	None	Borrow, Return	4	All Devices	0	0	0	0	0
Modify	5	3592	JA or JJ	None	Borrow, Return	5	All Devices	0	0	0	0	0

Figure 2-15 Pool Properties

The VTS takes into account all of the policies specified for a pool independently when determining the physical volumes in the pool that are eligible for reclamation. Each pool can have a different set of reclamation policies. If the dual copy function of APM is used, the reclamation policies need to be set up for both the primary and secondary pools.

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 standalone 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 it prior to upgrading B20 VTS with existing 3590 drives to add IBM 3592 tape drives. If Advanced Policy Management is installed, 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 3592s 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 you may want to remove the 3590 drives, so IBM has provided two options to hasten this migration.

With VTS Release 7.4 you can activate the new *Secure data erase* function using the pool properties screen. 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.

> **Important:** A physical volume becomes eligible for reclamation (and subsequent data erasure if requested) if *one* of the reclamation policy settings applies.

2.6.2 Selective dual copy

With Advanced Policy Management, storage administrators have the option to selectively create dual copies of logical volumes within a VTS. This function is also available in the Peer-to-Peer VTS environment. At the site or location where the second Distributed Library is located, logical volumes can also be duplexed:

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

- The copy feature is enabled by the management class setting at the LM where you define the secondary pool.
- Secondary and Primary pools can be intermixed:
  - A primary pool for one logical volume can be the secondary pool for another logical volume.
  - Multiple primary pools can use the same secondary pool.
- At rewind/unload time, the secondary pool assignment is determined and the copy of the logical volume is scheduled. The scheduling of the backup is determined by the premigration activity occurring in the VTS.
- The copy is created prior to the primary volume being migrated.

In a standard VTS, a logical volume could only exist as a single entity on a single Stacked volume. If the Stacked volume is damaged there is potential to lose access to a logical volume, the only ways of ensuring data availability is to use Peer-to-Peer VTS that would create two logical volumes in a separate 3494 and VTS or by using host software to duplex the logical volume.

Advanced Policy management also allows users to create a copy of a logical volume within a single VTS. The copy is created in second Physical pool ensuring physical cartridge separation. Control of Dual copy is via the Management class construct (see 4.4.3, “Creating management classes” on page 147). The second copy is created when the original volume is pre-migrated. Figure 2-16 summarizes the process flow.

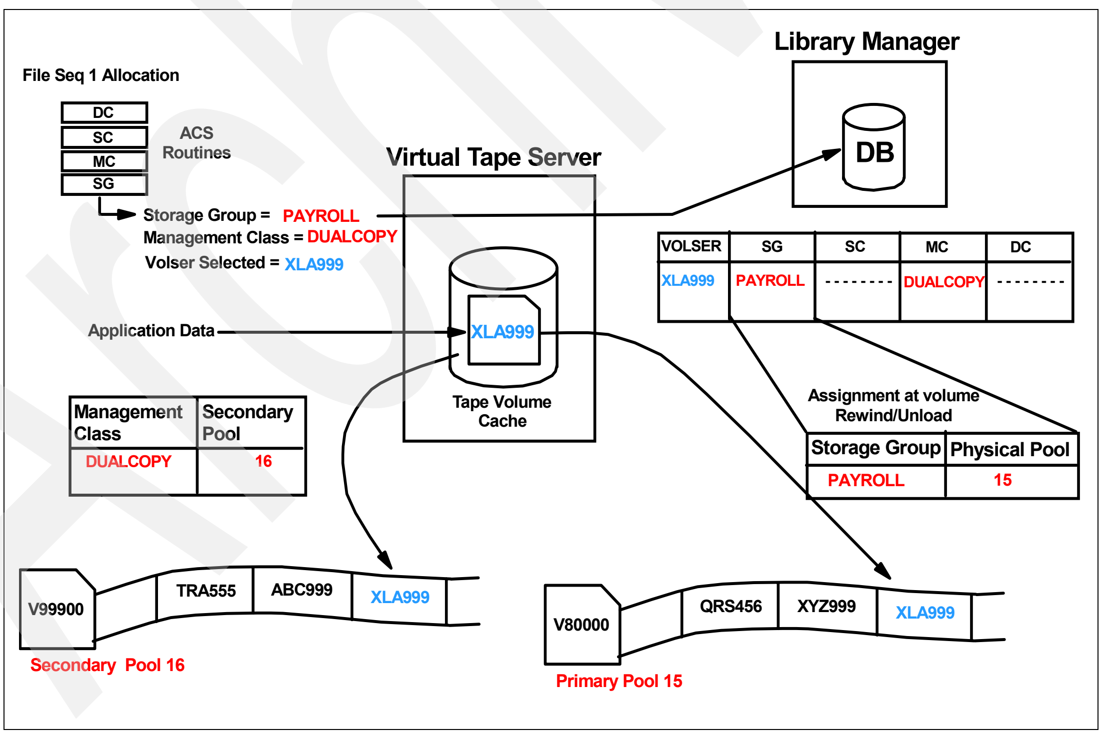

Figure 2-16 VTS dual copy process flow

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

If the primary volume is exported, the secondary copy will be deleted from the VTS; this means that when a critical volume is exported it resides in only one place. If this volume is imported back into the library the second copy is not re-created unless the volume is pre-migrated again.

A primary pool can be a secondary pool for another primary and vice versa. All logical volumes that create a second copy can exist in the same secondary pool regardless of their primary pool. Dual copies can also be created in a Peer-to-Peer VTS environment.

Note: The secondary volume is only used for recovery and as such cannot be addressed by the host system. Without Advanced Policy Management enabled, only a single copy of a logical volume can be created in a standalone VTS.

Logical volume database

Infrastructure changes have been introduced as part of the introduction of APM. One of the changes has been the creation of the Logical Volume Database:

- Its function is to track Advanced Policy Management information, such as:
  - Primary pool assignment
  - Backup pool assignment
  - Cache preference levels
- Where an error has been encountered on a logical volume that has been copied, it provides a locking mechanism used to coordinate the swap from the primary pool to the secondary pool so the recall is transparent to the user.

This database is maintained as a logical database within the Token Database which was added for Peer-to-Peer VTS. The Token Database is used by the VTS but, token data is only used for Peer-to-Peer VTS. The backup of the Token Database is performed at the same time as for the VTS logical volume Database. The backup is done to every physical volume when it has been filled.

Unlike the Physical Volume Database, after code upgrade, there is no initialization process performed on the Logical Volume Database. The database is initialized when new logical volume data is created. It is currently only used when Advanced Policy Management is installed.

For logical volumes that have been dual copied, the VTS maintains a primary copy and a backup copy in its database.

All normal logical volume operations such as mount, read/write, and unload are always done to the primary volume.

Secondary copy initiation

Use of the secondary copy is initiated when a recall failure occurs on the primary volume or for volumes flagged as read-only. If a recall error occurs:

1. The VTS checks for the existence of a secondary copy for the logical volume.
2. Flags in the Logical Volume Database are used to confirm integrity of the primary and secondary files.
3. At this point, the recall is tried again using data from the secondary file. If the recall is successful, then the mount is successful and VTS removes the dependency on the errant primary copy. The VTS also schedules the creation of a new primary copy.
4. If the recall fails again, the mount is failed.

Note: The secondary volume is only used for recovery and as such cannot be addressed by the host system. Also, without Advanced Policy Management enabled, only a single copy of a logical volume is available.

2.6.3 Tape volume cache management

The TVC in a VTS is managed such that it is generally full of virtual volumes. The virtual volumes in cache at any point in time, is determined by the cache management algorithms within the VTS. Virtual volumes are removed from the TVC primarily when space is needed to write a new virtual volume, recall a volume from physical tape or append to a cache resident volume. Volumes to be removed must also have been copied to a physical tape. The default selection of the volumes to remove is based on a 'least recently used' algorithm, so volumes copied to physical tape and that have been in cache the longest without access are removed first.

Through the use of the Advanced Functions (FC4000) or the Advanced Policy Management (FC4001- FC4004) features, you have the ability to influence the residency of virtual volumes in the TVC through the assignment of a cache preference level to a virtual volume.

Preference level 0

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

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

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

Preference level 1

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

When a volume is assigned preference level 1, the VTS adds it to the queue of volumes to copy to physical tape after a four minute time delay and after any volumes assigned to preference level 0. The four minute time delay is to prevent unnecessary copies to be performed when a volume is created, then quickly mounted and appended to again. When space is needed in the cache, the VTS will first see if there are any preference level 0 volumes that can be removed. If there are no preference level 0 volumes that can be removed, the VTS selects preference level 1 volumes to remove based on a 'least recently used' algorithm, so volumes copied to physical tape and that have been in cache the longest without access are removed first.

Preference level 1 is assigned to volumes as a default when the 4000-4004 features are not installed or where a scratch volume is created by a host that does not support Advanced Policy Management.

Note: When a volume is written from beginning of tape, it is initially assigned a preference level of 1. That preference level is overridden at volume close time based on the host set IART or storage class definition. This is to account for re-use of scratch volumes between hosts that are at a software level that supports cache management and those that do not.

The cache management described above is also supported for the PTP VTS for the I/O VTS. Figure 2-17 illustrates the effect of using Preference Levels on the TVC.

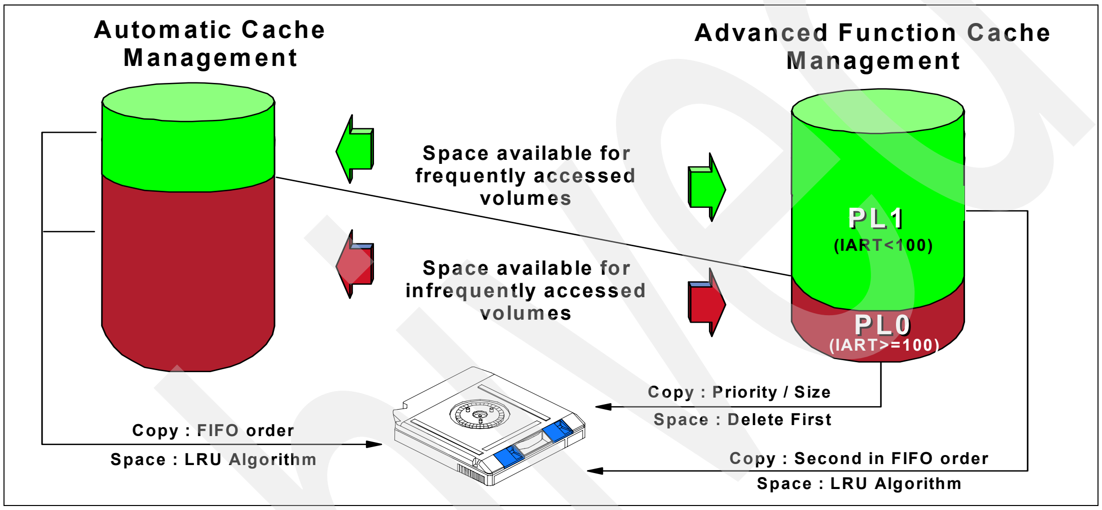

Figure 2-17 Effect of tape volume cache management on TVC

With either of the Advanced Function or Advanced Policy Management features, either cache management preference level 0 or 1 can be associated with a volume. The method for assigning the preference level is different for the difference features. Once a preference level has been assigned to a volume, that assignment is persistent until the volume is re-used for scratch and a new preference selection is made. For example, this means that a volume assigned a preference level of 0 will maintain that preference level when it has been subsequently recalled into cache.

With the introduction of the Advanced Functions feature, FC4000, you have the ability to assign a preference level to virtual volumes through storage class. It uses the Initial Access Response Time Seconds (IART) attribute of the storage class to determine whether a virtual volume (data set) should be assigned to preference level 0 or 1. The IART is an existing storage class attribute. It was originally added to specify the desired response time (in seconds) for an object using the Object Access Method (OAM).

See Table 2-3 for the values that can be assigned and the corresponding preference level. The host passes the IART value to the VTS based on the definition of the storage class assigned to the volume through the Automatic Class Selection (ACS) routines. With FC4000, the IART value is interpreted by the VTS and the volume is assigned the appropriate preference level.

If you wish to have a virtual volume preferred to remain in cache, assign a storage class to the volume whose IART value is 99 seconds or less. Conversely, if you wish to have a virtual volume be preferred out of cache, assign a storage class to the volume whose IART value is 100 seconds or more.

Table 2-3 .Cache preference group summary

Cache Level Preference	IART	Management technique
0	100 seconds or more	Preferential removal of volumes from cache, largest first.
1	Less than 100 seconds	Removal of volumes after group 0, least recently used first.

With the Advanced Policy Management (FC4001 - FC4004), the storage class name assigned a volume in the ACS routine is directly passed to the Library Manager. Figure 2-18 shows the process of passing a host defined storage class to the Library Manager. For non-DFSMS environments, you can assign a storage class to a range of logical volume (either when inserting them or afterward). See 4.3.4, “Define (insert) logical volumes” on page 130 for more details. At the library, you define one or more storage class names and pick preference level 0 or 1 for them. To be compatible with the IART method to set the preference level, the storage class definition also allows a Use IART selection to be made. Assuming that the Use IART selection has not been specified, the VTS sets the preference level for the volume based on the preference level 0 or 1 definition for the storage class assigned to the volume.

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

The primary difference between the two feature codes for cache management is that with FC4001 to FC4004, the definition has been moved outboard of the host. This includes the default cache management behavior, should you choose to. While this can be done, it is not recommended for z/OS hosts. Defining Advanced Policy Management functions for cache management is covered in 4.4.4, “Creating storage classes” on page 149.

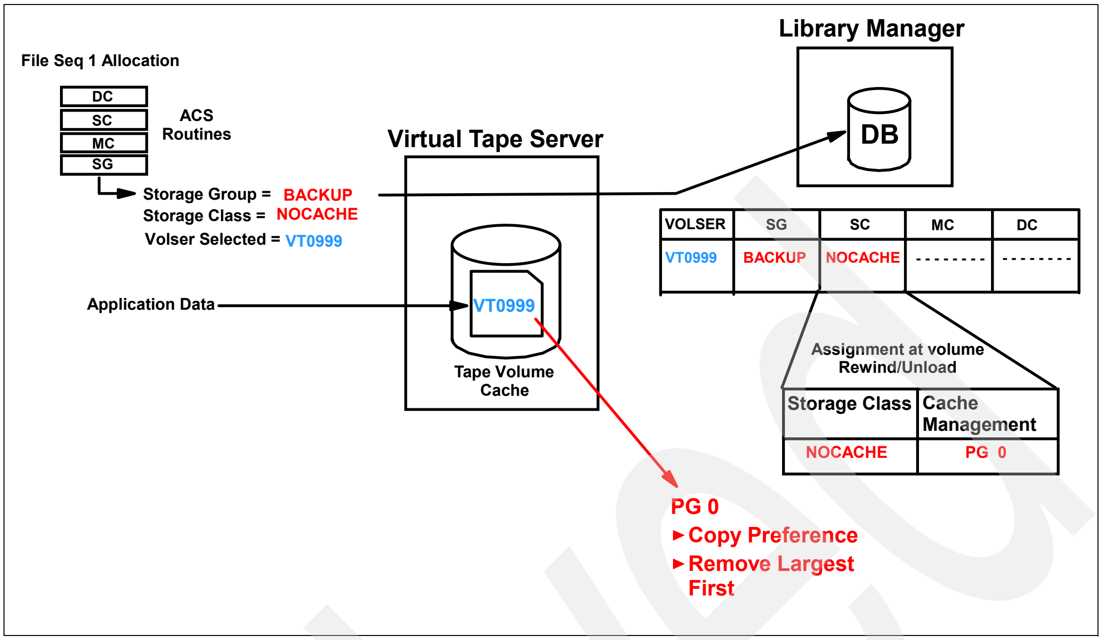

Figure 2-18 TVC management with Advanced Policy Management

If you are currently using the storage class IART value to set cache residency in the VTS and wish to migrate to the controls with FC4001-FC4004, we recommend that you define the same storage class names outboard at the Library Manager and specify Use IART for the preference level setting. If at a later time you wish to no longer use the host set IART values, you can modify the previously defined storage class names to explicitly specify the preference level.

2.6.4 Large Logical Volume sizes

With VTS Release 7.4, the default logical volume sizes of 400 and 800 MB are extended to 1000, 2000 and 4000 MB. The Library Manager honors the data class name passed down by the ACS routines and uses the logical volume size defined in the data class construct on the Library Manager. Figure 2-19 shows the process flow when using large logical volume sizes.

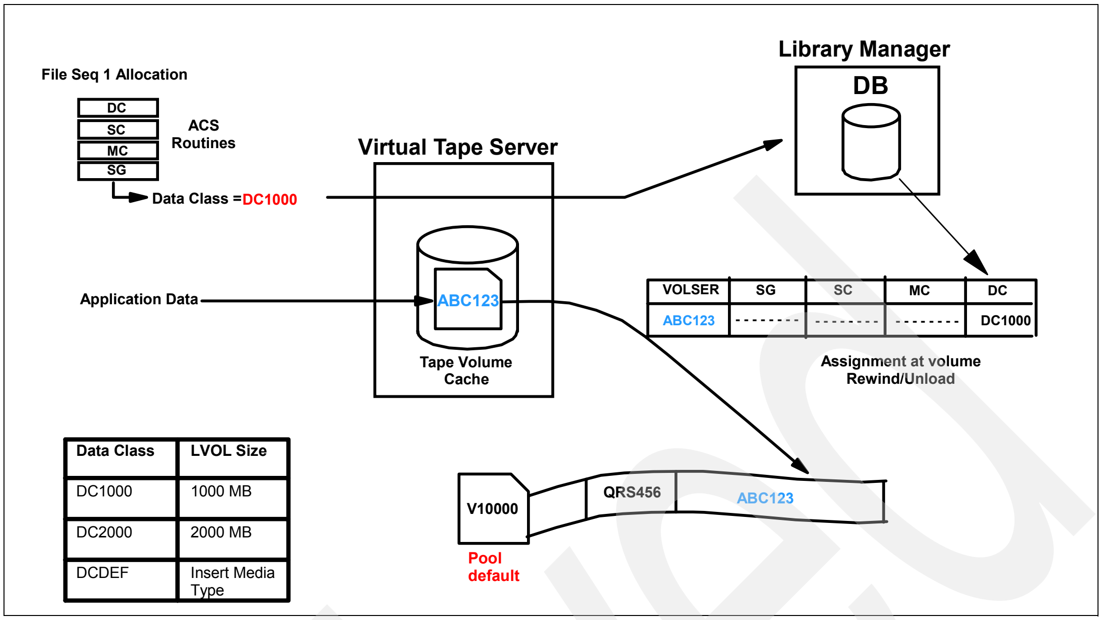

Figure 2-19 The new data class construct

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 the host closes the volume and opens another volume to continue writing.

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. The definition process is described in 4.3.4, “Define (insert) logical volumes” on page 130.

Note that you might have to define larger logical volume sizes to your applications as well, if an application calculates how much space is left on a volume to fill it to its full capacity. For example, DFSMShsm only knows that a MEDIA1 or MEDIA2 volume is mounted and will stop filling the volume at around 400 or 800 MB. To allow HSM to use the full capacity of a large logical volume, you should change the PERCENTFULL parameter from the usually used value of 97% to about 480 to 490% when using 4000 MB logical volumes.

2.6.5 PTP copy mode control

In a Peer-to-Peer VTS, there are three 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 a second copy is queued but not created and the application is allowed to end.
- **No copy:** Where the logical is only written to one VTS and no copy is created in the other VTS.

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.

For more information about IBM TotalStorage Peer-to-Peer VTS, refer to the *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115.

2.6.6 Reclamation

The rules that govern reclamation 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, e.g. 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.

2.6.7 Import/Export function

To move a set of logical volumes for vaulting, or to move volumes to another IBM VTS, an Import/Export function is available with the Advanced Function (FC4000) or Advanced Policy Management (FC4001-4004). The export process allows you to move logical volumes from the VTS to exported stacked volumes which are subsequently removed from the VTS.

These are the basic steps of the export process:

1. The host designates logical volumes to be exported from the VTS.
2. The VTS copies the logical volumes to a set of tape cartridges, *exported stacked volumes*, including information fully identifying the logical volumes exported.
3. The VTS deletes all internal references to the exported logical volumes.
4. The VTS puts the tape cartridges containing exported logical volumes in the export-hold category for later operator ejection and informs the host of the results of the export process.

The import process allows you to copy logical volumes from exported stacked volumes back into the VTS, where the VTS can be the originating machine or another VTS. These are the basic steps of the import process:

1. An operator inserts import cartridges (cartridges containing data to be imported) into the IBM TotalStorage Enterprise Automated Tape Library.
2. The host designates the stacked volumes or specific logical volumes to be imported into the VTS.
3. The VTS determines the logical volumes involved and copies them from the set of import cartridges onto cartridges already in the VTS. The VTS creates internal references to each imported logical volume.
4. The VTS informs the host of the results of the import process.

Exported stacked volumes created on IBM 3590 Model B1A tape drives can be imported into IBM 3590 Model E1A configurations; however, exported stacked volumes created on IBM 3590 Model E1A tape drives cannot be imported into IBM 3590 B1A configurations. The same applies to exporting H1A cartridges (you cannot import to a lower capacity cartridge). FC4001 contains Import/export. If you add this feature to an existing VTS to use Import/Export, there must be a Convenience I/O installed in the Library Manager for this feature to work.

When a logical volume is exported from a VTS that has FC4001 installed, the SMS constructs associated with that logical volume. If imported on to a non APM VTS the constructs are ignored. The policies can be changed at Import time by overriding the constructs and if the export volume is from a Non-APM VTS then constructs can be assigned to the logical volumes at import. Export will automatically select a 3590 K type cartridge unless there are only 3590 J type cartridges available or it calculates that the exported data will fit on a J cartridge.

Refer to Appendix B, “VTS Import / Export” on page 369 for more details.

2.6.8 Open Systems considerations

In an Open Systems environment, SMS constructs are not available from the host but Advanced Policy Management is made available to these environments. This allows controls such as Cache management, Volume Pooling and Selective Volume Dual copy to Open Systems hosts. The constructs can be set at insertion time, so a range of logical volumes can be set with a storage group, management class etc.

We recommend that you do not set these constructs at this stage for volumes from SMS hosts, as they will be overridden by information from the host.

- Static assignment of constructs is ignored by the host.
- Assignment is kept when the logical volume is returned to scratch.
- Non-MVS systems support multiple logical scratch pools.

> **Note:** Exploitation of APM in non-MVS environments can be achieved through static assignment of constructs to logical volumes. Host control is given through utilization of multiple scratch pools.

3

Preinstallation planning and sizing

In this chapter we document the key actions that need to take place before you install an IBM TotalStorage VTS. We provide the information you need to prepare for installation:

- Summarize possible hardware configurations of an IBM TotalStorage Enterprise Automated Tape Library (3494) including the IBM TotalStorage Virtual Tape Server (VTS).
- Describe the VTS attachments to IBM 3494 and 3584 tape libraries.
- Give you planning information for sizing the VTS configuration.
- Provide information for physical planning and placement.
- Explain which data are suitable for the VTS.
- List the software requirements.
- Describe switch and fabric support.
- Provide information for operator education and training.

In Appendix C, “VTS implementation step-by-step” on page 419”, we provide a detailed description of all tasks to implement a VTS.

3.1 Hardware configurations

The first generations of VTS had to be connected to the 3494 Tape Library with 3590 drives. In the meanwhile the support for 3592 drives was added, as well as the heterogeneous support of 3590 and 3592. With the recent availability of R7.4 we have the additional option to connect a model B10 or B20 VTS to a 3584 Tape Library with fibre connected 3592. These alternatives are described separately in the following sections.

3.1.1 VTS configuration options

Here we present a model comparison between Model B18 VTS (remember that this model is withdrawn from marketing), Model B10 VTS, and Model B20 VTS with their supported configurations and attachments (Table 3-1). More detailed descriptions of specific configurations are given in the following sections.

Table 3-1 VTS model comparison

Model	Model B18 VTS[1]	Model B10 VTS	Model B20 VTS
Library Support	3494	3494 and 3584	3494 and 3584
Number of VTS (3494/3584)	2 / no	2 / 8	2 / 8
Virtual Drives to z/OS	32/64/128 (2/4/8 Control Units)	64 (4 Control Units)	128/256 (8/16 Control Units)
Virtual Drives to SCSI	up to 32	up to 32/64	up to 64
3590 Tape Drives	3 to 12	4 to 6	6 to 12
3592 Tape Drives	no	4 to 12	4 to 12
3590 and 3592[2]	no	no	6 and 12
Host Attachment	0/2/4/8 ESCON 0/2/4 SCSI	0/2/4 ESCON 0/4/8 SCSI 0/2/4 FICON	8/16 ESCON 0/8 SCSI 0/4/8 FICON
ESCON/SCSI	up to 4/4	up to 4/8	up to 8/8
FICON/ESCON	no	2/4	4/4
Tape Volume Cache[3]	216 to 5184 GB	648/1296 GB	2592/5184 GB
Logical volumes	up to 250,000	up to 250,000	up to 500,000
Logical volume sizes (GB) compressed (3:1)	1.2 / 2.4	1.2 to 6.0	1.2 to 12.0
Common	Master Console for Service feature Peer-to-Peer VTS GDPS support ETL Specialist and ETL Expert 3590 EHPCT (K) cartridge support Advanced Policy Management		

Model	Model B18 VTS[1]	Model B10 VTS	Model B20 VTS
Notes: 1. The Model B18 VTS has been withdrawn from Marketing. 2. 3590-B1A, 3590-E1A, 3590-H1A, 3592-J1A,and 3592-E05 are supported as physical tape drives to the VTS. 3. Assuming 3:1 compression			

3.1.2 Model B10 VTS and Model B20 VTS configurations

In this section we describe the following configuration options, which are available with the Model B10 VTS and Model B20 VTS.

Host attachment

The host attachment options are shown in Table 3-2 through Table 3-5.

Table 3-2 Model B10 VTS ESCON/SCSI attachment configuration options

No. of ESCON channels	No. of SCSI buses	Attachment features
2	0	3412
4	0	3412 x 2
4	4	3412, 3418, 3422 x 2
0	8	3422 x 4

Table 3-3 Model B20 VTS ESCON/SCSI attachment configuration options

No. of ESCON channels	No. of SCSI buses	Attachment features
8	0	3412 x 4
16	0	3412 x 4, 3418 x 4
8	8	3412 x 2, 3418 x 2, 3422 x 4

Table 3-4 Model B10 VTS ESCON/FICON attachment configuration options

No. of ESCON channels	No. of FICON channels	Attachment features
2	2	3415 / 3416 x 2, 3412 x 1
0	2	3415 / 3416 x 2
0	4	3415 / 3416 x 4

Table 3-5 Model B20 VTS ESCON/FICON attachment configuration options

No. of ESCON channels	No. of FICON channels	Attachment features
0	4	3415 / 3416 x 4
4	4	3415 / 3416 x 4, 3412 x 2
8	4	3415 / 3416 x 4, 3412 x 2, 3418 x 2
0	8	3415 / 3416 x 8

Each ESCON channel in the VTS is capable of supporting 64 logical paths, providing up to 1024 logical paths for Model B20 VTS with sixteen ESCON channels and 256 logical paths for Model B10 VTS with four ESCON channels. Each logical path can address any of the 32, 64, 128 or 256 virtual devices in the Model B20 VTS.

Each FICON channel in the VTS can support up to 128 logical paths, providing up to 1024 logical paths for the Model B20 VTS with eight FICON channels. With a Model B10 VTS, 512 logical paths can be provided with 4 FICON channels. As with ESCON, each logical path can address any of the 64, 128 or 256 virtual devices in the Model B20 VTS.

Note: Intermixing of FICON and SCSI interfaces is not supported.

Cache sizes

The cache size options are shown in Table 3-6.

Table 3-6 IBM TotalStorage VTS Model B10 VTS and Model B20 VTS TVC sizes

Feature codes	Base capacity	Model B10 VTS	Model B20 VTS
3704	144 GB / 216 GB	648 GB	N/A
3705	288 GB / 432 GB	1296 GB	N/A
3705 x 2	576 GB / 864 GB	N/A	2592 GB
3705 x 4	1152 GB / 1728 GB	N/A	5184 GB

Performance Accelerator introduced with Model B18 as an optional feature is standard in Model B10 VTS and Model B20 VTS. Usable capacity is assuming 3:1 compression with the GPFS file system with Performance Accelerator.

Drives for the VTS

These are the drives you can choose:

- For the Model B10 VTS, four, five or six 3590-B1A/E1A/H1A or four to twelve 3592-J1A drives can be associated with VTS.
- For the Model B20 VTS:
  - Six to twelve 3590-B1A/E1A/H1A in one or two adjacent 3494-D12 frames can be associated with VTS or
  - Four to twelve 3592 drives in a dedicated 3494-D22 frame or
  - Four to twelve 3592 drives installed in one or more 3584-D22 frames,
  - A mix of four to twelve 3592 drives in a 3494-D22 frame and one to six 3590 drives in a 3494-D12 frame.

3.2 IBM 3494 Tape Library attachment

The VTS has to be connected to a 3494 Tape Library or a 3584 Tape Library with the 3953-L05 Library Manager. When connected to a 3494, the VTS may be attached to 3590 and/or 3592 tape drives. With a 3584 Tape Library, 3592 drives are supported exclusively. Next we describe the different configurations in more detail.

3.2.1 VTS with 3590 tape drives

The IBM TotalStorage Virtual Tape Server is housed in two 3494 frames: an IBM Model Bxx VTS frame and IBM D12 Frame(s). The D12 Frame is located anywhere in the 3494, whereas the Model Bxx VTS frame can be located at a distance of up to 14 m from the D12 Frame.

See Figure 3-1, where the VTS is located outside the 3494 and its associated D12 Frame is located as frame two. The configuration shown here is a two-frame 3494 with a Model Bxx VTS frame, which is the minimum configuration when a VTS is installed. The L1x Frame has the option of having additional non-VTS drives if required; the type and number of drives depend upon the L1x Frame model. If more than six IBM 3590 tape drives are attached to the VTS, a second D12 frame is required which must be installed adjacent to the first D12 frame.

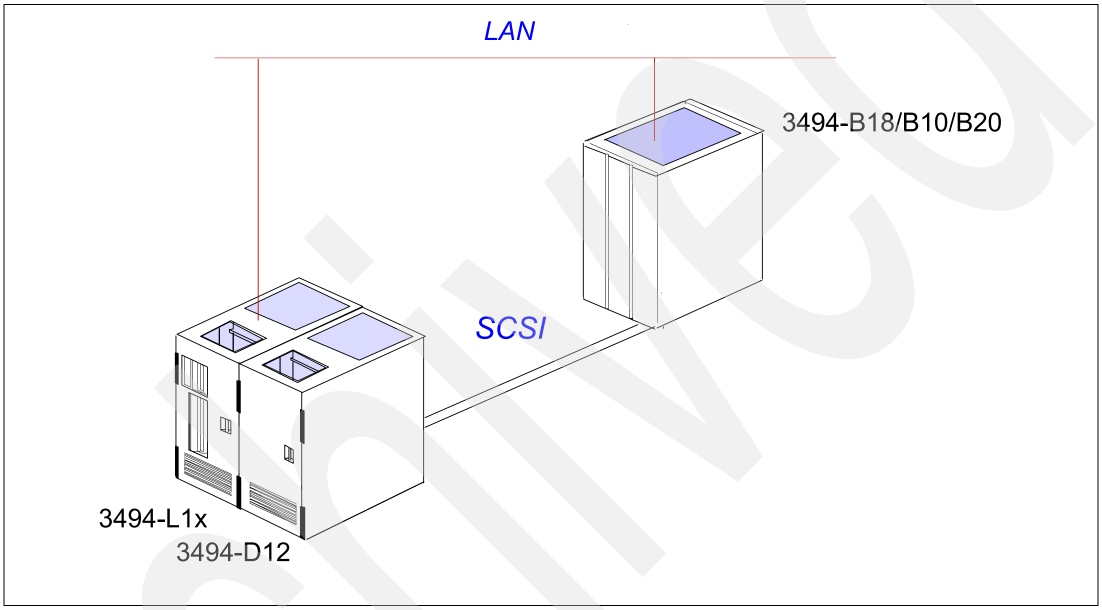

Figure 3-1 Minimum VTS configuration

The minimum configuration shown in Figure 3-1 can be enhanced to a library of up to 16 frames by adding other D1x Frames, IBM 3494-S10 frames and the D12 Frames for a second VTS in the same tape library.

Figure 3-2 shows a 3494 with an IBM 3494 Model L1x frame, two D12 Frames attached to two Model Bxx VTS subsystems, and a D14 Frame, which can include zero to four 3590 tape drives attached to the hosts through the IBM 3590-A60 controller and ESCON channels.

> **Important:** IBM 3590 tape drives are not supported in the IBM 3584 Tape Library. A VTS with IBM 3590 tape drives attached can only be installed in an IBM 3494 Tape Library.

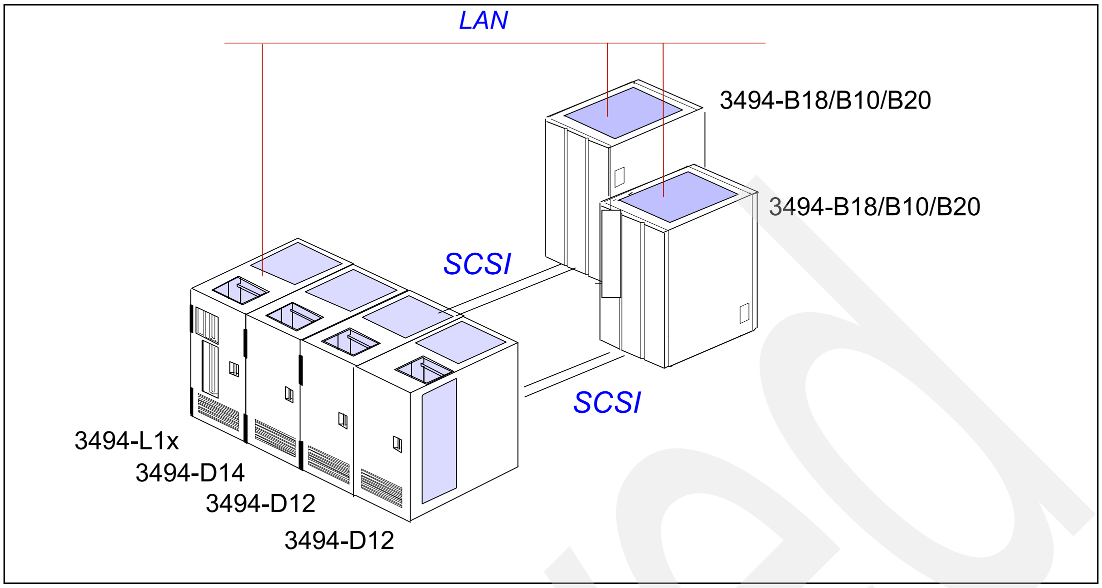

Figure 3-2 IBM TotalStorage Enterprise Tape Library with VTS and native tape drives

3.2.2 VTS with 3592 tape drives in a 3494 Tape 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.

This is not possible for a Model B10 VTS. Existing B10 VTS installations need to be upgraded to a Model B20 VTS to allow migration to 3592 tape drives. In a heterogeneous drive environment in the Model B20 VTS, the tape drives need to be installed in the same 3494 frames as in an Open Systems environment. 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.

In the IBM 3494, 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 are required, and in a HA1 environment the optional 3494 D22 dual PCC can be used to attached 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 3-3.

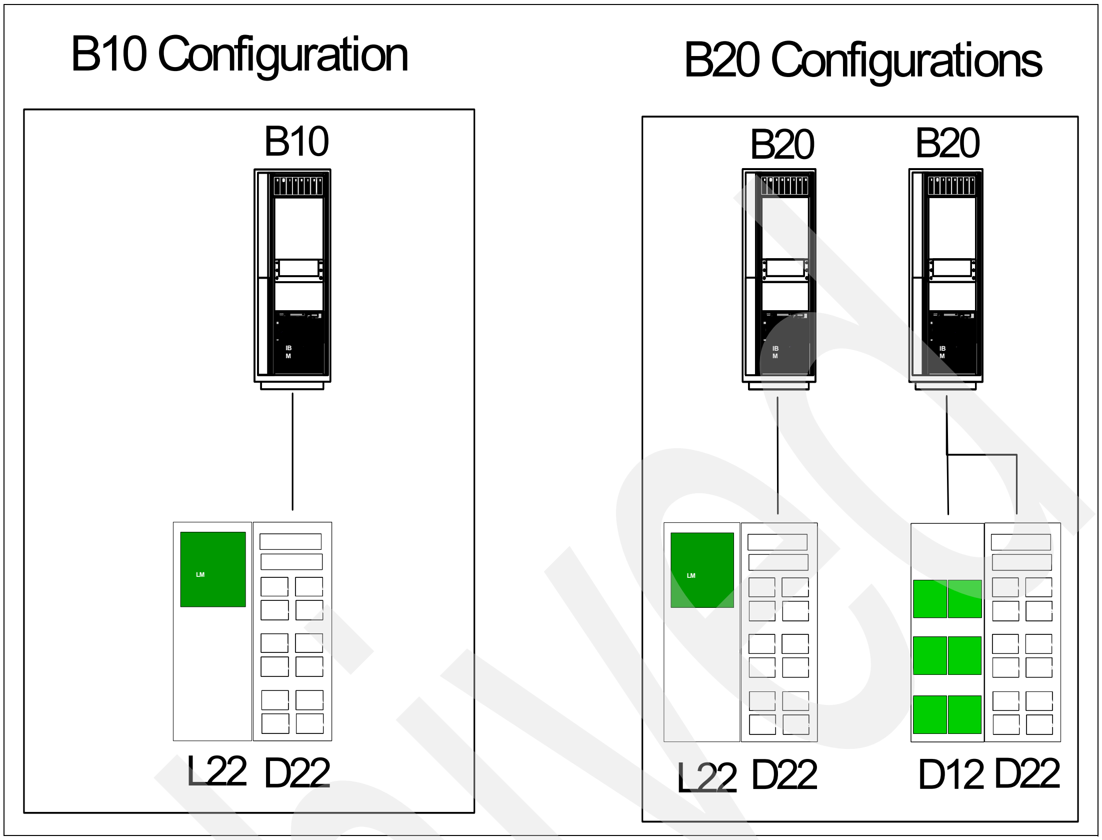

Figure 3-3 VTS Configurations with 3592 installed in a 3494 Tape Library

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 is supported as the maximum configuration in a heterogeneous environment.

The installation options for B20 VTS subsystems with only 3590 drives attached remain the same as before: Up to 12 drives in two IBM 3494-D12 frames can be attached. 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 329 for more details.

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 new 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 B20 VTS, which has only 3592 drives installed. This might be required when merging workloads using import/export, for example.

The configuration rules for IBM 3592 tape drives attached to a Model B10 or B20 VTS can be summarizes as follows:

- 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 6 drives is supported.
- A maximum of twelve 3592 drives are supported and have to be installed all in one D22 frame.
- A maximum of six 3590 drives are supported per D12 frame. All drives are SCSI attached.
- A maximum of two drive frames is supported for the B20 VTS.
  - When two D12 frames are installed, they must be adjacent to each other.
  - When one D12 and one D22 frame are installed, they do not need to be adjacent.
- A maximum of one drive frame is supported for a Model B10 VTS.

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.

The IBM 3584/3953 Tape Library configuration only supports IBM 3592 tape drives attached to the VTS. Up to twelve Fibre Channel drives are supported per VTS. These drives do not need to be installed in a dedicated frame but can be spread across all D12 frames of the 3584 Tape Library.

The VTS is connected to the IBM 3953-L05 Library Manager and, through switches installed in the IBM 3953-F05 base frame, to the BIM 3952 tape drives installed in the IBM 3584 Tape Library.

Both IBM 3592- J1A and IBM 3592-E05 models are supported behind the same VTS. When Model E05 drives are installed, they always operate in J1A emulation mode. You cannot replace existing Model J1A drives with E05 drives, but you can add new E05 drives to a VTS configuration which already has Model J1A drives installed. When intermixing Model J1A and E05 tape drives behind the same VTS, the total number of tape drives cannot exceed twelve drives.

3.2.3 Library configuration rules

Figure 3-5 shows a sample configuration of an IBM 3494 Tape Library which can contain up to two VTS systems, plus IBM 3592 and IBM 3590 tape drives attached to Open Systems hosts or to IBM 3590 or 3592 tape controllers.

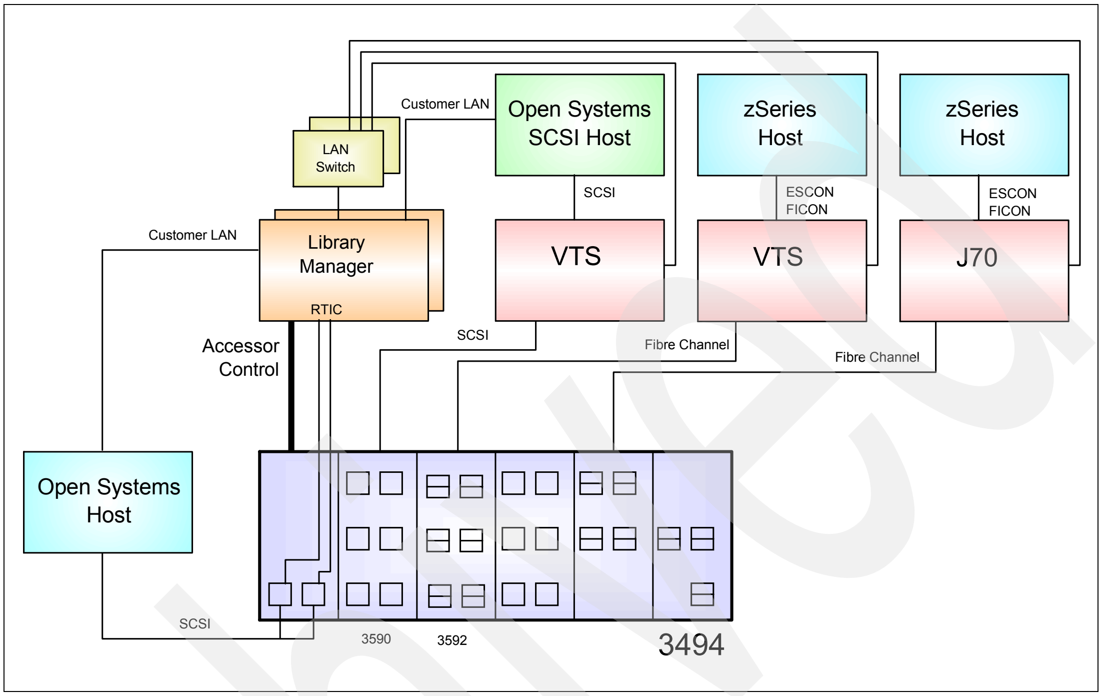

Figure 3-4 The logical libraries in a 3494 Tape Library

Physical 3494 library: This can contain up to three logical libraries with their own LIBID:

- All native, zSeries and/or Open Server attached tape drives
- VTS_1
- VTS_2

Logical library: The logical library is defined in SMS:

- The 3494 can be partitioned between two or more hosts by using different Library Manager volume categories and by using their own set of control data sets (TCDB, TMC, ACDS).
- It can be shared between different systems (LPARs) or sysplexes.

You can install a maximum of two VTS subsystems within a single 3494 Tape Library. The following configurations are possible:

- Model B10 VTS or Model B20 VTS
- Two Model B18 VTSs with 3590 drives only
- A mixture of Model B18 with 3590 drives only, Model B10 VTS, or Model B20 VTS

Cartridges can be stored in any storage cell within the 3494. The storage cells in the D12 Frame and D22 frame of a VTS subsystem are not dedicated to the cartridges used by that VTS subsystem. The Library Manager function is shared among all VTS and tape subsystems installed and manages all operations inside the 3494.

Optionally, you can install the IBM TotalStorage Enterprise High Availability Tape Frames (HA1 Frames), if your 3494 consists of 3, 4, 6, 8, 10, 12, or 16 frames, not counting the Model Bxx VTS frames. The HA1 Frames consists of two frames that are attached on the leftmost and rightmost sides of the 3494. The frame next to the L1x Frame is a service bay. The frame at the other end of the 3494 contains a second Library Manager and a second accessor in hot-standby mode.

Only one of the Library Managers can be active at any point in time. You can switch the active Library Manager to allow, for example, for concurrent Library Manager LIC update. See the redbook, *IBM TotalStorage Enterprise Tape: A Practical Guide,* SG24-4632, for details on switching to and from the primary Library Manager. Both Library Managers are synchronized permanently. If the active Library Manager fails, the second Library Manager will take over automatically.

The second accessor is in hot-standby mode. If the active accessor fails, the standby accessor will push it into the service bay next to the L1x Frame and will take over automatically.

For high mount performance, you can optionally activate the second accessor from standby to active mode. High availability and high reliability of an HA1 Frames are still maintained because each accessor can pick up activities of the other accessor should it require service. Figure 3-5 shows the configuration options for the HA1 Frames and IBM Magstar 3494 Tape Library Dual Active Accessor (DAA).

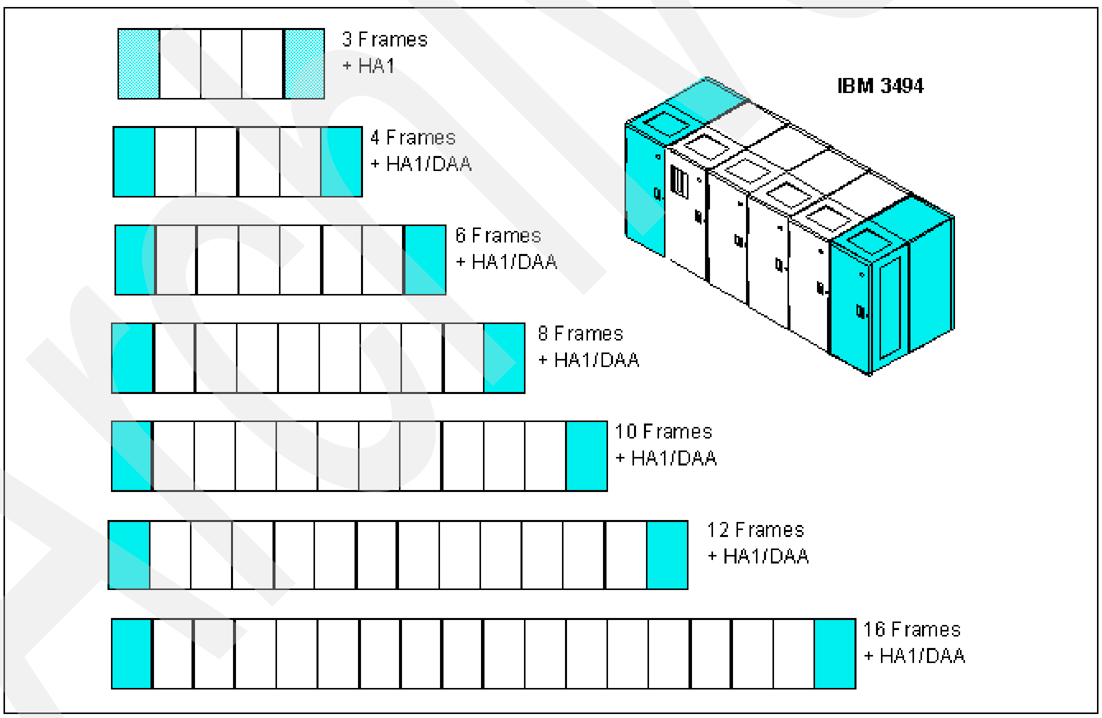

Figure 3-5 3494 configuration options

The 3494 shown in the upper-right corner of Figure 3-5 is the minimum configuration for the DAA feature. The minimum configuration for installation of the HA1 Frames is three library frames, if the DAA is not installed. Attached VTS frames are not considered library frames, because they do not account for the total physical length of the 3494. The minimum configuration for the DAA feature is four library frames. For planning and operating information concerning the IBM Magstar 3494 Tape Library Dual Active Accessor, refer to *IBM TotalStorage Enterprise Tape: A Practical Guide,* SG24-4632.

3.3 IBM 3584 Tape Library attachment

The 3953 Library Manager Model L05 provides attachment to the zSeries enterprise environment for the 3584 Tape Library. With this support, a VTS can attach to 3592 Tape Drives within a 3584 Tape Library and zSeries host have access to native 3592 drives via the J70 controller. Both base VTS and PtP VTS configurations are supported. Within a PtP VTS configuration, a mix of 3494 and 3584 tape libraries is allowed.

3.3.1 The 3953 Library Manager within a 3584 Tape Library

A *physical 3584 Tape Library* supports multiple logical libraries (as many as drives are installed if open server connected). Each logical library consists of:

- Tape drives
- Storage slots

The *Advanced Library Management System* (ALMS) is used to facilitate the definition and management of multiple logical libraries. It has to be used even if only one logical library is required.

A 3953 Library Manager can only be connected to a single 3584 logical library — which may be one out of many other logical libraries in this specific 3584 Tape Library. Within a 3953 Library Manager logical library partition, up to three logical libraries (identical to the logical libraries of a 3494) are supported:

- All native zSeries attached 3592 tape drives via J70 controller - *NO open server hosts*
- VTS_1
- VTS_2

The logical library is defined in SMS:

- A *logical library* itself can be partitioned between two or more hosts by using different Library Manager volume categories and by using their own set of control data sets (TCDB, TMC, ACDS).
- Or, it can be shared between different systems (LPARs) or sysplexes.

Open system hosts are connected directly to the 3584 Tape Library, as you can see in Figure 3-6.

The 3953 no longer needs to control the accessor and does not care how many accessors or grippers there are. Also, all physical motions, the opening and closing of doors, the cell scanning after doors are closed again, the inventory processing, the scanning of the I/O station and any error recovery are controlled by the *library controller* of the 3485 library.

The VTS subsystem connects to the 3953 Library Manager Model L05 which is installed in the IBM 3953-F05 base frame. It also connects to the 3592 Tape Drives in the 3584 Tape Library through the switches which are also installed in the IBM 3953-F05 base frame. The control path between the 3953 Library Manager and the automation in the 3584 Tape Library is through the VTS or the 3592 Tape Controller Model J70.

The base firmware microcode for the 3953 Library Manager is the same as for the current 3494 Library Manager. The Library Manager code base is modified to allow it to be separated from the physical 3494 Tape Library. Part of the modification is to support the SCSI Medium Change (SMC) command set to the 3584 Tape Library.

The existing 3584 concept of logical libraries is designed to segregate the physical devices associated with the VTSs and the 3592 Tape Controllers Model J70 managed by a 3953 Library Manager from the Open Systems drives. All VTS and 3592 Model J70 physical devices managed by the same 3953 Library Manager will reside in the same logical library. The logical library associated with the VTSs and J70s may be a portion of the library or the whole library. See Figure 3-6 for a sample configuration.

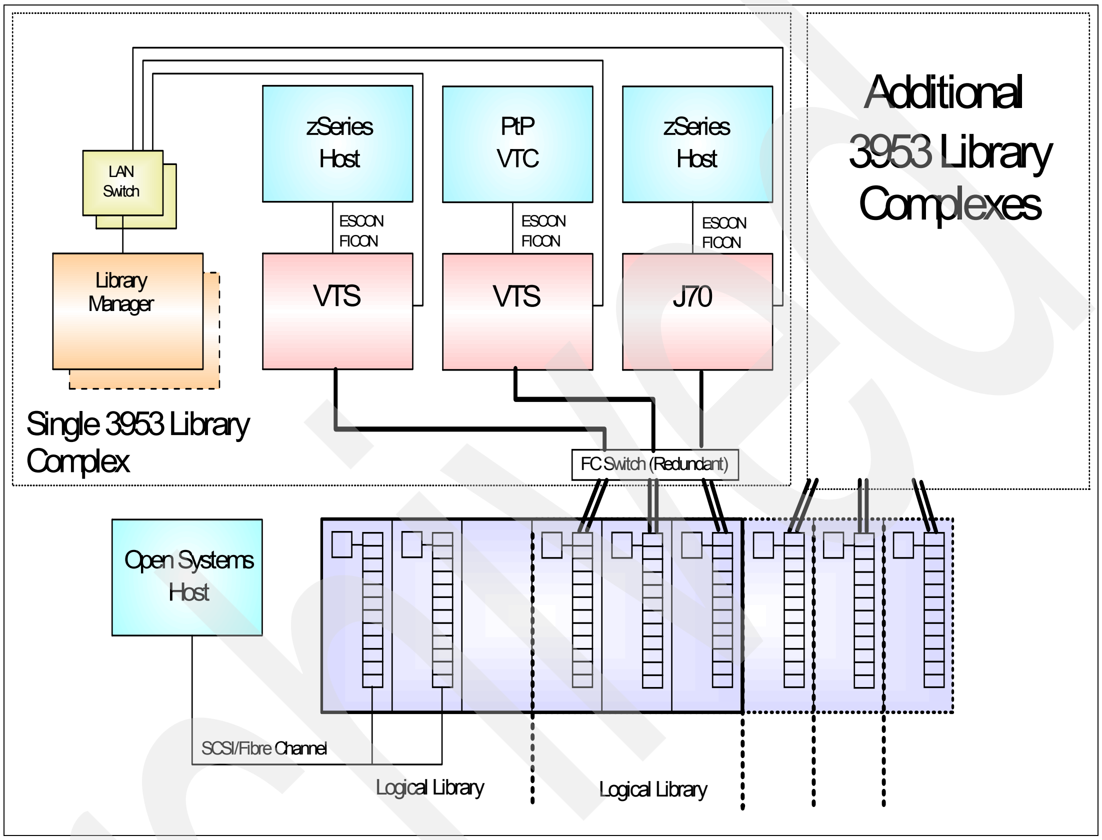

Figure 3-6 3953 Library Manager connected to 3584 Tape Library

With the introduction of the 3953 Library Manager, attachment of stand-alone drives in the 3584 Tape Library to Fibre Channel attached hosts remains the same as before its introduction. The host system Fibre Channels attach directly to the tape drives without the intermediary 3953 Library Manager. All direct attached drives will be logical libraries that are separate from the logical library associated with a 3953 Library Manager.

The physically collection of Library Managers, LAN switches and fibre channel switches for the VTS and J70 connections is what makes-up the 3953.

These units are spread across one or multiple frames which can comprise a complete 3953 system. The 3953 tape system and it's corresponding VTSs and or zSeries hosts comprise what we are calling a "Library Manager Partition", which aligns on the 3584 side of the Tape system with a single logical library partition.

As noted before, connections to both the VTS and J70s are exclusively fibre, while the VTS and J70 connections to their respective hosts can either be ESCON or FICON. On the 3584 side, each 3953 Tape System is cabled into a single logical library partition of the 3584, which can consist physically of a portion of a frame, cartridges and drives or collection of multiple frames housing several drives and hundreds of cartridges. The drives and cells are assigned using ALMS.

The 3953 Library Manager is supported in the following configurations:

- VTS Model B10 or B20 and PtP VTS, J70
- Up to 16 subsystems per 3953 Library Manager (of which two can be VTSs)
- Up to two 3953 Library Managers per 3584

> **Note:** Both the 3494 and 3584 tape libraries use the terms *logical library* and *partition* interchangeably.

3.3.2 Configuration example for a VTS in a 3584 Tape Library

As shown in Figure 3-7, a VTS is connected via the Library Manager in the 3953 frame to its 3592 drives installed in a 3584 Tape Library.

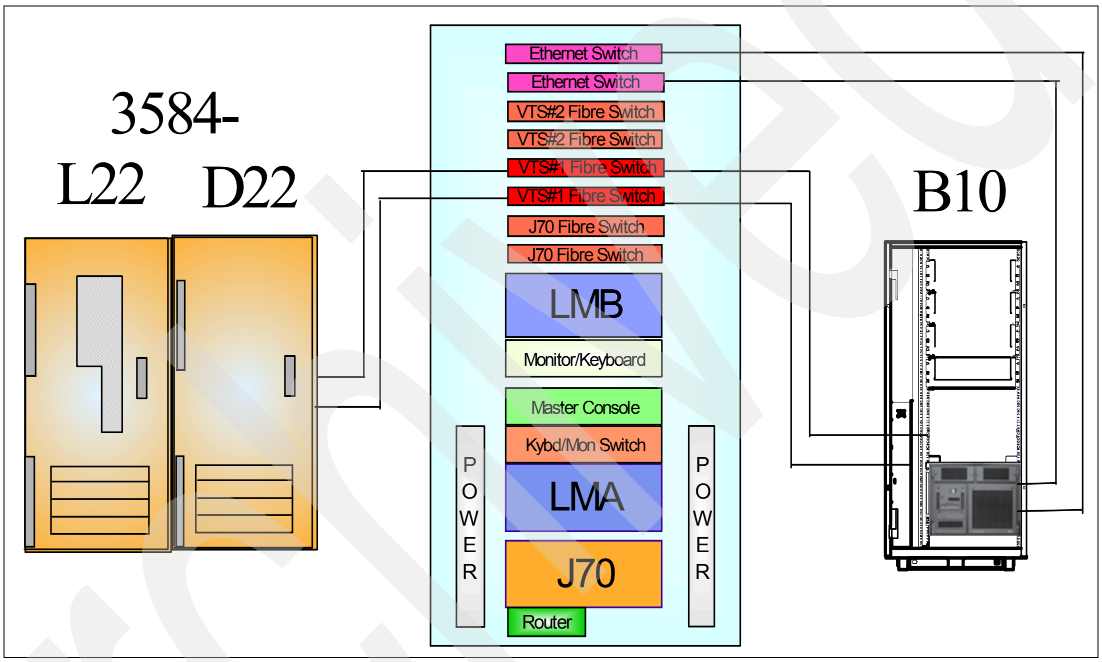

Figure 3-7 VTS connected to a 3584 Tape Library

The minimum 3584 Tape Library configuration for a VTS only environment where no other native drives are needed may consist of only the L22 frame with up to 12 3592 drives. Within a 3584 configuration, the VTS attached drives can be spread over multiple frames, they do not need to be installed in one frame only. If you need more storage slots than one frame provides, we recommend to install the VTS attached drives in more than one frame to further decrease the access time to your cartridges. For high availability, we recommend to configure a second power supply unit and the second Library Manager. You can find more information about the configuration rules and recommendations in the redbook, *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning, Implementation, and Monitoring,* SG24-6789.

3.4 The VTS as a logical library

When an IBM 3494 Tape library contains both an IBM VTS and native tape drives, the Library Manager splits the physical library into logical libraries, one for each VTS and one for all non-VTS tape drives. The non-VTS drives may be z/OS or open server connected. Each logical library appears to the host as a separate physical library. There may be another VTS attached to this tape library, which would look like a third logical library.

When a z/OS connected IBM 3584 Tape Library contains both an IBM VTS and native 3592 drives, which are connected to the host via a Model J70 controller, the Library Manager in the dedicated 3953 frame splits this logical library, which is a subset of the IBM 3584, into two logical library partitions: one for the VTS drives and one for the J70 connected drives. Each logical library partition looks to the host like a separate physical library. There may be another VTS attached to this logical library which would look like a third logical library partition.

Each logical library partition has its own full set of Library Manager categories that can overlap. The Library Manager maintains two sets of category tables and keeps the VTS and physical volumes separate.

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

3.4.1 Partitioning the VTS into multiple logical libraries

The virtual drives and virtual volumes in a VTS (3494 or 3584 attached) can be partitioned just like physical drives and real volumes in any other IBM 3494 Tape Library. Any virtual volume can go to any physical stacked volume. The VTS places no restrictions on the use and management of those resources. However, prior to Advanced Policy Management (APM), the physical stacked volumes managed by the IBM VTS could not be partitioned.

With this feature (APM) you have the ability to partition your stacked media in up to 32 separate pools. This is achieved by assigning a storage group to a defined range of stacked volumes prior to insert time. The steps to achieve this can be found in 4.3.6, “Physical cartridge insertion” on page 135.

Figure 3-8 and Figure 3-9 show the process of copying to stacked volumes without APM (All logical volumes regardless of the VTS logical library are stored on the same media) and with APM (all logical volumes are defined to separate stacked media pools).

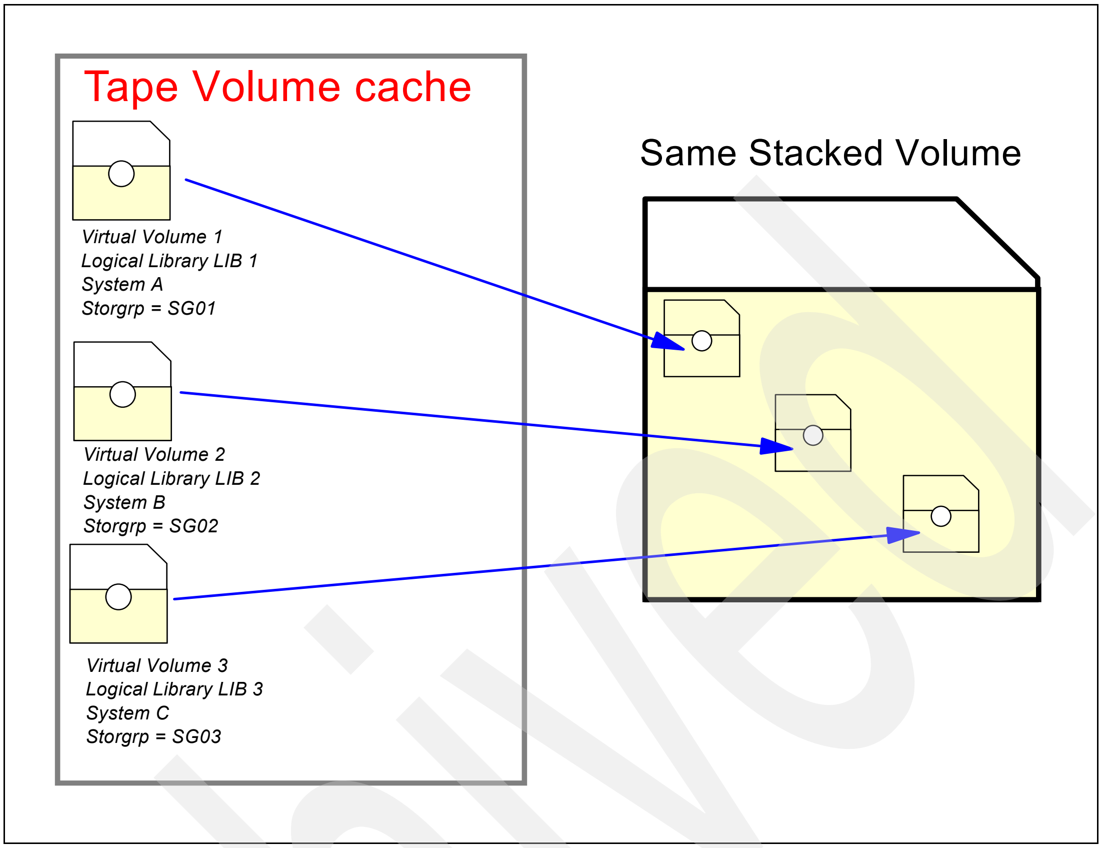

Figure 3-8 Copy process in a partitioned VTS library without APM and pooling

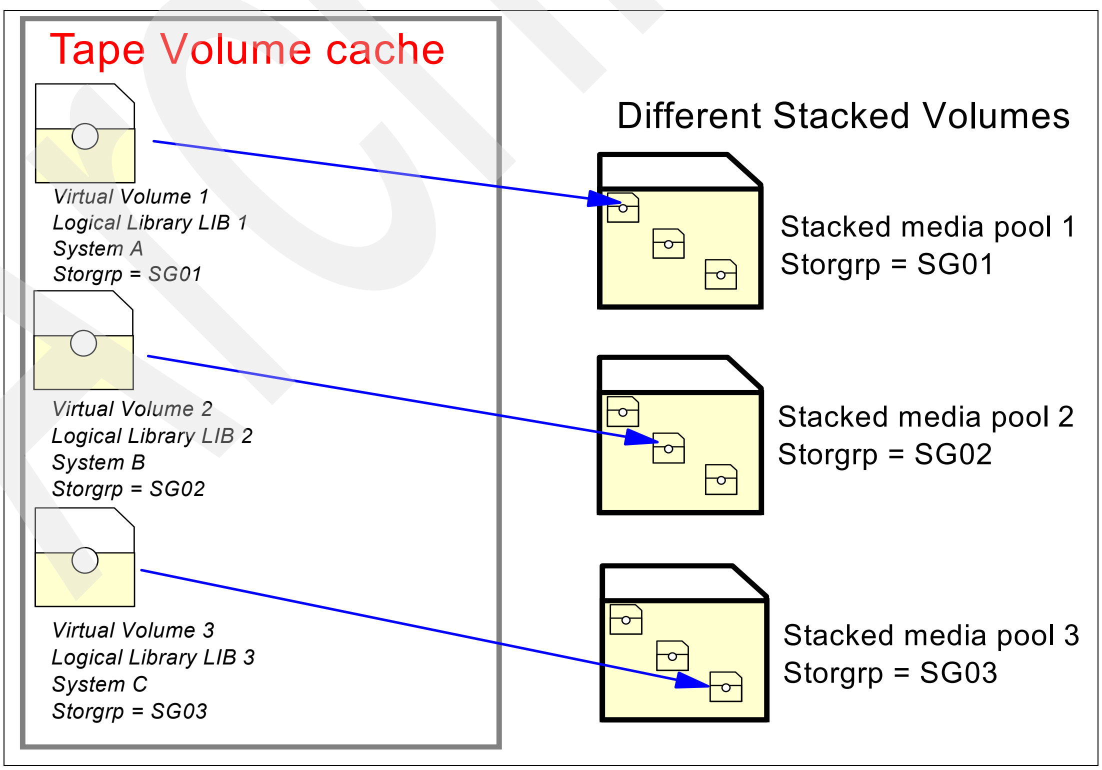

Figure 3-9 Copy process in a partitioned VTS library with APM and pooling

When the VTS itself is partitioned into two or more libraries, the categories used by each accessing system must be unique. For z/OS DFSMS, the default category numbers for all systems in SYS1.PARMLIB(DEVSUPxx) must be changed. The DEVSUPxx member is activated by an entry pointing to it from the IEASYSxx member of SYS1.PARMLIB. Each accessing system should use different category numbers for the cartridges and each partition will have its own range of logical volumes.

3.4.2 Sharing a VTS subsystem

An ESCON-attached VTS supports up to sixteen physical channels, each of which is capable of supporting 64 logical paths. A FICON-attached VTS supports up to eight physical channels, each of which is capable of supporting 128 logical paths. Both ESCON and FICON have a total of 1024 logical paths. Each logical path can address any of the 256 virtual devices in the VTS. We recommend that you use a Director when connecting the VTS to more than one system.

> **Note:** The maximum supported number of ESCON and FICON channels described above applies to Model B20 VTS.

The VTS places no limitations on the number of hosts that can use those channel paths, or the types of hosts, or their operating system environments. This is the same as for any tape technologies that are supported in IBM tape libraries.

An operating environment, through its implementation however, does impose limits. BTLS has a limit of eight because of its eight scratch pools. z/OS DFSMS (starting with Version 1.3) can support up to 32 systems or groups of systems.

A host running BTLS can be attached to a VTS at the same time as a host running system-managed tape.

Basically, anything that can be done with standard drives in an IBM 3494 can be done with the virtual drives in a VTS. Refer to the *Guide to Sharing and Partitioning IBM Tape Library Dataservers,* SG24-4409, for more information about sharing a 3494.

3.5 Tape analysis and sizing the VTS

In this section we document the process of using different tools to analyze your current tape environment and to size the VTS to meet your requirements. With this link you get access to a tools library which offers a bunch of jobs to analyze your current environment and a procedure to unload specific SMF records for a comprehensive sizing with *BatchMagic*, which has to be done by your IBM representative.

3.5.1 IBM Tape Tools

Most of the IBM Tape Tools are available for customers, but some, such as, for example, BatchMagic, are only available for IBM personnel. The tools that are generally available can be downloaded using the following link:

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

You will get a list of text, PDF, and EXE files. At first, open the OVERVIEW.PDF file to get a brief description of all the different tool jobs. All jobs are to be found in the IBMTOOLS.EXE file.

It is a “Self-Extracting Zip file” and, after having downloaded it to your PC, it will expand into four separate files:

- IBMJCL.BIN - JCL for current tape analysis tools.
- IBMCNTL.BIN - Parameters needed for job execution
- IBMLOAD.BIN - Load library for executable load modules
- IBMPAT.BIN - Data pattern library only needed if you will be running the QSAMDRVR utility.

There are two procedures which help to detect weak points in your current tape environment which could influence the overall performance of the VTS in a negative way. Those are bad blocksizes, i.e. smaller than 16 kilobytes and small compression ratios.

Tape Compression Analysis for VTS

By analyzing the MDR (Miscellaneous Data Records) from SYS1.LOGREC or the EREP history file, we can see how well current tape volumes are compressing.

The following job-stream has been created to help analyze these records. See installation procedure in member $$INDEX.

- EREPMDR - JCL to extract MDR records from EREP history file.
- TAPECOMP - Program that reads either SYS1.LOGREC or the EREP history file and produces reports on current compression ratios and MBs transferred per hour.

MDR Analysis for bad VTS Block Sizes

Again by analyzing the MDR (Miscellaneous Data Records) from SYS1.LOGREC or the EREP history file, we can identify tape volumes that are writing small blocks to the VTS and causing extended job runtimes.

The following jobstream has been created to help analyze these records. See installation procedure in member $$INDEX.

- EREPMDR - JCL to extract MDR records from EREP history file.
- BADBLKSZ - Program that reads either SYS1.LOGREC or the EREP history file, finds volumes writing small block sizes, then gets job-name and dsname from a TMS copy.

Data collection and extraction

To correctly size the VTS, you have to analyze your current workload. The SMF records that are required to perform the analysis are record types 14, 15, 21 and 30. One SMF type 94 record per hour is generated by z/OS per VTS and can be used to analyze an already operating VTS for capacity and performance planning. See 7.4, “Monitoring and evaluating VTS performance” on page 290 for more information on analyzing SMF type 94.

You should collect the stated SMF records for all z/OS systems that share the current tape configuration and may have data migrated to the VTS. The data collected should span one month (to cover any month-end processing peaks) or at least those days which are representing the peak load in your current tape environment. You need to check in SYS1.PARMLIB in member SMF to see whether the required records are being collected. If they are not being collected, arrange for collection.

The steps shown in Figure 3-10 are as follows:

Step 1: The TMS data and SMF data collection using the FORMCATS and SORTSMF procedure selects only the required tape processing related SMF records and the TMS catalog information.

Step 2: The files created are compressed by BMPACKT and BMPACKS procedures.

Step 3: Download the packed files (ZIPPED format) to your PC and send them per email to your IBM representative or burn them on a CD if the size is larger than allowed for transmission per email.

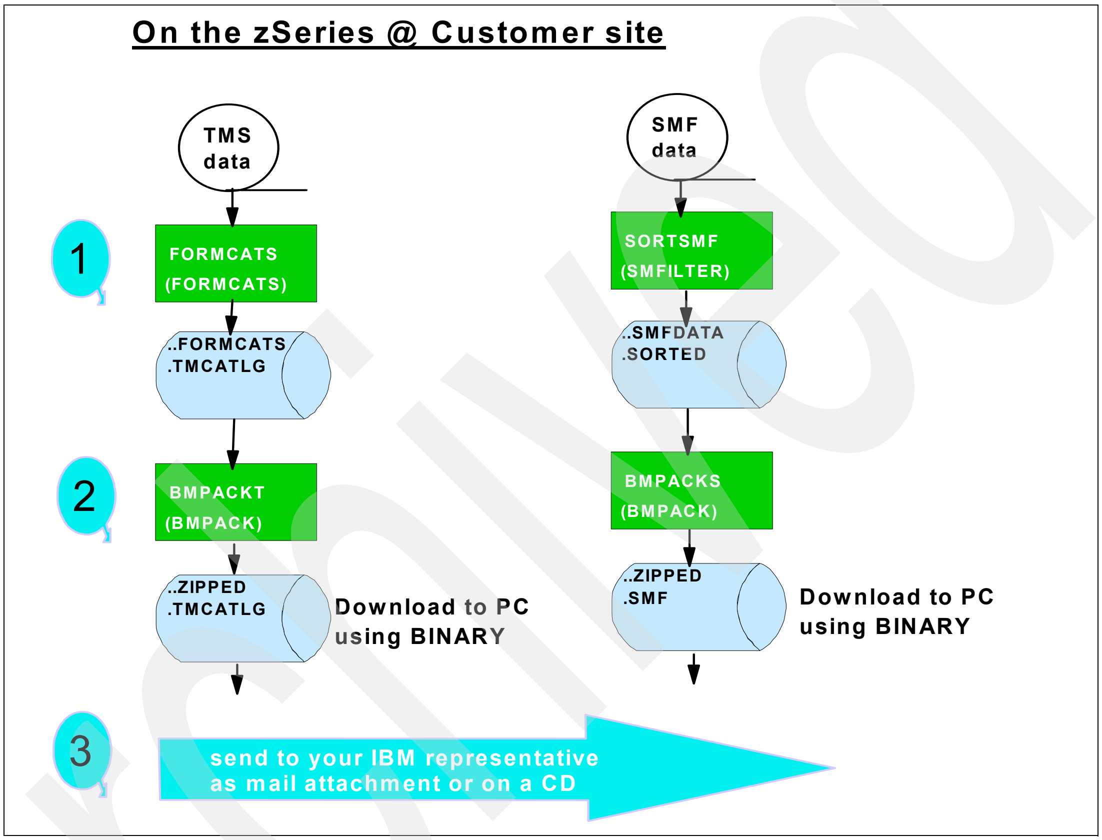

Figure 3-10 Unload process for TMS and SMF data

In addition to the extract file, the following information is useful for sizing the VTS:

- **Number of volumes in your current tape library:** This should include all your tapes (those in automated libraries, on shelves, and off site). If you provide the unloaded Tape Management Catalog data, you do not have to collect the number of volumes.
- **Criteria for identifying volumes:** Because volumes that go off site, that are used as backup, or that should be ignored are of particular interest, their identification is important. Identifiers such as high level qualifiers (HLQs), program names, or job names, must be documented.
- **Number and type of tape control units installed:** This information provides a good understanding of the current configuration and will help you identify the reasons for any apparent workload bottlenecks.
- **Number and type of tape devices installed:** This information will also help you identify the reasons for any apparent workload bottlenecks.
- **Number and type of host channels attached to tape subsystems:** This information will also help you identify the reasons for any apparent workload bottlenecks.

3.5.2 BatchMagic

This tool gives a comprehensive view of the current tape environment as well as predictive modeling on workloads and technologies. The general methodology behind this tool involves analyzing customer SMF type 14/15 21 and 30 records, as well as data extracted from the customer tape management system. The TMS data is required only if you want to make a precise forecast of the cartridges to be ordered based on the current cartridge utilization that is stored in the TMS catalog.

A BatchMagic “run” involves data extraction, grouping data into workloads and then targeting workloads to individual or multiple IBM tape technologies. Batch Magic examines Tape management System catalogs and projects cartridges required with new technology and it models the operation of a VTS and 3590 or 3592 drives and projects required resources. The reports from BatchMagic give a clear understanding of your current tape activities and even more important, make projections for a VTS solution together with its major components such as 3590/3592 drives, which cover your overall sustained and peak throughput requirements.

It is specifically for IBM internal and Business Partner use.

3.5.3 Suitable data for the VTS

The VTS appears to the host systems as two, four, eight, or sixteen 3490E strings with a total of 64, 128 or 256 devices attached respectively. Any data that can reside on a 3480/3490/3590/3592 or cartridges from other vendors can reside on the VTS today. However, processing characteristics of workloads differ, so some data is more suited for the VTS than other data. In this section we highlight some of the considerations to bear in mind when deciding which data to place in an enhanced VTS. When you are selecting the data to migrate to the VTS subsystem, consider the following recommendations:

- **Throughput:** The VTS has a finite bandwidth, as has any other device attached to a computer, so you should plan to put workloads into the VTS that fit within its capabilities. The tools and processes described in 3.5, “Tape analysis and sizing the VTS” on page 72 can provide you with the necessary information to select the correct workloads for the VTS.
- **Cartridge capacity utilization:** One of the key benefits of the VTS is its ability to fully utilize the huge capacity of the 3590 and 3592 cartridges and to manage that capacity effectively without host or user involvement. With enhanced ESCON /FICON and SCSI host attachments that have data compression, a virtual volume may contain up to 12 GB of data using the extended logical volume sizes via the new data classes. Keep in mind that SCSI is not configurable for a VTS attached to a 3584 Tape Library. Very large data sets gain little from the VTS's ability to stack data, so you may want to leave them on native TotalStorage 3590 cartridges or better migrate them to native TotalStorage 3592 cartridges.
- **Data location:** The Advanced Function feature with the Import/Export component feature of the VTS allows you to export logical volumes, for example, to store them in a disaster vault, or to import them into another VTS subsystem. You should consider, however, that exporting logical volumes involves writing them onto exported stacked volumes, which then are ejected from the tape library. If you need to export a large number of logical volumes frequently, you might consider placing them on native 3592 cartridges or writing them to a 3494 installed at the vault location. The Peer-to-Peer VTS offers the perfect solution for automatic duplicating all your tape data to a remote site.
- **Read access after write:** Often, one step of a job writes a tape volume and a subsequent step (or job) reads it. A major benefit can be gained in an enhanced VTS environment when this occurs, as the data can be read directly from the TVC, which effectively removes the rewind time, the robotics time, and load or thread times for the mount.

Figure 3-11 shows the effect on a job and drive assignment that a VTS can have as compared to a native drive. The figure is a freehand drawing — it is not to scale. It shows typical estimated elapsed times for elements that would make up the reading of data from a tape. When comparing the three time lines in Figure 3-11, notice that the VTS cache hit time line does not include robotics, load, or thread time at the beginning of the time line, neither does it include any rewind or unload time at the end.

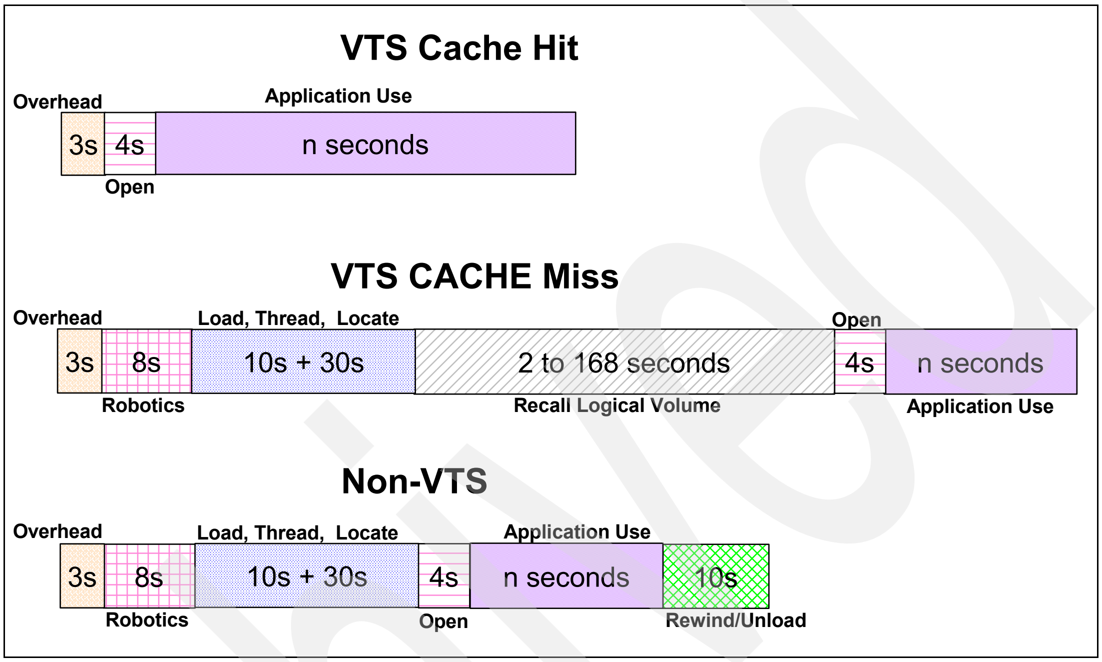

Figure 3-11 Tape processing time comparison (not to scale)

In this example, the VTS cache hit results in a savings in tape processing elapsed time of 40 seconds.

The time reduction in the tape processing has two effects. First, it reduces the elapsed time of the job processing the tape. Second, it frees up a drive earlier so the next job needing a tape drive can access it sooner, as there is no rewind or unload and robotics time after closing the data set.

When a job attempts to read a volume that is not in the TVC, the logical volume is recalled to the TVC. This is shown as the cache miss and includes additional time to move the data from the tape to the TVC. In this case the processing profile is similar to that of a non-VTS environment. Too much recall activity can negatively impact overall throughput of the VTS.

- **Scratch write mount times:** When a program issues a scratch mount to write data, the VTS uses information held in the TVC to simulate the tape being mounted. We describe the scratch mount process in Section 2.1.2, “Mounting a scratch tape” on page 13. The effect of using the stub data in the TVC is that the performance of scratch mounts is fast, because no physical mount is required. The performance for scratch mounts is the same as for TVC read hits, which are compared with non-VTS in Figure 3-11 on page 76.
- **Drive concurrency:** The B10 and B20 VTS appears to the host operating system as 64, 128 or 256 3490E drives. If you have periods of time during the day when your tape processing jobs are limited by drive availability, the VTS may help your processing considerably. The design of the VTS allows transparent access to multiple logical volumes on the same stacked volume, because access to the logical volumes is solely through the TVC. If you need access to more than one logical volume, it is provided without requiring any user involvement, unlike some alternatives such as stacking by using JCL.

- **Multifile volumes:** If you stack multiple files onto volumes today, by using JCL constructs or some other method, the reason you are stacking is most likely to better use cartridge capacity. Best utilization of cartridge capacity is, as already stated, a key objective of the VTS. Therefore, you should find that in many cases manual stacking of data sets onto volumes is no longer required. If you are planning a new application that would have used JCL to stack data sets onto a volume, the VTS makes this JCL step unnecessary.

  Multifile volumes moved to the VTS work without changing the stacking. However, the VTS recalls the complete logical volume to the TVC if the volume is not in cache, rather than moving each file as you access it. Therefore, in some cases, it can be advantageous to let the VTS do the stacking automatically for you. It can save you manual management overhead and, in some cases, host CPU cycles, host channel bandwidth, DASD space, or a combination of all of these.

For further discussion on VTS usage with specific types of data, see Chapter 5, “Data migration” on page 177.

3.6 Planning for volume and cartridge storage

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

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

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

3.6.1 Support for 500,000 logical volumes

With VTS R7.4 you can expand the number of logical volumes to 500,000 per VTS; that means 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 meta data systems must be re configured 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 plan for a data migration you have to consider the following:

- If the source VTS has FC4036 installed, then the target VTS <u>must</u> 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 data base backup additional 250 MB per cartridge are used.

3.6.2 Data compression

When writing data to a virtual volume, the host compression definition is honored. Compression is turned on or off by the JCL parameter DCB=TRTCH=COMP(or NOCOMP), the data class parameter COMPACTION=YES|NO, or the COMPACT=YES|NO definition in the DEVSUPxx PARMLIB member. The TRTCH parameter overrides the data class definition and both override the PARMLIB definition. We strongly recommend that you check your definitions to make sure that you specify compression when writing data to the VTS in order to achieve the optimum throughput. Note that the backward reads executed on a VTS do not degrade performance as they did on real IBM 3480 and IBM 3490 tape drives, when compression was added.

3.6.3 Size of virtual volumes

The effective size of virtual volumes can range from 1,200 MB up to 12,000 MB. This is the size of the virtual volumes after compression (3 : 1). You may see the number of multivolume cartridges (logical volumes) grow depending on the media size from which you are converting. For example, if you migrate data from native 3590 cartridges with 10 GB of uncompressed data and your target logical volume size is 400 MB, then the number of logical volumes needed will increase dramatically.

With R7.4 which provides beside the 400 MB CST emulated cartridges or 800 MB with ECCST emulated cartridges additional logical volume sizes of 1000 MB, 2000 MB and 4000 MB you have the choice to select the most appropriate size depending on the data you put in the IBM VTS.

The 400 MB CST emulated cartridges or 800 MB with ECCST emulated cartridges are still the only two types available for insertion.

> **Important:** In most applications, if you do not have a special need for using CST emulation, use the ECCST.

It is possible to mix both CST and ECCST emulated cartridges which you have inserted with the additional sizes you may select using the data class constructs inside the IBM VTS. The amount of data copied to the stacked cartridge is only the data that has been written. The choice between all available logical volume sizes does not affect the real space used in either the VTS TVC or the stacked volume, as you can see from Figure 3-12.

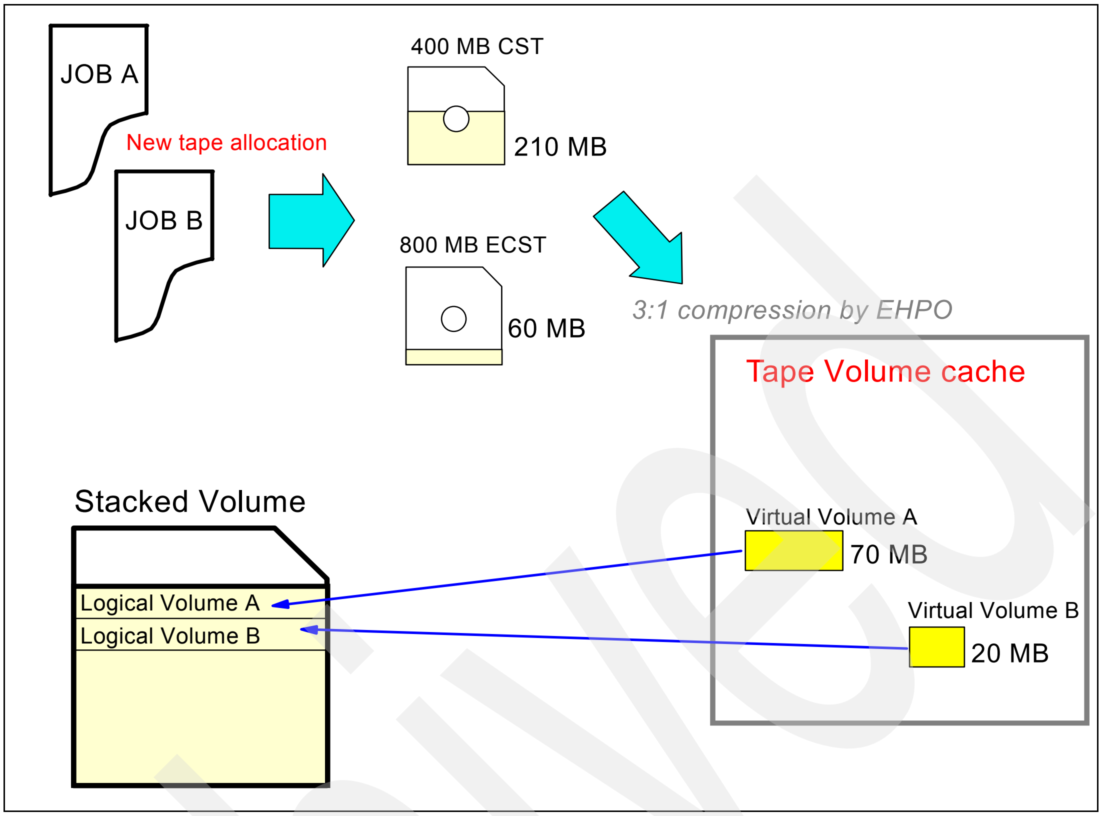

Figure 3-12 VTS space for CST and ECCST emulated cartridges

Sizing the number of logical volumes

The VTS B10 and B20 IBM TotalStorage provide compression for all data which is written to the TVC. Thus you will not see many changes regarding the number of logical volumes compared to the number of physical 3480 or 3490 cartridges that you are using.

You will need a larger number of logical volumes compared to the number of physical cartridges if you migrate from real volume sizes which contain a multiple of the 3480 or 3490 cartridges that you are using.

We recommend considering the following points before defining logical volumes to the VTS subsystem:

1. First insert five virtual volumes to make sure insert and the CBRUXENT exit work with the tape management system.
2. Then put in a few hundred virtual volumes for functional testing.
3. Set a reasonable scratch limit in ISMF for each library and media type and add enough virtual volumes (less than your target amount) so that when you go into production, you will soon hit your threshold. Track the CBR3660A message, which indicates that you are low on scratch, and validate your procedures for adding additional logical volumes (which eventually you may only rarely do).

   Refer to 4.3.6, “Physical cartridge insertion” on page 135 for a description of the insert process.

4. Here is a recommended formula to calculate the number of logical volumes needed:

Number of Virtual Volumes:

Vv = T + (Ax + Si)N

Where:

Vv = number of virtual volumes to be inserted.

T = the value used in the ISMF Library Define panel for the scratch threshold for the media type used (normally media2).

Ax = the average expiration period in days for the volumes destined for the VTS. This figure sets the expected duration of active volumes.

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

N = number of scratch volumes mounted each day by the applications using the VTS. This figure sets the expected frequency or requirements for new logical volumes.

Important: With the larger logical volume sizes, you have to repeat this calculation for each size you want to support by a data class definition.

- If you define more volumes than you need, you can delete the additional volumes. When you delete the volumes, however, you may leave gaps in the tape ranges. This will be difficult to manage because you have to keep track of which volumes you can insert later. Reinserting the volumes later has to be done on a volume-by-volume basis instead of in groups. See 6.7, "Ejecting logical volumes from the VTS" on page 248 for the procedures to delete logical volumes.
- Tie the volser to the volume's location, rather than to the data that resides on the volume. A volser of VT0001 tells you that the logical volume is VTS resident but does not tell you whether the data on the volume is test or production. A volser of TEST02 tells you that the data that resides on the tape is test data. It does not, however, tell you that the volume is a VTS volume and cannot be physically removed.

  Make sure that you can clearly distinguish between your logical volumes and the physical cartridges.

Tip: Take the time you need to decide which volsers and which media type to define to the IBM VTS.

The number of virtual volumes in the VTS inventory that have never been mounted is not reported in the SMF94 record. In general it is possible to use the tape management system to determine the number of virtual volumes in a given volser range that had never been mounted. As long as there are virtual volumes that have never been mounted, the demand for stacked cartridges will climb.

Sizing the number of stacked volumes

The number of stacked volumes that you need for your VTS subsystem depends on the number of logical volumes you define, the average amount of data you write on the volumes, and the average compression ratio achieved for the data. If Expired Volume Management is not implemented, logical volumes occupy space on the stacked cartridges even after they have been returned to scratch: Only when a logical volume is rewritten, the data on the stacked volume for this logical volume is eligible to be reclaimed.

The number of stacked volumes also depends on the reclaim threshold percentage that you have specified. The default is set to 10%, therefore, if a stacked volume contains less than 10% of active data, it is eligible for space reclamation.

You can calculate the number of stacked volumes required to support a given number of logical volumes as:

```
# of stacked volumes = (active data / stacked volume capacity) + # of empty stacked
volumes + minimum # of empty stacked volumes
```

Where:

- **Active data** = # of logical volumes * average logical volume size.

  This is the amount of uncompressed user data. Note that the number of logical volumes is the number you have defined to the VTS, not merely the number of private volumes.

- **Stacked volume capacity** = native cartridge capacity * compression ratio * stacked volume utilization.

  This is the amount of active data a stacked volume will hold, on average.

  The native cartridge capacity of a cartridge written by an IBM 3590 Model B1A tape drives on a standard length 3590 cartridge equals 10 Gbytes. The native cartridge capacity of a cartridge written by an IBM 3590 Model E1A tape drives on a standard length 3590 cartridge equals 20 Gbytes or double the capacity using the extended length media. The native cartridge capacity of a cartridge written by an 3592 tape drive on media type JA equals 300 Gbytes.

  Compression ratio is the average compression ratio for data you are moving into the VTS.

  If you do not know the compression ratio assume a conservative 2:1, unless the data is already in compressed format, in which case you assume 1:1. See 3.5, “Tape analysis and sizing the VTS” on page 72 for information on obtaining your compression ratio.

  Stacked volume utilization is calculated as:

  ((100% + reclaim threshold percentage)/100)/2

  We assume that the utilization of non-empty stacked volumes varies evenly between 100% and the reclaim threshold percentage specified for the VTS.

- **# of Empty stacked volumes** = active data * (buffer percentage/100) / stacked volume capacity

  The VTS needs empty stacked volumes to store new copies of logical volumes. It also needs empty stacked volumes to do reclamation and import/export of stacked volumes if the Advanced Function feature has been installed. Reclamation will eventually return low utilized stacked volumes back to scratch but you need to make sure the VTS has enough free storage to operate with. In addition, you should have additional free storage as a buffer against an overall increase in the amount of data in the VTS which may be due to an increase in the average logical volume size, or a decrease in the average compression ratio. We express the size of the free storage buffer as a percentage of active data and calculate the required number of empty stacked volumes based on it. The buffer should be at least 10% of your active data, preferably in the range of 15% to 25%.

- **Minimum # of empty stacked volumes**:

  At all times, you should have a minimum of 10 empty stacked volumes in the VTS. This minimum number of stacked 3592 volumes may need to be higher dependent on whether and how you are using Physical Volume Pooling. We recommend the following minimum numbers of empty stacked volumes as listed in Table 3-7.

Table 3-7 Minimum number of empty stacked volumes required

Physical Volume Pooling implementation	**Minimum # of Scratch stacked cartridges 3590 or 3592 JJ media**	**Minimum # of Scratch stacked cartridges 3592 JA media**
One physical pool or no pooling implemented	50	10
Physical Volume Pooling using Borrow/Return to dynamically acquire Scratch cartridges from the Common Scratch Pool (CSP)	50 + (2 per physical pool in use)	10 + (2 per physical pool in use)
Physical Volume Pooling using No borrow/Keep to explicitly assign Scratch stacked cartridges to a pool	5 per physical pool in use	5 per physical pool in use

A pool will hold onto the last two physical cartridges for 72 hours before returning them to the common scratch pool, so to prevent running the common scratch pool empty, a minimum of two per pool should be assumed in addition to the 50 or 10 minimum. If the volumes are explicitly assigned to a physical volume pool, we recommend five scratch cartridges as a minimum to account for a few volumes being in the filling state and a few needing to be reclaimed.

You should never let the number of empty stacked volumes fall below 10, as that will have an impact on VTS performance. Each time you add new logical volumes you must check that you have enough stacked volumes to support them. Do not simply add stacked volumes in the same ratio as you add logical volumes. Rather, take into account the characteristics (volume size and compression ratio) of the new workload you are moving into the VTS.

Here are some guidelines for the values of the formula mentioned above:

- If you do not have a chance to determine your current compression ratio based on LOGREC information, then use a value of 2.5 : 1, which we found as an average value in many tape environments.
- The reclaim threshold is 10% (the default), yielding a stacked volume utilization of 0.55.
- The free storage buffer is 20%; however, a minimum of 50 empty stacked volumes are included.

Example of stacked volume calculation for 3590 model H

Table 3-8 lists the parameters for the calculation example.

Table 3-8 Stacked volume calculation example parameters

Input Parameter	Value
Reclaim threshold percentage	10
Buffer percentage	20
Average logical volume size (MB)	250
Number of logical volumes	10,000
Expected compression ratio	3:1
Target backend devices	3590H
Target backend cartridge type	HPCT (“J”)
Number of pools	1
Pooling implementation	n/a

1. Determine the amount of active data:

   Average logical volume size * Number of logical volumes = Active data

   `For example: (250 * 10000) = 2,500,000 MB`

2. Determine the expected average utilization percentage:

   ((100 + Reclaim threshold percentage) / 100) /2

   `For example: ((100+10)/100)/2 = 0.55`

3. Determine the stacked (physical) volume utilization (MB):

   Target backend device on Target backend cartridge type = Native cartridge capacity

   `For example: 3590H1A written on HPCT = 30,000 MB`

   Native cartridge capacity * Expected compression ratio = Possible data on each cartridge

   `For example: 30,000 * 3:1 = 90,000 MB`

   Possible data on each cartridge * expected utilization percentage = stacked volume utilization

   `For example: 90,000 * 0.55 = 49,500 MB`

4. Determine the number of base cartridges:

   Active data / Possible data on each cartridge = base cartridges

   `For example: 2,500,000 / 49,500 = 50.505051`

5. Determine the number of scratch (empty) cartridges to be used:

   base cartridges * buffer percentage = scratch cartridges

   ```
   For example: 50.505051 * 0.20 = 10.10101
   ```

   Since this figure is below the minimum recommended scratch level of 50 cartridges, then the minimum is the value to be used.

6. Minimum required stacked volumes:

   base cartridges + scratch cartridges

   ```
   For example: 50.505051 + 50 = 101 (always round up)
   ```

This calculation can be used for the subsequent multiples of 10,000 logical volumes.

Important: With the larger logical volume sizes, you have to repeat this calculation for each number of logical volumes you have calculated before with the formula in “Sizing the number of logical volumes” on page 79 to get the final number of physical cartridges you have to order.

Example of stacked volume calculation for 3592 model J1A with Media JA

Table 3-9 lists the parameters for the calculation example.

Table 3-9 Stacked volume calculation example parameters

Input Parameter	Value
Reclaim threshold percentage	10
Buffer percentage	20
Average logical volume size (MB)	400
Number of logical volumes	100,000
Expected compression ratio	3:1
Target backend devices	3592-J1A
Target backend cartridge type	ECT (“JA”)
Number of pools	4
Pooling implementation	Borrow/return

1. Determine the amount of active data:

   Average logical volume size * Number of logical volumes = Active data

   ```
   For example: (400 * 100,000) = 40,000,000 MB
   ```

2. Determine the expected average utilization percentage:

   ((100 + Reclaim threshold percentage) / 100) /2

   ```
   For example: ((100+10)/100)/2 = 0.55
   ```

3. Determine the stacked (physical) volume utilization (MB):

   Target backend device on Target backend cartridge type = Native cartridge capacity

   `For example: 3592-J1A written on ECT = 300,000 MB`

   Native cartridge capacity * Expected compression ratio = Possible data on each cartridge

   `For example: 300,000 * 3:1 = 900,000 MB`

   Possible data on each cartridge * expected utilization percentage = stacked volume utilization

   `For example: 900,000 * 0.55 = 495,000 MB`

4. Determine the number of base cartridges:

   Active data / Possible data on each cartridge = base cartridges

   `For example: 40,000,000 / 495,000 = 80 (rounded)`

5. Determine the minimum number of empty stacked cartridges:

   10 + (2 per physical pool)

   `For example: 10 + 2 * 4 = 18`

6. Determine the number of scratch (empty) cartridges to be used:

   base cartridges * buffer percentage = scratch cartridges

   `For example: 80 * 0.20 = 16`

   Since this figure is below the minimum recommended scratch level of 18 cartridges, then the minimum is to be used.

7. Minimum required stacked volumes:

   base cartridges + scratch cartridges

`For example: 80 + 18 = 98 (always round up)`

Thus you may replace 100,000 3490 cartridges by 98 3592-J1A cartridges.

3.6.4 Selecting logical volume serial numbers

You need to give some thought to selecting the logical volume serial numbers that you use in the VTS. You will also want to keep in mind any local considerations for these ranges and any future growth.

If you are going to use (BLP-Bypass Label Processing) mounts outside the VTS, you can run into problems when using a volser range beginning with L (Lxxxxx). When you mount an NL or BLP scratch tape, an Lxxxxx volser is generated unless overridden by a tape management system or installation exit. If this Lxxxxx volser is later cataloged, it will conflict with any Lxxxxx volser that you have defined with the same xxxxx. You should therefore avoid using the Lxxxxx volsers for the logical volumes within the VTS.

Virtual volumes in the VTS are held in RAID-5 SSA disk arrays. This is the Tape Volume Cache (TVC). The VTS subsystem runs two file systems for storing data in and retrieving data from the TVC. One file system is for even volume serial numbers and the other for odd volume serial numbers. This can be seen in Figure 3-13.

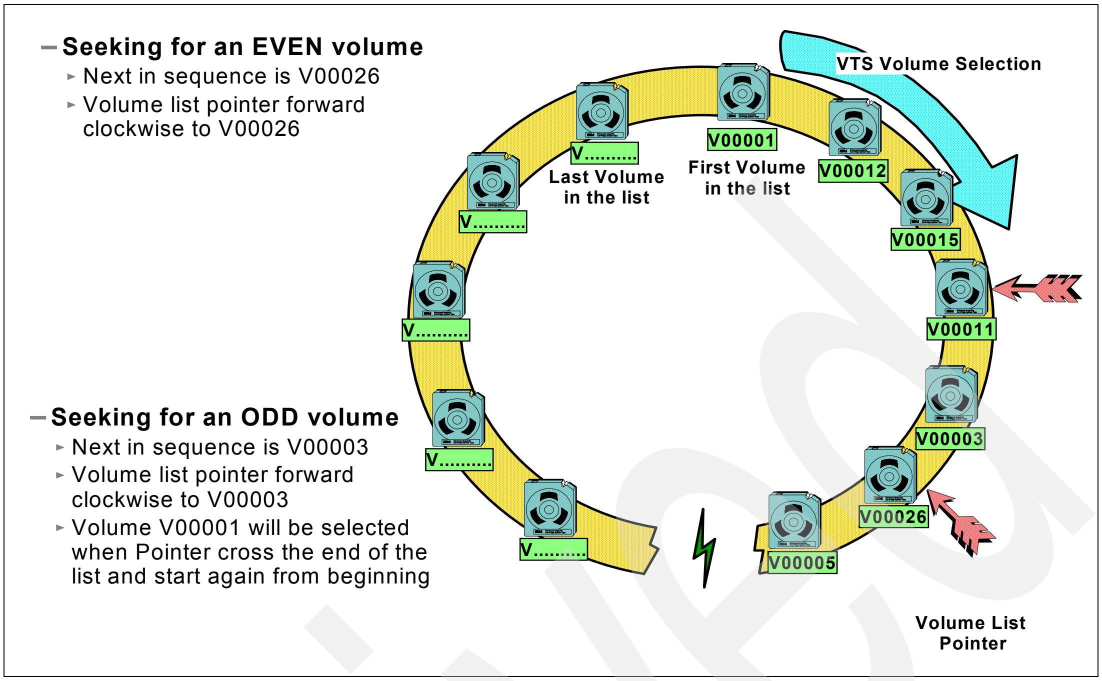

Figure 3-13 VTS scratch volume selection

Each of the two file systems maintains a pointer which designates the next volume (ordered by category order number) to be selected for a mount from category (therefore a non-specific mount). This assignment of odd or even volsers to a file system helps to balance workload and increase performance in conjunction with the method used by the Library Manager to select a volume from a category. At any given time when the Library Manager has the opportunity to make a volser selection, the number of odd volumes mounted is compared to the number of even volumes mounted and an odd or even volser is selected for approaching a balance.

Note: You may have noted that the VTS supports alphabetic characters for the logical volume serial numbers. All logical volume numbers are converted to ASCII when referenced in the VTS. Whether or not the logical volume's number is even or odd is then determined by whether or not bit 7 is on.

The Library Manager maintains an ordering by *time* for each volume entered into a category. When a logical volume is moved to a category, it is placed at the bottom of the ordered list, using the highest category order number. When a volume is requested from a category, the next in order, odd or even, volume will be selected based upon the balance of odd and even volumes mounted in the VTS.

As an example, assume that we have a VTS with the volsers listed in Table 3-10. Each volser has a category order number that the Library Manager uses to keep track of which logical volumes to mount next for a nonspecific request.

Table 3-10 Scratch tape order example

Volser	Category	Category order
V00012	0002	010
V00015	0002	012
V00011	0002	782
V00003	0002	783
V00026	0002	785

Now assume that we start out with the Library Manager pointer specifying volser V00011 as the next odd and V00026 as the next even to be selected.

When the first request comes in, the VTS selects an odd volser due to activity. This would result in volser V00011 being mounted in response to the scratch request.

Now another request comes in for a scratch mount, and the VTS chooses even. This would result in volser V00026 being mounted in response to the scratch request.

Now for our next scratch request, the VTS chooses odd. This results in volser V00003 being chosen. This logical volume is the next highest one in the category order that is odd.

For our last scratch request, the VTS chooses even. This results in volser V00012 being chosen assuming any subsequent volumes on the picture are odd. This logical volume is the next highest one in the category order that is even. The category order “wrapped” around to the low number because there were no more even logical volumes with a category order number higher than 785 (V00026).

Let's summarize some of the important factors here:

- Table 3-10 represents one scratch category, but there may be many such categories in a single library.
- Volumes are arranged in sequence of category order number, not by volser.
- The two pointers (even/odd) to the next volume move in ascending category order sequence.
- Volumes returned to scratch (reassigned to this category) are given a new, higher category order number.
- Volumes inserted into the library (and assigned to this category) are given a new, higher category order number.
- Category order numbers are virtually infinite.
- If volumes are assigned to a category, by insertion *or* return-to-scratch, faster than they are selected for use, the pointer will not start over at the beginning of the category order number list unless the Library Manager is restarted.
- Category order numbers are never reused. A volume returned to this category will be given a new (higher) number.

3.6.5 Extended Length Media Support

Support for the 3590 Extended High Performance Tape Cartridge stacked volumes (3590’K’ cartridges) in the VTS was introduced with VTS code level 2.23.xx.x and the Library Manager code 526.

This support allows for both the standard High Performance Tape Cartridge and the Extended High Performance Tape Cartridge to co-exist in the VTS. There is no external influence to force selection in a mixed environment, the cartridges are loaded without regard to the media type; the VTS will be aware of which media type it has loaded.

All 3590 drives must be Extended Length Media capable. The rear panel of the 3590 will have a 2X label near the model number sticker. If the drives need to be upgraded, use one of the following feature codes:

- 5780 EXTENDED MEDIA SUPPORT -(3590 B1A)
- 5781 EXTENDED MEDIA SUPPORT -(3590 E1A)

Note: If a 3494 Model L14 or D14 Frame contains one or more 3590 Model H1A drives, either a 3590 Model A60 controller must also be installed in that frame, or the drives must be attached to an A60 controller in an adjacent Model D14 frame.

The 3494-Dxx frame and the 3494-Lxx should also have FC9780 to identify that the frames are extended length media aware.

3.6.6 System managed tape considerations

In this section we discuss the most important factors to consider when using a IBM VTS in a system managed tape environment.

Integrated Cartridge Loader simulation

The IBM VTS appears to the host to be an IBM 3490E subsystem with an Integrated Cartridge Loader (ICL) installed but inactive. A cartridge loader's status is reported as ACTIVE in a Tape library if a media type has been assigned to the device and there are scratch volumes of that media type available in the library. You can use the z/OS `LIBRARY SETCL` command to load the emulated ICL with scratch cartridges of a specified media type.

In selecting a drive for a scratch mount, MVS allocation prefers ones with an active cartridge loader containing the media that matches the scratch request. If you wish to prefer one set of drives over another in a VTS, the SETCL command can be used for that purpose. However, since there is no real advantage in using one virtual drive over another, we recommend that you do not do so. There is no scratch mount performance advantage to assigning a media type to a virtual drive, so the normal recommendation is to assign the drives to NONE. This is because, when you use the fast-ready attribute for that scratch category, the new scratch volume is created directly in the TVC without staging the data from tape.

Define more than one VTS per storage group

When selecting a virtual tape drive for non-specific mounts, the host operating system doesn't have information about the following critical VTS resources and configurations:

- Number of available drives (virtual drives as well as physical drives)
- Number of available scratch cartridges
- Busy condition of VTS control unit or tape library accessor (such as reclamation activity)
- Size and degree of occupation of the TVC

“Balanced” allocation in this scenario is limited. Our recommendation is to define one storage group per VTS and decide how to balance the main workloads between the VTSs, based on the knowledge of the applications and criteria developed during the analysis phase. A second advantage of having each VTS in its own storage group is the operational flexibility: You can ensure that two copies of a data set are sent to a different VTS. If the VTSs are in one storage group, there is no way to guarantee that the allocation will hit a different VTS.

Support for cartridges with no external label

Although cartridges without an external label are allowed in a IBM 3494 Tape Library, their use is not recommended. Every tape library inventory operation would cause unlabeled cartridges to be ejected.

When the VTS is attached to an IBM 3584 Tape Library, unlabeled cartridges are not supported for the VTS.

Standard label

The IBM VTS always creates an IBM standard label for every new virtual volume regardless of the JCL specification. It works the same as for a real 3494. Non-standard and unlabeled virtual volumes are supported using bypass label processing.

TCDB size considerations

The TCDB size must be checked and possibly increased to fit the possible 500,000 virtual volumes one VTS is capable of holding. Refer to 4.5, “VTS software definition for zSeries” on page 153 for more information related to allocating the TCDB.

3.6.7 Basic Tape Library Support (BTLS) considerations

The considerations applied to system-managed tape are also valid when using a IBM VTS in a BTLS environment.

The emulated integrated cartridge loaders can be used, but they do not give any performance benefit in the same way as in an system-managed tape environment. BTLS deals with NL and NSL tapes as system-managed tape.

You should increase the size of the BTLS catalog to fit the additional 500,000 volumes that VTS can manage. To calculate the correct size of the BTLS catalog, refer to *Basic Tape Library Support User's Guide and Reference*, SC26-7016.

3.6.8 Tape Mount Management (TMM) considerations

TMM uses DFSMS to redirect selected data sets onto DASD as they are allocated and DFSMShsm to move those data sets back onto tape.

TMM enables installations to stack multiple compacted tape data sets automatically onto a tape, in most cases without having to identify and modify any JCL procedures. If you are already using TMM, you will have to decide which data will go to VTS and which data will stay with TMM.

After introducing the IBM VTS to your environment, you must review your actual TMM environment and make changes as required, to better fit VTS into your tape environment.

During the data collection for a Batch Magic analysis, TMM data should be identified to be included in the tape analysis.

3.7 Physical planning and placement

In this section we present information that your planning department needs to ensure that the site planned for the VTS is suitable. Refer to the *IBM 3494 Tape Library Dataserver Introduction and Planning Guide*, GA32-0279, for additional information.

3.7.1 Operational requirements

The physical dimensions of the VTS frame, which is located outside the 3494, differ from the dimensions of the other IBM 3494 frames. Table 3-11 lists the dimensions without service clearance.

Table 3-11 IBM 3494 Tape Library Frame dimensions in millimeters

Frame	Height	Width	Depth
Model B10 or B20 VTS	1800	724	1036
L1x Frame/D1x Frame	1800	725	1520
L22, D22, L52, D52 Frames	1800	725	1212

In addition to the dimensions listed in Table 3-11, a service clearance of 1016 mm is required in front of and behind the VTS. On both the left-hand and right-hand sides of the 3494, a service clearance of 762 mm is recommended.

Table 3-12 provides the VTS specifications for weight, heat output, airflow, power requirements, and cartridge slots.

Table 3-12 IBM TotalStorage VTS specifications

Frame	Weight (kg)	Heat output (kW)	Max. power (KVA)
B10	540	1.82	2
B20	757	3.64	4

Table 3-13 provides the specifications for weight of the 3584 Tape Library frames with and without drives and cartridges.

Table 3-13 IBM 3584 Frame weight

	Weight of base frame with one drive and 0 cartridges	Weight of base frame with 12 drives and max. cartridges
D52	274 kg	483 kg
L22	364 kg	534 kg
D22	270 kg	494 kg

For further information about the physical installation of VTSs, refer to the *IBM 3494 Tape Dataserver Library Introduction and Planning Guide*, GA32-0279 and to the *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Introduction and Planning Guide*, GA32-0472 as well as *IBM TotalStorage 3584 Tape Libraries Guide for Open Systems*, SG24-5946

3.7.2 Physical planning considerations

While physical planning is a customer responsibility, we recommend that you consider the following factors when you decide where to put your new VTS:

- Do you plan to add additional frames to the IBM 3494 or 3584 in the near future? Do you, for example, plan to install the High Availability option some time later? If so, allow enough space for the new frames when the IBM VTS is installed. A few extra hours spent planning up front can save you days relocating equipment later.

- Do you plan to relocate the computer room to another room, building, or site in the near future? If so, how much time does it take to relocate a VTS and its tape library? If the VTS is in production use before the relocation, how will the relocation affect your production schedule?
- Have you allowed for service clearances when determining where you will put the VTS? For additional information about physical specifications and service clearances, consult:
  – *IBM TotalStorage 3494 Tape Library Introduction and Planning Guide*, GA32-0279
  – *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280
  – *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Introduction and Planning Guide*, GA32-0472
  – *IBM TotalStorage 3584 Tape Libraries Guide for Open Systems*, SG24-5946
  – *IBM TotalStorage Enterprise Automated Tape Library (3494) Systems Assurance Product Review (SAPR) Guide*, SA01-005, which is available from your IBM representative or can be found at:

    `http://w3-1.ibm.com/support/assure/assur30i.nsf/PubAllNum/SA185?OpenDocument`
- Will your 3494 or 3584 be located in a remote area without operating staff, for example, in a disaster recovery vault? If it will be, you might consider installing the Remote Library Manager Console feature, which allows you to see intervention required messages or perform such actions as defining additional logical volumes without having to be at the site where the 3494 is installed. For additional information about the Remote Library Manager Console, refer to *IBM 3494 Tape Library Dataserver Introduction and Planning Guide*, GA32-0279 and IBM TotalStorage 3953 Tape Frame Model F05 and *Library Manager Model L05 Introduction and Planning Guide*, GA32-0472.

  In addition to the Remote Library Manager Console, you might consider implementing Simple Network Management Protocol (SNMP) traps for monitoring IBM 3494 events. For more information about SNMP, see 6.3.3, "Simple network management protocol (SNMP) traps" on page 219.

The considerations listed above are intended to help you think about your physical environment and the placement of the IBM VTS in that environment. It is easier to move equipment on paper than in real life when the equipment is fully loaded.

3.7.3 Power considerations

The stand-alone VTS frame needs its own power connection. VTSs are shipped with two primary control compartments and two power cards allow connection to two independent power sources. Other than that, the VTS imposes no additional environmental requirements over those of the 3494 and IBM 3590 or 3592 products.

3.7.4 IBM TotalStorage Master Console

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 IBM TotalStorage Master Console for transmission to IBM for review for preventive action. The IBM TotalStorage Master Console 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 IBM TotalStorage Master Console feature (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 System Service Representative (SSR) and is a required feature on the Model B10 and B20 VTS. It is an optional feature for Model B18 VTS. Figure 3-14 describes IBM TotalStorage Master Console connections.

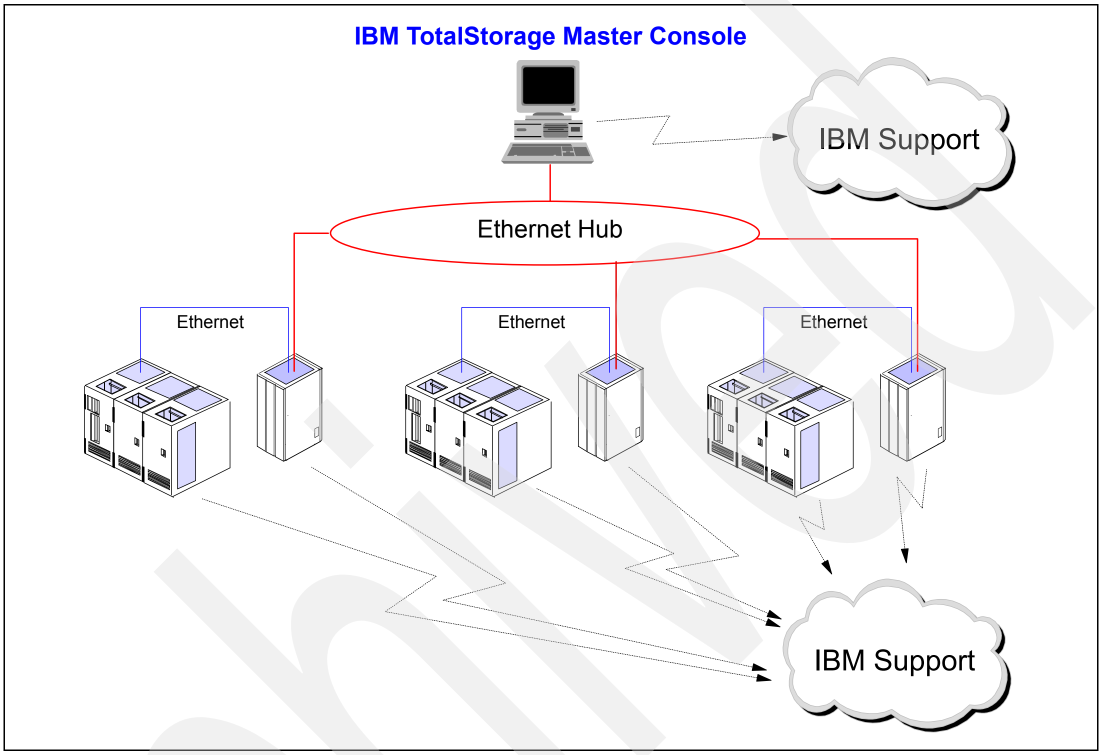

Figure 3-14 IBM TotalStorage Master Console connectivity

3.7.5 VTS and Library Manager microcode levels

We summarize the microcode levels for VTS and Library Manager and their major improvements in Table 3-14.

Table 3-14 VTS & 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	3584/3953 support, APM enhancements, VTS upgrade options

3.8 Software requirements

The VTS appears to the host operating system as two, four, eight, or sixteen fully configured IBM 3490E tape subsystems. The software levels required to support the VTS are the same as those for IBM 3490E tape subsystems, because the VTS appears to the host as an IBM 3490E tape subsystem in an 3494 or 3584 Tape Library.

3.8.1 Software requirements for ES/3090, ES/9000, S/390, or zSeries

Here we discuss the software requirements for ES/3090™, ES/9000®, S/390®, or zSeries. The IBM TotalStorage VTS supports z/OS, OS/390®, z/VM®, VM/ESA®, VSE/ESA™, and TPF.

z/OS support

Support is provided with z/OS V1R4 and above, which also includes the recommended support for the 3584 Tape Library with the 3953 Library Manager. For more information on the 3584 support provided, refer to APAR OA09751, OA10566, OA09753, and OA10566. EREP Version 3.5 plus PTFs is required as well.

z/OS software support for the Advanced Policy Management feature is provided as a small Programming Enhancement (SPE) for DFSMS at OS/390 V2R10 and z/OS V1R1 and above (PTFs are required) with base support available at z/OS V1R5.

Previous levels of z/OS and OS/390 already supported the VTS, APM, FICON attachment, and FICON cascading. Those software levels, however, are no longer supported.

z/VM and VM/ESA support

The VTS is supported with z/VM Version 3 Release 1 or later releases. APAR VM62710 required for native z/VM support with FICON.

VM/ESA Version 2 Release 3 or later releases already supported the VTS with APAR VM62710 required for native VM/ESA support with FICON.

EREP Version 3.5 plus PTFs is required as well.

Support for the 3584 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).

Note: No support is available for the VTS Advanced Function Feature or the Advanced Policy Management Features. Support for the IBM TotalStorage Peer-to-Peer Virtual Tape Server is available in an environment in which a zSeries host with full support of the PtP VTS is also attached to the Tape Library.

Transaction Processing Facility (TPF) support

TPF Version 4.1 with PTFs supports the VTS.

VSE/ESA support

Native support for the VTS is available with z/VSE™ V3.1.

As a guest under VM/ESA or z/VM, the VTS is supported with VSE/ESA Version 2 Release 3 or later plus Program Temporary Fixes (PTFs), DFSMS/VM Function Level 221 with PTFs, and EREP Version 3.5 with PTFs.

For FICON channel attachment of the VTS, either z/VSE V3.1 or later release or VSE/ESA Version 2 Release 5 or later are required together with EREP Version 3.5 plus PTFs.

Support for the 3584 Tape Library with the 3953 Library Manager is provided at VSE/ESA 2.6 and above.

z/VSE 3.1 supports all 3592 J1A configurations without APARs in all automation offerings, including 3494 and 3584.

Additional software

Functional exploitation of the VTS Advanced Functions feature including Import/Export is currently provided with z/OS V1R4. Additional software support for Import/Export is provided with:

- **DITTO/ESA:** DITTO/ESA provides a VTS Exported Volume Read Utility to:
  - Extract a logical volume from an Exported Stacked Volume and create a physical volume containing only the single logical volume.
  - Provide a list of logical volumes contained on an Exported Stacked Volume.

  This support is available through APAR PQ26582.
- **Tape Management System support:** Using DFSMSrmm (or independent software vendor tape management system), you can initiate the movement of logical volumes to stacked volumes or the copy of logical volumes from stacked volumes into a virtual tape server. Contact the independent software vendors for the tape management system regarding their support plans for support of the import/export capability.

3.8.2 Software requirements for SCSI host systems

IBM TotalStorage VTS supports pSeries, Sun, HP, and Windows® systems with required software levels and adapters as documented in Table 3-15. To attach VTS to these servers, the required number of FC3422 should also be ordered.

Table 3-15 IBM TotalStorage VTS SCSI host software and adapters support

Server	Host adapters	Host software	VTS features
RS/6000®, pSeries	2412 6204 6209 6207	AIX 4.3.2	9106 and 9201
Sun	X1062A X1065A X6541A	Solaris™ 2.6, 7 or 8	9211 and 9201
HP	A4800A	HP-UX 11.0	9210 and 9201
Windows NT®	AHA-2944UW	Windows NT Server 4.0 with SP6	9212 and 9201
Windows 2000	AHA-2944UW	Windows 2000 Build 2195	9212 and 9201

Server	Host adapters	Host software	VTS features
Notes: 2412 MCA Enhanced SCSI-2 Differential Fast/Wide Adapter/A (4-C) 6204 PCI Universal Ultra SCSI Adapter 6209 PCI SCSI-2 Fast/Wide Differential Adapter (4-B) 6207 PCI Differential Ultra SCSI Adapter (4-L) X1062A Sun SBus Differential Fast/Wide Intelligent SCSI-2 Host Adapter X1065A Sun SBus Ultra Differential Fast/Wide Intelligent SCSI-2 Host Adapter X6541A Sun PCI Dual-Channel Differential Ultra SCSI Host Adapter A4800A HP Fast/Wide SCSI-2 Adapter AHA-2944UW Adaptec PCI Differential Ultra SCSI Adapter			

RS/6000 and pSeries SCSI Host attachment

The VTS Open Systems Device Drivers (FC9201) provides device driver support for pSeries host systems attached to SCSI buses in the VTS. The Attach to RS/6000, pSeries (FC9106) should also be specified to indicate attachment to an pSeries host system. Application programs supporting library attached 3490E devices (3490 Model C1A, C2A or F1A) will operate without change with a VTS. APM and PtP are also supported.

Sun system attachment

The VTS Open Systems Device Drivers (FC9201) provides device driver support for a Sun server attached to SCSI buses in the VTS. The Sun Attachment feature (FC9211) should also be specified to indicate attachment to an Sun host system. Application programs supporting library attached 3490E devices (3490 Model C1A, C2A or F1A) will operate without change with a VTS. APM and PtP are also supported.

HP Server attachment

The VTS Open Systems Device Drivers (FC9201) provides device driver support for an HP server attached to SCSI buses in the VTS. The HP-UX Attachment feature (FC9210) should also be specified to indicate attachment to an HP host system. VTS is supported HP 9000 L-class, N-class or V-class servers running HP-UX 11.0. Attachment to V-class server requires the In line SCSI Terminator (FC9798). Application programs supporting library attached 3490E devices (3490 Model C1A, C2A or F1A) will operate without change with a VTS. APM and PtP are also supported.

Microsoft Windows attachment

The VTS Open Systems Device Drivers (FC9201) provides device driver support for an Intel®-compatible server attached to SCSI buses in the VTS. For Microsoft® attachment, the Attached to MS Windows feature (FC9212) should also be specified to indicate attachment to MS Windows 2000 or MS Windows NT. Application programs supporting library attached 3490E devices (3490 Model C1A, C2A or F1A) will operate without change with a VTS. APM and PtP are also supported.

3.9 Planning for Advanced Policy Management exploitation

Preparing the exploitation of the multiple functions offered by APM needs careful planning. Investigate first, which of the functions are of high importance for your tape production and avoid to interact more than required the self-managing VTS solution.

3.9.1 Prerequisites

APM is implemented in the VTS LIC and exploited in z/OS host software. Non-MVS hosts can use the functions by only using the definitions which are done at the Library Manager level. For details see 4.4, “Implementing Advanced Policy Management” on page 142.

Software support

There are two modes within which an APM enabled VTS can run: Exploitation or Toleration.

Toleration support is provided by:

- PTFs for DFSMS v1.5 or higher (APAR OW54056) but this release with FMID HDZ11E0 is not supported anymore.

Exploitation support is provided by:

- SPE for OS/390 V2R10 (APAR OW54054)
- SPE for z/OS V1R1 or higher (APAR OW54054)
- SPE for z/OS.e V1R1 or higher (APAR OW54054)
- z/OS V1R5 or later (include in the base code)
- z/OS.e V1R5 or later (include in the base code)

Hardware prerequisites

APM is supported since Library Manager LIC 527 and VTS LIC 2.26. To upgrade, the following code levels and configurations are required:

Minimum hardware requirements:

- B10/B18/B20/PTP VTS
- Four physical tape drives
- 216 GB cache
- For a B18 FC5236, Performance Accelerator
- VTS LIC level 2.23.29.0 or later
- LM LIC level 526.07 or later

> **Note:** The Library Manager LIC can be installed first while still running with a prior level of VTS LIC.
>
> Once you have upgraded to the new level 2.26, regardless of whether of not you have you have installed FC4001, IBM strongly recommends that you do not back down to a previous level of VTS LIC.

> **Attention:** If you have upgraded to the VTS LIC that supports APM and have written new data, you cannot revert to a previous level of LIC. This is due to the changes that have occurred to the database, which form part of the APM functionality.

The upgrading to the new LIC levels may have implications for customers who do disaster recovery tests with vendors or import/export volumes between multiple VTS’s. These are some important considerations:

- All locations in a disaster recovery test will need to ensure that their machines are all running at the VTS level of 2.26 and LM level 527 to ensure compatibility between all system environments and the VTSs running APM.
- VTS’s running level 2.22.xx.x, or later, will be able to import volumes exported from a VTS running LIC level 2.26.

- VTS's running level 2.26 will be able to import volumes exported from a VTS running LIC level 2.26. If there is the requirement to scratch and override data of logical volumes after they are expired, you have to group those logical volumes to prevent excessive reclamation activity.

You can check whether your VTS is APM capable through the ETL Specialist. Figure 3-15 shows the ETL Specialist panel, which is displayed when you select **VTS Status** from the **Monitor VTS** work items.

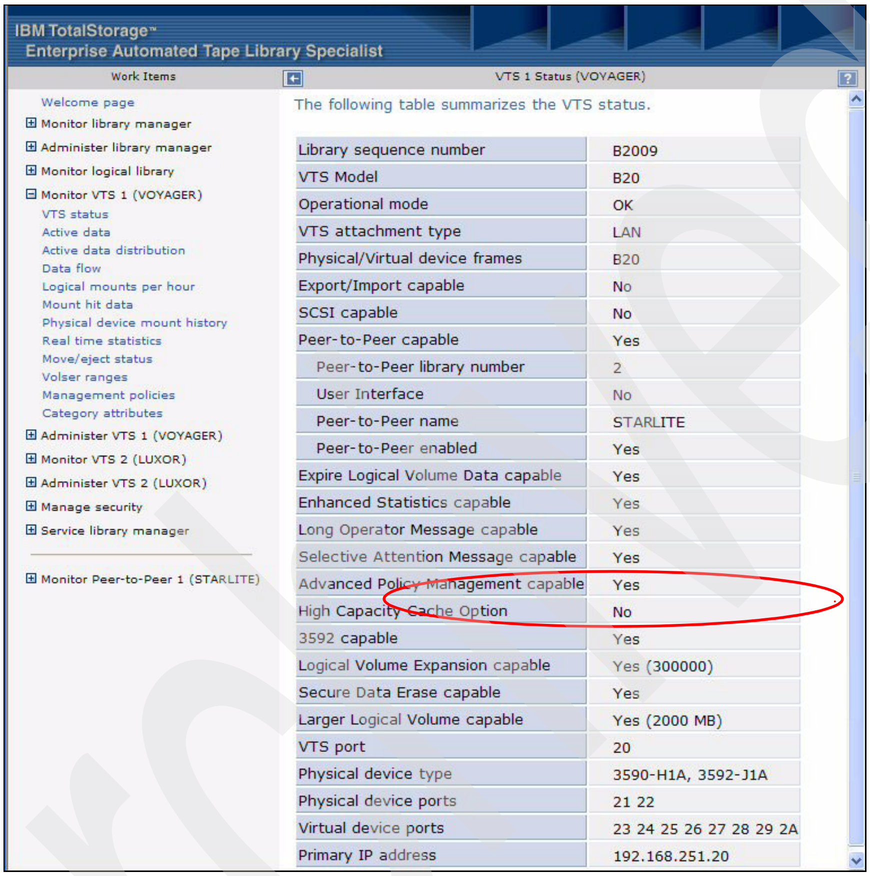

Figure 3-15 VTS Status panel

3.9.2 Cache Management exploitation

Prior to this feature, there was no way to influence cache residency. As a result, all data written to the TVC was premigrated using a FIFO (First-In-First-Out) method. The mount response time performance for specific volume requests in a VTS subsystem vary greatly depending upon residency of the requested logical volume in the Tape Volume Cache or the need for the logical volume to be recalled from a stacked volume.

You may want the VTS to manage the residency of logical volumes in the Tape Volume Cache to maximize the probability of the needed volumes being resident in the Tape Volume Cache. At a minimum, you want to have some influence over which logical volumes reside in the Tape Volume Cache and for how long.

3.9.3 APM pooling considerations

Pooling 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 3-16 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 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.

Changes for copying data to tape

With APM enabled, the VTS will attempt to chain logical volumes together that are part of the same pool. These chains are based upon LRU Age and Cache Preference Levels. These chains could also contain primary and backup copies of logical volumes. The VTS will continue to manage the chains and determine the most appropriate size depending on a number of factors. These can include:

- Export logical volumes which are eligible to premigrate.
- Old resident logical volumes which are eligible to premigrate.
- The VTS is idle.
- The VTS is busy and must premigrate data to free space within the TVC.

Figure 3-16 shows the process for the post 527 and 2.26 LIC levels with Advanced Policy Management enabled and volume pooling being used to direct logical volumes to different stacked volume pools.

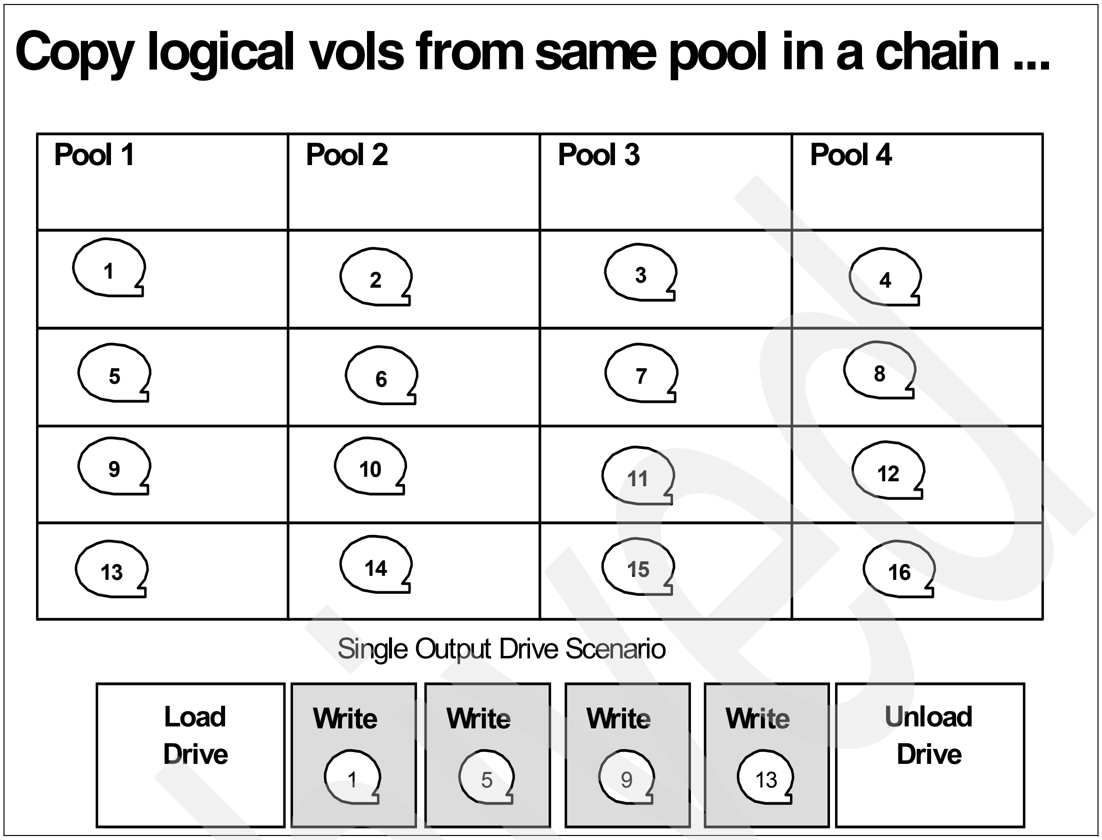

Figure 3-16 Copy process with APM enabled and volume pooling in use

3.9.4 Selective Dual Copy

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

If you activate Dual Copy for a group of data or a specific pool, consider that all tasks and properties 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

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

3.9.5 Larger logical volumes

With the exploitation of larger logical volume sizes less logical volumes are needed for large datasets which exceed the 800 MB barrier of today. If you need five logical volumes each with 800 MB today for a specific large dataset you can store this onto one logical volume with 4000 MB now. In case of a recall, there is a need of only one recall action and one physical tape mount instead of five recalls (and in worst case five physical mounts) which interrupts the processing of the job. If there is no spanning of this large volume. The probability of more spanning on the physical volume is increased as you select larger volume sizes, as follows:

- 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 datasets that span five or more logical volumes.

Furthermore, the cache premigration thresholds needed to be changed because more space is needed for 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 premigration 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 impacts.

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 dataset on this volume for “Write” then the size will remain as it was.

3.9.6 Secure Data Erase

As soon as you have decided to implement Secure Data Erase for a *limited group of data separated on a dedicated pool*, the number of additional reclamation tasks plus data erase tasks will increase. Less physical drives 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 trigger the physical volume to be erased, an almost full physical cartridge will be first reclaimed and second erased.

We highly recommend to group logical volumes that require secure erase after they are expired, in such a way that almost no unnecessary reclamation and subsequently erasure operations take place. Pooling by expiration date may help to reduce unnecessary reclamation. Although proper grouping reduces the amount of reclamation that needs to be done, it will not eliminate the erasure step.

3.9.7 Considerations for cache management

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 you want to be candidates for expedited removal from cache.

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 Batch Magic to allow a simulation showing the improvements in read hits by using cache management.

3.10 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.

Distances supported are shown in Table 3-16.

Table 3-16 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

3.10.1 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, standalone 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 & Inrange), CISCO and Brocade) but you may mix models of one vendor. See the switch Web pages for specific intermix combinations supported.

3.10.2 FICON transfer modes

Two types of FICON channel extenders are available working with either:

- 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 standalone and Peer-to-Peer VTS for USD-X and the Edge Router from CNT (now McData) and for CN2000 from CIENA.

3.10.3 Factors that affect distance

Fibre Channel distances depend on many factors, which include:

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

Native shortwave FC transmitters have a maximum distance of 500m with 50micron 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 300m.Native longwave FC transmitters have a maximum distance of 10km 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 100km. 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 relationship is very vendor specific. There should be a match between the buffer credit capability of the nodes at either end of the extenders. A host bus adapter (HBA) with a buffer credit of 64 communicating with a switch port with only eight buffer credits would be able to read at full performance over 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 2km in order to maintain full performance.

Buffer credits within the switches and directors have a large part to play in the distance equation. The buffer credits in the sending and receiving nodes heavily influences the throughput that is attained within the Fibre Channel. Fibre Channel architecture is based on a flow control that ensure 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 *Introduction to SAN Distance Solutions*, SG24-6408.

3.11 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 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?

3.11.1 Adding a VTS to an existing 3494 or 3584

If you are adding a VTS to an existing 3494 containing a VTS, 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:

- 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:

- Software choices and how they affect the IBM VTS
- Disaster recovery considerations

3.11.2 New VTS and new 3494 Tape Library or new 3584 Tape Library with 3953 LM

If you are putting in a new 3494 or 3584/3953 with a VTS, the education and training for your staff should include the VTS training items listed in 3.11.1, “Adding a VTS to an existing 3494 or 3584” on page 104, 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*, GA32-0279
- *IBM TotalStorage 3494 Tape Library Operator’s Guide*, GA32-0280
- *IBM TotalStorage Tape Library Guide for Open Systems*, SG24-5946
- Appendix C, “VTS implementation step-by-step” on page 419
- *IBM TotalStorage Enterprise Automated Tape Library (3494) Systems Assurance Product Review (SAPR) Guide*, for additional information about education and training. (See “Other resources” on page 519 for bibliographic information for this publication.)

You also may contact your IBM marketing representative for further information about education and training for the 3494 and VTS.

3.11.3 Migration of an existing VTS

Model B18 VTS and Model B10 VTS can be upgraded to Model B20 VTS

For details on these changes, refer to Chapter 8, “Upgrade scenarios” on page 329. We recommend that you also review Chapter 6, “Operating the VTS” on page 211, as well as 7.4, “Monitoring and evaluating VTS performance” on page 290, for additional information about operational changes.

3.11.4 Sharing the VTS with other systems

In addition to the material in Section 3.4.2, “Sharing a VTS subsystem” on page 72, refer to the *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409.

3.11.5 Implementation services

In this section we describe the services available.

IBM Global 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 lifecycle, including these:

- **Assessment:**

  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:**

  Planning assists you 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 (VTS) 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 (VTS) 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:**

  Support-Line 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 world-wide e-business leader. IBM Global Services offers comprehensive, integrated industry solutions that will deliver ongoing value to help keep you competitive in today's globalized 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.

For availability in your country, check with your IBM System Service Representative (SSR).

4

Implementation

In this chapter we describe how to implement the VTS. Because the VTS works with the 3494's Library Manager, it must be installed together with a 3494. When we speak about the VTS installed in a 3584, then you also need a 3953 Library Manager for each VTS.

You can install a VTS together with your existing 3494 or 3584, or you can install a new 3494 or 3584/3953 and a VTS at the same time. Either way, you must have a 3494 or a 3584 if you want to take advantage of the unique performance and functions that a VTS offers.

VTS software implementation:

We only discuss the specific issues for the VTS in this chapter.

If you are new to the IBM 3494 or IBM 3584 Tape Library, we strongly recommend for you to consult the following publications for the tape library implementation guidelines that apply to VTS, as well as the non-VTS libraries that are not included here:

- *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632
- *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning, Implementation, and Monitoring,* SG24-6789.

4.1 VTS implementation considerations

The VTS implementation can logically be described under the following areas.

HARDWARE I/O configuration definition: This area relates to the system generation, and consists of such processes as ESCON channel attachment to a host, HCD/IOCP definition, and power preparation. Because the host operating system sees the VTS as a 3494, the HCD and IOCP definition for VTS virtual drives are essentially the same as the HCD and IOCP definition for drives resident in a 3494 without VTS.

VTS definition from the Library Manager: Here you define volumes, policies and categories using Library Manager operations. With the addition of FC4001 - FC4004 (Advanced Policy Management or APM), you also define data class, management class, storage class, and storage group construct names and policies. These Library Manager operations are performed by using the operator panels for the VTS definitions. If you already have an existing 3494, you benefit from the transparent operational capability of the Library Manager console.

VTS software definition for zSeries: Here you define the new tape library to the individual operating system. The definition process is the same for VTS virtual drives as it is for 3494 resident drives. If you are using z/OS DFSMS and SMS tape, you update DFSMS ACS (Automatic Class Selection) routines, Object Access Method (OAM), and your tape management system during this phase. With the addition of FC4001 - FC4004 (Advanced Policy Management (APM)) you also define data class, management class, storage class, and storage group construct names, and selection policies, which are in turn passed to the VTS for actioning. Installation of VTS in Open System environment is documented in 4.8, “VTS Open System attachment” on page 167.

The tasks we outline in this section are meant to give you an overall concept of the tasks required for VTS installation. The task lists we provide are not all-inclusive, and they are not intended to replace any current documentation for the VTS. We provide a simple checklist for the recommended process.

4.1.1 Background and supplementary materials

Whether you are installing a VTS in an existing or new 3494, review the following items:

- 3494 specific hardware items:
  - Review the *IBM TotalStorage Enterprise Automated Tape Library (3494) Systems Assurance Product Review (SAPR) Guide*, including checklists. (See “Other resources” on page 519 for bibliographic information for this publication.)
  - Review the *IBM 3494 Tape Library Dataserver Introduction and Planning Guide*, GA32-0279.

    The latest versions of this document can be found at:

    http://www.storage.ibm.com/hardsoft/tape/pubs/prodpubs.html
  - Complete systems assurance for the 3494.
- 3494 specific software items:
  - Review software support based on platform and installed features:
    - z/OS (see 3.8.1, “Software requirements for ES/3090, ES/9000, S/390, or zSeries” on page 93
    - VM/ESA (see 3.8.1, “Software requirements for ES/3090, ES/9000, S/390, or zSeries” on page 93)

  - VSE/ESA (see 3.8.1, “Software requirements for ES/3090, ES/9000, S/390, or zSeries” on page 93)
- Review software support based on storage management products:
  - z/OS DFSMS and OAM (see the *Object Access Method Planning, Installation and Storage Guide for Tape Libraries*, SC26-0427, and *Implementing System Managed Storage*, SC26-3123).
  - BTLS (see the *Basic Tape Library Support User's Guide and Reference*, SC26-7016 and Chapter 3.6.7, “Basic Tape Library Support (BTLS) considerations” on page 89).
  - DFSMS/VM RMS (see the *VM/ESA DFSMS/VM Removable Media Services User's Guide and Reference* and 4.6, “VM/VSE considerations” on page 160).
- Review software support based on the tape management system and based on the VTS functions you are planning to use, such as the Advanced Function feature Import/Export component. For DFSMSrmm, see 5.2, “DFSMSrmm and other tape management systems” on page 186. For other tape management systems, contact the independent software vendor for your product.
- Review *IBM Magstar 3494 Tape Library and 3490E PSP bucket* for applicable maintenance (available on IBMLINK, or you can contact your IBM representative, who can assist you in obtaining PSP buckets).
  - Include PSP buckets for 3490E.

- VTS-specific hardware items:
  - Review the *IBM TotalStorage Enterprise Automated Tape Library (3494) Systems Assurance Product Review (SAPR) Guide*. (See “Other resources” on page 519 for bibliographic information for this publication.)
  - Review the *IBM TotalStorage 3494 Tape Library Introduction and Planning Guide*, GA32-0279.
  - Review the *IBM TotalStorage 3494 Tape Library Operator’s Guide*, GA32-0280.
  - Review the VTS and EVTS QBUCKET (available on IBMLINK, or you can contact your IBM representative to obtain them).
  - Complete systems assurance for the VTS.
- Decide on the volser ranges for VTS-owned physical cartridges and logical volumes (see 6.4.1, “Defining volser ranges for stacked volumes” on page 222). Order these physical cartridges with sufficient lead time to coincide with the expected installation.
- Review Chapter 3, “Preinstallation planning and sizing” on page 57.

4.1.2 Overview of installation tasks

The tasks outlined in this section are specific to the simultaneous installation of a 3494 and a VTS. If you are installing a VTS in an existing 3494, some of these tasks may not apply to you:

- Hardware items (completed by your IBM service representative):
  - Install the 3494 and Library Manager.
  - Install any native drives that will not be VTS controlled:
    - IBM TotalStorage 3590
    - IBM 3490E
  - Install the Model Bxx VTS Frame and D12 Frame(s).

- Host-specific activities (define drives to the host):
  - z/OS
  - z/VM and VM/ESA
  - z/VSE and VSE/ESA
- Software-specific items:
  - Apply maintenance for the IBM 3494 or IBM 3584/3953.
  - Verify or update exits for the tape management system (if applicable) and define logical volumes to it.
  - Update the storage management system software:
    - z/OS DFSMS
    - SYS1.PARMLIB members
    - OAM parameters
    - Define the TCDB
    - Define or update SMS classes and groups
    - Define or update ACS routines
    - Install applicable OAM exits
  - BTLS:
    - Follow tasks in the BTLS user's guide.
    - Install and update exits.
    - Apply maintenance.
    - Update JCL with new unit parameter.
  - DFSMS/VM RMS:
    - Consult the *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141, and Chapter 4.6, “VM/VSE considerations” on page 160.
- VTS-specific installation items:
  - Define VTS environment, using the Library Manager.
  - Define the volser ranges for VTS-owned physical and logical volumes to the Library Manager (see 6.5, “Inserting logical volumes into the VTS” on page 241).
  - Insert VTS-owned physical volumes in the 3494 (see 6.4.4, “Inserting stacked volumes in the VTS” on page 229).

These tasks will be discussed further, including the recommended order of events, later in this chapter.

Once your VTS is installed in the 3494, carry out these post installation tasks:

- Schedule and complete operator training.
- Schedule and complete storage administrator training.

Route data to the VTS (see 5.7.1, “Moving data into the VTS: phased method” on page 197) through:

- SMS tape, using ACS routines
- BTLS, using the unit parameter
- DFSMS/VM RMS

Analyze your current environment and move selected existing data to the VTS to:

- Reduce racking and floor space requirements
- Ensure the readability of data requiring long term storage
- Allow frequently mounted volumes to benefit from an automated environment

4.2 Hardware I/O configuration definition

From a host perspective, the VTS subsystem looks like four, eight, or sixteen IBM 3490E tape control units, each with sixteen devices attached through ESCON or FICON channels. Virtual 3490E tape drives are defined just like a physical IBM 3490-A10 controller with sixteen addresses through the HCD, or through the IOCP and the MVSCP or HCPRIO macro. However, you should not define a preferred path for the VTS virtual devices.

Before you can use the VTS, you need to define it to the zSeries system and processor through HCD. Because the virtual tape drives of the VTS are library resident, you must define them through the HCD dialog, specifying LIBRARY=YES. As for the VTS drives resident in the physical library, there is no definition in IOCP, MVSCP and HCPRIO.

For BTLS-managed libraries, HCD is not required, because BTLS does not require the LIBRARY=YES definition. If you are using BTLS to control the library functions of the VTS, you can use either HCD or IOCP/MVSCP.

You may need to update missing interrupt handler (MIH) values as well, as described inSet values for Missing Interrupt Handler (MIH)

After defining the VTS to a system by whichever method, you should check that the devices can be brought online. Also plan to update the expected IPL configuration so that this is automatic and coincides with the implementation via the Library Manager and host software.

4.2.1 Logical path considerations: ESCON

The VTS can be attached to zSeries hosts and logical partitions through up to 16 ESCON channels. Each channel attachment provides 64 logical paths, which makes up to 1024 logical paths in a sixteen channel configuration.

Depending on the configuration of the VTS, you will have 32, 64, 128 or 256 virtual devices. You need to define one control unit (CU) macro in the HCD dialog for every 16 virtual devices.

The Virtual Tape Server (VTS) provides the following internal connectivity:

- Model B10 VTS provides 4 control units.
- Model B20 VTS provides 8 or 16 control units.

The VTS provides the following ESCON connectivity, and you should plan so that these limits will not exceeded:

- Up to 1024 device numbers per ESCON channel
- Up to 64 logical path per ESCON port
- Up to 128 logical path per CU for ESCON

Figure 4-1 shows how logical paths are established in the VTS. The VTS ESCON ports are connected to all control units; and all control units are connected to all ESCON ports.

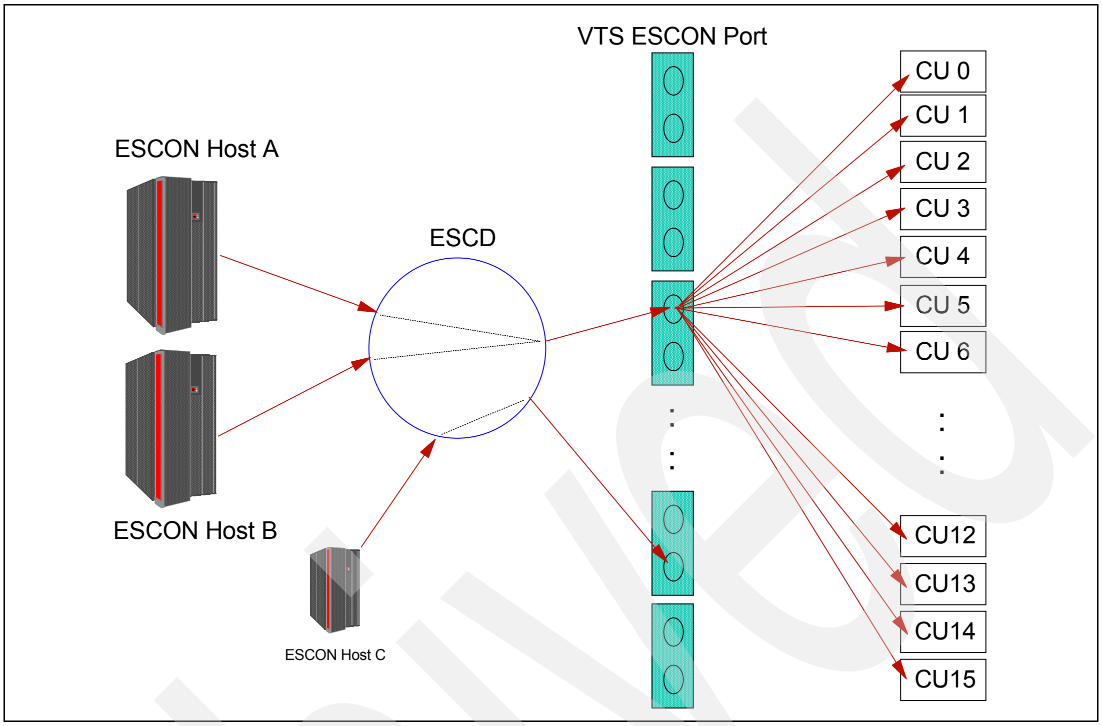

Figure 4-1 Logical path establishment

The ESCON link architecture limits the accessible devices to 1024 per ESCON channel.

The VTS may have up to 256 devices defined, therefore this limit does not apply.

Up to 64 logical paths per ESCON port

In each ESCON port, the VTS establishes the link between host and control units.

Number of logical paths = number of control units x number of systems to this port.

- If you have 2 CUs with 32 virtual tape drives, up to 32 systems can access this port.
- If you have 4 CUs with 64 virtual tape drives, up to 16 systems can access this port.
- If you have 8 CUs with 128 virtual tape drives, up to 8 systems can access this port.
- If you have 16 CUs with 256 virtual tape drives, up to 4 systems can access this port.

If more systems should be connected to VTS with given number of CUs, define them to access another ESCON port in the VTS as ESCON Host C in Figure 4-1.

Up to 128 logical paths per CU

Control unit images in the VTS can establish ESCON logical paths up to a total of 128 paths.

Number of logical paths = number of systems to this CU x number of paths.

- If you define 4 paths to a CU, up to 32 systems can access to this CU.
- If you define 8 paths to a CU, up to 16 systems can access to this CU.

If more systems should be connected a control unit, adjust the system definition 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 with low throughput requirements or those that require fewer virtual devices.

Total logical path that VTS provides

Each ESCON channel in the VTS provides 64 logical paths, which makes up a total of 128 logical paths in a two-channel configuration. Therefore:

- In a 4-channel VTS, 256 logical paths are provided.
- In an 8-channel VTS, 512 logical paths are provided.
- In a 16-channel VTS, 1024 logical paths are provided.

Number of logical paths = number of hosts x number of CUs x number of channels.

The calculated number of logical paths must not exceed theoretical maximum logical paths. The number of channels are channel paths from all hosts.

This formula assumes all hosts access all control units in VTS with same channel paths. If the calculated number of logical paths exceeded, adjustment is required.

4.2.2 Logical path considerations: FICON host

The FICON link architecture is superior to ESCON in that it relieves the addressable device limitation of 1024 per ESCON channel. With FICON, the limitation has been increased to 16384 in either FICON bridge mode (FCV) or FICON native (FC) mode.

The VTS may have up to 256 devices defined, therefore this limit does not apply.

Up to 128 logical paths per FICON port

In each FICON port, the VTS establishes the link between host and control units.

Number of logical paths = number of control units x number of system to this port.

- If you have 4 CUs with 64 virtual tape drives, up to 32 systems can access this port.
- If you have 8 CUs with 128 virtual tape drives, up to 16 systems can access this port.
- If you have 16 CUs with 256 virtual tape drives, up to 8 systems can access this port.

Up to 128 logical paths per CU

Control unit images in the VTS can establish FICON logical paths up to a total of 128 paths.

Number of logical paths = number of systems to this CU x number of paths.

- If you define 4 paths to a CU, up to 32 systems can access to this CU.
- If you define 8 paths to a CU, up to 16 systems can access to this CU.

If more systems need to be connected to a control unit, adjust the system definition 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 with low throughput requirements or those that require fewer virtual devices.

Figure 4-2 details the FICON/ESCON intermix for a VTS.

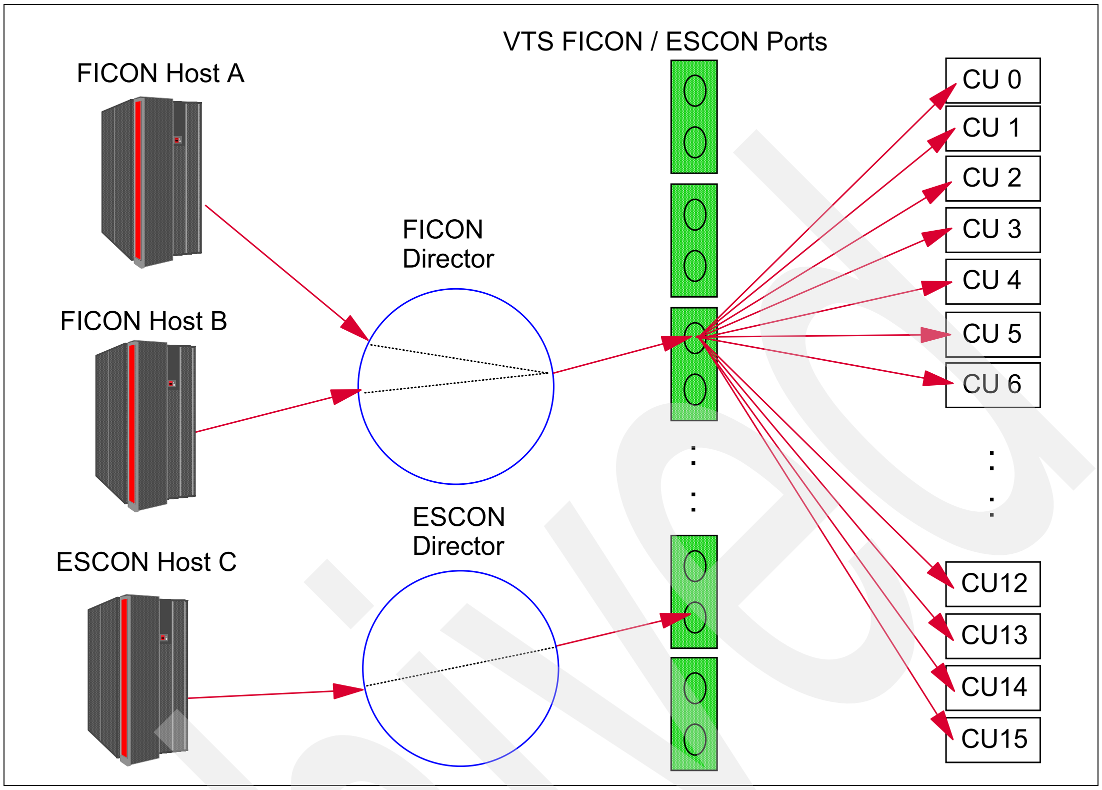

Figure 4-2 FICON / ESCON intermix on a VTS

Total logical path that VTS provides

Each FICON channel in the VTS provides 128 logical paths, which makes up a total of 256 logical paths in a two-channel configuration. Therefore:

- In a 4-channel VTS, 512 logical paths are provided.
- In an 8-channel VTS, 1024 logical paths are provided.

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

The calculated number of logical paths must not exceed theoretical maximum logical paths. The number of channels are channel paths from all hosts.

This formula assumes all hosts access all control units in VTS with same channel paths. If the calculated number of logical paths exceeded, adjustment will be required.

The redbook, *FICON Native Implementation and Reference Guide*, SG24-6266, covers the planning and implementation of FICON channels, operating in FICON native (FC) mode for the IBM zSeries 900 and 9672 Generation 5 (G5) and Generation 6 (G6) processors. It also discusses the FICON and Fibre Channel architectures, terminology, and supported topologies.

Note: A reduction in the number of physical paths will reduce the throughput capability of the VTS and the number of available logical paths. A reduction in control units will reduce the number of virtual devices available for any individual host.

4.2.3 Defining devices through IOCP

In the following IOCP examples, the virtual drives have been configured by using the CUADD parameter in the CNTLUNIT macro statements. Although this method is not used in most installations, the definitions can be used to get an idea of how the logical control unit and device addresses have to be defined in HCD.

IOCP statements for VTS and ESCON director

In this section we discuss the steps required to configure, a VTS attached through IBM 9032-3 ESCON director. If you are installing ESCON directors for the first time, the ESCON directors themselves also have to be defined in the IOCP or HCD input/output definition file (IODF). For information about defining the ESCON directors, refer to:

- *S/390 Input/Output Configuration Program User's Guide and ESCON CTC Reference*, GC38-0401
- *z/OS Hardware Configuration Definition User's Guide*, SC33-7988

A CNTLUNIT and an IODEVICE macro must be coded for each ESCON director.

Figure 4-3 shows a four-channel-path and four-control-unit configuration with a single ESCON director.

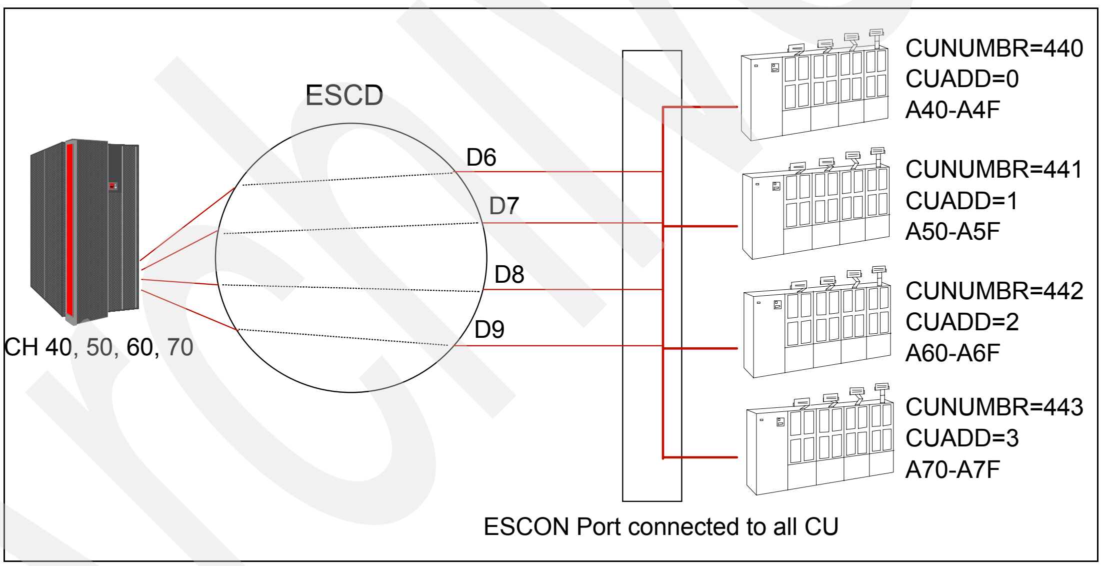

Figure 4-3 VTS definition with an ESCON director

The sample input deck for the configuration in Figure 4-3 is shown in Example 4-1.

Example 4-1 IOCP for VTS with ESCON director

```
Virtual Tape Server
CHPID PATH = 40,TYPE=CNC,SWITCH=01
CHPID PATH = 50,TYPE=CNC,SWITCH=01
CHPID PATH = 60,TYPE=CNC,SWITCH=01
CHPID PATH = 70,TYPE=CNC,SWITCH=01
                :
CNTLUNIT CUNUMBR=440,PATH=(40,50,60,70),UNIT=3490,          X
UNITADD=((00,16)),LINK=(D6,D7,D8,D9),CUADD=0
```

```
IODEVICE ADDRESS=(A40,16),UNIT=3490,CUNUMBR=(440),            X
UNITADD=00
CNTLUNIT CUNUMBR=441 PATH=(40,50,60,70),UNIT=3490,           X
UNITADD=((00,16)),LINK=(D6,D7,D8,D9),CUADD=1
IODEVICE ADDRESS=(A50,16),UNIT=3490,CUNUMBR=(441),            X
UNITADD=00
CNTLUNIT CUNUMBR=442,PATH=(40,50,60,70),UNIT=3490,           X
UNITADD=((00,16)),LINK=(D6,D7,D8,D9),CUADD=2
IODEVICE ADDRESS=(A60,16),UNIT=3490,CUNUMBR=(442),            X
UNITADD=00
CNTLUNIT  CUNUMBR=443,PATH=(40,50,60,70),UNIT=3490,           X
UNITADD=((00,16)),LINK=(D6,D7,D8,D9),CUADD=3
IODEVICE ADDRESS=(A70,16),UNIT=3490,CUNUMBR=(443),            X
UNITADD=00
```

IOCP statements for VTS without ESCON director

Figure 4-4 shows a four-control-unit configuration without ESCON directors.

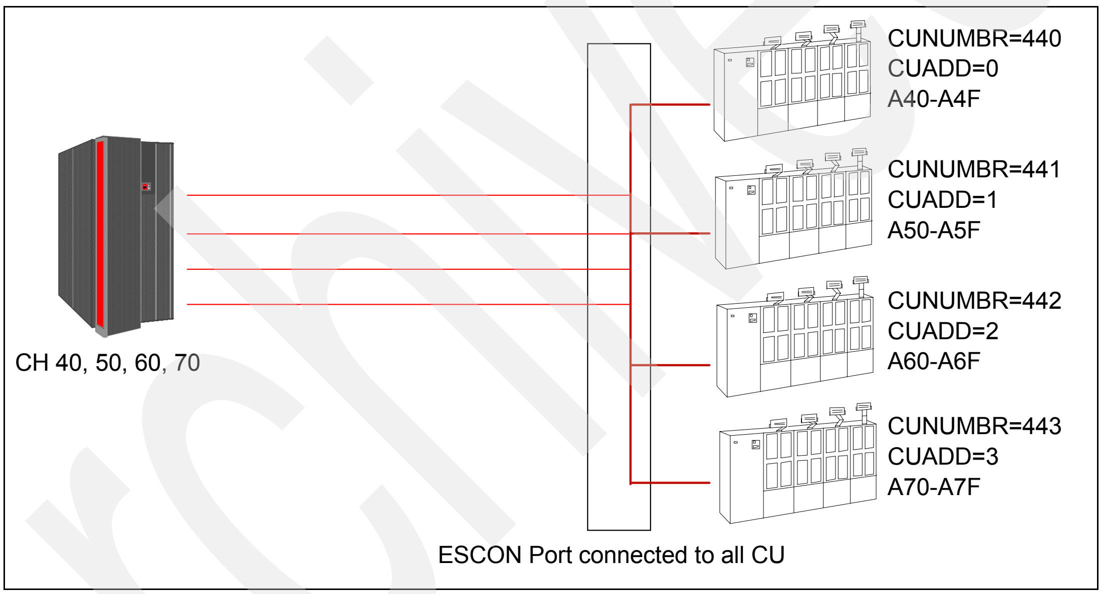

Figure 4-4 VTS definition without an ESCON director

The sample input deck for the configuration in Figure 4-4 is listed in Example 4-2.

Example 4-2 IOCP for VTS without ESCON director

```
Virtual Tape Server IOCP
CHPID PATH = 40,TYPE=CNC
CHPID PATH = 50,TYPE=CNC
CHPID PATH = 60,TYPE=CNC
CHPID PATH = 70,TYPE=CNC
              :
CNTLUNIT CUNUMBR=440,PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=0,
LIBRARY-ID=11111,LIBPORT=01
IODEVICE ADDRESS=(A40,16),UNIT=3490,CUNUMBR=(440),
UNITADD=00
CNTLUNIT CUNUMBR=441 PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=1,
LIBRARY-ID=11111,LIBPORT=02
```

```
IODEVICE ADDRESS=(A50,16),UNIT=3490,CUNUMBR=(441),
UNITADD=00
CNTLUNIT CUNUMBR=442,PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=2,
LIBRARY-ID=11111,LIBPORT=03
IODEVICE ADDRESS=(A60,16),UNIT=3490,CUNUMBR=(442),
UNITADD=00
CNTLUNIT CUNUMBR=443 PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=3,
LIBRARY-ID=11111,LIBPORT=04
IODEVICE ADDRESS=(A70,16),UNIT=3490,CUNUMBR=(443),
UNITADD=00
```

IOCP statements for VTS and a FICON director

Figure 4-5 shows a four-control-unit configuration with a FICON director.

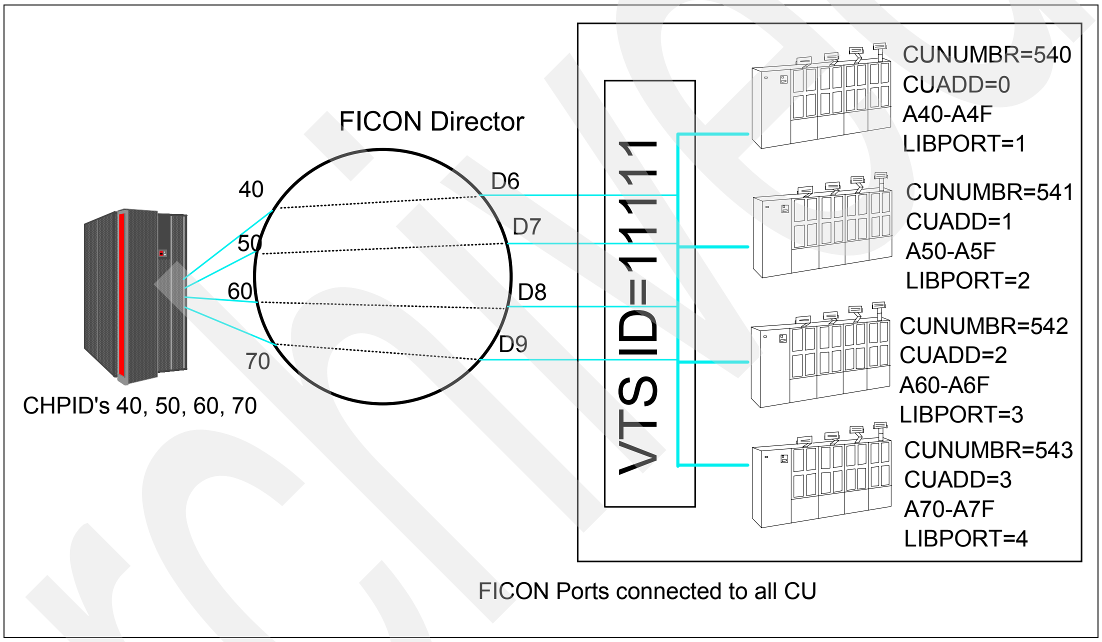

Figure 4-5 VTS with FICON director connections

As you can see in Figure 4-5, the z/OS host can address all of the drives, through all of the channels defined. The sample input deck for this configuration is listed in Example 4-3.

Example 4-3 IOCP for VTS with FICON director

```
Virtual Tape Server IOCP
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
              :
CNTLUNIT CUNUMBR=540,PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=0
LINK=(D6,D7,D8,D9),LIBRARY-ID=11111,LIBPORT=01
IODEVICE ADDRESS=(A40,16),UNIT=3490,CUNUMBR=(540),
UNITADD=00
CNTLUNIT CUNUMBR=541 PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=1
LINK=(D6,D7,D8,D9),LIBRARY-ID=11111,LIBPORT=02
```

```
IODEVICE ADDRESS=(A50,16),UNIT=3490,CUNUMBR=(541),
UNITADD=00
CNTLUNIT CUNUMBR=542,PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=2
LINK=(D6,D7,D8,D9),LIBRARY-ID=11111,LIBPORT=03
IODEVICE ADDRESS=(A60,16),UNIT=3490,CUNUMBR=(542),
UNITADD=00
CNTLUNIT CUNUMBR=543 PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=3
LINK=(D6,D7,D8,D9),LIBRARY-ID=11111,LIBPORT=04
IODEVICE ADDRESS=(A70,16),UNIT=3490,CUNUMBR=(543),
UNITADD=00
```

IOCP statements for VTS without a FICON director

This definition (Example 4-4) is essentially the same as the point to point ESCON configuration expect for the channel type (FICON = FC, ESCON = CNC).

Example 4-4 IOCP for VTS without FICON director

```
Virtual Tape Server IOCP
CHPID PATH = 40,TYPE=FC
CHPID PATH = 50,TYPE=FC
CHPID PATH = 60,TYPE=FC
CHPID PATH = 70,TYPE=FC
              :
CNTLUNIT CUNUMBR=540,PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=0,
LIBRARY-ID=11111,LIBPORT=01
IODEVICE ADDRESS=(A40,16),UNIT=3490,CUNUMBR=(540),
UNITADD=00
CNTLUNIT CUNUMBR=541 PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=1,
LIBRARY-ID=11111,LIBPORT=02
IODEVICE ADDRESS=(A50,16),UNIT=3490,CUNUMBR=(541),
UNITADD=00
CNTLUNIT CUNUMBR=542,PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=2,
LIBRARY-ID=11111,LIBPORT=03
IODEVICE ADDRESS=(A60,16),UNIT=3490,CUNUMBR=(542),
UNITADD=00
CNTLUNIT CUNUMBR=543 PATH=(40,50,60,70),UNIT=3490,
UNITADD=((00,16)),CUADD=3,
LIBRARY-ID=11111,LIBPORT=04
IODEVICE ADDRESS=(A70,16),UNIT=3490,CUNUMBR=(543),
UNITADD=00
```

4.2.4 Defining devices through HCD

In this section we describe the process of defining the VTS through HCD panels. The most important points to observe are as follows:

- HCD definitions are required for SMS tape.
- Two, four, eight, or sixteen 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.

- To access control units in the VTS properly, it is necessary to define the correct number of control units using logical unit address (CUADD) = 0 to 15 and associated strings, 16 drives each.
- Mixed device types are not allowed in the same device pool in a tape library.

Figure 4-6 and Figure 4-7 show the two important screens for specifying a tape control unit.

```
------------------------ Add Control Unit ---------------------
CBDPCU10

Specify or revise the following values.
Control unit number  . . . . 0440  +

Control unit type  . . . . . 3490         +

Serial number  . . . . . . . __________
Description  . . . . . . . . ________________________________

Connected to switches  . . . 01  01  01  01  __  __  __  __  +
Ports  . . . . . . . . . . . D6  D7  D8  D9  __  __  __  __  +

If connected to a switch, select whether to have CHPIDs/link
addresses and unit address range proposed.

Auto-assign  . . . . . . . . 2   1.  Yes
                                 2.  No

F1=Help     F2=Split     F4=Prompt   F5=Reset    F9=Swap    F12=Cancel
```

Figure 4-6 Adding the first VTS Control Unit through HCD, Part 1

Specify the control unit number and the type (here, 3490) as shown in Figure 4-6, then press Enter. The screen 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   50   60   70   __   __   __   __  +
Link address . . . . . . D6   D7   D8   D9   __   __   __   __  +

Unit address . . . . . . 00  __   __   __   __   __   __   __  +
Number of units  . . . . 16  ___  ___  ___  ___  ___  ___  ___

Logical address  . . . . 0   + (same as CUADD)

Protocol . . . . . . . . __   + (D,S or S4)
I/O concurrency level  . 2  + (1, 2 or 3)


 F1=Help     F2=Split    F4=Prompt   F5=Reset    F9=Swap    F12=Cancel
```

Figure 4-7 Adding the first VTS Control Unit through HCD, Part 2

Note that if the VTS is not being attached through ESCON directors, the link address fields would be blank. You should define the second to sixteenth VTS tape control units, specifying the logical unit address (CUADD) = 1 to F in the CBDPCU12 panel as shown in Figure 4-7. The number of drives to be specified for an IBM VTS control unit is always 32, 64, 128 or 256, but you must specify these in two or four strings, with each string associated with one of the previously defined control units.

To define the VTS 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 shown in Figure 4-8 is displayed.

```
------------------------------- Add Device ---------------------------
CBDPDV10

Specify or revise the following values.
Device number  . . . . . . . . 0A40  (0000 - FFFF)
Number of devices  . . . . . . 16__
Device type  . . . . . . . . . 3490_________  +

Serial number  . . . . . . . . __________
Description  . . . . . . . . . ________________________________

Connected to CUs . . 0440  ____  ____  ____  ____  ____  ____  ____  +


F1=Help     F2=Split    F4=Prompt   F5=Reset    F9=Swap    F12=Cancel
```

Figure 4-8 Adding the first 16 drives through HCD, Part 1

After entering the required information and specifying to which processors and operating systems the devices are connected, the screen in Figure 4-9 is displayed, where you can update the device parameters.

```
CBDPDV13          Define Device Parameters / Features          Row 1 of 6
Command ===> ___________________________________________ Scroll ===> PAGE

Specify or revise the values below.
Configuration ID . : AB          MVS operating system
Device number  . . : 0A40        Number of devices  :16
Device type  . . . : 3490

Parameter/  Value    P Req.  Description
Feature
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
AUTOSWITCH  No               Device is automatically switchable
LIBRARY-ID  12345            5 digit library serial number
LIBPORT-ID  02               2 digit library string ID (port number)
ALTCTRL     No               Separate physical control unit path
SHARABLE    yes.             Device is Sharable
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 Adding the first 16 VTS drives through HCD, Part 2

Notes: If you are defining drives that are installed in a system-managed IBM 3494 VTS Library, you must specify LIBRARY=YES.

If more than one zSeries system will be 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.

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.

4.2.5 HCD support for library and port IDs

LIBRARY-ID and LIBPORT-ID are optional parameters provided by HCD.

These fields allow HCD to provide the library configuration information that is normally obtained from the device at IPL time. For devices that are available during IPL, the HCD information is redundant. However, if the library is unavailable during IPL or in OFFLINE status, the HCD information allows the library device to be brought online (when it subsequently becomes available to the system) without reactivating the IODF.

.

Important: We strongly recommend to specify LIBRARY-ID and LIBPORT-ID; it may protect you from needing additional IPLs each time you change your I/O configuration.

If you provide the LIBRARY-ID and LIBPORT-ID, they must match the IDs provided by the hardware. 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. For example, in a five frame plus two Model B20 VTS configuration, as shown in Figure 4-10, you have the frames in the following order:

- L12 Frame with two 3590 tape drives - Open Systems attached
- D22 Frame providing VTS physical drives
- D22Frame providing VTS physical drives
- D14 Frame providing four native 3590 drives (and a controller)
- D14 Frame providing four native 3590 drives (and a controller)
- B20 Frame providing first VTS
- B20 Frame providing second VTS

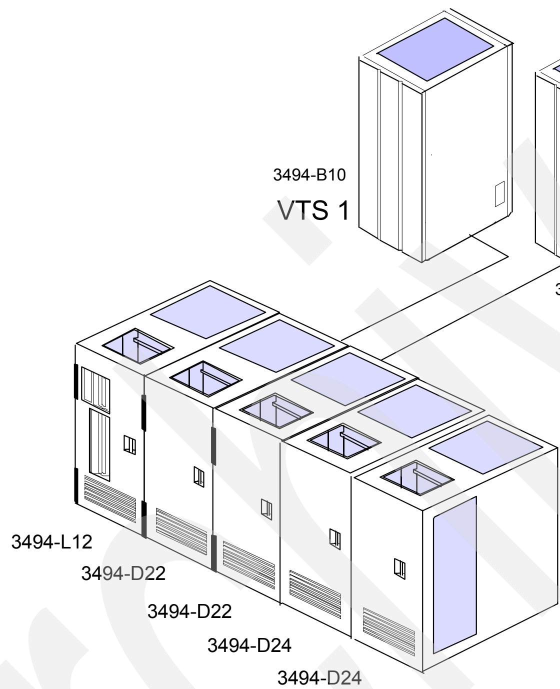

Figure 4-10 Sample configuration showing LIBRARY-ID and LIBPORT-ID

You need to specify the following LIBRARY-IDs and LIBPORT-IDs:

- L12 LIBRARY-ID = none, LIBPORT-ID = none
- D22 none
- D22 none
- D24 LIBRARY-ID = 12345, LIBPORT-ID = 01
- D24 LIBRARY-ID = 12345, LIBPORT-ID = 02
- B20 LIBRARY-ID = 54321, LIBPORT-ID = 01 to 04, CUADD = 0 to 3
- B20 LIBRARY-ID = 98765, LIBPORT-ID = 01 to 08, CUADD = 0 to 7

Note:
If 64 virtual drives are defined:
LIBPORT-ID = 01 to 04 and CUADD = 0 to 3 are used.
If 128 virtual drives are defined:
LIBPORT-ID = 01 to 08 and CUADD = 0 to 7 are used.
If 256 virtual drives are defined:
LIBPORT-ID = x’01’ to x’10’ and CUADD = 0 to F are used.

When attaching the VTS to an IBM 3584/3953 configuration, the LIBRARY-IDs and LIBPORT-IDs for the VTS are the same as when attached to an IBM 3494 Tape Library, because the VTS has its own LIBRARY-ID, and the LIBPORT-ID for the VTS always starts with 01.

In an existing installation you can use the DEVSERV QTAPE system command to find out what to specify. All tape drives connected to a control unit have the same LIBPORT-ID. Therefore you have to issue the DS QT command only once per control unit. 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 4-11 details the output of a DS QT system command:

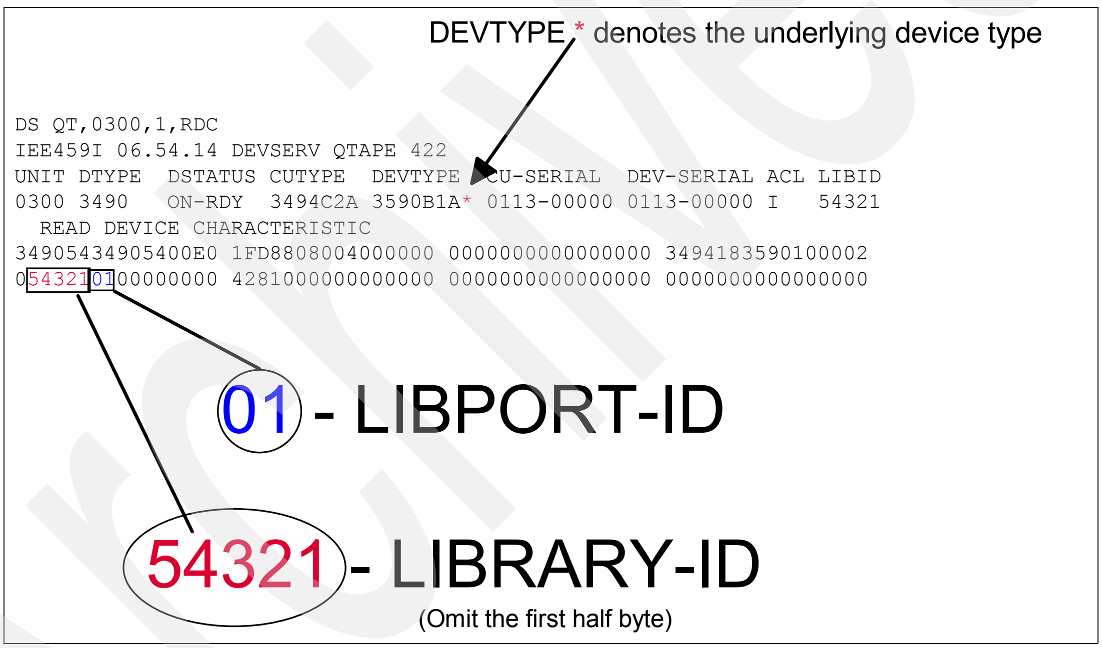

Figure 4-11 Sample DEVSERV QT command output

4.2.6 Activate the I/O configuration

There are differences in the process of activating the IODF for a new tape library or for changes made to an existing tape library that cause changes in the device addresses, if you want to avoid an IPL. As an alternative to the procedures described below, you can always IPL the system.

Installing a new tape library

If you are installing a VTS for the first time, from a host software definition point of view this is the installation of a new library. When you are activating the IODF for a new tape library, the following steps are required to get the tape library or VTS ONLINE without IPLing your systems:

1. Activate the IODF

2. MVS VARY ONLINE the devices in the library. This will creat4e some of the control blocks. You will see the message

   ```
   IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED
   ```

3. Do the final ACTIVATE. This is required to build the Eligible Device Table (EDT) for MVS Allocation.

You can check details using the DEVSERV QTAPE command, which provides information on Unit Control Block (UCB), UCB prefix, UCB common extension, Device Class Extension (DCE), and Read Device Characteristics (RDC) data and Read Configuration Data (RCD) data, which are data buffers acquired directly from the device.

> **Note:** If you are just adding additional device addresses to an existing VTS, you can use the same process as for a new tape library.

Modifications to an existing tape library

When you are modifying an existing tape library so that existing device addresses are to be changed, follow this process flow:

1. Activate an IODF deleting all devices from the library.
2. Activate an IODF that defines all of the devices of the modified library.
3. MVS VARY ONLINE the devices in the library. This will create some of the control blocks. You will see the message

   ```
   IEA437I TAPE LIBRARY DEVICE(ddd), ACTIVATE IODF=xx, IS REQUIRED
   ```

4. Do the final ACTIVATE.

Alternatively, you can use DS QL,nnnnn,DELETE (where nnnnn is the LIBID) to delete the library's dynamic control blocks. If you have IODEF with LIBID and LIBPORT coded already, proceed as follows:

1. Use QLIB LIST to display that the INACTIVE control blocks have been deleted.
2. Use ACTIVATE IODF to redefine the devices.
3. Use QLIB LIST to display that the ACTIVE control blocks are properly defined.

IF LIBID and LIBPORT are not coded, perform the following steps:

1. MVS Vary online the devices in the library. This will create some control blocks, and you will see the message:

   ```
   IEA437I TAPE LIBRARY DEVICE(DDD), ACTIVATE IODF=XX, IS REQUIRED
   ```

2. Use ACTIVATE IODF to redefine the devices.
3. Use QLIB LIST to verify that the ACTIVE control blocks are properly defined.

4.2.7 Display and control your settings

The DISPLAY SMS command can be used to display and check the settings within DEVSUPxx member for the scratch categories.

Note that the scratch count of MEDIA2 does not necessarily match the number of scratch volumes of your tape management system when you use the **Expire Hold** function in the VTS. OAM displays the scratch count it receives from the Library Manager. Example 4-5 shows the sample output of a DISPLAY SMS,LIBRARY command.

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
----------------------------------------------------------------------
Library supports outboard policy management.
Library supports import/export.
Convenience I/O station installed.
Convenience I/O station in Output mode.
Convenience I/O station Empty.
Bulk input/output not configured.
```

The DEVSERV QUERY LIBRARY or DS QL command should always be used to query your Library configuration before and after an activate of your IODF when you do changes for the library.

Figure 4-12 shows an example of how you would list all libraries using the DS QL command. Those LIBIDs marked with the asterisks 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 4-12 DEVSERV QLIB,LIST

Figure 4-13 shows a detailed list of one single Library using the DS QL,libid,DETAIL command. Check that no duplicate port IDs are listed and that each port has 16 devices — this is the correct output for a VTS.

```
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 4-13 Sample output of the DEVSERV QLIB,libid,DETAILS command

You can display the syntax with DS QL,?. For a complete description of the QLIB command, see Appendix D, "DEVSERV QLIB command" on page 443 or refer to the manual, *z/OS V1R7.0 MVS System Commands*, SA22-7627.

4.2.8 Set values for Missing Interrupt Handler (MIH)

An MIH time of 45 minutes is recommended for all the VTS drives. The 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 IBM 3590 or IBM 3592 tape drives that the VTS manages.

You should specify the MIH values in the PARMLIB member IECIOSxx,. Alternatively, you can use the z/OS operator command SETIOS; this setting will be available until changed or the system is initialized.

To check the current MIH time, use:

```
D IOS,MIH,TIME=TAPE
```

To specify MIH values for emulated 3490E tape drives, use:

```
SETIOS MIH,DEV=(0A40-0A7F),TIME=45:00
```

You can specify an MIH value in the IECIOSxx PARMLIB member as follows:

```
MIH DEV=(0A40-0A7F),TIME=45:00
```

To display the new settings, use:

```
D IOS,MIH,DEV=0A40
```

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. For more details, see the z/OS *MVS Initialization and Tuning Reference*, SA22-7592.

When specifying time intervals, consider the following facts:

- 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.

Refer to the specific device's reference manuals to determine if the device supports self-describing MIH time values.

4.3 VTS definition from the Library Manager

After your IBM service representative installs the VTS, you must define the environment so that you can use high-performance VTS functions. All VTS definition procedures are performed from the Library Manager console. For further details on operating an IBM 3494 Tape Library, refer to the *IBM TotalStorage 3494 Tape Library Operator's Guide,* GA32-0280. For details on operating the IBM 3584 Tape Library, refer to the *IBM TotalStorage 3584 Operator Guide,* GA32-0468 and to the *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473.

The procedure of defining the VTS from the Library Manager includes defining initial volser ranges and management policies for the VTS operation and the 3590 cartridge insertion for VTS stacked volume preparation. Perform the following tasks at the Library Manager:

- Define stacked volume ranges. See 4.3.3, "Define stacked volume ranges" on page 129.
- Define (insert) logical volumes. See 4.3.4, "Define (insert) logical volumes" on page 130.
- Define Fast Ready categories. See 4.3.5, "Define Fast Ready categories" on page 133.
- Define VTS management policies. See "Reclamation and reconciliation" on page 137.
- Define Inhibit reclaim schedule.
- Set reclaim policies.
- Define Free storage threshold.
- Define cleaning schedule. See 4.3.8, "Cleaning schedule pop-up window" on page 142.
- SNMP Traps definition. See 4.3.9, "Simple network management protocol traps" on page 142.

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 zSeries implementation of the VTS the Library Manager definitions and the zSeries software definitions are intertwined.

4.3.1 Navigating the Library Manager

Although a physical tape library logically defined as described in Figure 4-10 on page 122 looks like two separate libraries to the host, the libraries share the same Library Manager and database. Therefore, the volsers of both the virtual and physical volumes in the same physical tape library must be unique. Before you can insert physical 3590 or 3592 cartridges for VTS use into a library, you must define them on the Library Manager (LM). You must also define a beginning set of volser ranges for your Virtual Volumes.

Most of the tasks can be performed either through the Library Manager console or through the Enterprise Tape Library (ETL) or the VTS Specialist, the web-based browser interfaces to the Library Manager. Wherever the task can be performed through the ETL or VTS Specialist, we show these panels in the subsequent chapters. A complete description on the Library Manager console screens can be found in the *IBM TotalStorage 3494 Tape Library Operator's Guide,* GA32-0280 and the *IBM TotalStorage 3584 Operator Guide,* GA32-0468 and to the *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473, depending on which tape library the VTS is attached to.

4.3.2 Check the VTS library sequence number

You need to ensure that the library sequence number defined through the Interactive Storage Management Facility (ISMF) is the same as the library-id given to the VTS by the IBM service representative during the teach operation. Check the library sequence number on the *Operational Status* panel on the ETL Specialist *Monitor Library Manager* selection. You can scroll down until you see the field, *Library Sequence Numbers NNNNN*, as shown in Figure 4-14.

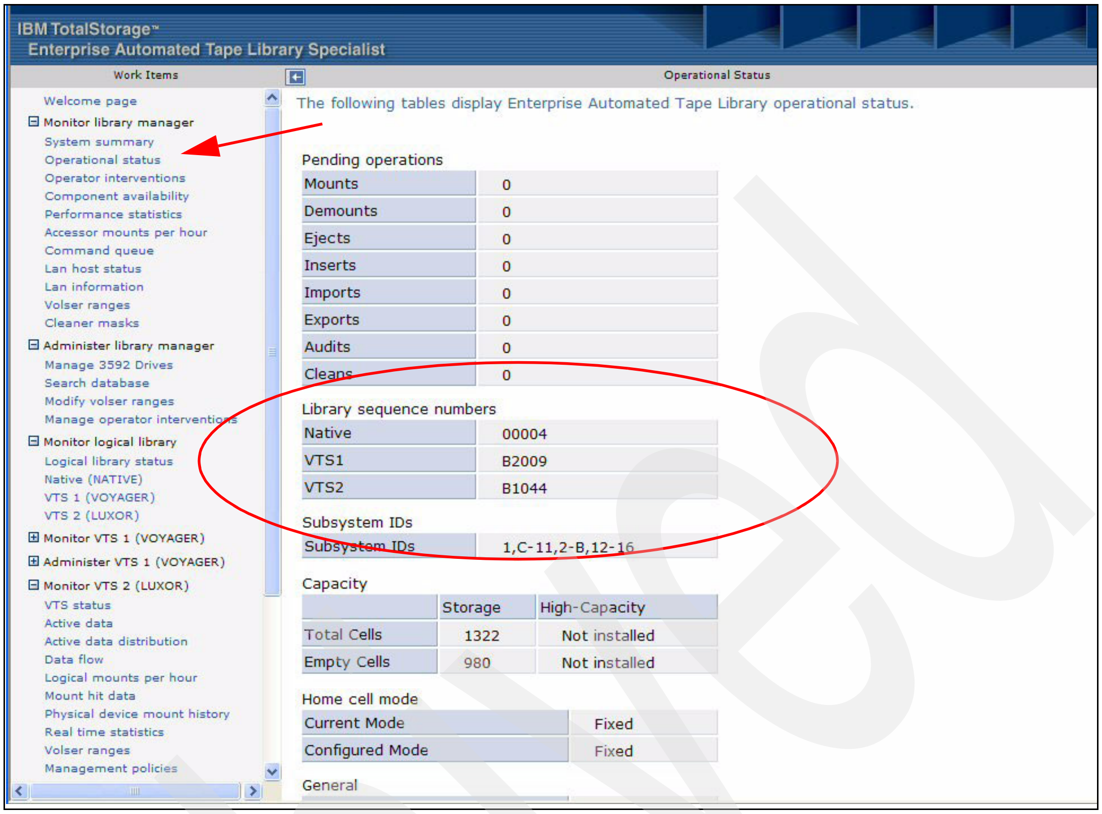

Figure 4-14 Operational Status screen

In the *Operational Status* window, library sequence numbers show the unique identification numbers for each non-VTS and VTS logical tape library within the physical 3494 or the IBM 3584 Tape Library Manager partition. These numbers are assigned at the plant of manufacture and set during the teach operation by the IBM service representative.

4.3.3 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 tape library. We highly recommended that the range be unique to all hosts within the customer environment. You can access the **Modify Volser** panel either from the **Administer Library Manager** or from the **Administer VTS** work items as indicated by the arrows in Figure 4-15.

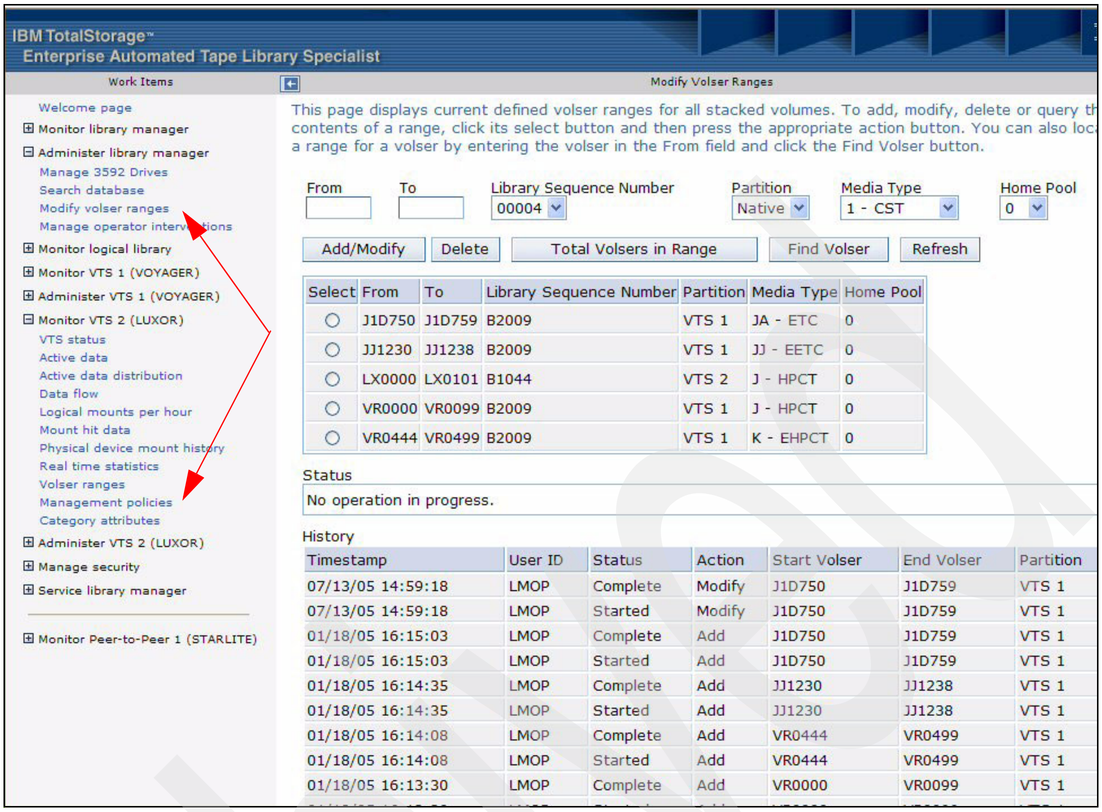

Figure 4-15 Modify Volser Ranges window

With Advanced Policy Management (FC4001-4004), you can also assign your stacked volumes (and / or ranges) to specific volume pools. Details on how this is done are shown in Figure 4-24 on page 146.

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 one for which there are available cartridges. See 6.4.1, “Defining volser ranges for stacked volumes” on page 222.

4.3.4 Define (insert) logical volumes

You insert logical volumes in the VTS 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. The maximum number of logical volumes in one IBM 3494 configuration is 1,000,000 when two VTSs are included.

Important: You should only define 10,000 volumes at a time and then define the next 10,000 when the first 10,000 volumes have finished with the insert processing.

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. For more detailed information how to eject logical volumes, refer to 6.7, “Ejecting logical volumes from the VTS” on page 248.

Use the **Administer VTS** work item **Manage logical volumes** to insert your logical volumes as shown in Figure 4-16.

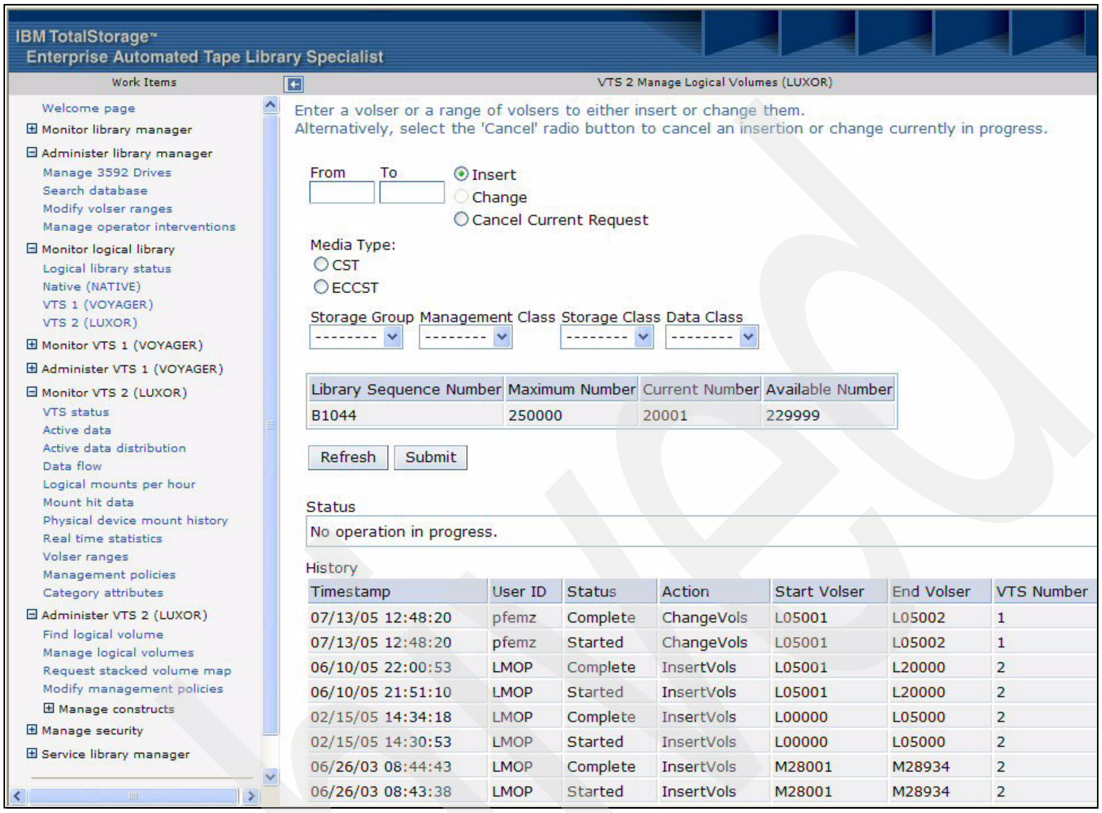

Figure 4-16 Manage logical volumes window

The logical volumes inserted here should already have been defined to the 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.

Should you have the additional feature, Advanced Policy Management (APM: FC4001-4004), and are operating the VTS in an 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.

See 6.5, “Inserting logical volumes into the VTS” on page 241 for details on how to determine the amount of logical volumes and on how to perform the action.

The information from the Library Manager database is uploaded to the attached hosts before host applications can access the tape cartridges inside the 3494. 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.

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.

A 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, we recommend that a GRS ring be used 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.

You may see within a SYSPLEX where all systems are part of a GRS ring or in GRS star mode — mostly a JES2 complex — that the insert process is very fast. A JES3 complex is often a plex beyond SYSPLEX boundaries. Because a GRS ring does not communicate with the other hosts, you may need to use other ISV products for serialization. In this case entry processing may take much longer.

Tip: We recommend to 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 above may appear multiple times until an owning host processes the volumes for entry.

Note: When you are inserting many logical volumes into the VTS, we recommend that you only add volumes 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 where 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 the CBRUXENT installation exit on all systems which are not the primary owners of specific volumes, then these volumes may have their category updated by a subsequent process.

4.3.5 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 4-17 shows the Define Fast Ready Categories window. Note that this definition cannot be done from the ETL Specialist.

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 the 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, which are described in Appendix E, “Library Manager volume categories” on page 451.

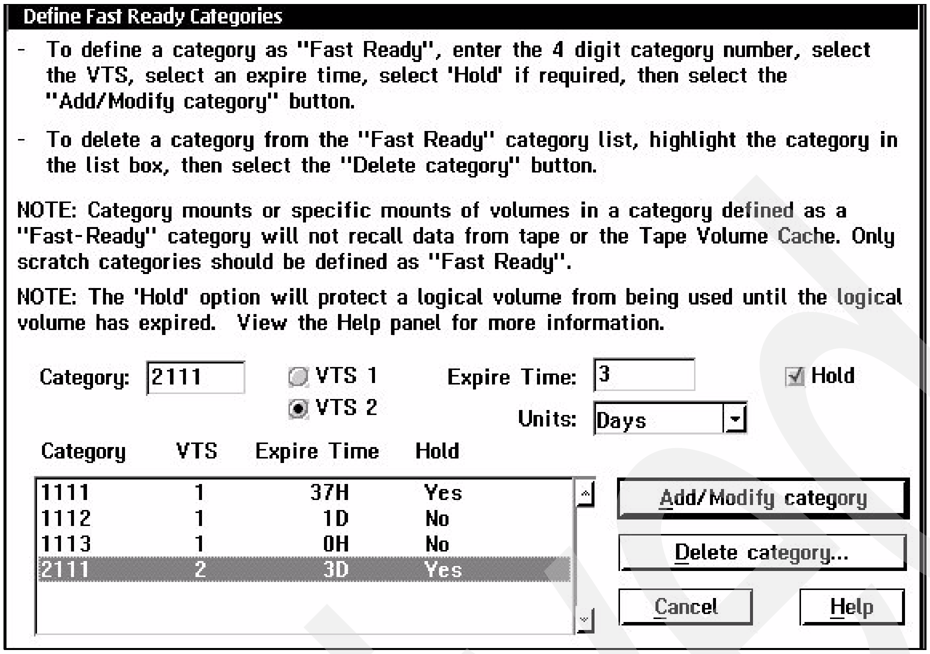

Figure 4-17 Set VTS category attributes pop-up window

In this panel you 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 z/OS DFSMS. These scratch mount categories are 0001 and 0002, respectively if you use the z/OS DFSMS default values. They are ‘0031’ and ‘0032’ in Figure 4-17 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. See 6.6, “Expired volume management” on page 245.

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 E, “Library Manager volume categories” on page 451 for the scratch mount category for each software platform. In addition to the z/OS DFSMS default value for the scratch mount category, you can define your own scratch category to the 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 “cleanup” activity once a category is to be removed from service.

4.3.6 Physical cartridge insertion

At this point, insert your cartridges in the tape library. As this is the first time you insert stacked volumes for your VTS, you may have a large number of cartridges to insert. The easiest way to insert larger numbers of cartridges is to open the frame doors and insert them in empty cells 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 6.4.4, “Inserting stacked volumes in the VTS” on page 229 for information about inserting cartridges in the library.

If you are attaching the VTS to an IBM 3584 Tape Library Manager partition, you must first make sure that the cartridges are assigned to the proper 3584 logical library, as the cartridge insert process requires an additional step. Figure 4-18 sketches the different phases of insert processing of physical cartridges in an IBM 3584 Tape Library.

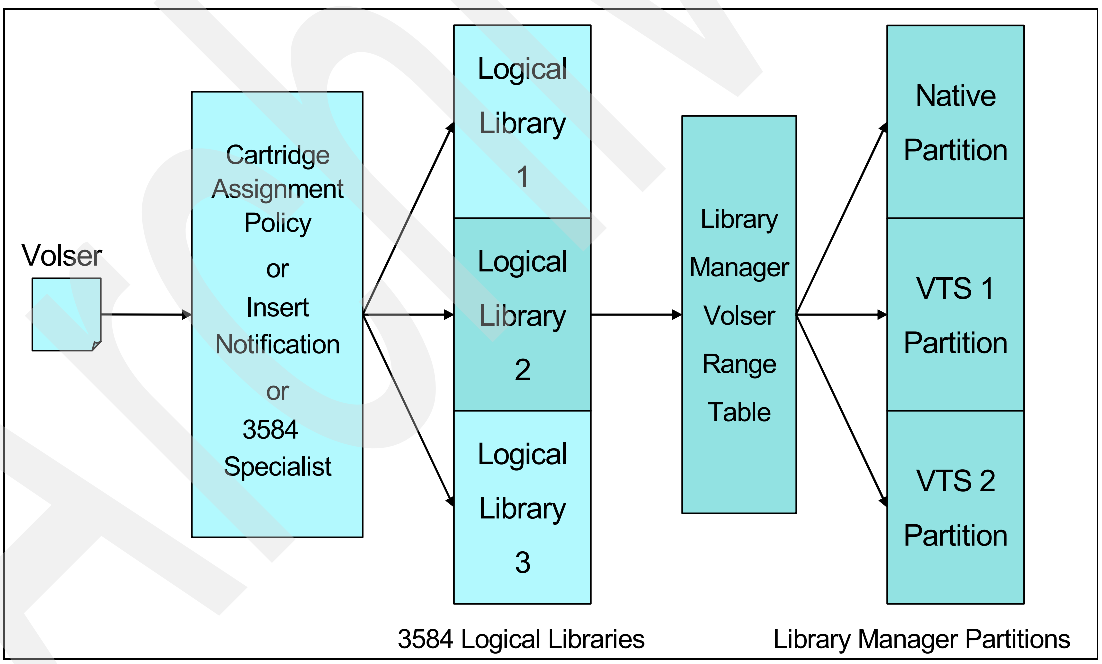

Figure 4-18 IBM 3584 Logical Libraries and cartridge assignment

There are two ways to insert physical cartridges:

- Insert volumes directly into storage cells by opening library doors:
  - The 3584 Cartridge Assignment Policy defines which volumes are assigned to which logical library. See Figure 4-19.
  - Library Manager is notified after doors are closed and 3584 performs inventory. Library Manager updates its database accordingly.

– 3584 Web Specialist can be used for unassigned volumes.

- Insert volumes via Convenience I/O Station (Insert notification is required to be enabled on 3584):
  – The 3584 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 3584 Specialist can be used to assign unassigned volumes.

Figure 4-19 shows the *Cartridge Assignment Policy* screen of the 3584 Tape Library Specialist that allows you to define the CAPs.

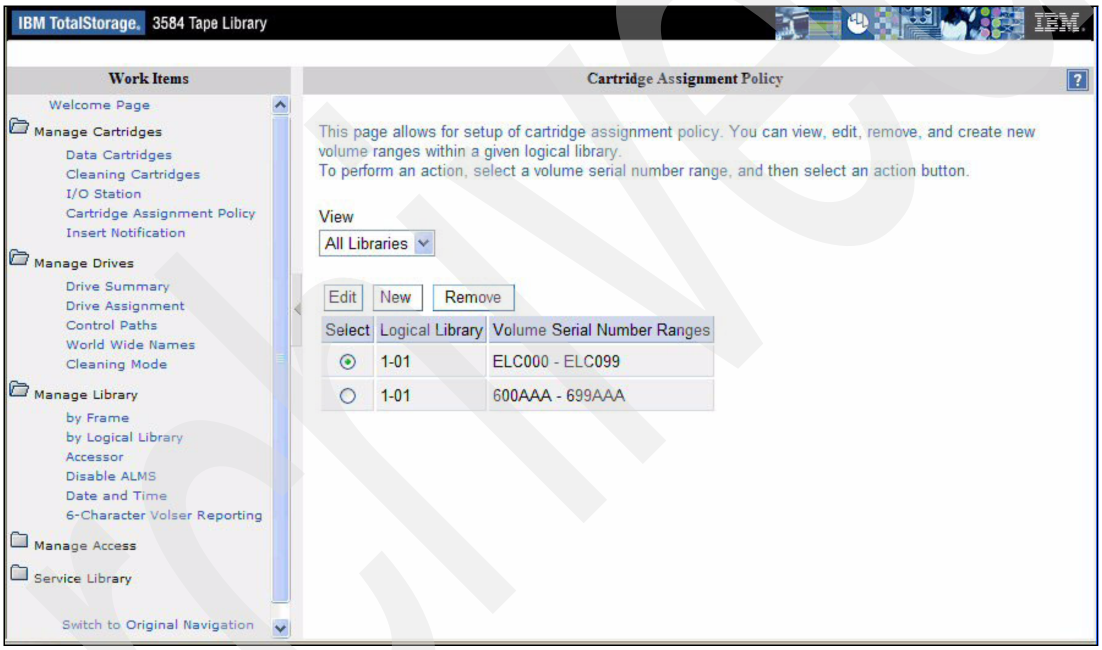

Figure 4-19 3584 Cartridge Assignment Policy

For more information on the setup of the IBM 3584 Tape Library, 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 6.4.1, “Defining volser ranges for stacked volumes” on page 222 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 4.3.3, "Define stacked volume ranges" on page 129 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: Install one or two cleaner volumes in the tape library for each type of (3490E/3590) tape subsystem 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.

4.3.7 Define VTS management policies

The **Set VTS Management Policies** window allows you to 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 and policies through the Library Manager console or the ETL Specialist.

Over time, more and more logical volume copies on a stacked volume become obsolete and the stacked volume contains less and less active data. The storage management software monitors the amount of active data on stacked volumes. It marks the cartridge eligible for reclamation when the percentage set by the Library Manager Reclaim Threshold Percentage value is met. During reclamation, the active data is copied to another stacked volume, leaving the source volume to be used as a scratch stacked volume by the storage management software in the 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 or using the ETL Specialist.

When Advanced Policy Management is installed, you can define additional reclamation policies. For details, refer to 4.4.1, “Defining stacked volume pool properties” on page 143.

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.

There is a panel on the Library Manager that allows you to monitor the amount of active data on stacked volumes and helps to 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 “VTS Active Data Distribution window”, Figure 7-18 on page 306. Alternatively, you can use the respective ETL Specialist Panel by selecting **Active Data** from the **Monitor VTS x** work items.

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.

Reclamation is enabled or disabled according to a 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 task/s 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 4.3.7, “Define VTS management policies” on page 137.

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.

The term, *access to,* includes any borrowing capability. Borrowing is described in "Borrowing and returning" on page 44. 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 4-1 summarizes the new thresholds.

Priority	Condition	Reclaim Schedule Honored	Active Data Threshold % Honored	Number of Concurrent Reclaims
1	Priority Move that benefits a Pool in Scratch Panic State	No	No	1, regardless of idle drives
2	Pool in Scratch Panic 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

Table 4-1 Reclamation Priority Table

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 as 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.

- **Day of week:**

  The choices are **Sunday** through **Saturday** or **Every day**. If Every day is selected, the Start time and Duration entered will apply to every day of the week.

- **Start Hour and Minute:**

  These are the start hour and minute for the inhibition. A 24-hour clock is used, where 00 in the hour field indicates midnight.

- **Duration Hours and Minutes:**

  These are the number of hours and minutes during which the inhibit reclaim should remain in effect. Up to 167 hours and 59 minutes (7 days minus 1 minute) can be specified. By specifying the maximum, it will essentially always inhibit reclaim.

- **Check boxes:**

  These indicate to which VTS to apply the schedule.

Add an entry to the inhibit reclaim schedule by selecting a day of week, a start time, and the duration, and select the **Add** button.

Modify an entry by pressing the **Modify** button.

Delete an entry by highlighting it in the list box and selecting the **Delete** button.

Figure 4-20 shows the ETL Specialist Panel that is displayed when you select **Management Policies** from the **Administer VTS** work items.

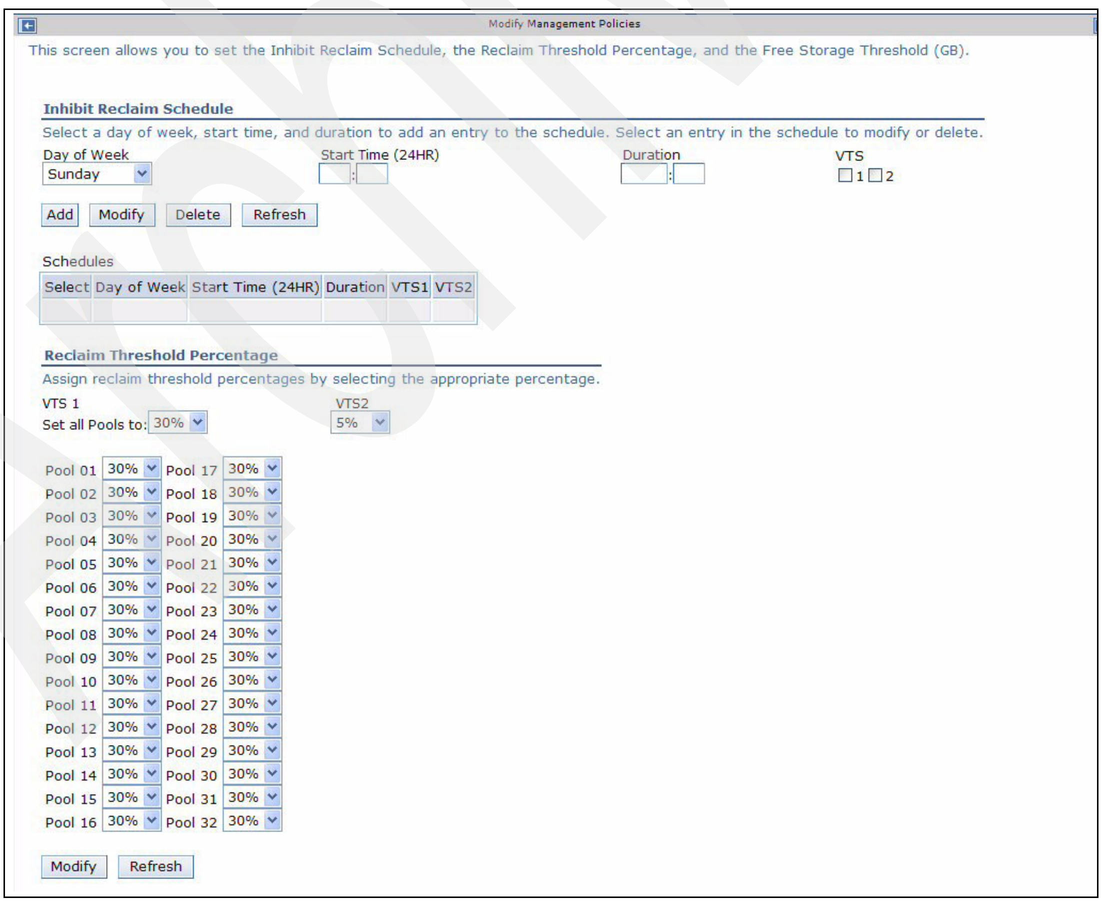

Figure 4-20 Set VTS management policies pop-up window pre LIC 527

VTS x - Set all Pools to If the VTS is Advanced Policy Management enabled, then use this push button to define the reclamation thresholds for individual pools.

Modify Saves all the changes made to the inhibit reclaim schedule and the free-storage threshold.

Refresh Refreshes the display.

Reclaim Threshold Percentage

The *Reclaim Threshold Percentage* (Figure 4-20) 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 that 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 (FF03 for pre 527 LIC or 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.

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 7.5.7, “Active data distribution” on page 305 for information about displaying the panel.

> **Note:** You can see the effect that reconciliation or reclamation has when they run by looking at the VTSSTATS report. See Figure 7-8 on page 295 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, the 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.

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 that the Free Storage Threshold percentage can only be defined from the Library Manager console and not from the ETL Specialist panels.

4.3.8 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 3590 tape drives is set correctly. These specifications can only be made from the Library Manager console from the window shown in Figure 4-21.

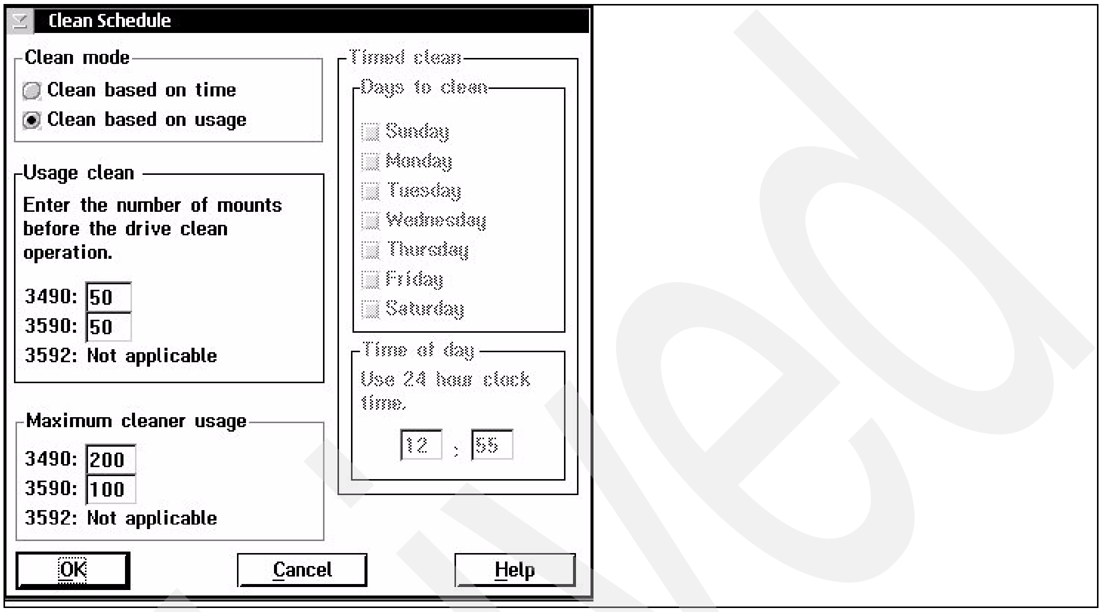

Figure 4-21 Library Manager Clean Schedule window

4.3.9 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.

4.4 Implementing Advanced Policy Management

Advanced Policy Management was introduced in 2002 and is supported on all VTS and Library Management LIC levels since then. If you are in doubt of whether APM is enabled on your VTS, you can check this on the ETL Specialist **VTS Status** window, which is displayed when you select **VTS Status** from the **Monitor VTSx** work items as shown in Figure 3-15 on page 97.

APM requires definitions at the Library Manager, which are typically done using the ETL Specialist, therefore we show the ETL Specialist panels required. You can also make these definitions from the 3494 or 3953 Library Manager console. For details on using the Library Manager panels, refer to the *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280, or to the *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473.

4.4.1 Defining stacked volume pool properties

For implementation of physical volume pooling, you must define the stacked volume pool properties. Physical volume pooling is required, for example, when you upgrade an existing Model B20 VTS to IBM 3592 tape drives. In addition, you can specify reclaim parameters for the physical cartridges of a specific pool.

To define or modify the pool properties, select **Modify Storage Pool Properties** from the **Administer VTS** work items. The panel shown in Figure 4-22 is displayed.

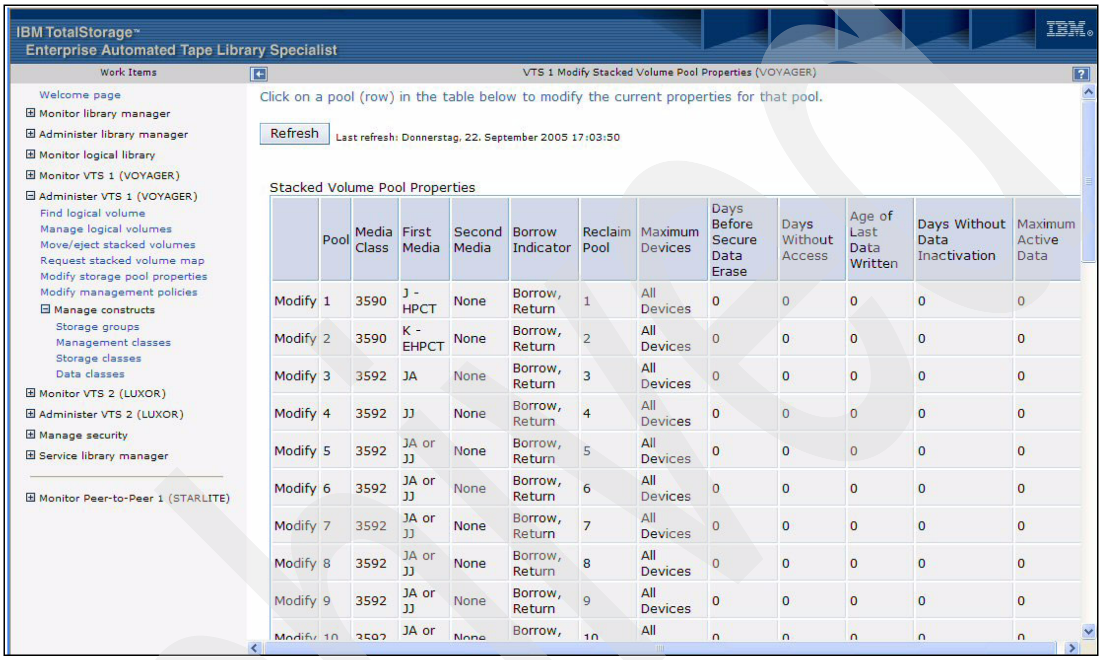

Figure 4-22 Stacked Volume Pool Properties window

The Stacked Volume Pool Properties window (Figure 4-22) allows you to modify pool properties. A list of the 32 general purpose stacked volume pools is produced, and for each of the pools, the following columns are displayed:

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.

Media 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.

First 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*.

Second Media
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*.

Borrow Indicator

Defines whether and how the pool is populated with scratch cartridges. *Borrow, Return* enables borrowing from the Common Scratch Pool When volumes become scratch, they are returned to the CSP. *Borrow, Keep* enables borrowing from the CSP. When volumes become scratch, they remain in the borrowing pool and are not returned. *No Borrow, Return* does not allow borrowing from the CSP. When any volumes become scratch, they are returned to the CSP (this setting can be used to empty pools). *No Borrow, Keep* does not allow borrowing from the CSP. When any volumes become scratch, they remain in the pool and are not returned.

Reclaim Pool

Lists the pool to which active logical volumes will be assigned when stacked volumes of this pool are reclaimed by the VTS. The stacked volume itself is treated according to the Borrow Indicator.

Maximum Devices

Defines the maximum number of physical tape drive that this pool can use for pre-migration.

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 3-15 on page 97.

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.

Each of the reclamation policies is taken into account independently of each other. A physical cartridge is reclaimed if it meets one of the reclamation criteria defined.

If you click the **Modify** button in the first column of the Modify Pool Properties panel, the screen shown in Figure 4-23 is displayed, from which you can update the pool properties. After you have completed your modifications, click **OK** to activate the new definitions or click **Cancel** to invalidate your changes. Both selections will return you to the Modify Pool Properties window.

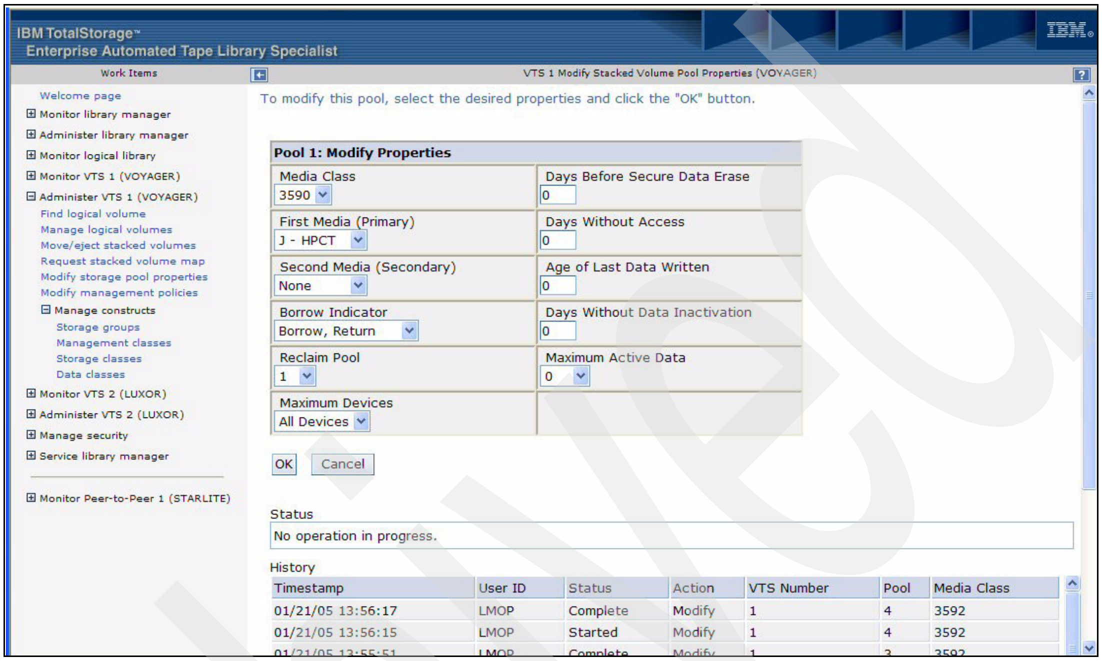

Figure 4-23 Modify Stacked Volume Pool Properties update window

The History table at the bottom of the panel shown in Figure 4-23 serves is displayed on all panels of the ETL and VTS Specialist that allow modification of parameters. It is also displayed on those panels of the Library Manager console that allow modifications of library or VTS settings, and it serves multiple purposes:

- 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.

As you can see, the last modification have been done from the Library Manager console as indicated by User ID LMOP.

4.4.2 Defining storage groups

The Storage Groups panel (Figure 4-24) is displayed when you select **Storage groups** from the **Manage constructs** pull-down of the **Administer VTS** work items.

When a logical volumes is assigned a storage group at the host, it will subsequently be written onto a physical cartridge of the pool specified for this storage group. This does not include logical volume pooling through multiple logical scratch pools, but it allows separation of groups of logical volumes onto different sets of physical cartridges.

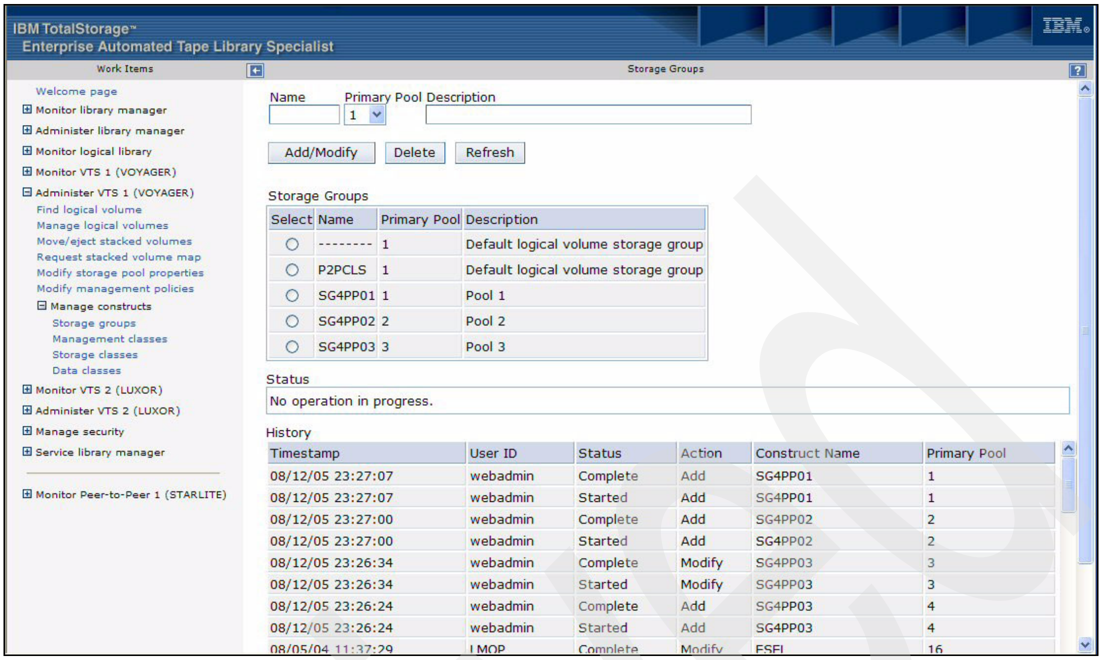

Figure 4-24 Manage Storage Groups panel

Through the Manage Storage Groups panel (Figure 4-24), you can:

- Add a storage group.
- Modify an existing storage group.
- Delete 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.
2. Select a **Primary Pool** and enter a short description.
3. Click 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.

If the host ACS routines assign a storage group that has not been defined, a new storage group will be defined at the Library Manager using the same specifications as the default storage group, which is usually pointing to Pool 1.

- You can update the default storage group to point to a different pool.
- The default storage group, identified by eight dashes (--------), cannot be deleted.
- Up to 256 storage groups, including the default, can be defined.

Tip: In an environment with multiple pools in use, we recommend not to assign a storage group to Pool 1. If you then encounter physical cartridges and logical volumes in Pool 1, you have an indication that your storage group definitions on the host and on the Library Manager might not be matching.

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, BatchMagic, in consultation with your IBM Storage Specialist.

4.4.3 Creating management classes

The Management Class panel (Figure 4-25) is displayed when you select **Management classes** from the **Manage constructs** pull-down of the **Administer VTS** work items.

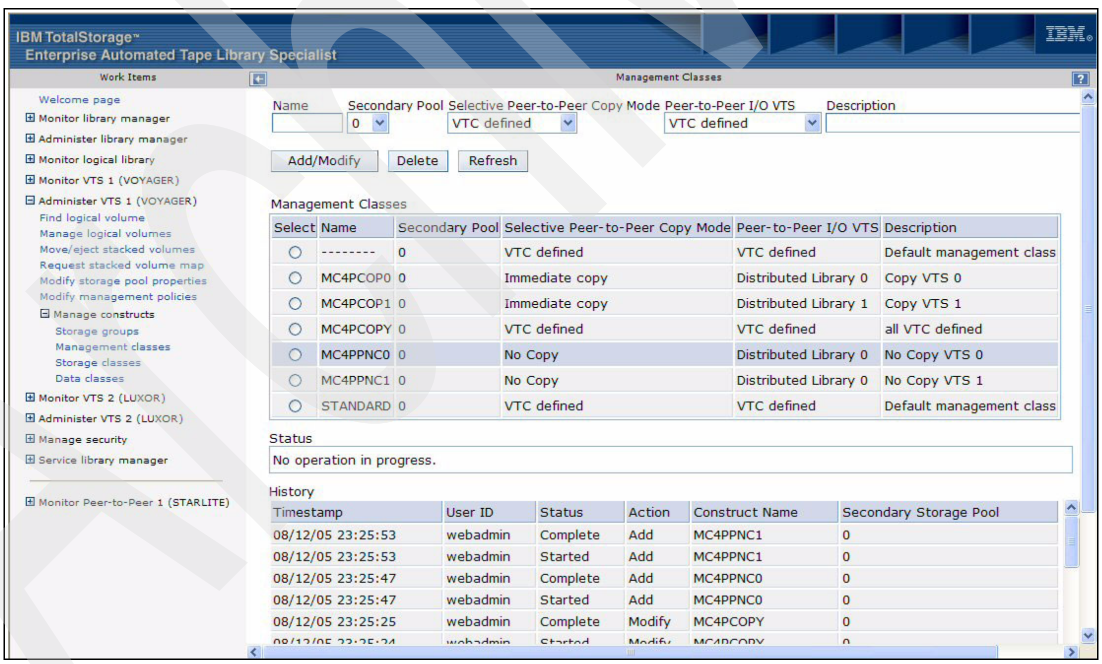

Figure 4-25 Manage Management Classes panel

The Manage Management Classes panel (Figure 4-25) allows you view and manage management classes. Through this panel, you can:

- Add a management class.
- Modify an existing management class.
- Delete a management class.

To add a management class, do the following steps:

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 the dual copy option. Select one of the following options:
   - Specify the secondary pool number (1-32). This will determine in which physical pool within the same VTS the copy of the logical volumes will reside.
   - If ‘00’ is selected, no secondary copy will be made. This is the recommended setting if you are using a Peer-to-Peer VTS.
3. Specify the selective Peer-to-Peer Mode:
   - Select **Immediate** to create the secondary copy at the same time as the first copy. This setting will over-ride the IBM System Service Representative (SSR) 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 SSR setting at the AX0/VTC setting as defined by the SSR.
   - Select **No Copy** if you want the logical volume only to be written to a single VTS of the PtP VTS configuration.
4. Specify the **Peer-to-Peer I/O VTS**:
   - Select **VTC Defined** to use the AX0/VTC setting as defined by the SSR.
   - 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

If the host ACS routines assign a management class that has not been defined, a new management class will be defined at the Library Manager using the same specifications as the default management classes.

- 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 BatchMagic which has been updated to accommodate the new functions.

4.4.4 Creating storage classes

The Storage Class panel (Figure 4-26) is displayed when you select **Storage classes** from the **Manage constructs** pull-down of the **Administer VTS** work items.

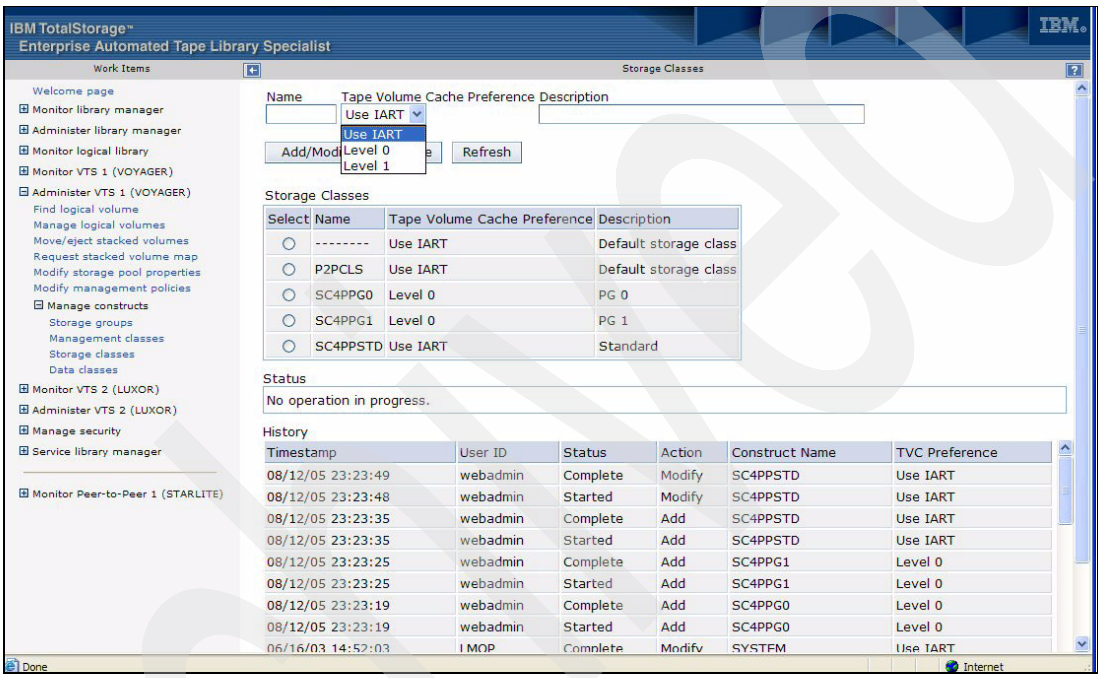

Figure 4-26 Storage Classes panel

The Storage Classes panel (Figure 4-26) allows you to view and manage storage classes. Through this panel, you can:

- Add a storage class.
- Modify an existing storage class.
- Delete a storage class.

To add a storage class, do the following steps:

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 within the storage class construct names.
2. Determine the **Tape Volume Cache Preference**. Three options are available:
   - **Use IART:**

     Instructs the VTS to honor host supplied Initial Access Response Time (IART) value defined in the storage class construct on the host.
   - **Level 0:**

     Instructs the VTS to remove volumes from the tape volume cache (TVC) as soon as they are copied to tape, but only if space is needed in the cache. Remove largest first.

Small volumes assigned to level 0 may also be removed from cache, if the VTS workload allows sufficient processor cycles for the background task to do so. This is done to avoid potential delays caused by the removal of a large amount of very small volumes during peak periods, should space be required in the TVC.

Level 1:

Instructs the 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.

If the host ACS routines assign a storage class that has not been defined, a new storage class will be defined at the Library Manager using the same specifications as the default storage classes.

- The default storage class, identified by eight dashes (--------), cannot be deleted.
- Up to 256 storage classes, including the default, can be defined.

4.4.5 Creating data classes

The Manage Data Classes panel (Figure 4-27) 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.

From a z/OS environment perspective, for SMS managed tape, the DFSMS data class defines:

- Media type parameters
- Recording technology parameters
- Compaction parameters
- Performance Scaling
- Segmentation

For the VTS, only the Media type, Recording technology, and Compaction parameters are used. Through the Media type parameter, the use of larger logical volume sizes can be requested.

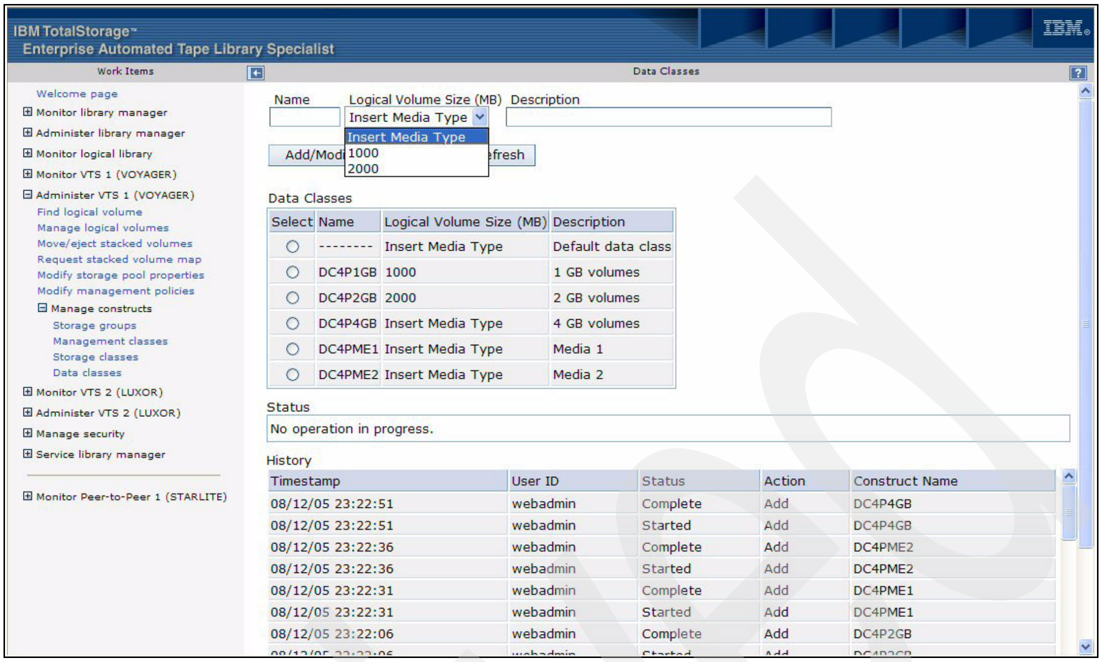

Figure 4-27 Manage Data Classes panel

To add a data class, do the following steps:

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 (MB) value:
   - Select **Insert Media Type** to use the media type that the logical volume had defined when it was inserted into the VTS. This can either be 400 MB for CST or 800 MB for ECCST.
   - Select **1000** to use logical volumes of 1000 MB of size.
   - Select **2000** to use logical volumes of 2000 MB of size.
   - Select **4000** to use logical volumes of 4000 MB of size.

   Larger logical volume support for 1000 MB volumes is included in the VTS R7.4 LIC level if APM is installed. Support for 2000 or 4000 MB logical volumes has to be enabled by the CE based on whether the hardware requirements are met.

   Each time, a logical volume is mounted in response on a scratch mount request, the logical volume size is determined based on the data class assigned. Re-writing a logical volume from loadpoint in after a specific mount request does however not change the current size of that logical volume.
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.

If the host ACS routines assign a data class that has not been defined, a new data class will be defined at the Library Manager using the same specifications as the default data classes.

- The default data class, identified by eight dashes (--------), cannot be deleted.
- Up to 256 data classes, including the default, can be defined.

Important: In a z/OS environment, data class constructs can be used even if system-managed tape is not implemented. The data class ACS routine is processed whether or not the allocation is subsequently going to a system-managed tape library.

4.4.6 Implementing Advanced Policy Management for non-z/OS hosts

Advanced Policy Management and its constructs are exploited only in DFSMS host environments where OAM has knowledge of the construct names and dynamically assigns and resets them. z/VM, z/VSE, Open Systems, and other hosts do not have knowledge of the construct names and therefore cannot change them. In addition, non-z/OS hosts use multiple LM categories for scratch volumes and therefore can use multiple logical scratch pools on the Library Manager as shown in Table 4-2.

Table 4-2 Scratch pools and Library Manager volume categories

Host SW	LM Scratch Categories	# scratch pools	LM Private Cat.
VM (+ VM/VSE)	X’0080’ - X’008F’	16	X’FFFF’
BTLS	X’0FF2’ - X’0FF8’, X’0FFF’	8	X’FFFF
Native VSE	X’00A0’ - X’00BF’	32	X’FFFF
Unassigned (can be used by Open Systems)	X’012F’ - X’0FF1’ X’0FF9’ - X’0FFE’ X’F00F’ - X’FEFF’		

As the hosts do not know about constructs, they ignore static construct assignment, and the assignment is kept even when the logical volume is returned to scratch. Static assignment means that at insert time of logical volumes, they are assigned construct names. Construct names can also be assigned later.

Select **Manage Logical Volumes** from the **Administer VTSx** work items to obtain the panel shown in Figure 4-2, which allows you to insert logical volumes or to change existing logical volumes.

Note: In a z/OS environment, OAM controls the construct assignment and will reset any static assignment made at the Library Manager level using the ETL Specialist or the LM console.

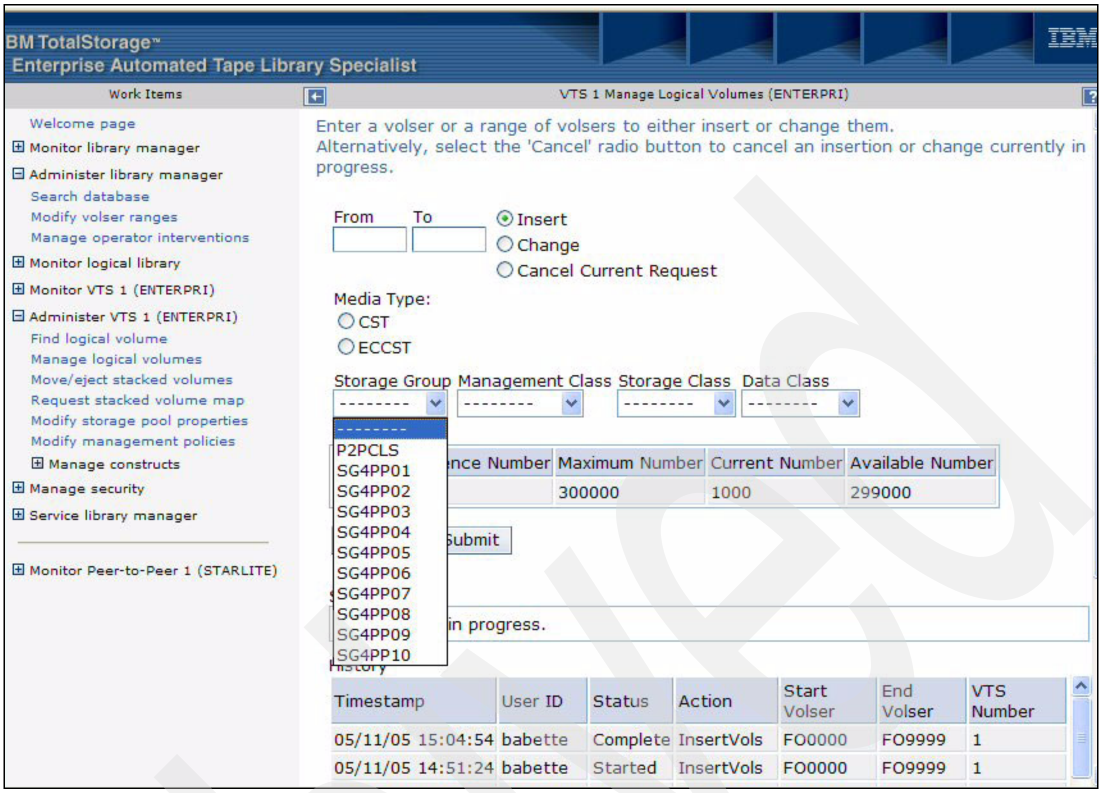

Figure 4-28 Manage Logical Volumes panel

To implement APM, first define your pools and constructs as described earlier in this chapter. Then insert your logical volumes in groups and assign the required construct names. We recommend that you define groups of logical volumes with the same construct names assigned and, during insert processing, direct them to different LM volume categories so that all volumes in one LM volume category have identical constructs assigned.

Host control is given through utilization of the appropriate scratch pool. By requesting a scratch mount from a specific scratch category, the actions defined for the constructs assigned to the logical volumes in this category will be executed at Rewind/Unload of the logical volume.

You should make sure that, when the logical volume is returned to scratch, it is returned to the correct LM volume category.

4.5 VTS software definition for zSeries

In this section, we describe the software definition considerations for implementing the VTS in z/OS, VM/ESA, and VSE/ESA environments. From a software point of view, the VTS is same as the IBM 3494 Tape Library with IBM 3490E tape drives.

The VTS must be defined as a new tape library with emulated IBM 3490E tape drives from the host system. Refer to *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632, for more information on defining this configuration.

The software support levels required to support an IBM 3494 or 3584 configuration with VTS are explained in 3.8, “Software requirements” on page 93.

4.5.1 DFSMS system-managed tape

To use the VTS in a z/OS environment, you can use either z/OS DFSMS or BTLS. During hardware installation of the VTS, the customer engineer assigns a five-character tape library sequence number to the VTS library. You use this sequence number to define a new tape library to SMS through the ISMF Library application. See 4.3.2, “Check the VTS library sequence number” on page 128.

DFSMS definition overview

In this section we describe the VTS in the DFSMS environment. You should read this in conjunction with the corresponding section in *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632.

When you install the VTS in a DFSMS environment, you must 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 the new installation of a real 3494 because the host recognizes the VTS as a 3494. One restriction on the name of this VTS library is that it must not begin with a ‘V’. To use the VTS, create a storage group and dedicate it to the VTS logical library, to allow the VTS logical tape library virtual drives to be allocated by the ACS routines.

To direct the host allocations to the VTS, you must define new storage groups that contain only VTS libraries. You should not intermix VTS and non-VTS libraries in the same storage group.

There is an overall definition procedure for a first tape library installation documented in this manual. For more information on this procedure, see “First time VTS library installation” on page 419. You should read this in conjunction with the corresponding section in *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632.

If you have a system-managed tape environment with more than one IBM tape library, the TCDB size will increase considerably if you have to define all 500,000 virtual volumes of one VTS to the Library Manager. A good solution is to split the TCDB, using a general VOLCAT for the real volumes and one or more user VOLCATs pointing to a unique first character identifier for each of the volsers of the logical volumes. In this way, you easily separate the volume catalog information of the VTS from the other libraries. An improvement in catalog performance is a further advantage of this recommendation. Refer to *z/OS Object Access Method Planning, Installation, and Storage Guide for Tape Libraries*, SC26-0427, for more information.

> **Important:** The TCDB is a key element in the operation of the VTS. We suggest that you back up the TCDB like any other catalogs.

The use and modification of OAM exits is applicable under the following conditions:

- You already have a DFSMS OAM environment.
- You have already installed the OAM installation exits and are using them.
- You are adding a VTS to an existing 3494.
- You are installing a VTS together with a new 3494.

The OAM installation exits are described in *z/OS Object Access Method Planning, Installation, and Storage Guide for Tape Libraries*, SC26-0427.

If you use the VTS in a single SMS environment, define a new storage group for each VTS. If you plan to use storage groups that include more than one VTS, refer to 3.6.6, “System managed tape considerations” on page 88 for a discussion of implications and possible drawbacks.

You must change your ACS routines to assign the newly defined storage groups in order to write data to the VTS. The modified ACS routine allocates data to a VTS-owned virtual volume on a VTS-owned virtual drive. You decide how to assign volumes for data that is written to the VTS, based on data set name, job name, or any of the other acceptable variables that DFSMS allows.

DFSMSrmm is an optional tape management software component of z/OS DFSMS. DFSMSrmm provides all installation-wide exits, CBRUXENT, CBRUXCUA, CBRUEXJC, and CBRUXVNL, used by the OAM library control system (LCS) to perform management of the IBM tape libraries. You can define and use DFSMSrmm for a VTS in the 3494 the same way as for the 3494 without VTS.

Refer to Appendix B, “VTS Import / Export” on page 369 for additional tape management system considerations when using the Advanced Function feature with the Import/Export function.

For details about DFSMSrmm, refer to the *DFSMSrmm Primer*, SG24-5983.

For other tape management systems, you need to ask the vendor of your tape management system (TMS) for their version of the exits that communicate with OAM.

4.5.2 Defining the VTS to the tape management system

Because the VTS emulates logical 3490E tape drives, the procedure for defining the VTS to your tape management system and other software products such as z/OS DFSMS, is the same as that for the 3494 with real IBM 3490E tape drives.

Additional considerations apply if you plan to use the Advanced Function feature with the Import/Export function. We discuss this topic in Appendix B, “VTS Import / Export” on page 369.

4.5.3 JES3 environments

Because the VTS emulates IBM 3490E tape drives, VTS virtual 3490E drives can be used in a JES3 environment. The VTS has the same limitations related to the JES3 Dynamic Support Programs (DSP) as native 3490E tape drives in a 3494. The VTS is not supported by DSPs.

A customer might want a JES3-managed tape library if it:

- Already has or plans to install a JES3 environment
- Wants to share the tape library among multiple hosts
- Does not have defined switchable tape-drive support in a Parallel Sysplex® environment
- Does not have a defined DFSMSplex environment

Following are the considerations to take into account in a JES3 environment:

UNITNAME: Because one tape library has only 3490E drives and the other tape library has only virtual 3490E drives, you must never specify a virtual tape drive as an alternate name for another tape library specific unit name. From an operational point of view, separate the tape allocations to the VTS from other real tape drives in the same 3494.

DFSMS: If you want to define the DFSMS-managed VTS drive, you define the JES3-managed drive as usual. To define the VTS as a JES3-managed drive, the following setups are required in the JES3 initialization deck:

1. Define JES3-managed drives in the VTS through DEVICE statements
2. Set JES3 device names through the SETNAME statement

3. Define which device names are subsets of other device names through the HWSNAME statement

Refer to *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632, for detailed information about the DEVICE, SETNAME, and HWSNAME statements to define in the JES3 initialization deck.

For a working example of how to define the JES3 initialization deck and esoteric names for a mixed configuration, native 3490E drives, native 3590 drives and virtual 3490E drives, refer to “JES3 sample initialization deck definition” on page 156.

JES3 can continue to manage tape devices outside an IBM 3494 or 3584 as long as those drives do not belong to the same generic or esoteric unit types as tape drives inside the tape library. For example, you *must not* have JES3-managed 3490E drives outside the library while there are physical IBM 3490E drives inside the library. However, you can have JES3-managed IBM 3480 or IBM 3490 base drives (non-3490E) outside the library while the drives inside the library are all IBM 3490E.

The IBM 3490 base drives (non-3490E) are identical to IBM 3480 drives as far as zSeries and JES3 are concerned. Therefore you cannot have IBM 3490 base drives (not 3490E) inside the 3494 Tape Library and JES3-managed IBM 3480 or 3490 base (not 3490E) drives outside it at the same time.

JES3 sample initialization deck definition

In this section we present a working example of a JES3 initialization deck and of an esoteric names definition in the HCD for a four frame plus one Model B20 VTS configuration.

> **Note:** You must not specify any other esoteric names for either physical or virtual devices inside a 3494 Tape Library or VTS. Therefore do not specify esoterics such as 3480, 3480X, SYS3480R, or 3490. Also, never use esoterics such as TAPE or CART.

Assuming the frames are in the following order:

- L10 Frame with two 3490E tape drives. Library ID = 12853 device numbers 520 to 521
- D14 Frame providing four native 3590 drives. device numbers 510 to 513
- S10 Frame cartridge storage only
- D12 Frame providing VTS physical drives
- Model B20 VTS frame providing VTS with 32 virtual drives. Library ID = 60286 device numbers 5A0 to 5AF and 5B0 to 5BF

These are the SETNAME rules that apply to defining a VTS for JES3:

- The complex-wide library name must be included in all statements.
- A library-specific name must be included for XTYPEs within the referenced library.
- The complex-wide device type name must be included for all XTYPEs of the corresponding device type in the complex.
- A library-specific device type name must be included for the XTYPE associated with the devices within the library.
- Do not specify generic (for example, 3480, 3480X, 3490) or esoteric (for example, TAPE, CART) unit names for library devices.

These are the general HWSNAME rules:

- Alternates must be a subset of or equal to the primary.
- List all names as the primary once, even if no alternates are possible.
- Do not list names as alternates which span multiple libraries unless the primary does.

For the implementation of a VTS, this means that:

- The complex-wide library name LDGW3495 must include all other LDGxxxxx names as alternates.
- The library-specific name LDGxxxxx must include all LDG names for the corresponding library as alternates.
- The complex-wide device type name LDG3490/LDG3490E must include all library-specific device type names.

Figure 4-29 shows the dependencies explained above:

<table>
<tr><td>510-513</td><td>520-521</td><td>5A0-5BF</td></tr>
<tr><td>LDB12853</td><td>LDE12853</td><td>LDE60286</td></tr>
<tr><td>LDG3591</td><td colspan="2">LDG3490E</td></tr>
<tr><td colspan="2">LDG12853</td><td>LDG60286</td></tr>
<tr><td colspan="3">LDGW3495</td></tr>
</table>

Figure 4-29 JES3 library SETNAMES

For the given configuration, the JES3 definition should be:

```
/* Devices (native 3590) 510 to 513
DEVICE,XTYPE=(ROBI,CA),XUNIT=(510,SY1,,ON)
DEVICE,XTYPE=(ROBI,CA),XUNIT=(511,SY1,,ON)
DEVICE,XTYPE=(ROBI,CA),XUNIT=(512,SY1,,ON)
DEVICE,XTYPE=(ROBI,CA),XUNIT=(513,SY1,,ON)
/*
/* Devices (native 3490E) 520 and 521
DEVICE,XTYPE=(ROB3490,CA),XUNIT=(520,SY1,,ON)
DEVICE,XTYPE=(ROB3490,CA),XUNIT=(521,SY1,,ON)
/*
/* Devices (virtual 3490E) 5A0 to 5AF and 5B0 to 5BF
DEVICE,XTYPE=(ATLVTS,CA),XUNIT=(5A0,SY1,,ON)
DEVICE,XTYPE=(ATLVTS,CA),XUNIT=(5A1,SY1,,ON)
                :                 :
                :                 :
                :                 :
DEVICE,XTYPE=(ATLVTS,CA),XUNIT=(5BE,SY1,,ON)
DEVICE,XTYPE=(ATLVTS,CA),XUNIT=(5BF,SY1,,ON)
```

```
/*
/* This SETNAME is for native 3590 drives
SETNAME,XTYPE=ROBI,NAMES=(LDGW3495,LDG12853,LDG3591,LDG12583)
/*
/* This SETNAME is for native 3490E drives
SETNAME,XTYPE=ROB3490,NAMES=(LDGW3495,LDG12853,LDG3490E,LDE12853)
/*
/* This SETNAME is for virtual 3490E drives
SETNAME,XTYPE=ATLVTS,NAMES=(LDGW3495,LDG60286,LDG3490E,LDE60286)
/*
/* Sysplex Wide Library Name
HWSNAME,TYPE=(LDGW3495,LDG12853,LDG60286,
              LDG3591,LDG3490E,LDB12853,LDE12853,LDE60286)
/* Library Specific Library Name
HWSNAME,TYPE=(LDG12853,LDB12853,LDE12853,LDG3591)
HWSNAME,TYPE=(LDG60286,LDE60286)
/* Sysplex Wide Device Name
HWSNAME,TYPE=(LDG3591,LDB12853)
HWSNAME,TYPE=(LDG3490E,LDE12853,LDG60286,LDE60286)
/* Library Specific Device Name
HWSNAME,TYPE=(LDB12853,LDG3591)
HWSNAME,TYPE=(LDE12853)
HWSNAME,TYPE=(LDE60286,LDG60286)
/*
```

The esoteric names definition in the HCD should be:

```
 / Esoteric  VIO    Token        State
 _ LDGW3495  No     ____         510-513, 520-521 5A0-5BF
 _ LDG12853  No     ____         510-513, 520-521
 _ LDG60286  No     ____         5A0-5BF
 _ LDG3591   No     ____         510-513
 _ LDG3490E  No     ____         520-521, 5A0-5BF
 _ LDB12853  No     ____         510-513
 _ LDE12853  No     ____         520-521
 _ LDE60286  No     ____         5A0-5BF
```

These are the HCD Library ID and Port ID values:

```
520-521 Library ID = 12853,  Port ID = 01
510-513 Library ID = 12853,  Port ID = 02
5A0-5AF Library ID = 60286,  Port ID = 01
5B0-5BF Library-ID = 60286,  Port-ID = 02
```

Basic Tape Library Support

If you implement the VTS under BTLS, you can define the tape drives through the hardware HCD dialog, or you can use IOCP and MVSCP. You also have to define a new logical tape library to BTLS that contains all virtual drives associated with this VTS logical library. To direct tape allocations to the VTS, define a new esoteric unit name and use it in the DD statement for the newly allocated tape data set. For BTLS-managed libraries, HCD is not required because BTLS does not require the LIBRARY=YES definition.

Because of the tape library definitions, BTLS knows that the drive allocated resides in a VTS tape library. For BTLS implementation details, refer to the *Basic Tape Library Support User's Guide and Reference*, SC26-7016. For additional practical information, refer to *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632.

> **Note:** There is no support available in BTLS for APM or the Advanced Function feature with the Import/Export operations.

JES3 provides BTLS support as follows:

- Because the tape subsystems in the 3494 must not be defined in the JES3 initialization deck, they are not managed by JES3.
- zSeries performs all IBM TotalStorage Enterprise Automated Tape Library tape drive allocations with the allocation assist function of BTLS. BTLS is responsible for communication with the Library Manager in the 3494.
- BTLS functions in the JES3 environment are identical to the BTLS functions in the JES2 environment.
- JES3 tape DSPs or JES3 tape commands for tape drives inside an 3494 are not supported.

4.5.4 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 looks like a separate library to the host. A VTS can be shared by multiple zSeries, VM, and VSE guest systems in the same way that a physical tape library can be shared.

> **Note:** If you are sharing a 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 3.8, "Software requirements" on page 93.

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.

4.5.5 Partitioning the VTS between multiple hosts

Partitioning is the solution of choice 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 test 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 z/OS DFSMS Environment" in the *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409. Also see 3.4.2, "Sharing a VTS subsystem" on page 72.

4.6 VM/VSE considerations

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)
- Native VSE/ESA

4.6.1 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 TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632 and *Lights Out! Advanced Tape Automation Using VM/ESA*, GG24-4347.

Figure 4-30 shows the VM/ESA and z/VM native support for the IBM VTS.

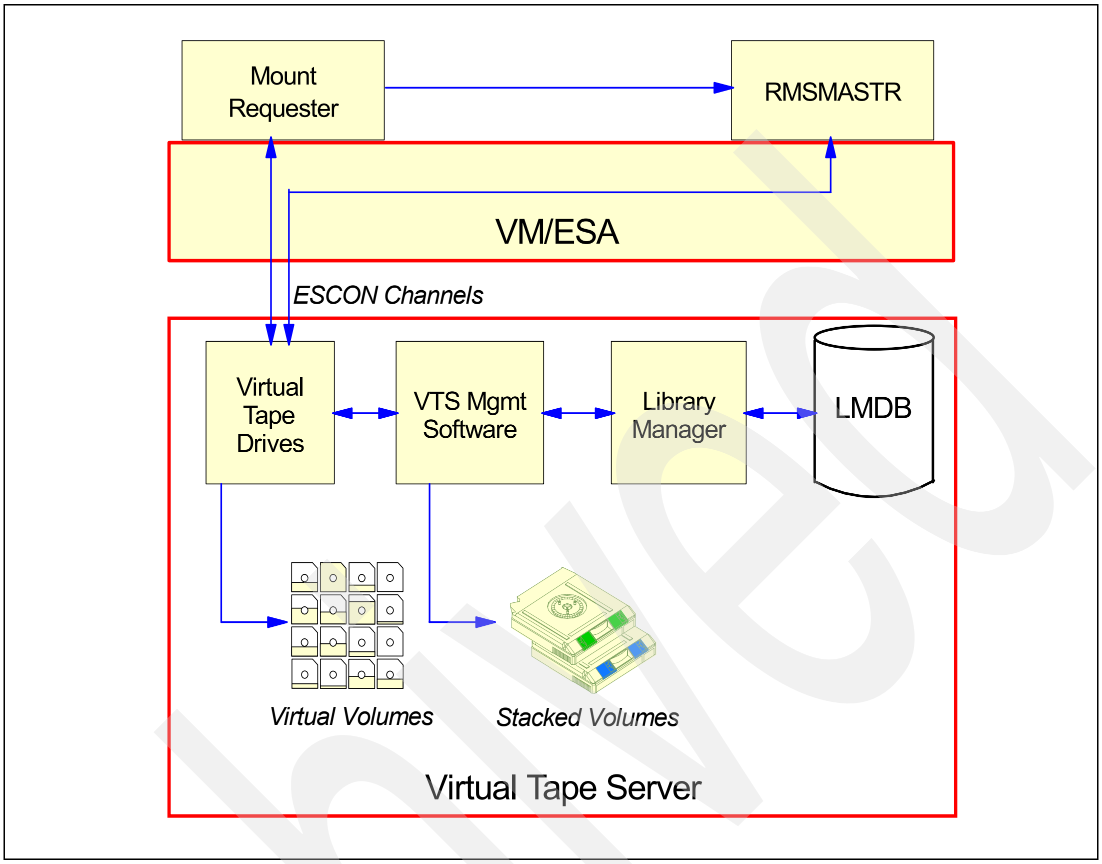

Figure 4-30 VTS in a native VM/ESA environment using DFSMS/VM

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 tape 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, additional considerations apply. Refer to the *Guide to Sharing and Partitioning IBM Tape Library Dataservers,* SG24-4409, for further information.

4.6.2 VM/ESA and z/VM guest support

In the following sections, we discuss two host environments that allow you to use an IBM 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, “DFSMS system-managed tape” on page 154 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 guests

VSE/ESA must run as a guest system under VM/ESA to use a VTS library. 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 4.6.3, “VSE/ESA as a VM/ESA guest using a VSE guest server (VGS)” on page 163. For further information, 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.

4.6.3 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 4-31 you can see the flow and connections of a IBM VTS in a VSE/ESA environment under VM.

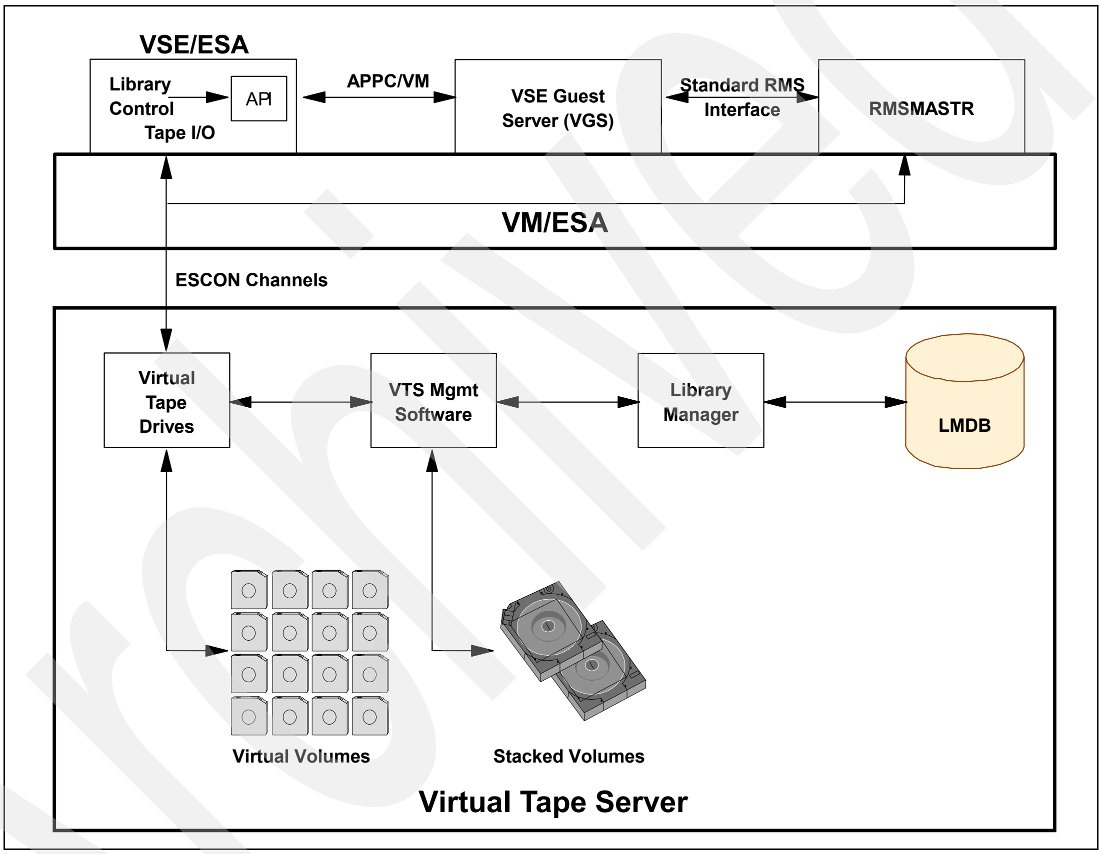

Figure 4-31 VTS in a VSE/ESA environment as a VM guest

Tape management

As with the VM/ESA native environment (see 4.6.1, "VM/ESA and z/VM native support using DFSMS/VM" on page 160), 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 4-32 shows the case of CA-DYNAM/T VSE.

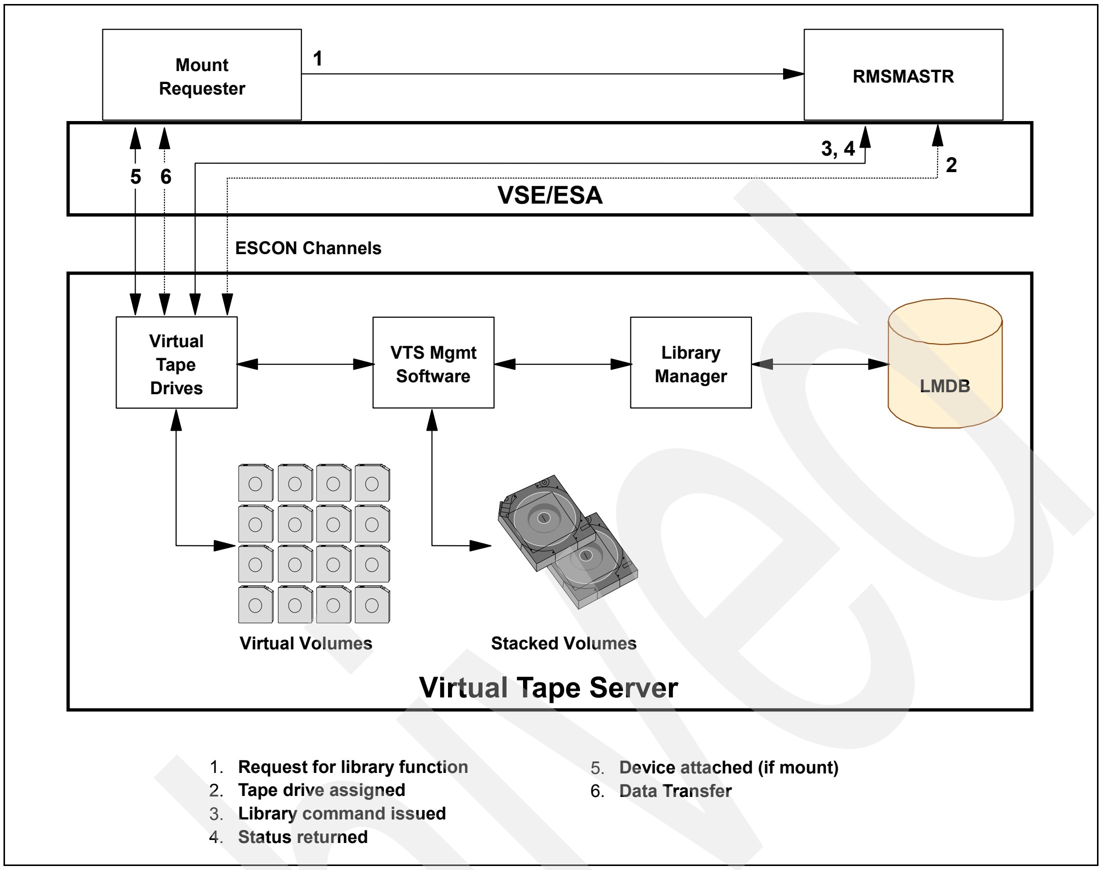

Figure 4-32 VTS in a VSE/ESA environment as a VM guest (no VGS)

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.

4.6.4 Native z/VSE

Native support is provided at z/VSE Version 3.1 which supports all IBM 3592-J1A configurations (J70 and VTS) without APARs in all automation offerings including IBM 3494 and IBM 3584/3953 configurations.

> **Restrictions for VM and VSE:**
>
> - No support is available for the VTS Advanced Function Feature or the Advanced Policy Management Features.
> - Support for the IBM TotalStorage Peer-to-Peer Virtual Tape Server is available in an environment in which a z/OS host with full support of the Peer-to-Peer VTS is also attached to the IBM 3494 or 3584/3953 configuration.

z/VSE supports the IBM 3584/3953 natively through its Tape Library Support (TLS). In addition to the old tape library support, a function has been added to allow the tape library to be supported via the /390 channel command interface commands thus eliminating any XPCC/APPC communication protocol as required with the old interface. The external interface (LIBSERV JCL and LIBSERV macro) remain unchanged.

Define library support

First you must define which type of support you are using by specifying the SYS ATL statement. You can define:

TLS Tape Library Support, which provides full support

VSE LCDD, which does not support 3592 and therefore not the IBM 3584

VM VM Guest Support

For native support under VSE, select TLS.

Define tape libraries

You must define your tape library or libraries. This is done though a batch job as shown in Example 4-6.

Example 4-6 Define libraries

```
* $$ JOB JNM=TLSDEF,CLASS=0,DISP=D
* $$ LST CLASS=A
// JOB TLSDEF
// EXEC LIBR,PARM='MSHP'
   ACCESS S=IJSYSRS.SYSLIB
   CATALOG TLSDEF.PROC REPLACE=YES
    LIBRARY_ID TAPELIB1 SCRDEF=SCRATCH00 INSERT=SCRATCH00            --- default library
    LIBRARY_ID TAPELIB2                * SECOND LIB DEF
    DEVICE_LIST TAPELIB1 460:463       *  DRIVES  460 TO 463
    DEVICE_LIST TAPELIB2 580:582       *  DRIVES  580 TO 582
    QUERY_INV_LISTS LIB=TLSINV         * MASTER INVENTORY FILES
    MANAGE_INV_LISTS LIB=TLSMAN        * MANAGE FROM MASTER
 /+
```

LIBSERV

The communication from the host to the Library Manager is through the LIBSERV JCL or macro interface. Example 4-7 shows a sample job using LIBSERV to mount volume 123456 for write on device address 480 and, in a second step, to release the drive again.

Example 4-7 Sample LIBSERV JCL

```
$$ JOB JNM=BACKUP,CLASS=0,DISP=D
$$ JOB BACKUP
// ASSGN SYS005,480
// LIBSERV MOUNT,UNIT=480,VOL=123456/W
// EXEC LIBR
BACKUP S=IJSYSRS.SYSLIB TAPE=480
/*
// LIBSERV RELEASE,UNIT=480
/&
$$ EOJ
```

LIBSERV provides the following functions:

Query all libraries for a volume: LIBSERV AQUERY,VOL=123456

Mount from category: LIBSERV CMOUNT,UNIT=480,SRCCAT=SCRATCH01

Query count of volumes: LIBSERV CQUERY,LIB=TAPELIB1,SRCCAT=SCRATCH01

Query device : LIBSERV DQUERY,UNIT=480

Query inventory of library: LIBSERV IQUERY,LIB=TAPELIB1,SRCCAT=SCRATCH01

Query library: LIBSERV LQUERY,LIB=TAPELIB1

Manage inventory: LIBSERV MINVENT,MEMNAME=ALL,TGTCAT=SCRATCH01

Change category: LIBSERV SETVCAT,VOL=123456,TGTCAT=SCRATCH01

Query library for a volume: LIBSERV SQUERY,VOL=123456,LIB=TAPELIB1

For additional information, refer to the *z/VSE System Administration Guide*, SC33-8224, and to the *z/VSE System Macros Reference*, SC33-8230.

4.7 Transaction processing facility (TPF) considerations

This section describes the support for an IBM VTS in a native transaction processing facility environment. The TPF control program and several new and modified TPF E-type programs support the IBM VTS. The support is limited to a command-based interface.

Here are some points to consider when implementing an IBM VTS in a TPF environment:

- Reserving a tape category does not prevent another host from using that category. It is the user's responsibility to monitor the use of reserved categories.
- Automatic insert processing is not provided within TPF.
- There is currently no IBM tape management system for TPF.

Library Manager interface

The TPF operator's only interface to the IBM VTS is a new TPF functional message ZTPLF. The various ZTPLF functions provided allow the operator to manipulate the tapes in the library as operational procedures require. These functions include Reserve, Release, Move, Query, Load, Unload and Fill. Refer to the *IBM TotalStorage Enterprise Tape: A Practical Guide,* SG24-4632 for additional information on these ZTPLF functions

Control data sets

The TPF host does not keep a record of the volumes in the IBM VTS Tape Library or manage the tape volumes within it. You can use the QUERY command to obtain information about the tape volumes held in the IBM 3494 Tape Library.

SIM and MIM presentation

SIMs and MIMs report hardware-related problems to the operating system. Refer to the *Statistical Analysis and Reporting System User Guide*, which you can access on the Web at:

http://www-1.ibm.com/support/docview.wss?uid=ssg1S7000247

SIM and MIM are presented in TPF through CEFR0354I, CEFR0355W, CEFR0356W, CEFR0357E, CEFR0347W, CDFR0348W, and CDFR0349E messages, as well as EREP reports.

4.8 VTS Open System attachment

The IBM TotalStorage VTS supports a variety of IBM and Open Systems hosts including IBM pSeries, RS/6000, Sun, HP, and Intel-compatible servers running Microsoft Windows NT and Windows 2000 when attached to an IBM 3494 Tape Library.

Restriction: When the VTS is installed in an IBM 3584/3953 configuration, Open Systems hosts cannot attach to the VTS because they are not supported by the IBM 3953 Library Manager.

From a customer viewpoint, or from an operating system perspective, the VTS appears as a 3494 with 32, 64, 128 or 256 3490E tape drives and up to 250,000 cartridges. Both the Model B10 VTS and Model B20 VTS provide access up to 64 virtual tape drives through the SCSI Attachment feature (FC3422). Maximum number of virtual tape drives available to SCSI hosts is dependent on the VTS models and number of FC3422. These virtual tape drives will remain accessible to ESCON hosts, if available.

The VTS can be shared by zSeries and Open Systems hosts to eliminate redundant tape drive attachments, reduce operational costs and allow increased data automation. Dedicated use of a VTS among multiple SCSI hosts is also possible to configure.

Virtual volumes are sharable between zSeries and other platforms provided software support is in place to handle the ASCII to EBCDIC conversion of data and that the owning system is clearly established. No secondary system should be allowed to use these volumes for output unless the owning systems metadata is kept up to date.

4.8.1 SAN attachment configuration

The VTS does not support native Fibre Channel, but it supports the Fibre Channel hosts through IBM 2108 SAN Data Gateway, which interconnects Fibre Channel and SCSI. The VTS can be connected to a Storage Area Network (SAN) with this attachment configuration. An IBM 2109 Fibre Channel Switch can also be connected to provide redundant path or configuration flexibility. Table 4-3 describes the supported servers for SAN configuration for the VTS.

Table 4-3 VTS SAN supported servers

Server	IBM 2108 SAN Data Gateway	IBM 2109 SAN Switch
pSeries, RS/6000	Yes	Yes
Sun	Yes	Yes
HP	No	No
Windows NT	Yes	Yes
Windows 2000	Yes	Yes

For additional information on technical requirements for supported servers, operating systems, Fibre Channel host bus adapters, and configuration, refer to these Web sites:

http://www.storage.ibm.com/hardsoft/products/sangateway/gatewayspec.htm
http://www.storage.ibm.com/hardsoft/products/fcswitch/fcswitch.htm

Figure 4-33 shows how a VTS can be attached to the various Open Systems with various configurations. For software requirements and supported adapters, refer to 3.8, "Software requirements" on page 93.

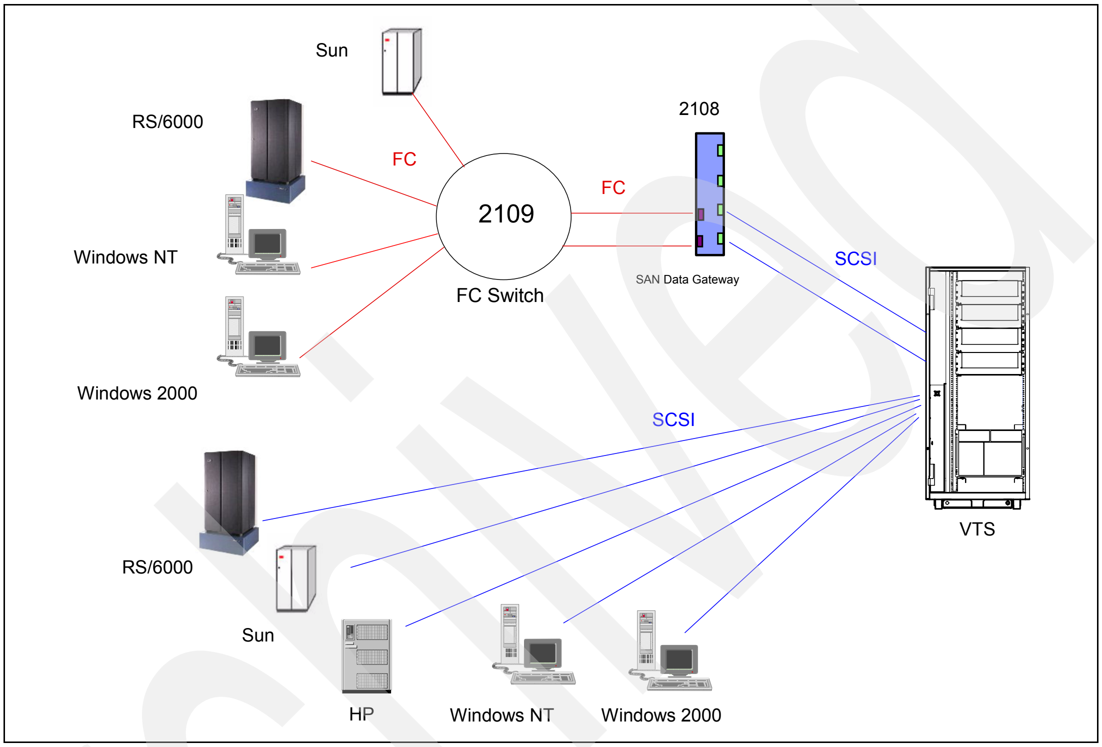

Figure 4-33 VTS attachment to Open Systems

For detail implementation of IBM 2108 SAN Data Gateway and IBM 2109 Fibre Channel Switch, refer to *IBM Tape Solutions for Storage Area Networks and FICON*, SG24-5474.

4.8.2 IBM TotalStorage VTS Open Systems design

VTS virtual drives are emulated 3490E tape drives. This emulation is used for both ESCON and SCSI attachments. Behavior on the SCSI bus is based on that of a 3490E Model C2A. Application programs written to support library attached 3490E tape drives should operate without change.

The maximum number of virtual drives available with a single VTS remains unchanged with the installation of the SCSI Host Attachment feature: that is, FC3422 does not add virtual devices. The number of tape drives which may be addressed as SCSI devices depends upon the number of FC3422 on a particular VTS Model. Up to 256 virtual drives can be defined and used in a VTS. Any virtual drive can be assigned to any host or may be switched between hosts based on usage requirements. A maximum of 64 virtual drives are available to the SCSI hosts of a single VTS with four SCSI Host Attachment features. One SCSI Host Attachment feature supports a maximum of 16 virtual drives. This maximum of 16 virtual drives per SCSI Host Attachment feature can be accessed over one or both SCSI interfaces of that SCSI Host Attachment feature.

The SCSI Host Attachment is a feature (FC3422) for a VTS which permits Open Systems host attachment through industry-standard SCSI busses. Each SCSI Host Attachment (FC3422) provides two SCSI bus attachments and includes data compression compatible with ESCON High Performance option (EHPO).

Each feature supports a maximum of 16 virtual devices. The feature supports attachment to the SCSI-2 standard interface on selected pSeries, RS/6000, Sun, HP, or Windows servers. Inclusion of the SCSI Host Attachment feature requires the VTS Open Systems Device Drivers (FC9201) as a prerequisite. Either an ESCON host system must be attached or the *Call Home* capability must be enabled so that service alerts can be sent.

The 3490E drive commands, status and data are passed on the SCSI bus and communications between the host and Library Manager are performed over a separate LAN or serial connection, bypassing the VTS. The conceptual drawing in Figure 4-34 shows the components and connections used by this feature.

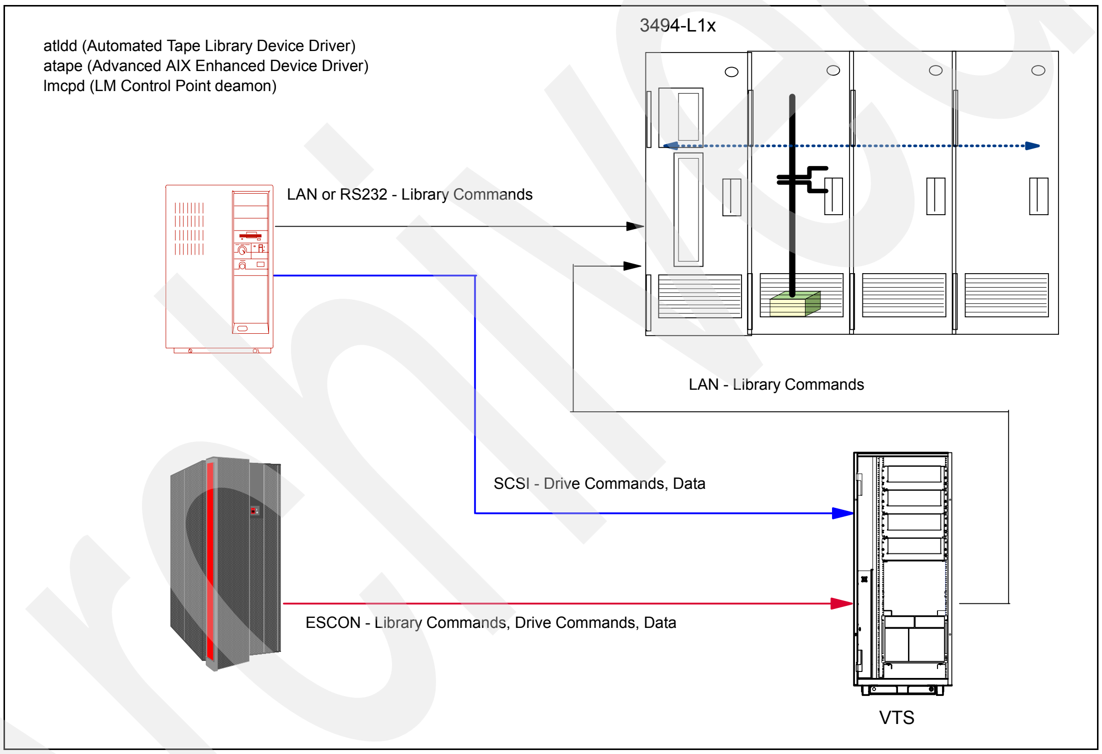

Figure 4-34 VTS with mixed ESCON and SCSI host attachments

4.8.3 Open Systems implementation

The VTS implementation for AIX, Solaris, HP-UX, Windows NT, and Windows 2000 are similar to implementing real IBM 3490-C2A Tape Subsystems with SCSI interfaces into an 3494. Detailed information about implementing a VTS in each platform is provided in the *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632.

For recommendations on using Tivoli Storage Manager, see 5.3, “Tivoli Storage Manager” on page 188 and the redbook, *Tivoli Storage Manager Version 4.2: Technical Guide,* SG24-6277.

VTS Open System Systems Device Drivers (FC9201) provides the tape drive device driver and library device driver. Open Systems device drivers are described in *SCSI Tape Drive, Medium Changer, and Library Device Drivers User's Guide,* GC35-0154. Any changes for VTS support will be documented in a revision of this publication.

With the addition of Advanced Policy Management (FC4001-4004), you have the ability to deploy volume affinity. This of course means that if you are sharing a VTS with multiple, unlike hosts (zSeries, Sun, Windows 2000) you may wish to segregate your logical volumes by host. In this way, the only stacked media used for each host will be that defined outboard of each host.

4.8.4 SCSI bus

The SCSI Host Attachment feature (FC3422) supports the SCSI-2 (Fast/Wide Differential) 20 MB/s standard interface, as well as the 40 MB/s Differential Wide/Ultra™ SCSI interface. The maximum cable length for a SCSI bus having a VTS SCSI connection is 20 meters. The SCSI Host Attachment feature on the VTS controller must be terminated externally by using the provided terminators. Adapter cards are fitted with connection blocks to permit daisy-chain connections or external terminators.

A daisy-chain connection of multiple VTS SCSI adapters requires reconfiguration of the adapter nodes to provide unique starting target IDs. This is because all busses have a default configuration with a starting target ID of 8. For recommended definitions and to avoid conflicts, see 4.8.5, “SCSI target addressing” on page 171.

Consider the following rules to configure virtual tape drives for the SCSI host:

- Each SCSI adapter provides two SCSI bus.
- Each SCSI bus can be configured to address any even number of up to 16 contiguous logical devices, starting with an even device number.
- Each SCSI adapter is restricted to accessing no more than 16 logical devices. The sum of the logical devices configured on both SCSI buses on an adapter card cannot exceed 16.

Figure 4-35 shows how SCSI buses are connected to virtual devices.

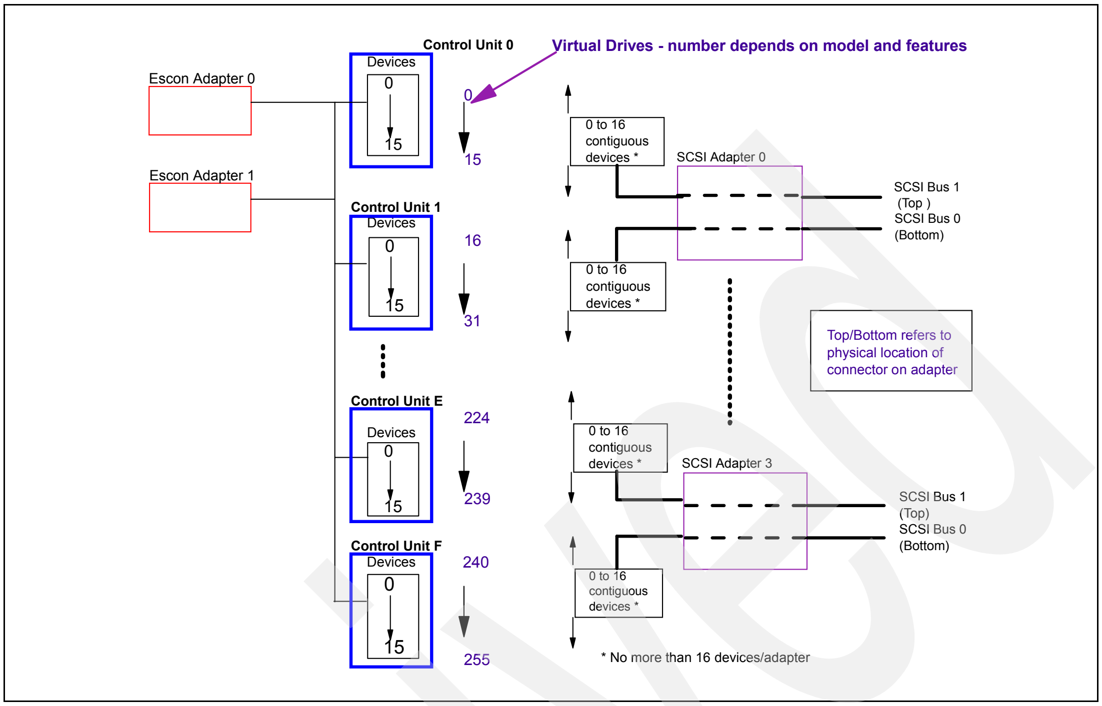

Figure 4-35 SCSI bus connections to virtual devices

4.8.5 SCSI target addressing

The SCSI target addresses to be used on each bus are configurable through a service panel by an IBM service representative upon installation of the VTS with a SCSI Host Attachment feature or installation of the SCSI Host Attachment feature in the field. Each SCSI Host Attachment feature provides two SCSI interfaces. One, two or four adapter cards may be installed in the VTS controller.

Each SCSI Host Attachment feature is restricted to accessing no more than 16 virtual drives. Each bus may be configured to address any even number of up to 16 contiguous virtual tape drives in a given control unit, starting with an even device number. Note that internally, virtual tape drives are numbered from 0 to 255. The target addresses for each bus of a SCSI Host Attachment feature must be contiguous (see Figure 4-36).

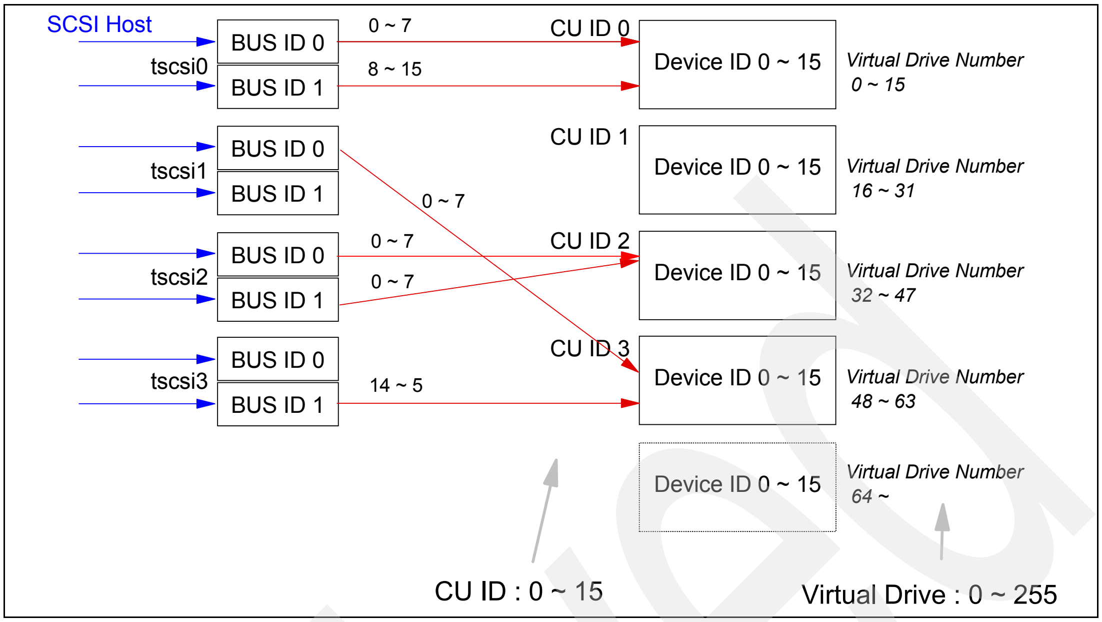

Figure 4-36 SCSI target configuration

With the example above, if we set the SCSI adapter as 3, bus 1 base SCSI target address as 4, starting logical control unit as 3, starting device as 14, then eight devices are attached as shown in Table 4-4.

Table 4-4 Virtual devices attached to SCSI3 bus1

SCSI adapter	Bus on adapter	CU ID	Device number	Virtual drive number	SCSI target address	LUN
3	1	3	14	62	4	0
3	1	3	15	63	4	1
3	1	4	0	64	5	0
3	1	4	1	65	5	1
3	1	4	2	66	6	0
3	1	4	3	67	6	1
3	1	4	4	68	7	0
3	1	4	5	69	7	1

The division of the virtual tape drives between virtual control units is done for ESCON addressing purposes. Therefore, device 0 for one virtual control unit is contiguous to device 15 of the previous (lower-numbered) virtual control unit. There may be an overlap between the virtual drives accessed on each SCSI bus and/or SCSI adapter.

Hosts are expected to use the SCSI `Reserve Unit` and `Release Unit` commands to prevent use of a single virtual tape drive by multiple initiators or contention with ESCON-attached hosts. The VTS virtual drives can be shared between ESCON and SCSI hosts if the ESCON `assign/unassign` and SCSI `reserve/release` functions are properly implemented. To keep it simple, we recommend to dedicate virtual drives to SCSI hosts if possible.

The adapter cards use one SCSI target address for every two virtual drives configured, starting with the specified base address for the bus. LUN (logical unit number) 0 and 1 are used for virtual drives at each target address. No target address on the adapter may exceed 15, even in the maximum configuration with 16 virtual drives using SCSI target address 8 to 15. Each bus is independently configurable through service terminal. Figure 4-37 shows emulated 3490E virtual tape drives, being the default configuration for one SCSI attachment to the VTS.

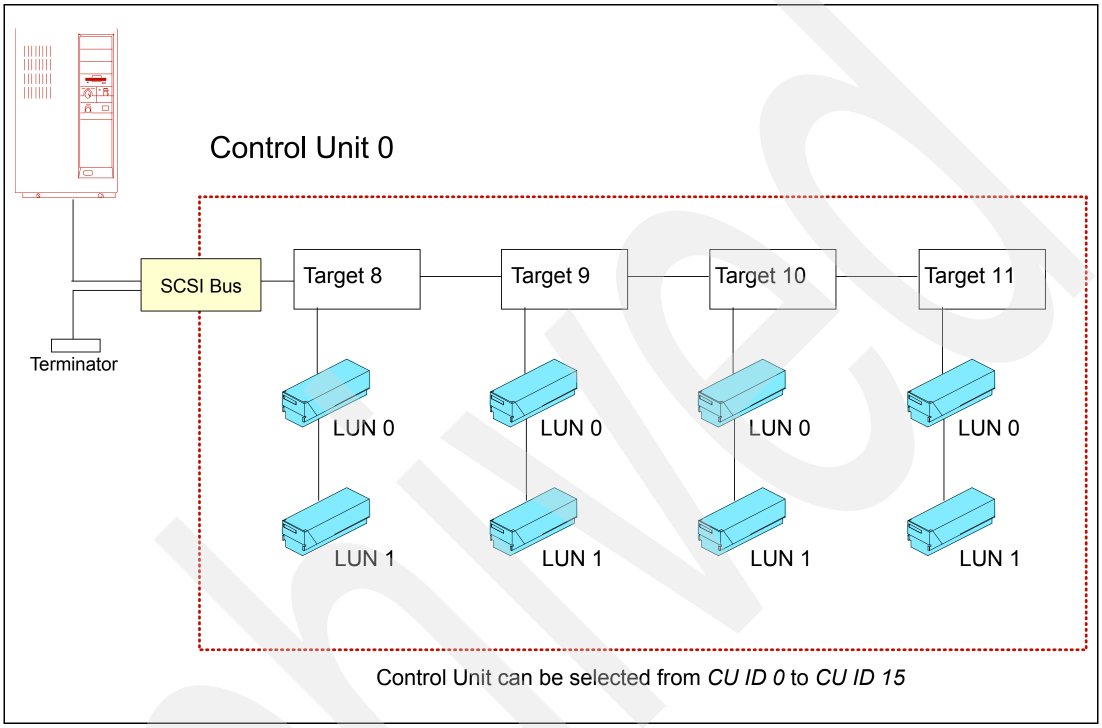

Figure 4-37 Emulated eight 3490E to SCSI host

The initial SCSI device addressing for each bus of a SCSI Host Attachment feature is:

- Eight virtual tape drives to a SCSI bus
- SCSI target address ranges are 8, 9, 10, 11

Table 4-5 shows the initial assignments of target addresses for the SCSI Host Attachment feature. These initial assignments may be changed from a service panel by an SSR. The default setting, for all SCSI virtual tape drives, is *non-shared*. If you would like to share the same virtual drives between SCSI interface 0 and 1 or share between a SCSI host and a zSeries, they must be configured by the SSR.

Table 4-5 SCSI target addresses initial setting

SCSI adapter	Bus on adapter	CU ID	Device number range	Virtual drive number range	SCSI target address range	LUN range
0	0	0	0-7	0-7	8-11	0-1
0	1	0	8-15	8-15	8-11	0-1
1	0	1	0-7	16-23	8-11	0-1

SCSI adapter	Bus on adapter	CU ID	Device number range	Virtual drive number range	SCSI target address range	LUN range
1	1	1	8-15	24-31	8-11	0-1
2	0	2	0-7	32-39	8-11	0-1
2	1	2	8-15	40-47	8-11	0-1
3	0	3	0-7	48-55	8-11	0-1
3	1	3	8-15	56-63	8-11	0-1

Multiple SCSI attached VTS subsystems on the same SCSI bus (daisy chaining) will require reassignment of the initial target address of one or more subsystems to avoid conflicts. Figure 4-38 shows daisy-chained control units providing 16 virtual tape drives to a SCSI host which has one SCSI adapter. In this configuration, two control units may be within one VTS or one control unit in one VTS, the other control unit in another VTS. In this figure, the tape drives in the second control unit changed their starting target address as 12 to avoid conflicts. Daisy chaining of multiple SCSI buses requires FC5004 to connect the two VTS SCSI ports.

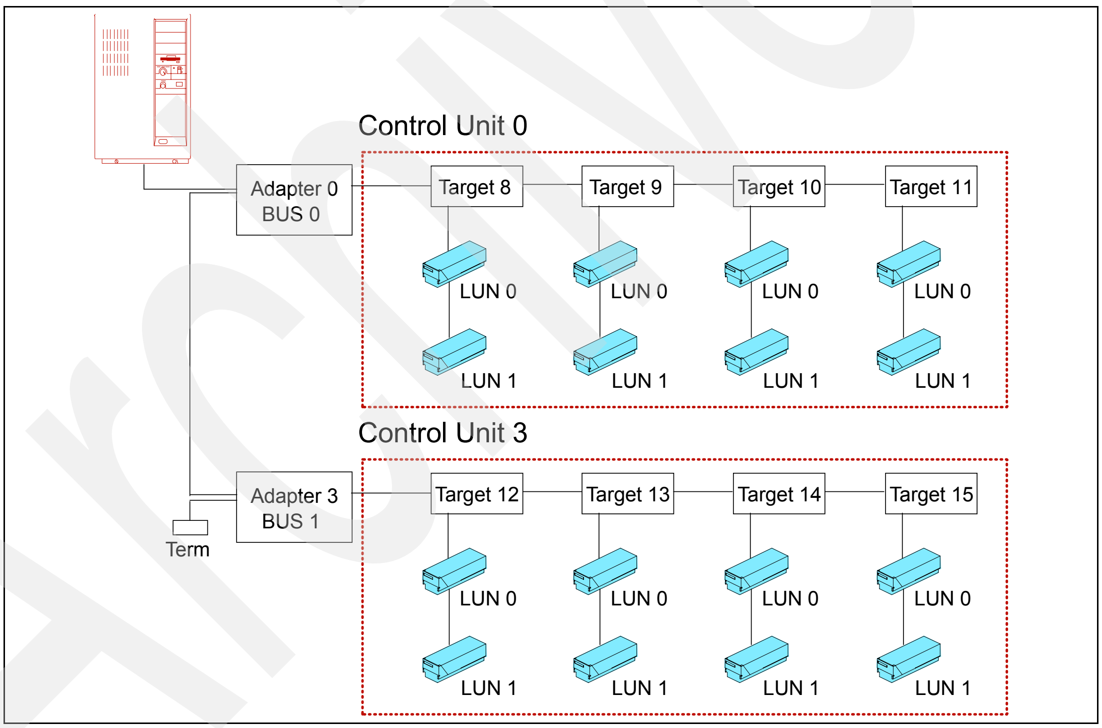

Figure 4-38 Daisy-chained Control Units providing 16 virtual tape drives to a SCSI Bus

In an environment where the total available virtual tape drives is 64, if the first 48 virtual devices are assigned to an ESCON host, a possible definition for 16 SCSI is shown in Table 4-6.

Table 4-6 SCSI target addresses sample

CU ID	Device IDs	Virtual drive number	SCSI target address	LUN
3	0-1	48-49	8	0-1
3	2-3	50-51	9	0-1
3	4-5	52-53	10	0-1
3	6-7	54-55	11	0-1
3	8-9	56-57	12	0-1
3	10-11	58-59	13	0-1
3	12-13	60-61	14	0-1
3	14-15	62-63	15	0-1

For a detailed discussion on tape library sharing between zSeries and SCSI hosts, including considerations about drive sharing, volume sharing and related software implementation steps, refer to the *Guide to Sharing and Partitioning IBM Tape Library Dataservers,* SG24-4409.

4.8.6 SCSI drive emulation

The VTS appears to the SCSI host as up to 64 IBM 3490E drives. In fact, one or more IBM 3490 Model C2A Tape Subsystems each with two SCSI adapters are emulated. The IBM 3490 Model C2A Tape Subsystem contains a single control unit that can access two IBM 3490E tape drives. Before a host application can access a tape drive, it must first address the tape drive. Figure 4-39 illustrates the components used to address IBM 3490E tape drives in an IBM 3490 Model C2A Tape Subsystem.

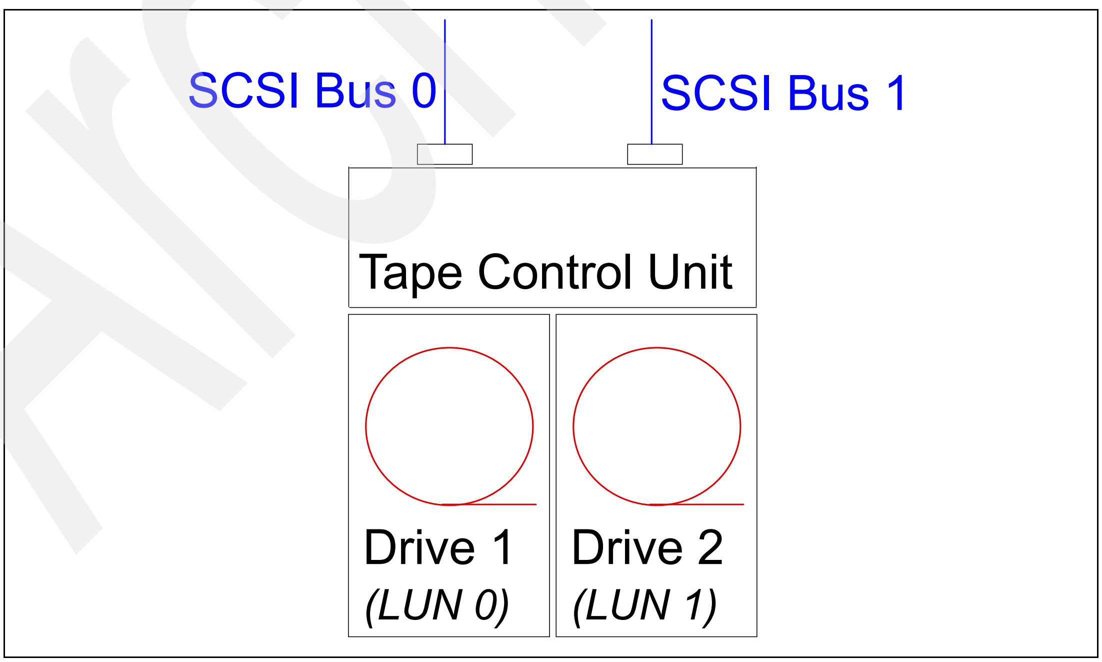

Figure 4-39 IBM 3490 Model C2A Tape Subsystem with two IBM 3490E tape drives

There are three components used to access a drive:

- The SCSI bus interface that the control unit resides on
- The target address of the control unit on the bus
- The address of the drive on the control unit

The format for addressing a SCSI drive is:

SCSI Address Format: SCSI Interface address / Target ID address / LUN address

Where:

- SCSI interface refers to either VTS SCSI adapter bus 0 or 1.
- Target ID refers to the target address of the control unit of the emulated IBM 3490-C2A on the SCSI bus. The valid IDs are 0 - 15.
- LUN refers to the drive you are addressing, address 0 is the first emulated 3490E tape drive, address 1 represents the second emulated 3490E tape drive.

The VTS SCSI host attachment feature provides two SCSI interfaces per adapter. The default for the virtual drive is non-shared. This means that one virtual drive can only be accessed by one SCSI interface. If the customer chooses to have the virtual drives defined as shared then either one or both of the interfaces can be used to access a drive as shown in Figure 4-40.

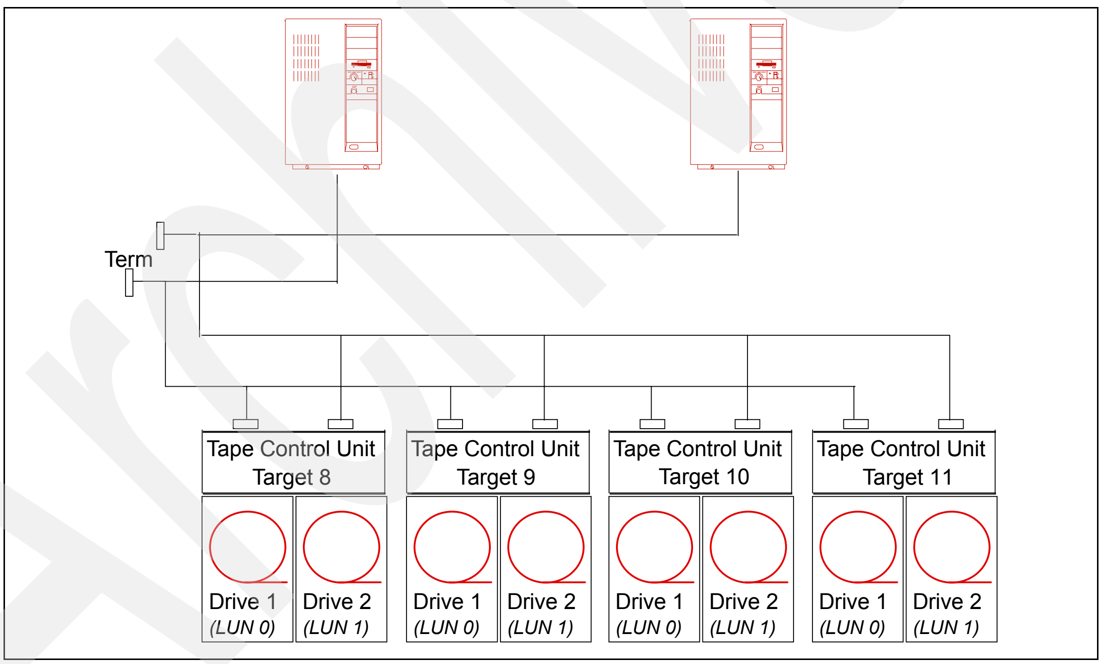

Figure 4-40 Control Unit with two hosts connected to share drives

5

Data migration

The VTS can eliminate a number of bottlenecks that you might have in your tape environment. If you need more drives, the VTS enables you to address up to 256 virtual drives. If you are wasting space on tape cartridges, the VTS can help by stacking the logical volumes and filling the real cartridge completely. The use of virtualization results in almost instantaneous scratch mounts at the host and a reduced need of operations staff.

The IBM VTS can also help you in scratch management, a task that must be performed daily to keep production going. By defining up to 500,000 virtual volumes in a VTS or 1,000,000 in an IBM 3494 Tape Library or an IBM 3584 Tape Library Manager partition with two VTSs, you can use many more scratch volumes than you ever used before, so that the scratch recycle operation can be performed far less often.

In this chapter we provide you with alternatives for and guidance to migrate your data into the VTS, but first we cover product-specific considerations for several applications that use tape such as DFSMShsm.

5.1 DFSMShsm

The z/OS DFSMS hierarchical storage manager (DFSMShsm) has its own functions that allow full utilization of the tape cartridge. So at first sight, using the IBM VTS in a DFSMShsm environment does not bring any special advantage from a capacity point of view, because you create virtual volumes, which, once copied, fill a stacked volume completely. You can achieve the same results from writing directly to the physical tape drive, leaving VTS storage for other applications that cannot use the full cartridge capacity and are therefore better candidates for VTS management.

Even though HSM is an application capable of using the full cartridge capacity, there are reasons to consider using the VTS instead of native physical drives for DFSMShsm data. For example, when writing ML2 data onto a cartridge with an uncompressed capacity of 300 GB, chances are higher that a recall request needs exactly this cartridge which is currently being written to by a space management task. This incident is known as *recall takeaway*.

The effects of recall takeaway can be a real disadvantage when writing ML2 data onto native, high capacity cartridges, because the space management task closes the ML2 cartridge after the user data set is written, when a recall is waiting for this volume. The space management task then requests a new scratch tape, and the partially filled ML2 tape is not reused to append any more data to it. This can result in many cartridges being only partially full and becoming candidates for subsequent recycles. To avoid this effect with, you might have to keep your user data longer on DASD Primary or ML1 volumes. In turn, this means that you could move your data to ML2 sooner when using a VTS instead of native drives, so that you can have smaller Primary and ML1 disk pools.

There are also other reasons to direct HSM data into a VTS. The limited amount of native drives often found in smaller customer environments limits the number of DFSMShsm tasks that can run concurrently. With the large number of up 256 virtual drives that a single VTS offers, you can dedicate a larger number of virtual drives to each HSM function and allow for higher throughput during your limited backup and space management window.

The VTS became a well suited tape storage device for DFSMShsm since the latest VTS generation — for example, the model B20 VTS — is more powerful than the first generations and has large tape volume cache sizes to allow for long periods of peak throughput.

DFSMShsm data is well suited for the VTS, given appropriate tailoring of those parameters that can affect DFSMShsm performance. The subsequent sections describe this tailoring in more detail.

5.1.1 Volume and data set sizes

The size of user data sets is very important when you choose between a VTS or native drives like 3592. DFSMShsm uses only Single File Format to write to tape cartridges.

z/OS supported data set sizes

Different data set sizes are supported for disk and tape data sets, based on the data set organization and the number of volumes, a single data set can span:

- DASD data sets are limited to 59 volumes, except for PDS and PDSE data sets, which are limited to one volume.
- A data set on a VIO simulated device is limited to 65 535 tracks and to one volume.
- Tape data sets are limited to 255 volumes.

Table 5-1 lists the maximum data set sizes supported in z/OS environments.

Table 5-1 Maximum volume sizes in z/OS

Storage medium	Max. volume size	Max. number of volumes	Max. data set size
DASD: 2105 or DS8000	65,520 CYL = 54 GB	59	3.18 TB
Tape: IBM 3592	300 GB x 2.5 Compression	255	191.25 TB
Tape: VTS with large logical volume support	4000 MB x 2.5 Compression	255	2.55 TB

DFSMShsm supported data set sizes

Single file format, as used by DFSMShsm, reduces I/O and system serialization because only one label is required for each connected set (as opposed to multiple file format tapes that require a label for each data set). The standard-label tape data set that is associated with the connected set can span up to the allocation limit of 255 tapes. This standard-label tape data set is called the DFSMShsm tape data set. Each user data set is written in 16K logical blocks to the DFSMShsm tape data set.

Important: A single DFSMShsm user data set can span up to **40** tapes. This is the limit for Migration, Backup, and Recycle.

After DFSMShsm writes a user data set to tape, it checks the volume count for the DFSMShsm tape data set. If the volume count is greater than 215, the DFSMShsm tape data set is closed, and the currently mounted tape is marked full and is de-allocated. The number 215 is used so that a data set spanning **40** tapes fits within the 255-volume allocation limit. DFSMShsm selects another tape, and then starts a different DFSMShsm tape data set.

Data set spanning can be reduced using the SETSYS TAPESPANSIZE command.

VTS large logical volume support

Before VTS Release 7.4, the maximum logical volume size in the VTS was 800 MB. With a maximum of 40 volumes supported and assuming a compression ratio of 2.5 : 1, the maximum user data set size in a VTS was:

```
800 MB x 2.5 x 40 = 80GB
```

Since Release 7.4, logical volume sizes 400, 8000, 1000, 2000, and 4000 MB are supported in the VTS. Let us assume that you have very large z/FS of 300 GB. Such a data set does not fit on 40 volumes of 800 MB each, but it would fit on the #4000 MB large virtual volumes, as proven in the following example:

```
4000 MB x 2.5 x 40 = 400 GB
```

Any single user data set larger than 400 GB is a candidate for native 3592. The compression factor is 2.5:1.

> **Important:** HSM can have more than one address space on one LPAR (MASH = Multi Address Space Support) or have a different setup on different LPARs in your PARMLIB member ARCCMDxx and separated via ONLYIF statements. One HSM address space can have a MIGUNIT(3590-1) and the other address space a MIGUNIT(VTS). The same is true for BUUNIT. The HSM instance which has the 3592 as Migration or Backup Unit can run Space Management or Auto Backup only for that storage group (SG) where all your large data sets like z/FS reside. The other HSM instance would migrate and back up all the smaller data sets onto VTS. You can use a command like F DFSMS2,BACKDS or F DFHSM2,BACKVOL(SG2) to issue the command to the second address space of HSM:
>
> - SETSYS TAPEMIGRATION(ML2TAPE(TAPE(unittype))
> - SETSYS RECYCLEOUTPUT(MIGRATION(unittype))
> - SETSYS BACKUP(TAPE(unittype))
> - SETSYS RECYCLEOUTPUT(BACKUP(unittype))

Migration to a different logical volume size

To make sure that DFSMShsm starts using larger data sets, you must mark full any empty or partially filled tapes written using the previous logical volume size. To identify these tapes, issue the following DFSMShsm command:

LIST TTOC SELECT(NOTFULL)

Each tape identified as being empty or partially filled must be marked full using the following DFSMShsm command:

DELVOL volser MIGRATION(MARKFULL)

and/or

DELVOL volser BACKUP(MARKFULL)

As DFSMShsm migrates data and creates backup copies, it prefers to add to an existing migration/backup volume. As the volume nears full, it handles spanning of data sets as described below in 5.1.2, "VTS implementation considerations" on page 181. If a data set spans across DFSMShsm volumes, it becomes a "connected set" in DFSMShsm terms.

A key point, however, is that whether the data set spans or not, DFSMShsm uses FEOV processing to get the next volume mounted. So, the system believes that the volume is part of a multi-volume set regardless of whether DFSMShsm identifies it as a connected set. Because of the EOV processing, the newly mounted DFSMShsm volume will use the same DC and other SMS constructs as the previous volume.

With the DFSMShsm SETSYS PARTIALTAPE MARKFULL option, DFSMShsm will mark the last output tape full, even though it has not reached its physical capacity. By marking the last volume full, the next time processing starts, DFSMShsm will use a new volume, starting a new multi-volume set and allowing use of new DC and other SMS constructs. If the volume is not marked full, the existing multi-volume set continues to grow and use the old constructs.

This discussion is relevant to APM and the implementation of different logical volume sizes. If all volumes have been marked full, you can simply update your ACS routines to assign a new DC and other SMS constructs and, from then on, each new migration or backup volume will use the new size.

5.1.2 VTS implementation considerations

If you direct HSM data into the VTS, we recommend that you modify your HSM mount wait timer to be 12 minutes. This allows for the possibility of extra time needed on specific mounts for the VTS to stage the data back in to cache.

Consider using large logical volumes like 4000 MB for backup but short logical volumes for migration. You need to select another SMS DATACLAS for backup than for migration via your ACS routine.

Consider tailoring the following ARCCMDxx SETSYS parameters:

```
SETSYS TAPEUTILIZATION(LIBRARYBACKUP(PERCENTFULL(97)))
SETSYS TAPEUTILIZATION(LIBRARYMIGRATION(PERCENTFULL(40)))
```

Using a high value for backup tapes helps you minimize the number of virtual volumes. Using a lower value for migration tapes speeds up recall processing, because having smaller logical volumes reduces the delay for the recall operation, while still allowing the use of MEDIA2 logical volumes.

The introduction of 5 sizes of logical volumes such as 400 MB, 800 MB, 1000 MB, 2000 MB, or 4000 MB bytes of uncompressed data is not known to DFSMShsm. Your installation can select one of the sizes by assigning a Dataclass via the ACS routines. You must tune the PERCENTFULL value for your requirements. The value of 97% for backup is correct for 800 MB logical tape volumes, and you may try 485% for 4000 MB logical volume sizes. We recommend the following values:

For backup:

Percentage = x MB x 97 / 800 MB, where x is your logical volume size in MB

For migration:

Percentage = x MB x 40 / 800 MB, where x is your logical volume size in MB

See also the text of 3584 support APAR OA09751.

Other applications may have a similar existing TAPECAPACITY-type specification or a PERCENTFULL-type specification enabling applications to write beyond the default volume sizes for MEDIA1 (cartridge system tape) and MEDIA2 (enhanced capacity cartridge system tape).

If multiple VTS libraries are eligible for a request, consideration should also be given so that the same logical volume size is then used for the request across all libraries. When displaying the volumes through your tape management system, the tape management system may continue to display the volume capacity based on the default volume size for the media type with the volume usage (or a similar parameter) showing how much data has actually been written to the volume reflecting its larger capacity.

The default volume size is overridden at the library through the Library Manager data class policy specification and is assigned or re-assigned when the volume is mounted for a scratch mount. For further discussion of this enhancement (including the larger logical volume sizes support) and other new VTS enhancements, refer to the 3494 Tape Library Operator Guide (GA32-0449).

```
SETSYS SELECTVOLUME(SCRATCH)
SETSYS PARTIALTAPE(MARKFULL)
SETSYS TAPEDELETION(SCRATCHTAPE)
```

Using a global scratch pool, you benefit from a fast mount time by using the fast-ready attribute for the scratch category (see Chapter 4, “Implementation” on page 107). The MARKFULL parameter does not mean a waste of space using VTS, because the stacked volume contains only the written data of each logical volume copied and the same applies to the TVC.

```
SETSYS TAPESPANSIZE(4000)
```

Explanation: TAPESPANSIZE(nnnn) is an optional parameter used to reduce the occurrences of data sets spanning migration or backup tape volumes. For nnnn, specify a value between 0 and 9999 in units of megabytes (MB). This value represents the maximum number of megabytes of tape (ML2 or backup) that DFSMShsm may leave unused while it tries to eliminate spanning of data sets. To state this differently, this value is the minimum size of a data set that is allowed to span tape volumes. Data sets whose size is less than the value do not normally span volumes. Only those data sets whose size is greater than or equal to the specified value are allowed to span volumes.

This parameter offers a trade-off: to reduce the occurrences of a user data set spanning tapes in exchange for some unused media at the end of each cartridge. The amount of unused media can vary from 0 to nnnn physical megabytes, but roughly averages 50% of the customer's median data set size. For example, if you specify 4000 MB and your customer's median-sized data set is 2 MB, then on average only 1 MB of media is unused per cartridge.

Installations that currently experience an excessive number of spanning data sets may want to consider specifying a larger value in the SETSYS TAPESPANSIZE command.

A larger absolute value is needed to represent the same amount of unused capacity on a percentage basis when the tape has a larger total capacity. For example, if you allow 2% of unused tape to avoid tape spanning for a 3590-Hxx device using enhanced media, this would be a TAPESPANSIZE of 1200 MB. To allow 2% unused tape for a 3592 Model J device, this would be a TAPESPANSIZE of 6000 MB. The 2% unused tape for a 4000 MB logical volume on a VTS would result in a TAPESPANSIZE of 60 MB.

Using a high value reduces tape spanning. This will reduce the number of virtual volumes that need to be recalled to satisfy DFSMShsm recall/recover requests. With VTS, you can be very generous with the value, as no space is wasted. For example, TAPESPANSIZE of 4000 would mean that any data set less than 4000 MB will not fit on the remaining space of a virtual volume will be started on a fresh new virtual volume.

DFSMShsm recommends a value of 4000 MB for all IBM 3490 and 3590 tape cartridges.

Note: The maximum value is 9999 MB and the default value is 500 MB.

5.1.3 DFSMShsm task related considerations

To better understand the use of the DFSMShsm with VTS, we now summarize the different DFSMShsm functions that use tapes and we analyze the benefit that VTS gives.

Backups of DFSMShsm control data sets

The backup of DFSMShsm control data sets (CDSs) can easily be done in a IBM VTS, exploiting the benefit of using virtual volumes instead of physical volumes, which might otherwise be under utilized.

ABARS

The purpose of the DFSMShsm ABARS function is disaster recovery. Therefore, ABARS should not be directed to a IBM VTS.

However, if the IBM VTS is installed off-site, the virtual volumes created reside in the VTS, with no strong need to eject them. The ABARS **SETSYS ABARSTAPES(STACK)** parameter is not recommended in a VTS environment. It would cause ABARS to write all output files to a single output volume. **SETSYS ABARSTAPES(NOSTACK)** should be specified instead.

ABARS function can be used to export data from the IBM VTS, as detailed in 5.7.6, “Moving data out of the VTS” on page 203.

Volume dumps

Using VTS as output for the DFSMShsm AUTODUMP function, do not specify the **DEFINE DUMPCLASS(dclass STACK(nn)) or BACKVOL SG(sgname)|VOLUMES(volser) DUMP(dclass STACK(10))** parameters. These parameters were introduced to force DFSMShsm to use the capacity of native TotalStorage physical cartridges. If used with VTS, they cause unnecessary multivolume files. We therefore recommend using the default value, which is **NOSTACK**.

Migrate/recall (DFSMShsm Level 2)

The only problem that you need to consider when using a VTS as DFSMShsm level 2 is the recall process throughput of the IBM VTS. Consider how many recall tasks are started at the same time and compare that number with the number of physical drives of your IBM VTS. For example, if your installation often has more than 10 tape recall tasks at one time, you probably need twelve 3590 back-end drives to satisfy this throughput request, because all migrated data sets may already have been copied into stacked volumes and may not be in the TVC.

Backup and recovery

Unlike the DFSMShsm RECALL operation, RECOVERY usually has lower frequency in an DFSMShsm environment. Therefore, using VTS for DFSMShsm backup and recovery functions benefits you without impacting DFSMShsm performance. However, be careful to review your DFSMShsm performance requirements before moving DFSMShsm BACKUP to the IBM VTS.

TAPECOPY

The DFSMShsm TAPECOPY function requires that original and target tape volumes are of the same media type and use the same recording technology. Using a VTS without the EHPO feature as target for the TAPECOPY operation can cause problems because VTS virtual volumes are specified with an IDRC compaction ratio of 1:1 and can have a capacity smaller than native input volumes.

Use Table 5-2 to tailor your TAPECOPY environment.

Table 5-2 TAPECOPY utilization

ORIGINAL volume unit name	ALTERNATE volume unit name	Percent full that should be defined (assuming IDRC 2:1)
VTS w/o EHPO (CST)	3490E (CST), VTS (CST)	97%
VTS w/o EHPO (ECCST)	3490E (ECCST), VTS (ECCST)	97%
3490E (CST)	VTS w/o EHPO (CST)	45%

ORIGINAL volume unit name	ALTERNATE volume unit name	Percent full that should be defined (assuming IDRC 2:1)
3490E (ECCST)	VTS w/o EHPO (ECCST)	45%
VTS w EHPO (CST)	3490E (CST)	85%
VTS w EHPO (ECCST)	3490E w EHPO (ECCST)	85%

For example, if you are planning to put DFSMShsm alternate copies into a IBM VTS, a tape capacity of 45% may not be enough for the input non-VTS ECCST cartridges. TAPECOPY will fail if the (virtual) output cartridge encounters EOV before the input volume has been copied completely.

On the other hand, using VTS logical volumes as original and 3490E native as TAPECOPY target may cause EOV at the alternate volume because of the higher IBMLZ1 compression seen on the virtual drive compared to the IDRC compression on the native drive.

For special situations where copying from standard to enhanced capacity media is needed, the following patch command can be used:

PATCH .MCVT.+4F3 BITS(.....1..)

Tapecopy will then allow copying from standard to enhanced capacity media if the input and output tape units are exactly the same geometry, that is, D/T3590E —> D/T3590E, etc. D/T3590B —> D/T3590E, etc. ,would not be allowed because of the device geometry mismatch.

Normally you should avoid this practice since it will produce more unused space on tapes that are marked full and resultant higher media costs. The DFSMShsm default is to disallow copying from standard to enhanced capacity tape media.

DUPLEX TAPE

DFSMShsm uses the same logic and has the same restrictions: Both output tapes have to be of the exact same size and unit type. We recommend to use a P2P VTS and let the hardware do the duplex rather the DFSMShsm software function. This also helps to easier manage the disaster side. You can use GDPS and switch to the remote DASD side and the tape volser itself has not to be changed. No TAPEREPL or SETSYS DISASTERMODE commands are needed.

When using the 3592 Model **J1A** for duplexed migration output, performance is degraded because of the back-to-back SYNCDEV operations done for the original and the alternate tapes. APAR OA09928 provides a patch allowing syncs on the alternate tape to be disabled. The performance improvement varies with data set size, with the greatest improvements seen for the smaller data sets. Performance improvements can be quite substantial.

See OW45264 for a similar duplexed recycle patch.

RECYCLE

The DFSMShsm RECYCLE function reduces the number of logical volumes inside the VTS, but when started can cause bottlenecks in the IBM VTS recall process. If you have an IBM VTS with three physical drives, use a maximum of two concurrent DFSMShsm RECYCLE tasks. If you have an IBM VTS with six physical drives, use no more than five concurrent DFSMShsm RECYCLE tasks.

Select the RECYCLEPERCENT carefully, bearing in mind that:

- What you will free up is logical volumes residing on a stacked volume with hundreds of other logical volumes.
- The space occupied by the logical volume will be freed up only if and when the logical volume is used (overwritten) again, unless you are using Expired Volume Management.
- To RECYCLE, the VTS has to load the input volumes into the TVC.

We recommend to use a RECYCLEPERCENT value depending on the logical volume size, for example:

```
5 for 1000 - 4000 MB volumes
10 for 400 to 5800 MB volumes
```

Using RECYCLE SELECT (INCLUDE(RANGE(nnnnn:mmmmm))) or RECYCLE SELECT (EXCLUDE(RANGE(nnnnn:mmmmm))) for RECYCLE input can be helpful while selecting and migrating data to and from a VTS. Its immediate purpose is to enable you to set up volume ranges for different media types and different emulation types, like VTS logical volumes and 3490-emulated cartridges. There are no special data set names for RECYCLEOUTPUT, although you must code your ACS routines or BTLS options if using BTLS, to route RECYCLEOUTPUT to the library, using the &UNIT variable.

Refer to the IBM Redbook, *DFSMShsm Primer*, SG24-5272, for more information on implementing DFSMShsm.

5.1.4 DFSMShsm AUDIT considerations

In nearly all instances in which the DFSMShsm AUDIT MEDIACONTROLS function is performed, it is for tapes in “failed create” status. These are tapes in the process of being written to when DFSMShsm experienced an outage. As a result, the final extension to the Tape Table-of-Contents record (TTOC) has not been written to the OCDS, causing DFSMShsm's inventory of data sets residing on that tape to be incomplete. Since the missing inventory information is limited to, at most, the last 33 data sets written to the tape, AUDIT MEDIACONTROLS processing has the capability of initiating the audit at the last known position on the tape, rather than reading through the entire tape.

Prior to APAR OA04419, this “fast path” is taken only when the final TTOC data set entry recorded in the OCDS remains valid. If the final data set recorded in the TTOC has been expired, recalled, or otherwise deleted, AUDIT MEDIACONTROLS will audit the entire contents of the tape, an exceedingly time-consuming process for high-capacity tapes. While a full audit of a high-capacity tape can take many hours, the “fast path” typically requires only a few minutes to complete.

APAR OA04419 addresses the dependency of AUDIT MEDIACONTROLS on positioning information for the last data set recorded in the TTOC. The APAR changes the process by which DFSMShsm migration and backup tapes are created, saving positioning information in the base TTOC, itself. AUDIT MEDIACONTROLS is also changed by the APAR to take advantage of this additional positioning information such that the audit need not be performed in such close proximity to the time of the error and can successfully take the “fast path” for the vast majority of tapes in “failed create” status.

Note that OA04419 will affect only those tapes created or added to after the APAR is applied, as it relies on information recorded when the tape is being written to. Tapes last written to prior to the installation of the APAR will continue to take the “fast path” only when the last data set entry in the TTOC is valid. As before, tapes not in “failed create” status will be audited in their entirety.

The APAR OA04419 does also help to accelerate logical tapes, especially when using large 4000 MB volume sizes. AUDITs run faster on VTS volumes since they hold less data than native tape cartridges such as 3590H or 3592.

5.2 DFSMSrmm and other tape management systems

No changes are required to any tape management system to support basic VTS, you need only review the retention and movement criteria for the data in the VTS. However, if you plan to use the Advanced Function feature with the Import/Export component, you must have the appropriate tape management system software levels installed and implemented to control and manage vaulting and expiration of logical volumes residing on exported stacked volumes outside of the VTS. DFSMSrmm keeps track of logical volumes outside of the VTS in a *CONTAINER* field of the specific volume record. Refer to Appendix B, “VTS Import / Export” on page 369 for detailed information about Advanced Function feature.

When you direct allocations inside the VTS, the Vital Record Specifications (VRSs) or vault rules should tell the tape management system that the data set will never be moved outside the library if you do not want to use the export function.

Stacked volumes cannot be used by the host; they are managed exclusively by the VTS. Do not allow any host to either implicitly or explicitly address the VTS stacked volumes.

Exported stacked volumes, residing outside the VTS, also cannot be used by the host. However, the tape management system needs to know about these volumes in order to track the location and movement.

To indicate that the stacked volser range is reserved and cannot be used by any host system, define the volsers of the stacked volumes to RMM (see Figure 5-1).

A good method is to add racks and volsers of stacked cartridges to RMM and put them in a *loan location*. When exporting logical volumes, you can manually assign vault locations to the exported stacked volumes to keep track of their location. This can be done using the command:

```
RMM CHANGEVOLUME LOCATION(newloc)
```

Use the following PARMLIB parameter, assuming that VT is the prefix of your stacked VTS cartridges:

```
REJECT ANYUSE(VT*)
```

This will cause RMM to deny any attempt to read or write those volumes on native drives.

You can use the following statement:

```
REJECT OUTPUT(VT*)
```

This will cause RMM to deny any attempt to write on those volumes on native drives, but will not prevent reading operations. This option is recommended if you plan to read exported stacked volumes outside the VTS to recover logical volumes using a special utility function.

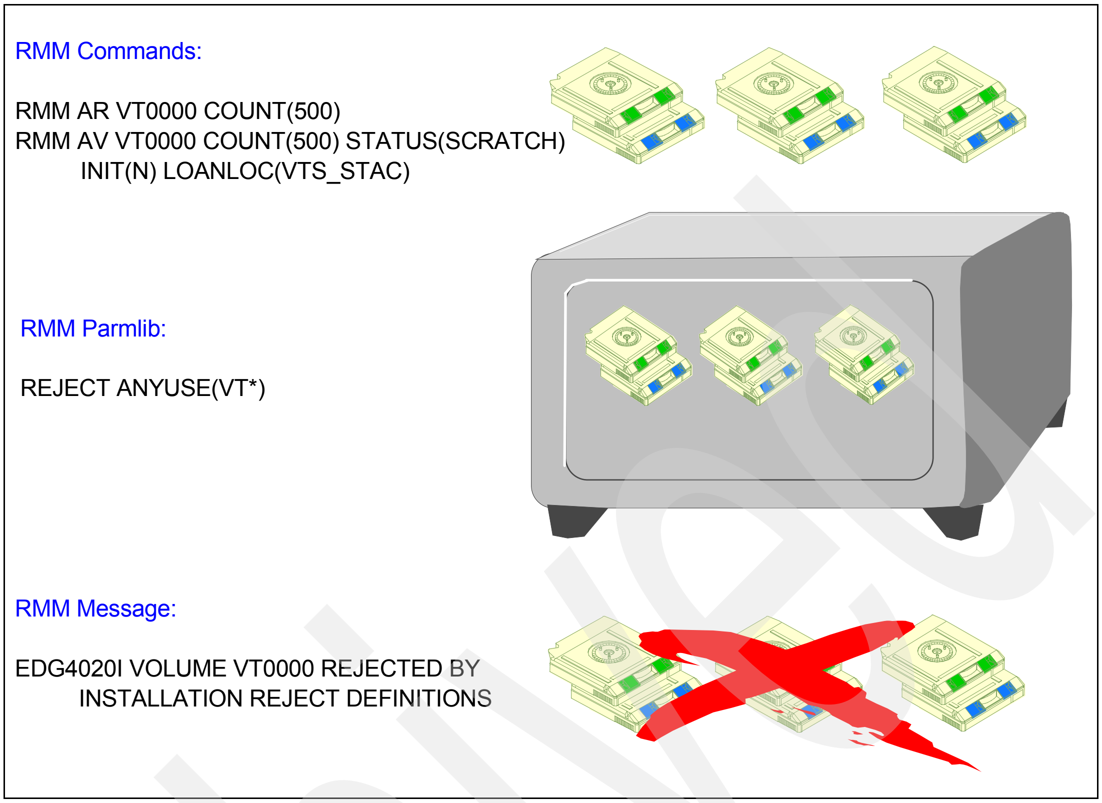

Figure 5-1 Defining stacked volumes to RMM

You do not need to explicitly define the virtual volumes to RMM. During entry processing, the active RMM automatically records information about each volume in its control data set. RMM uses the defaults you specified in ISMF for the library entry values if there is no existing RMM entry for an inserted volume; set the default entry status to SCRATCH.

When adding 500,000 virtual volumes, the size of the RMM CDS and the amount of secondary space available must be checked. RMM uses 1 MB for every 1000 volumes defined in its CDS; an additional 500,000 volumes would need 500 MB of space. However, we do not recommend that you add all volumes initially. Refer to 4.3.4, “Define (insert) logical volumes” on page 130.

To increase the size of the RMM CDS, you have to stop RMM activities, back up the CDS, then reallocate a new CDS with a larger size and restore the CDS from the backup copy. To calculate the correct size of the RMM CDS, refer to the *z/OS DFSMSrmm Implementation and Customization Guide*, SC26-7404.

Other tape management systems, such as BrightStor, CA-1 Tape Management Copycat Utility (BrightStor CA-1 Copycat) and BrightStor CA-Dynam/TLMS Tape Management Copycat Utility (BrightStor CA-Dynam/TLMS Copycat) need to reformat their database to add more volumes.

Additionally, some tape management systems do not allow the specification of tape volumes with alphanumeric characters or require user modifications to do so. Refer to the proper product documentation for this operation.

In both RMM and the other tape management systems, the virtual volumes do not have to be initialized. The first time a volser is used, VTS marks the virtual volume with VOL1, HDR1 and a tape mark, as if it had been done by EDGINERS or IEHINITT.

The virtual volumes, once used, appear to the tape management system as normal CST or ECCST cartridges with IDRC compaction used. The information on media errors (temporary and permanent) for virtual volumes is always null, because the errors for the real media (physical stacked volumes) are registered and managed by the IBM VTS without host knowledge.

5.3 Tivoli Storage Manager

With Tivoli Storage Manager there are significant new features and performance improvements over previous versions of ADSM. With support for SANs, Tivoli Storage Manager protects over 35 operating platforms from a single management console, making it the best choice for customer's heterogeneous server environments. Tivoli Storage Manager provides data availability, records retention and enterprise management for data on multiple vendor platforms in a networked environment.

The VTS provides ESCON and FICON channel attachment, as well as attachment through SCSI : a new approach to storage management. If you plan to use Tivoli Storage Manager with VTS, it is important to understand how these products work together.

Customers who use Tivoli Storage Manager with 3490s or tape solutions that have slower overall performance than VTS might be very satisfied with the performance of Tivoli Storage Manager with VTS. Other customers who use 3590 or 3592 today and go to VTS to solve drive availability or sharing problems or to eliminate dedicated libraries for Tivoli Storage Manager, might be less satisfied with performance of Tivoli Storage Manager and VTS compared to the connection to native drives. In general, Tivoli Storage Manager operations will perform better with native 3590 or 3592 drives than with VTS, and you should understand the trade-off between the performance and the VTS functionality.

5.3.1 Recommendations for VTS usage

Tivoli Storage Manager, like DFSMShsm, can automatically fill a native 3590 cartridge. It can use the tape up to the EOV, independent of the media type. If you plan to store Tivoli Storage Manager data into the VTS, consider the following suggestions for placing data on your VTS.

Use VTS for Tivoli Storage Manager archiving for archiving and backup of large files or databases for which you do not have a high performance requirement during backup and restore. VTS is ideal for Tivoli Storage Manager archive or long term storage because archive data is not frequently retrieved. Archives and restores for large files should see less impact from the staging overhead. Small files, such as individual files on file servers, can see performance impacts from the VTS staging. (If a volume is not in cache, the entire volume must be staged before any restore can be done).

Set TSM reclamation off by setting the reclamation threshold to 100 percent. Tivoli Storage Manager, like DFSMShsm, has a reclamation function to consolidate valid data from tapes with a low valid data percentage onto scratch tapes so that tapes can be freed up for reuse. TSM Reclamation with VTS can be slower because all volumes have to be staged to the cache. You should periodically set TSM reclamation on, by setting the threshold to a lower value to regain the use of VTS volumes with small amount of valid data that will not expire for a longer period of time. TSM reclamation should be scheduled for non-peak hours.

Use collocation to reduce the number of VTS volumes required for a full restore. Tivoli Storage Manager has a collocation function to group Tivoli Storage Manager client data onto a minimum set of tapes to provide a faster restore and to provide separation of client data onto different physical tapes. Collocation with VTS will not minimize the physical tapes used but will minimize the number of logical volumes used. Collocation with VTS can improve restore time for large amounts of data. VTS will not ensure physical tape separation when collocation is used because different logical volumes can reside on the same physical tape.

Use VTS for Tivoli Storage Manager database backups that are to be used for recovery from local media and use VTS at a recovery site or native drives for backups that are to be used for recovery from off-site media. Tivoli Storage Manager requires a different tape for every backup of the Tivoli Storage Manager database, therefore a large number of logical volumes with less data is created. Using the VTS you don't have to worry about the unused capacity of logical volumes.

Use VTS for backups of primary pool, noting that similar considerations to database backups apply to copy storage pools. If only one copy pool is used for local backups, then that storage pool should not be in the VTS, because there can be no guarantee that data in the copy storage pools are on different physical volumes. If storage pools for local and off-site backups are used, the copy storage pools for local backups can be in the VTS. The copy storage pools for off-site backups should use native drives or a VTS at the recovery site.

Use VTS in server-to-server configurations for multiple Tivoli Storage Manager server implementations. If you are using TSM server-to-server configuration, the data from your remote TSM servers are stored as virtual volumes, which appear as sequential media volumes on the source server and which are actually stored as archive files on a target server. These are ideal candidates for a VTS.

5.3.2 Recommendations for native drives

Use native drives for data that will be used for frequent individual file restores or have a requirement for high performance for backup and restore without any delays because of staging activity. Tivoli Storage Manager uses EXPORT to move data from one Tivoli Storage Manager server to another. This requires that both servers have compatible devices for the EXPORT media. Native drives should be used for Tivoli Storage Manager EXPORT unless you have the advanced function IMPORT/EXPORT at both VTSs.

5.3.3 Tivoli Storage Manager parameter settings

The setting for the following parameters can affect the performance of Tivoli Storage Manager with VTS.

MAXSCRATCH (storage pool definition): As for DFSMShsm, Tivoli Storage Manager should use a scratch pool for tapes, because you do not have to predefine tapes to Tivoli Storage Manager and you could benefit from the faster VTS scratch mounts.

MOUNTLimit (device class definition): With VTS, you have up to 64 virtual drives available. The number of drives available for Tivoli Storage Manager use can probably be increased, taking into account VTS performance. Set MOUNTLimit high enough so that the number of available drives does not limit the performance of Tivoli Storage Manager tape operations.

MOUNTRetention (device class definition): When storing data in the VTS, you can set this parameter to zero, because you have a greater chance of finding the virtual volume still in the TVC when Tivoli Storage Manager needs it. This avoids the need to keep the virtual volume mounted and frees a virtual drive for other users.

MAXCAPacity (device class definition): Using this parameter, you can tailor the size of the data written in a virtual volume. Having smaller virtual volumes can speed up recall processing. Using the full capacity of the virtual volume can limit the number of volumes used by Tivoli Storage Manager.

BAckup DB (database back up): Use SCRATCH=YES to use tapes from the TMS scratch pool and benefit from the faster VTS scratch mounts.

For details on setting up Tivoli Storage Manager, refer to the Tivoli Storage Manager Administrators Guide that can be found at:

http://www.tivoli.com/support/public/Prodman/public_manuals/td/TD_PROD_LIST.html

5.4 DFSMSdss

DFSMSdss full volume dumps are a good use of the IBM VTS. You have to plan carefully, however, and make sure you can achieve the required throughput: a DFSMSdss full volume physical dump can easily provide a data transfer rate of 10 MB/sec and higher for a single job. However, with today's VTS throughput capabilities of up to more than 400 MB/sec, the VTS throughput capabilities are unlikely to be a limiting factor. In the past, the data rate was often limited by the bandwidth of the DASD subsystem as the weakest part in the chain.

With VTS, you fill the stacked cartridge completely without changing JCL, using multiple virtual volumes. VTS then moves the virtual volumes created onto a stacked volume.

5.4.1 Full volume dumps

The only problem you might experience when using VTS for the DSS volume dumps is related to the size of the virtual volumes. If a single dump does not fit onto five logical volumes, you can either use an SMS DATACLAS specification, **Volume Count nn**, to enable more than five volumes. A better way is if you have VTS release 7.4 installed and can choose a 4000 MB logical volume via your SMS DATACLAS. This prevents unneeded multi-volume files.

Using the *COMPRESS* keyword of the *DUMP* command, you obtain a software compression of the data at the host level. As data is compressed at the VTS before being written into the tape volume cache, host compression is not required, unless channel utilization is high already.

5.4.2 Stand-Alone Services

Stand-Alone Services of DFSMSdss provide a stand-alone restore function that enables you to restore vital system packs without needing to rely on an zSeries environment.

Stand-Alone Services supports the 3494 and the Virtual Tape Server. It enables you to restore from native as well as virtual tape volumes in a VTS. Stand-Alone Services lets you specify the input volumes on the RESTORE command and sends the necessary mount requests to the Library Manager.

You can initial program load (IPL) the Stand-Alone Services core image from a virtual tape device and use it to restore dump data sets from virtual tape volumes.

Stand-Alone Services is provided as a replacement to the previous DFDSS V2.5 and DFSMS/MVS® V1 stand-alone functions. The installation procedure for Stand-Alone Services retains, rather than replaces, the existing stand-alone restore program so you do not have to immediately change your recovery procedures. We recommend that you implement the procedures as soon as you can and start using the enhanced Stand-Alone Services.

To use Stand-Alone Services, create a stand-alone core image suitable for IPL, using the new BUILDSA command of DFSMSdss.

Use the steps listed below to IPL the Stand-Alone Services program from a virtual device and to restore a dump data set from virtual volumes. Refer to 6.10.3, “Stand-alone support” on page 262 on how to use Library Manager menus to set a device in stand-alone mode.

- Ensure that the virtual devices you will be using are offline to other host systems. Tape drives to be used for stand-alone operations must remain offline to other systems.
- Set in stand-alone mode the virtual device from which you will load the Stand-Alone Services program by using the **Setup Stand-alone Device** window on the Library Manager console. Select **Mount a single volume** and specify the virtual volume that contains the core image for IPL.
- Load the Stand-Alone Services program from the device you just set in stand-alone mode. As part of this process, select the operator console and specify the input device for entering Stand-Alone Services commands.
- When the IPL is complete enter Stand-Alone Services RESTORE command from the specified input device. Example 5-1 shows a group of statements for use:

Example 5-1 Restore command

```
RESTORE FROMDEV(TAPE) FROMADDR(0A40) TOADDR(0900) -
   NOVERIFY TAPEVOL((L00001),(L00002))
```

  L00001 and L00002 are virtual volumes that contain the dump data set to be restored, 0A40 is the virtual device used for reading source volumes L00001 and L00002 and 0900 is the device address of the DASD target volume to be restored.

  Stand-Alone Services requests the Library Manager to mount the source volumes in the order in which they are specified on the TAPEVOL parameter. It automatically unloads each volume, then requests the Library Manager to demount it and to mount the next volume.

- When the restore is complete, unload and demount the IPL volume from the virtual device by using the Library Manager’s **Setup Stand-alone Device** window.
- Take the IPL device out of stand-alone mode by using the Library Manager's **Reset Stand-alone Device** window.

Stand-Alone Services issues the necessary mount and demount orders to the library. If you are using another stand-alone restore program which does not support the mounting of library resident volumes, you would have to set the source device in stand-alone mode and manually instruct the Library Manager to mount the volumes using the **Setup Stand-alone Device** window.

For details on how to use Stand-Alone Services, refer to *z/OS DFSMSdss Storage Administration Reference SC35-0424*

5.5 Object Access Method (OAM)

Tape cartridges provide a low-cost storage medium for storing primary and/or backup copies of OAM objects.

Allowing objects to be stored on tape volumes in conjunction with DASD and optical media provides flexibility and more efficiency within the storage management facility.

OAM stores objects on a IBM VTS as in a normal IBM 3494 or 3584 Tape Library, with up to 256 virtual drives and many virtual volumes available.

Consider using the TAPEPERCENTFULL parameter with object tape data, because the retrieve time of an OAM object is important. The recall time for smaller logical volumes can be reduced considerably.

There are also new functional enhancements associated with the Virtual Tape Server (VTS), one of which includes support for larger logical volume sizes. Application configuration related changes may be needed to fully utilize the larger logical volumes sizes defined at the library. For example, with OAM's object tape support, the existing SETOAM TAPECAPACITY keyword can be used to fill the larger logical volumes. However, before using this support with OAM's object tape support, install existing field APAR OA08963.

Virtual volumes in a IBM VTS can contain primary or backup copies of OAM objects, addressing either OBJECT or OBJECT BACKUP storage groups. An advisable solution is to address VTS with the OBJECT storage group and other non-VTS tape devices with the OBJECT BACKUP storage group, or use VTS Import/Export because usually the volumes in the object backup storage group are moved off-site for disaster recovery purposes.

A virtual volume can contain multiple OAM objects, separated by a buffer space. To optimize the use of VTS storing OAM object data, consider the following suggestions:

- Review the MOUNTWAITTIME parameter when using VTS to store OAM object tape data. The default (5 minutes) should probably be increased. Twelve minutes is a better number in case you have to recall a logical volume to read object data and there are other recall requests queued at the time. The VTS may need to stage the data back in to cache and this accounts for the extra mount time.
- Review the MAXTAPERETRIEVETASKS and MAXTAPESTORETASKS parameters when using VTS, because you have more virtual tape drives available.
- There are also other parameters, such as DEMOUNTWAITTIME, TAPEPERCENTFULL and TAPEFULLTHRESHOLD, that may need to be reviewed when using VTS to store OAM data.

OAM does support the VTS High Capacity Cache Option (FC4020 and FC4021). This VTS does, however, have a limited number of logical volumes available to it due to the fact that all volumes reside in cache. The benefit to OAM is the fast retrieval of both scratch and private data that is often needed for OAM object processing.

5.6 Database backups

Using a IBM VTS as output confers several advantages on database backups. In this section we provide a detailed description of these benefits for database products, such as DB2®.

5.6.1 DB2 data

DB2 uses tapes for two purposes: for storing archive logs and for storing image copies. Either one may be created in multiple copies, to be stored both on-site for local recovery and off-site for disaster recovery purposes. To use DB2 tape data with VTS, we recommend the approaches described in the following sections.

Archive logs

DB2 keeps track of database changes in its active log. The active log uses up to 31 DASD data sets (up to 62 with dual logging) in this way: When a data set becomes full, DB2 switches to the next one and automatically offloads the full active log to an archive log.

Archive logs are sequential data sets that are allocated on either DASD or tape. When archiving to tape, a scratch tape volume is requested each time.

Archive logs contain unique information necessary for DB2 data recovery. Therefore, to ensure DB2 recovery, customers usually make backups of archive logs. You can use general backup facilities or DB2's dual archive logging function.

When creating a dual copy of the archive log, usually one is local and the other is for disaster recovery. The local copy can be written to DASD, then moved to tape, using Tape Mount Management (TMM). The other copy can be written directly to tape and then moved to an off-site location.

With VTS, you can write the local archive log directly inside the IBM VTS. Avoiding the use of TMM saves DASD space, saves DFSMShsm CPU cycles and simplifies the process. The disaster-recovery copy must be created on non-VTS tape drives, so that it can be moved off-site.

The size of an archive log data set varies from 150 MB to 1 GB. The size of a virtual volume on a VTS without EHPO can be up to 4000 MB, so be sure that your archive log can fit in only one virtual volume. This is because it is suggested to use a single volume when unloading an archive log to tape. The size of a virtual volume on a VTS can be up to 12,000 MB, assuming a 3:1 compression ratio.

Tailoring the size and the number of active log DASD data sets allows you to obtain an archive log on tape whose size does not exceed the virtual volume size.

Limiting data set size may increase the frequency of offload operations and reduce the amount of active log data on DASD. However, this should not be a problem because VTS does not require manual operation and archive logs will stay in the TVC for some time and be available for fast recovery.

One form of DB2 recovery is backward recovery, typically done after a processing failure, where DB2 backs out uncommitted changes to resources. When doing so, DB2 processes log records in reverse order, from the latest back toward the oldest.

If the application being recovered has a large data set and makes only a few commits, you probably need to read the old archive logs that are on tape. When archive logs are on tape, DB2 uses read-backward channel commands to read the log records. Read-backward is a slow operation on tape cartridges processed on real IBM 3480 (if IDRC enabled) and IBM 3490 tape drives. On a VTS it is only about 20% slower than a normal I/O because data is retrieved from the TVC, so the tape drive characteristics are replaced by the random access disk characteristics.

Another benefit VTS can provide to DB2 operations is the availability of up to 256 virtual drives, because DB2 often needs a large number of drives concurrently to perform recovery or backup functions.

Having two VTSs, one on-site and one off-site, makes dual logging very easy, exploiting the benefit of the IBM VTS.

Figure 5-2 shows an example of where to put tape data.

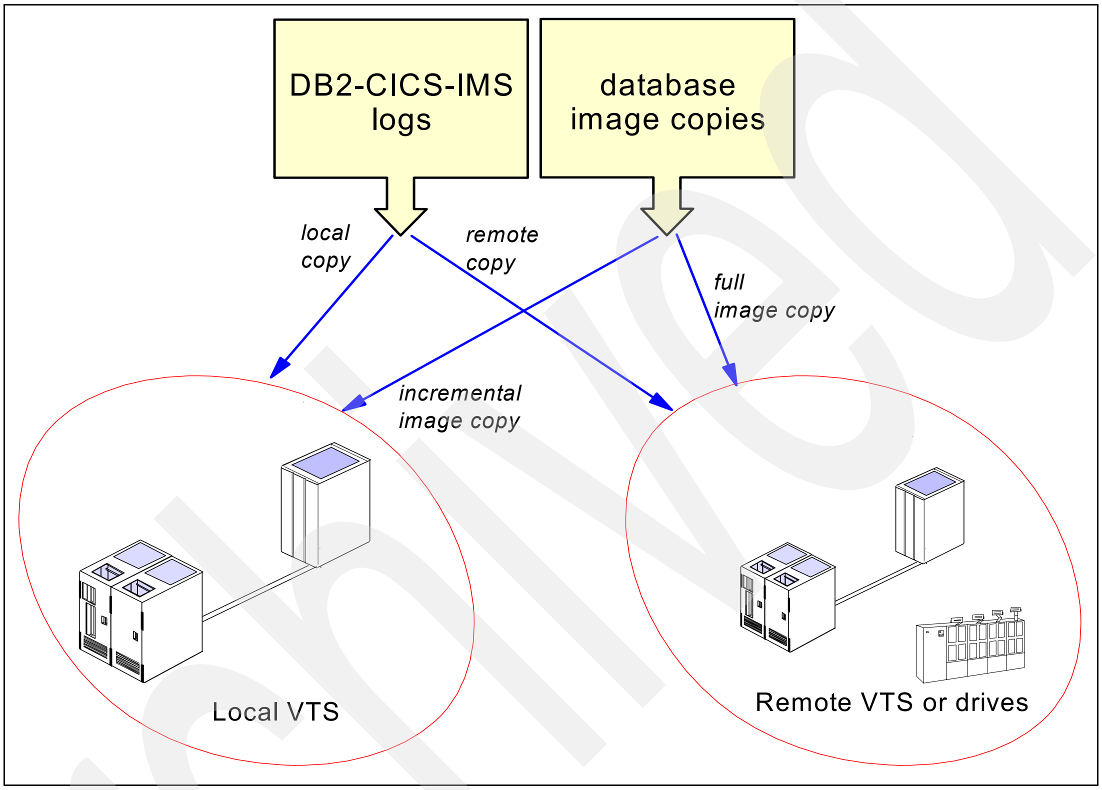

Figure 5-2 How to distribute database tape data among local and remote sites

Image copies

Image copies are backup copies of table spaces in a DB2 database. DB2 can create both full and incremental image copies. A full image copy contains an image of the whole table space at the time the copy was taken. An incremental image copy contains only those pages of a table space that have changed since the last full image copy was taken. Incremental image copies are typically taken daily, whereas full image copies are typically taken weekly.

DB2 provides the option for multiple image copies. You can create up to four identical image copies of a table space, one pair for local recovery use and one pair for off-site storage.

The size of the table spaces to be copied varies from a few megabytes to several gigabytes. The VTS solution is the best for small and medium sized table spaces, because you need a higher bandwidth for large table spaces.

When a database is recovered from image copies, a full image copy and the subsequent incremental image copies need to be allocated at the same time. This can potentially tie up many tape drives and, in smaller installations, can prevent other work from being run. With VTS, with its 256 virtual drives, this is not an issue.

The large number of tape drives is important also for creating DB2 image copies. Having more drives available allows you to run multiple copies concurrently and use the MERGECOPY DB2 utility without impact. An advisable solution is to run a full image copy of the DB2 databases once a week outside the VTS and run the incremental image copies daily using VTS. (The smaller incremental copy fits better with the VTS volume sizes.)

5.6.2 CICS and IMS

Like DB2, both CICS® and IMS use tapes to store logs and image copies of databases.

CICS is only a data communication product, whereas IMS has both the data communication and the database function (IMS-DL/1). CICS uses the same DL/1 database function to store its data.

CICS journals and IMS logs

CICS keeps track of database changes in its journal data sets. IMS keeps track of database changes in its online log data sets. Once these data sets become full, both CICS and IMS offload the logs to tape.

CICS and IMS logs are sequential data sets. When offloading these logs to tape, you must request a scratch volume every time.

The logs contain information necessary to recover databases and usually those logs are offloaded, as with DB2, in two copies, one local and one remote. You can write one local copy and then create the second for disaster recovery purposes later, or you can create the two copies in the same job stream.

With VTS, you can create the local copy directly on VTS virtual volumes, then copy those volumes to non-VTS tape drives, or to a remote IBM VTS.

Having a local copy of the logs written inside the VTS allows you faster recovery, because the data will stay in the tape volume cache for some time.

When recovering a database, you can complete backout operations in significantly less time with the VTS, because, when reading logs from tape, IMS uses the slow *read backward* operation (100 KB/s) on real tape drives. With the VTS, the same operation is much faster, because the data is read from TVC. Lab measurements do not see much difference between read forward and read backward in a VTS; both perform much better than on physical drives. The reason is not just that the data is in the tape volume cache, but the VTS code also fully buffers the records in the reverse order that they are on the volume when in read backwards mode.

Another benefit VTS gives to recovery operations is the availability of up to 256 virtual drives, allowing you to mount several logs concurrently and therefore to back out the database to be recovered faster.

The IMS change accumulation utility is used to accumulate changes to a group of databases from several IMS logs. This implies the use of many input logs that will be merged into an output accumulation log. Using the VTS, you can use more tape drives for this function.

Image copies

Image copies are backup copies of the IMS databases. IMS can create only full image copies. To create an image copy of a database, use a batch utility, copying one or more databases to tape.

With the VTS you do not have to stack multiple small image copies to fill a tape cartridge. Using one virtual volume per database does not waste space, because the VTS then groups these copies into a stacked volume.

IMS, unlike DB2, has a batch function that works with databases through an IMS batch region. If you do not use logs when running an IMS batch region, then in order to recover the database, you must use an image copy taken before running the batch job. Otherwise, you can use logs and checkpoints, which allows you to restart from a consistent database image taken during the batch execution processing. Using VTS, you can access these image copies and logs at a higher speed.

The VTS volume stacking function is the best solution for every database backup, because it is transparent to the application and does not require any JCL procedure change.

5.6.3 Batch data

Other applications that write to tape and benefit from using the VTS include:

- VSAM REPRO
- IEBGENER / IEBCOPY / ICETOOL
- DSS data set COPY or DUMP
- DFSMSrmm Tape Copy Tool
- Any other tape copy utility

The amount of data from these applications can be huge if your environment does not use TMM or if you do not have DFSMShsm installed. All such data benefit from using the IBM VTS for output.

With VTS, the application can write one file per volume, using only part of the volume capacity and VTS takes care of completely filling the stacked cartridge for you, without JCL changes.

The only thing you must remember is that, if you need to move the data off-site, you must address a device outside the local VTS, or use other techniques to copy VTS data on other movable tapes, as described in 5.7.6, “Moving data out of the VTS” on page 203.

5.7 Moving data into and out of VTS

In this section, we discuss techniques for moving data in and out of VTS. You can start using the VTS by moving data into it. The best method depends on the application you want to manage with the IBM VTS. We describe two methods:

Phased method: This method consists of starting to use VTS with new allocations. The migration of data takes longer, but it can be more controlled and flexible.

Quick method: Use this method when you want to move existing data into the VTS. It is called *quick* because it swiftly puts all data you want to move under VTS control.

Some hints on how to move data out of the VTS are provided in 5.7.6, “Moving data out of the VTS” on page 203. However, the VTS is a closed-storage method, so you must be careful in selecting data to move into it. You do not want to store a large amount of data in the VTS that will need to be moved back out.

A way of moving all data belonging to one VTS to another VTS by using physical cartridge movement is discussed in 5.8, “Physical cartridge movement” on page 205.

5.7.1 Moving data into the VTS: phased method

The data movement techniques outlined here depend more on changes in parameters, routines, or procedures than on overt data movement.

A phased method, independent of the data you are moving into the IBM VTS, is illustrated in Figure 5-3.

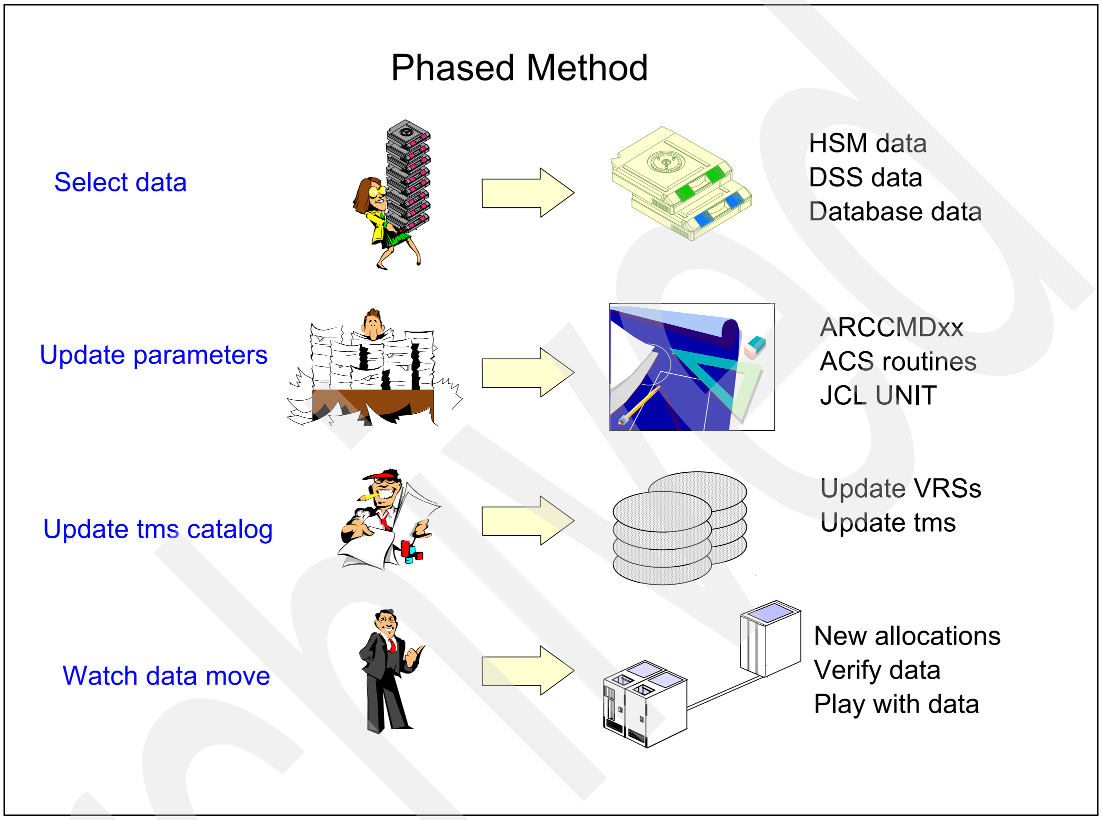

Figure 5-3 Sample process flow to move data in the VTS: phased method

Select the data

If you select DFSMShsm-owned data, you can group your data as listed here:

- Migration data (DFSMShsm level 2)
- Backup copies (user data, CDS data, or both)
- Dump copies
- All of the above

You can select data based on data set name, by application, or by any other variable that you can use in the ACS routines. You can also select data based on type, such as SMF data, or DASD DUMP data.

Update the applicable parameters

If you select DFSMShsm-owned data, review the ARCCMDxx member according to the previous recommendations in this chapter and update the following definitions:

- The data class ACS routines (if you are using data class)
- The storage class ACS routines (required)
- The storage group ACS routines (required)
- For BTLS, the unit parameter in the JCL

For DFSMSdss, update the following definitions:

- The data class ACS routines (if you are using data class)
- The storage class ACS routines (required)
- The storage group ACS routines (required)
- For BTLS, the unit parameter in the JCL

If you use database data, such as logs or image copy, direct new allocations into the VTS by updating the following definitions:

- The data class ACS routines (if you are using data class)
- The storage class ACS routines (required)
- The storage group ACS routines (required)

If you are using BTLS with the VTS to store database data, make the following updates:

- For DB2 archive log, update the archive log data set parameters in the DB2 DSNZPARM initialization member.
- For IMS logs, tailor the appropriate IMS PROCLIB member.
- For CICS logs, update your journal archive procedures.

For data other than DFSMShsm and DFSMSdss, if you are using SMS tape, update the ACS routines to include the data you want to move. You decide what you filter for and how you write the ACS routines. If you are using BTLS, update the UNIT parameter in the JCL to reflect the applicable unit for the VTS.

Update the tape management system

Even though you are not overtly copying data in this option, be sure to update the TMS catalog to reflect the changes that you expect. Check the retention rules and limits and update accordingly. Refer to 5.2, “DFSMSrmm and other tape management systems” on page 186 for more information.

Watch the data move to the VTS

Data movement using this option does not involve overt actions, such as COPY, RECYCLE, or DUMP. When you activate the ACS routines containing the code for the VTS, all new data allocations for the data you have selected are written to the IBM VTS. You simply verify that data is going where you expect it to go and you add code to the ACS routines to manage more data as you see fit.

You can select data types that create large quantities of data, like SMF records or DASD DUMPS and you can also select data types that create many very small data sets. By observing how the VTS handles each type of data, you become familiar with the VTS, its functions and capabilities.

5.7.2 Moving data into the VTS: quick method

The steps outlined in this section involve overt actions on your part to move data into the IBM VTS. As with the techniques outlined in 5.7.1, “Moving data into the VTS: phased method” on page 197, you choose the data you want to move to the IBM VTS. A quick method, independent of the data you are moving in the IBM VTS, is illustrated in Figure 5-4.

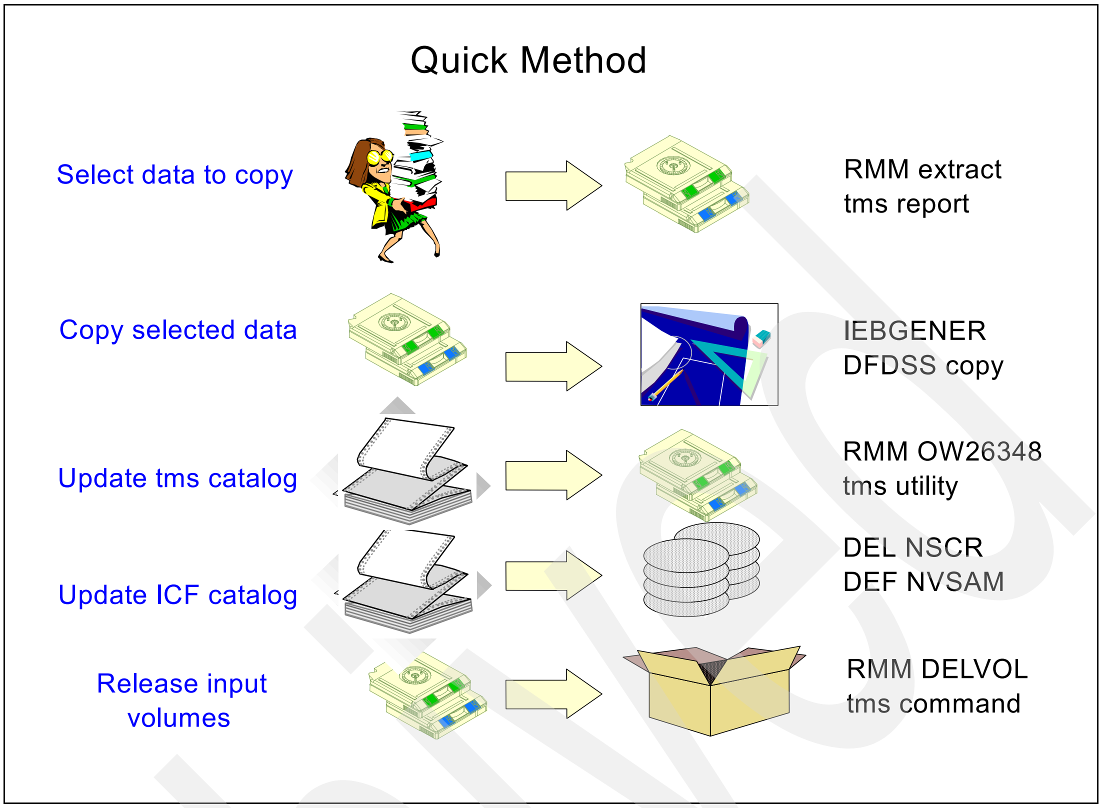

Figure 5-4 Sample process flow to move data in the VTS: quick method

Select the data to copy

The data you select influences all subsequent steps in this process. If you select DFSMShsm-owned data, the process for moving the data to the VTS will be different from the process that you use for DFSMSdss data. You can select data based on the data's attributes, as when using expiration date as a basis for your choice. For example, you could select data that you keep for 7 years. The choice is yours, because it is your data and your environment and you understand it better than anyone else. Probably the best method of selecting data to copy in the VTS is based on the data set name, application by application.

You need to be aware that some applications have knowledge of the volser where the data is stored. There are special considerations for these applications. If you change the volser that the data is on, the application will have no way of knowing where the data resides. For more information about this topic, see 5.7.4, "Considerations for static volsers" on page 202.

An easy method is to gain information from the TMS database. Some reports can give you details on the data you actually have in the tape shop, helping to select the input volumes.

If you are using DFSMSrmm, you can easily acquire data from an RMM EXTRACT file which is normally created as part of the regular housekeeping. Then with a REXX EXEC you extract the information needed, such as data set name, volser and file sequence of the input volumes.

Copy data with a tape copy utility

If you are using SMS tape, update the ACS routines to include the data that you are copying to the VTS. If you are using BTLS, update the UNIT parameter in the JCL to include the valid unit for the VTS.

After selecting data, the next step is to create a job stream that copies the data with the right utility from tape to tape.

If you selected DFSMShsm-owned data, use DFSMShsm recycle to move the data to the VTS. Use a COPYDUMP job to move DFSMSdss data to the VTS.

The utility to use depends on the data selected. In most cases, it is sequential data that can be copied using the IEBGENER utility, DITTO/ESA. If you have DFSORT™ or a similar utility, ICEGENER and ICETOOL can probably give better performance.

You must use a specific utility when the input data is in a special format, for example, DFSMSdss dump data. DFSMSdss uses a 64 KB block size and only the proper DSS utility, like COPYDUMP, can copy with that blocksize. Also take care when copying multifile multivolume chains, you may wish to separate these files since there is no penalty when they are in a VTS.

Update TMS with the correct retention information

When the copy operation has been successful, update the tape management system catalog. This data must be updated on the output volume:

- File sequence number
- Creation date and time
- Last read and last write date
- Jobname

Optionally, you can also update:

- Stepname
- DDname
- Account number
- Device number

In RMM, this can be easily done with a CHANGEDATASET command with special authority to update O/C/EOV recorded fields (OW26348).

Update the ICF catalog with the correct output volume

The next step is to uncatalog the input data sets (if they were cataloged) and recatalog the output data sets with the new volume information.

Release the input volume for SCRATCH processing

This final step needs to be done after you are sure the data has been correctly copied. You also need to verify that the retention and catalog information is correct.

Using this quick method sequence, you can copy every kind of tape data, including GDGs, without modifying the generation number.

In an RMM environment, you can use a REXX CLIST and RMM commands, listing data from the input volumes and then using the RMM REXX variables with the CD command to update the output. Afterward call IDCAMS to update the ICF catalog. When the operation completes and all errors have been corrected, use the RMM DELETEVOLUME command to release the input volumes.

Refer to the *z/OS DFSMSrmm Guide and Reference*, SC26-7404, for more information about RMM commands and REXX variables.

If you are using a TMS other than RMM, refer to the appropriate product functions to obtain the same results.

The first approach when migrating data inside the VTS can be easier with products like DFSMShsm or Tivoli Storage Manager. If you are planning to put DFSMShsm or Tivoli Storage Manager data in the VTS, see Sections 5.1, “DFSMShsm” on page 178 and 5.3, “Tivoli Storage Manager” on page 188.

With DFSMShsm, you can change the ARCCMDxx tape device definitions addressing an esoteric name with VTS virtual drives (in a BTLS environment) or changing SMS ACS routines to direct DFSMShsm data in the VTS. The DFSMShsm RECYCLE command can help speed the movement of the data.

A similar process can be used with Tivoli Storage Manager, changing the device class definitions for the selected data to put in the VTS and then invoking the space reclamation process.

If you are moving DB2 data into the VTS, be sure that when copying the data, the DB2 catalog is also updated with the new volume information. You can use the DB2 MERGECOPY utility to speed up processing, using VTS virtual volumes as output.

5.7.3 Products to simplify the task

You may want to consider using a product designed to copy data from one medium to another. The first choice is the IBM enhancement for DFSMSrmm called Tape Copy Tool (see Table 5-3). The Tape Copy Tool function of the internal IBM ADDONS package is designed to copying all types of MVS tape data sets from one or more volumes or volume sets to a new tape volume or tape volume set. Any tape media supported by DFSMSrmm is supported by this tool. The input tape media can be different than the output tape media. This tool should not be used to copy tape data sets owned by DFSMShsm (Hierarchical Storage Manager) or TSM (Tivoli Storage Manager) or similar program products, because the information of the new volser(s) is not automatically updated in their databases.

The DFSMSrmm Copy Tool cannot be used when you have another Tape Management System than DFSMSrmm. You need to choose another tape copy tool from Table 5-3.

Be sure to consider these factors when evaluating a tape copy product:

- Interaction with your tape management system
- Degree of automation of the process
- Speed and efficiency of the copy operation
- Flexibility in using the product for other functions such as duplicate tape creation
- Ease of use
- Ability to create a pull list for any manual tape mounts
- Ability to handle multivolume data sets
- Ability to handle volume size changes whether from small to large or large to small
- Ability to review the list of data sets before submission
- Audit trail of data sets already copied
- Ability to handle failures during the copy operation such as input volume media failures
- Flexibility in being able to filter the data sets by wild cards or other criteria such as expiration or creation date

Table 5-3 lists some common tape copy products. You can choose one of these products or perhaps use your own utility for tape copy. Certainly you do not need any of these products, but a tape copy product will make your job easier if you have many tapes to move into the VTS.

Table 5-3 Tape copy products

	Vendor	Web address for more information
DFSMSrmm enhancement Tape Copy Tool	IBM	http://w3.itso.ibm.com/redpieces/abstracts/zg246751.html
Beta 55	Beta Systems Software AG	http://www.betasystems.com
BrightStor CA-1 Copycat	Computer Associates International, Inc.	http://www.cai.com
BrightStor CA-Dynam/TLMS Copycat	Computer Associates International, Inc.	http://www.cai.com
CARTS-xx or Tape/Copy	Technologic Software Concepts, Inc.	http://www.technologic.com/
FATAR	Innovation Data Processing	http://www.innovationdp.com/
TapeSaver	OpenTech Systems, Inc.	http://www.opentechsystems.com/Tape-Copy.htm
TelTape	Cartagena Software Ltd.	http://www.cartagena.com
Zela	Software Engineering of America	http://www.seasoft.com/

In addition to using one of these products, consider using IBM Global Services to assist you in planning and moving the data into the VTS. For more information about these services, see 3.11.5, “Implementation services” on page 105.

5.7.4 Considerations for static volsers

Some applications have knowledge of the volser of the volume where the data is stored. One example of an application that has this knowledge is DFSMShsm. When moving data for these applications, you have two choices. You can utilize instructions from the application author to copy the data or you can copy the data to a volume with the same volser. For assistance with DFSMShsm tapes, see 5.7.2, “Moving data into the VTS: quick method” on page 198.

The preferred method for moving this type of data is to utilize instructions from the application author. If, however, you must copy the data to a volume with the same volser, take these points into consideration:

- If the source tape is system-managed tape, you cannot have two volumes with the same volser.
- The source and target media may not be the exact same size.
- You cannot mount two volumes with the same volser at the same time.

This is not the preferred method for moving data to the VTS. This method applies only if you have to maintain the data-set-to-volser relationship and it has limitations and weaknesses.

- Copy the non-VTS tape volumes to DASD or other tape volumes.[1]
- If the source volume is 3494-resident, eject the cartridge from the 3494, using the LIBRARY EJECT command or the ISMF EJECT line operator command from the ISMF panel.
- Delete the ejected volume from the tape management system.
- Define the volser range, including the once-duplicated number, to the VTS.
- Update ACS routines so that the data is directed to the VTS. For BTLS, update the UNIT parameter in the JCL.
- Create a job to copy the data currently on DASD to the VTS.[2]
- Run the copy job.
- Update the tape management system records and any other catalog structures.

5.7.5 Combining methods to move data into the VTS

You will most likely want to use a combination of the phased and quick methods for moving data into the VTS. One approach is to classify your data as static or dynamic.

Static data is information that is going to be around for a long time. This data can only be moved into the VTS with the quick method. You have to decide how much of this data is to be moved into the VTS. One way to make this decision is to examine expiration dates. You can then set a future time when all volumes or a subset is copied into the VTS. There may be no reason to copy volumes that are going to expire in two months. By letting these volumes go to scratch status, you will save yourself some work.

Dynamic data is of temporary nature. Full volume backups and log tapes are one example. These volumes typically have a very short expiration period. You can move this type of data with the phased method. There is no reason to copy these volumes if they are going to expire soon.

5.7.6 Moving data out of the VTS

There are many reasons why you would want to move data out of the VTS. The most common reason would be for disaster recovery or data interchange. You can move data out of the VTS in three ways (see Table 5-4).

Table 5-4 Methods of moving data out of the VTS

	Import/Export	**Host copy**	**ABARS**
Source copy retained	No	Yes	Optional
Compatibility	With another VTS	Depends on target	Depends on target (must have DFSMShsm)
Target aggregation (volser selection)	Yes	Maybe	Yes
Requirements	VTS	Standard tools or product	DFSMShsm

[1] The size of the DASD buffer needed to perform this step could be huge. Be sure you have sufficient space.
[2] Be sure that the size of the non-VTS volume previously copied to disk is compatible with the VTS volume sizes, 400 or 800 MB (times the compression ratio if EHPO installed).

	Import/Export	Host copy	ABARS
Data stacked on output media	Yes	Possibly	Yes
Source aggregation multiple sources	No	No	Yes
Outboard data movement	Yes	No	No

You can use the Import/Export feature of the VTS to move data out of a VTS. This data can only be brought into the same or another VTS. This method is useful if you are relying on the availability of another VTS available at the recovery site. It can also be used for data interchange between two sites that both have a VTS. For more information about the Import/Export feature, refer to Appendix B, “VTS Import / Export” on page 369.

With DITTO/ESA for MVS you have the alternative to create copies of exported logical volumes outside the normal Export/Import procedure. For more information about the new DITTO/ESA function, refer to , “DITTO VTS functions” on page 414.

The second method of moving data out of the VTS is to use a host-based tool to copy the data from the VTS to the target. With this method the data is reprocessed by the host and copied to another medium.

This method is described in 5.7.1, “Moving data into the VTS: phased method” on page 197. The only difference is that you need to address the VTS as input and the non-VTS drives as output.

The third method is to copy the data with the DFSMShsm ABARS function. This method is described next.

ABARS is the command-driven DFSMShsm function that backs up a user-defined group (called an *aggregate group*) of data sets (usually for recovery purposes) at another computer site or at the same site. ABARS can be used to back up and recover both SMS- and non-SMS-managed data, on DASD and on tape.

Using the DFSMShsm ABARS function, group the data you want to move outside the VTS, then start addressing other tape drives outside the IBM VTS. In this way, you obtain an exportable copy of the data that can be put in an off-site location.

Here is an overview of the process.

Create selection data set

Before you can run aggregate backup, you have to create one or more selection data sets. The selection data set lists the names of the data sets to be processed during aggregate backup.

You can identify the data set names in a single selection data set, or you can divide the names among as many as five selection data sets. You can specify six types of data set lists in a selection data set. The type you specify determines which data sets are backed up and how they are recovered.

An INCLUDE data set list is a list of data sets to be copied by aggregate backup to a tape data file where they can be transported to the recovery site and recovered by aggregate recovery. The list can contain fully qualified data set names or partially qualified names with place holders. DFSMShsm expands the list to fully qualified data set names.

Using a selection data set with the names of the data sets you want to export from the IBM VTS, you obtain a list of files on logical volumes that the ABARS function copy to non-VTS drives.

Define aggregate group

You have to define an aggregate group and related management class to specify exactly which data sets are to be backed up.

You define the aggregate group and management class used for aggregate backup to DFSMS through ISMF panels.

The aggregate group lists the selection data set names, instruction data set name and additional control information used by aggregate backup in determining which data sets are to be backed up.

Execute ABACKUP VERIFY command

You have two options with the ABACKUP command. You can choose to verify the contents of the aggregate backup without actually backing up any data sets, which is the same as performing a test run of aggregate backup.

An example of the ABACKUP command is:

```
HSEND ABACKUP agname VERIFY UNIT(non_VTS_unit) PROCESSONLY(USERTAPE)
```

With the PROCESSONLY(USERTAPE) keyword, only tape data sets are processed. In this way you can be sure that only the input data from VTS logical volumes is used.

Execute ABACKUP EXECUTE command

When you are ready, you can perform the actual backup, using this command:

```
HSEND ABACKUP agname EXECUTE UNIT(non_VTS_unit) PROCESSONLY(USERTAPE)
```

When you issue the ABACKUP command with the execute option, the following tape files are created for later use as input to aggregate recovery:

- **Data file:** Contains copies of the data sets that have been backed up.
- **Control file:** Contains control information needed by aggregate recovery to verify or recover the application's data sets.
- **Instruction/activity log file:** Contains the instruction data set, which is optional.

At the end of this process, you obtain an exportable copy of the VTS data, which can be used for disaster recovery and stored off-site using other physical tapes.

For more information about using the DFSMShsm ABARS function, refer to the *DFSMShsm Storage Administration Guide*, SC35-0421.

5.8 Physical cartridge movement

Figure 5-5 shows the source and target configuration when you migrate all data belonging to one VTS to another by using physical cartridge movement. Just as a sample the source VTS is a Model B18 VTS and the target VTS is a Model B20 VTS.

This method cannot be used to move the workload from a Model B18 VTS with EHPO to a Model B18 VTS without EHPO or from an IBM 3590 Enhanced Model based VTS to an IBM 3590 Base Model based VTS. This method requires close cooperation between the customer and the responsible IBM System Service Representative (SSR) during the planning and execution of the physical cartridge movement.

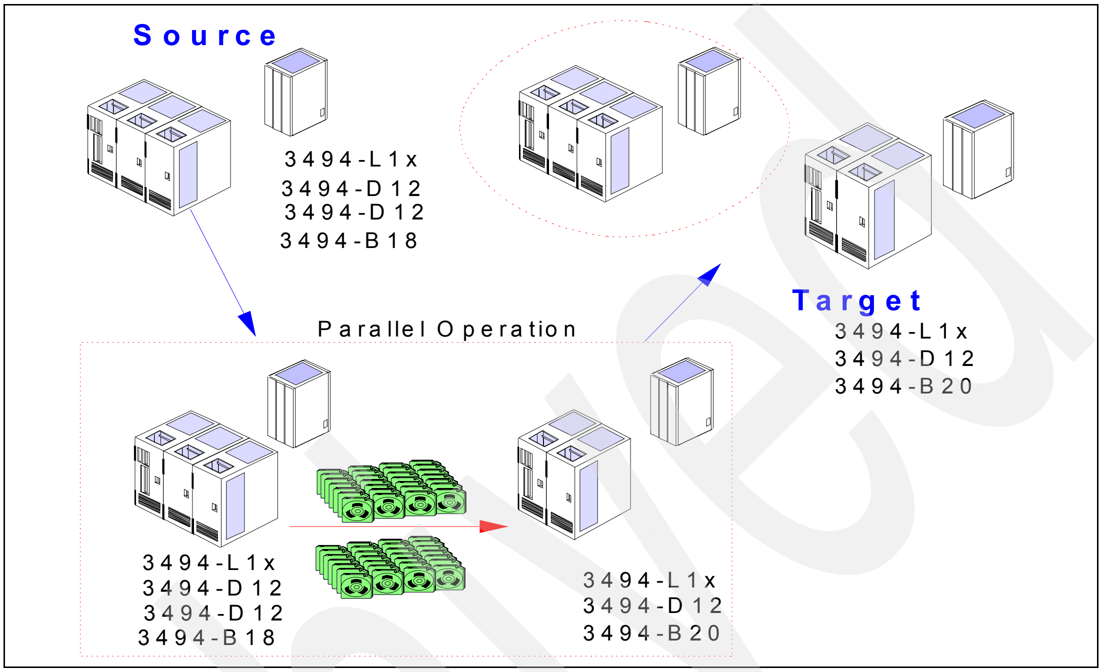

Figure 5-5 Physical cartridge movement

These are the steps to be executed:

Install a target Virtual Tape Server, for example, a Model B20 VTS:

- At least one D12 Frame, one L1x Frame Control Unit frame and a Model B20 VTS frame.
- Define the HCD definitions for the new library.
- Define the new library through ISMF.
- Define or update SMS constructs for the new 'logical' library, especially storage groups.
- Prepare ACS routine changes.
- If applicable, update JES3 definitions.
- Vary the new Model B20 VTS online.

Some of the subsequent tasks are done by the IBM service representative using specialized tools.

Move the content of the Source VTS to the Target VTS:

- Initialize concurrent fast migrate to drain the TVC of the Source VTS.
- Vary the Source VTS offline.
- Copy all TVC resident logical volumes to stacked cartridges using Fast Migrate function.
- Copy TVC fragments and VTS database to service cartridges (3590).
- Copy Library Manager database to diskette.
- Move and load stacked cartridges into the new library.

Note: This might involve physical transportation of data if the new VTS is in a different location. The duration and difficulty of this step should not be underestimated.

- Update the new Library Manager database with data from the old library.

- Restore the VTS database and the TVC fragments onto new Model B20 VTS.
- Reinventory the library for physical cartridges, keeping the logical volumes.
- Vary the Target Model B20 online.
- Update TCDB and tape management system CDS regarding location of the logical volumes if needed.

Redirect the workload to the Target Model B20 VTS library, using nondisruptive adjustments in the DFSMS ACS routines and constructs.

5.9 Migration scenarios when installing 3592 drives into a VTS

You may want to migrate your data to the new technology tape drives when you install 3592 drives as backend drives into an existing VTS where you already have 3590B/E/H drives installed. Or you want to use the new drives for new data.

We decided not to document these scenarios in this redbook. Instead, we refer you to the White Paper (17 pages), which is available on the Internet at:

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

Please ask your SSR if you do not have access to this document.

5.10 Summary

The VTS is an excellent choice for storing most tape data, with the possibility to handle off-site backups and interchange data that need to be taken out of the library using the Advanced Function feature. The VTS utilizes cartridge capacity fully, is easy to use, enables you to run more concurrent tape applications and in many cases offers very fast mount times by eliminating physical mounts entirely.

Native tape drives, particularly TotalStorage 3590 or 3592 drives, or a TMM implementation may in some cases offer advantages over the VTS. Remember, however, that directing tape data to different media on a selective basis requires more tuning effort from the storage administrator than an “all-in-VTS” solution. If you have native tapes drives or have implemented TMM and want to keep using them for storing part of your tape data, consider the following points when you decide when to use each:

- 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. The VTS offers better response times for scratch mounts than native tapes. Write-intensive functions, such as backup and dump, benefit particularly from fast VTS scratch mounts.

- The VTS offers better response times for private mounts than native tapes, provided the referenced volume resides in the VTS TVC. Intermediate data sets that are used by several jobs or job steps in succession benefit from this and are good candidates for the VTS.

- Applications that require very high tape bandwidth and fill up high capacity cartridges may be more effective if kept on native IBM 3590 drives that offer higher megabytes/second throughput rates than the VTS.

- With VTS, we recommend that you write only one data set per volume if you are not going to use all the data sets on the volume at once. When you have more than one data set on a logical volume, the VTS recalls the whole volume in the tape volume cache even if you need only one of the data sets. This lengthens mount times and fills the tape volume cache unnecessarily.
- Small data sets, when stored one per logical volume, quickly increase the number of logical volumes you have to insert in the VTS. This increases the number of entries in the Library Manager database, in the VTS database, in your tape management system catalog and in the SMS volume catalog. If you have already implemented TMM, you will need to evaluate the costs of managing these small data sets in the VTS. The maximum of 250,000 logical volumes per VTS or 500,000 logical volumes per IBM 3494 Tape Library will have to be taken into consideration also.
- To protect your data against media failures, do not store a data set and its backup in the same VTS because you have no control over which logical volumes end up on which stacked volumes. For example, a data set on a DFSMShsm ML2 volume and its backup on an DFSMShsm backup volume may end up on the same stacked volume. With APM, dual copy can further protect your data.
- If you regularly need to append a data set daily, native tape drives may offer better performance, because a logical volume has to be recalled into the TVC before it can be appended. This will, of course, depend on the size of your TVC and the activity that has taken place since the last append.
- As a general rule we recommend that you use ECCST emulated logical volumes, so that you can maximize the storage capability of your VTS. Remember that you do not waste stacked volume capacity if not writing a logical volume full.
- Use of small logical volumes improves recall performance. Functions that can benefit from small logical volumes are, DFSMShsm RECALL, Tivoli Storage Manager RESTORE and OAM object RETRIEVE. Remember, each of these products provides a command or parameter that enables you to limit the amount of data that they write on tape volumes; you can therefore achieve the performance of CST emulated volumes even when using ECCST emulated volumes.
- Some functions restrict the media type you can use. For example, DFSMShsm TAPECOPY requires that the input and output volumes are of the same media type. Therefore your physical cartridges may dictate the VTS logical volume type you can use.

Table 5-5 summarizes our considerations regarding placement of different types of tape data sets. When VTS is not an obvious first choice, we indicate an alternative.

Table 5-5 Tape data placement

Data type	Device	Comments
Batch		
Small data sets	1. VTS 2. TMM	Prefer VTS for small data sets as long as the number of logical volumes is not a constraint. Each VTS can have a maximum of 500,000 logical volumes. TMM is also a possibility especially if you are already using it.
Medium and large data sets	1. VTS	These go well in the VTS. VTS does not waste cartridge capacity like native 3590 and requires less cartridge shelf space in the library than native 3490.

Data type	Device	Comments
Very large data sets	1. 3590/3592 2. VTS	Native 3590 or 3592 offers better performance for very large batch data sets. Use 3590 or 3592 if you can do it without wasting too much cartridge capacity.
Intermediate data sets	1. VTS	The VTS is well suited for data sets that pass information from one job or job step to the next. With the VTS, you avoid mount delays if the data set is referenced before it is deleted from the VTS TVC.
DFSMSdss		
Physical dumps	1. 3590/3592 2. VTS	Prefer native 3590 or 3592 drives for performance reasons. Concurrent full-volume physical dumps with OPT(4) quickly use more bandwidth than a VTS can offer. VTS may be an option for physical dumps without OPT(4). You need to design your jobs so as to stack multiple dumps on one volume to utilize 3590 or 3592 cartridge capacity.
Logical dumps	1. VTS	VTS is good for logical dumps which require less bandwidth than physical dumps. With VTS it is easy to utilize cartridge capacity. You also benefit from fast virtual scratch mounts.
DFSMShsm		
ML2	1. 3590/3592 2. VTS	DFSMShsm can use the full capacity of native drives. Recalls from 3590 or 3592 are generally faster.
Backup	1. VTS	Although recovery from native 3590 or 3592 is generally faster, it is less of a concern because you should only seldom need to recover backups. By using the VTS, you can benefit from fast scratch mounts during backup.
CDS backup	1. VTS 2. TMM	With the VTS you can benefit from fast scratch mounts. If you are already utilizing TMM, you may want to keep them in TMM.
Volume dumps	1. 3590/3592 2. VTS	With the prerequisite z/OS DFSMS maintenance, DFSMShsm can stack multiple dumps on one volume. With the VTS you can also stack many dumps on one physical cartridge. You also benefit from fast scratch mounts.
ABARS	1. Off-site	Typically needs to go off-site
Tape copy	1. Off-site	Typically needs to go off-site

Data type	Device	Comments
Tivoli Storage Manager		
Tape storage pools	1. 3590/3592 2. VTS	Tivoli Storage Manager can use the full capacity of native 3590 or 3592. Restores from native is generally faster.
Storage pool backups, incremental	1. VTS	Good for VTS use. You benefit from fast scratch mounts and can fully use cartridge capacity. If the primary storage pool is on native 3590, the backup cannot end up on the same volume as the primary data.
Storage pool backups, full	1. Off-site 2. VTS	Often need to go off-site.
Database backups, incremental	1. VTS 2. TMM	These are typically small, so they should work fine in the VTS. Could stay in TMM if you already are utilizing it.
Database backups, full	1. Off-site 2. VTS	Good for the VTS, if backup must go off-site, use Advanced Function feature. The VTS does not waste cartridge capacity, as native tapes do and gives you fast scratch mounts.
OAM Objects		
OBJECT storage groups	1. 3590/3592 2. VTS	OAM can use the full capacity of native 3590 or 3592. Retrieval from native is generally faster.
OBJECT BACKUP storage groups	1. Off-site 2. VTS	Often need to go off-site. When the primary object is on native drives, there is no danger that it and its backup will end up on the same physical volume.
Database Systems		
Archive logs and journals	1. VTS	These go well in the VTS. You benefit from fast scratch mounts.
Dual archive logs and journals	1. Off-site	These typically go off-site.
Local image copies, incremental	1. VTS 2. TMM	The VTS should be fine for these data sets. TMM should work if you are already utilizing it.
Local image copies, full	1. VTS	The VTS is good here.

Note: The Device column indicates the primary and optionally a secondary, recommendation for storing the data type. The key to the Device column is:

VTS Use VTS.
3590 Use native 3590 drives, if available.
TMM Refers generally to an implementation where data is initially written on DASD and later migrated to tape with DFSMShsm.
Off-site Use preferably native 3590 or 3490 drives (stand-alone or in the 3494), use a VTS at the remote site, if available, or use the Advanced Function feature Import/Export component.

6

Operating the VTS

In this chapter we describe operational considerations unique to the IBM TotalStorage Virtual Tape Server (VTS).

For general guidance on how to operate an IBM TotalStorage Enterprise Tape Library (3494), refer to the *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280, and the *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632.

For general guidance to operating the 3584 and 3953, refer to the *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning, Implementation, and Monitoring,* SG24-6789. For guidance on how to operate the 3953, refer to the *IBM TotalStorage 3953 Tape Frame Model F05 and Library Manager Model L05 Operator Guide,* GA32-0473. For guidance on how to operate the 3584, see the *IBM TotalStorage 3584 Tape Library Operator Guide,* GA32-0468.

Note: We use the term *Library* to refer to the 3494 Tape Library and the 3584/3953 Tape Library system. In most cases, what applies to the 3494 also applies to the 3584/3953. When specific differences occur, we explain what this difference is and what the specifics are for each library.

We also use the term *tape drive* to refer to both the 3590 models and 3592 models of tape drives. Where there are differences, we specify these.

6.1 Power control

In this section we discuss the power on sequence for the library and the VTS. We also discuss the differences between a VTS on a 3494 Tape Library and a 3953/3584 Tape Library configuration.

6.1.1 Powering on the VTS as part of a 3494

The VTS is powered on as part of the 3494 power-on sequence. The Unit Emergency switch on the Model B18 VTS, Model B10 VTS, and/or Model B20 VTS frame must be set to the ON position. The VTSs are always the last frame to power on in the sequence. Powering on the VTS can be done independently of the Library as part of certain service functions.

Once the VTS has been powered on, ensure that all virtual drives are varied online to the appropriate host.

6.1.2 Powering off the VTS as part of a 3494

The VTS is powered off as part of the 3494 power-off sequence. You should not power off a VTS independent of the rest of the 3494 where it resides by using the VTS Unit Emergency Switch.

Ensure that all virtual devices are varied offline from the host before you power off the tape library that contains a VTS. Note that the VTS will not go offline if there is a *Reclamation* task currently active. Make sure you *Inhibit Reclamation* processing at least three hours before you are planning to power down the VTS. See 4.3.7, “Define VTS management policies” on page 137.

In a library with a VTS, the system can take up to 20 minutes to shut down. When the shutdown completes, the 3494 (including the VTS, if installed) is powered off.

To service the VTS, the IBM service representative can take a VTS offline from the Library Manager and shutdown and restart the VTS controller independent of the rest of the 3494 by using the Library Manager's Service Mode menus. This will allow other hosts, either zSeries or Open Systems attached through their respective native tape drives, to continue to operate independent of the 3494 VTS.

Attention: Using the Unit Emergency switch for immediate power off can cause database problems or check disk (CHKDSK) problems with the Library Manager or VTS controller. The VTS will not be powered off by the 3494 Unit Emergency switch, but the associated 3590 tape drives will have power removed.

6.1.3 Power control when part of a 3584/3953 system

When the 3494 VTS is part of a 3584/3953 system, there is no power sequence connecting the 3494, the 3953, and the 3584. Each machine must be powered on or off by its own unique operation methods.

Powering off a complete library system

The process of powering off a complete library system would include:

- Varying offline all virtual and native devices on all hosts (zSeries or Open Systems)
- Perform a soft shutdown of the Master Console
- Shutting down the active 3953 Library Manager

- Selecting “Shutdown Computer for Power Off” on 3953 Library Manager screen (this will issue a quiesce and subsequent shutdown of all attached VTSs and 3592-J70 control units attached to this 3953)
- Shutting down the standby 3953 Library Manager (if installed)
- Powering off Library Manager (and standby manager installed
- Powering off 3584 Tape Library

Powering on a complete library system

The process of powering on a complete library system would include:

- Powering on the 3584 Tape Library
- Switching on the Library Manager (and standby Library Manager if installed)
- Switching on the Master Console
- Switching on the 3592-J70 Controllers (if installed)
- Switching on the KVM switch
- Switching on the VTS

Further detailed information can be found in the *IBM TotalStorage 3953 Library Manager Model L05 Operator Guide*.

6.2 Operational states and modes

In this section we discuss the operational states and modes of both the 3494 Tape Library and the 3953/3584 Tape Library configuration. The Library Manager in the 3494 and the 3953-L05 Library Manager are very similar in operation. In fact, most of the microcode running on the 3953-L05 is the same as for the 3494 Library Manager. The difference is in the fact that the 3953-L05 does not directly control the accessor in the library.

6.2.1 3494 Tape Library

The installation of the VTS does not bring any changes to the operational modes and states of the 3494. For a detailed explanation of the modes and states, refer to the *IBM TotalStorage 3494 Tape Library Operator's Guide,* GA32-0280.

The operational modes are:

- Auto
- Pause
- Manual

The operational states are:

- Library Manager initialization
- Initialization complete
- Online
- Offline
- Shutdown pending
- Library Manager switchover in progress (only the 3494 with HA1 Frames installed)
- Accessor switchover in progress (only the 3494 with HA1 Frames installed)
- Dual active accessor enabled/disabled (only the 3494 with dual active accessor feature)
- Dual active accessor transitioning (only the 3494 with dual active accessor feature)

6.2.2 3584/3953 Tape Library

Some library functions that are controlled by the IBM 3494 Library Manager are not relevant in the IBM 3584 environment. Library functions are managed entirely within the IBM 3584 Tape Library, and the IBM 3953 Library Manager is not involved. Hence some of the displays or information at the operator console are no longer used:

- IBM 3584 does not support High Capacity I/O, which is assigning some block of library cartridge slots to be used for bulk input or ejection of cartridges. This means that any display of this is removed in the IBM 3953 Library Manager.
- Dual accessors in the IBM 3584 Tape Library are always in dual active mode. The activity and any accessor failures are managed by the library intelligence without involvement from the Library Manager. Individual accessor mounts per hour are not displayed since the Library Manager has no visibility of which accessor performed the mount.
- The Library Manager is not involved in and has no visibility of how a mount/demount is performed. Hence the IBM 3953 Library Manager does not display Home Cell mode: floating or fixed; these are functions handled within the IBM 3584 Tape Library.
- In the IBM 3494 Tape Library you can assign a default media type during the Teach process which will be used when the media type is not determined from the barcode or VOLSER ranges. The default media type is meaningless for an IBM 3953 Library Manager since the IBM 3584 Tape Library handles all aspects of label reading.
- The IBM 3953 Library Manager does not support ARTIC cards, which are the means by which the tape controllers and VTSs are attached to the Library Manager inside the IBM 3494. Any display relating to ARTIC cards is removed.
- The IBM 3584 Tape Library handles the availability of the Convenience I/O station and the Library Manager does not display I/O station availability or whether the door is open or closed. However, the Library Manager will display whether the I/O station is in Import or Insert mode.
- Any reference to a frame number in the Library Manager panels is removed since the Library Manager has no knowledge of a physical location within the IBM 3584, and does not manage the physical location within the library. Instead, a subsystem number, based on the IP address of the subsystem, is displayed.
- The volume state of Unreadable does not exist in the IBM 3953 Library Manager. If a volume barcode is unreadable, it is handled by the IBM 3584 Tape Library; it is not added to the IBM 3953 Library Managers database.
- Any references to a cell (Rack, Column, Row) will be changed to reflect the element address. Because of Advanced Library Management System (ALMS), which is required on the IBM 3584 when the IBM 3953 Library Manager is attached, the element address displayed does not relate to a fixed position in the physical library. The IBM 3584’s operator panel or Web interface needs to be used to locate a cartridge’s physical location by Frame, Column, and Row.
- Cleaning is not scheduled through the Library Manager according to day of the week or throughput. The IBM 3592 tape drive indicates when a clean is required and the IBM 3584 Tape Library responds automatically and handles all aspects of the clean. Also, the drive indicates when a cleaning cartridge has reached the end of its life and is automatically ejected by the IBM 3584. Note that the cleaning cartridge has a record on the tape of how many cleaning cycles it has completed, so cleaning cartridges cannot be reinserted after they are ejected.

Operational modes and states

Operational states describe the status of the Library Manager and are mostly unchanged from the IBM 3494, except where previously the state may have resulted from a function no longer controlled by the Library Manager — for example, the dual accessor transitioning or switchover. Operational modes in the IBM 3494 describe accessor movement. The three operational modes are:

- **Auto:** In this mode the library is automatically controlled by the IBM 3584 MCP which receives simple SCSI commands from attached hosts.
- **Pause:** This mode allows the operator to open the tape library door and access the interior of the IBM 3584. While the library is in Pause mode, requests are queued until the library returns to Auto. From Pause mode you can change the library to Auto or Manual.
- **Manual:** In this mode of operation the operator, under the direction of the Library Manager, locates and moves tape cartridges to and from storage cells and tape drives manually. This mode allows data to be retrieved when normal operations are interrupted by unexpected conditions. You can move into this mode only when the IBM 3584 is in Pause mode.

These modes are relevant in the IBM 3953 Library Manager. However, accessor movements are not functions of the Library Manager in the IBM 3953, and you cannot move between Auto and Pause states using the Library Manager operator console; you must use the IBM 3584 operator panel to put the library into Pause mode before, for example, moving to Manual mode of operation. See Figure 6-1 for a view of the 3584 Operator Panel.

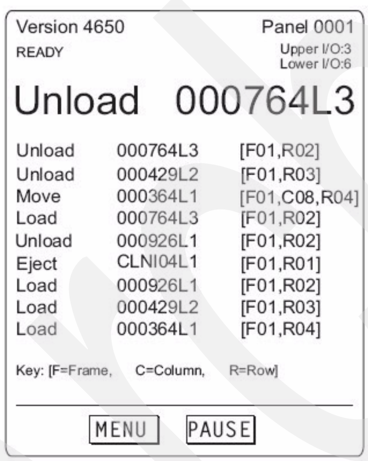

Figure 6-1 Operator panel on the 3584 Tape Library

6.2.3 The informational states of the 3494

The informational states of the 3494 Tape Library are:

- Degraded operation
- Safety interlocks open
- Bar code or vision system not operational
- Intervention required
- Library Manager check-1 condition
- All storage cells full
- Out of cleaner volumes
- Dual write disabled
- Dual Library Manager status (only the 3494 with HA1 Frames installed)
- Accessor status (only the 3494 with HA1 Frames installed)

The modes and states apply to the 3494 as a whole; you cannot set the mode or state of logical libraries independent of each other, with the exception of taking a VTS offline to the Library Manager. In this case only the single VTS is offline — any other defined logical library will still operate as normal.

Note in particular that the Manual mode is supported by the VTS. In Manual mode the Library Manager console provides the necessary information for the operator to manually mount stacked volumes to satisfy mount requests from the VTS controller.

6.2.4 The informational states of the 3584/3953

Informational states of the 3584 Tape Library and the 3953 Library Manager differ little when compared to the 3494 (Figure 6-2). The main difference is where to view information about the state. For instance, the 3953 Library Manager has no specific information about the state of the accessors in the 3584. If the accessors are Paused, then the Library Manager is Paused.

If a whole or part of an accessor is disabled, then the Library Manager cannot display this information. Specific hardware errors on the 3584 get posted, but the 3584 Specialist is the place to view the status of and manage the library. See the *IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0472 for details.

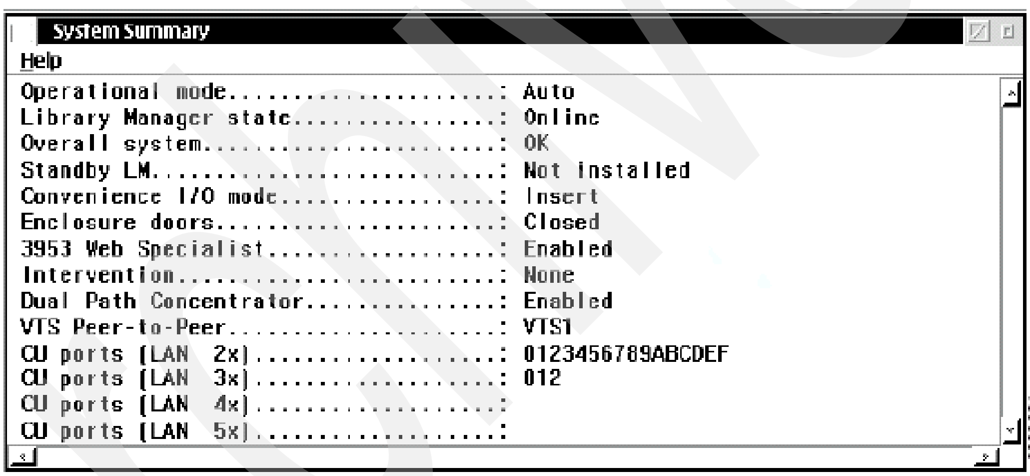

Figure 6-2 System summary screen of the 3953 Library Manager

6.3 Conditions that require intervention

The intervention-required facility of the Library Manager is used to report error and status conditions specific to the VTS and to indicate that assistance is required. When the Library Manager determines that an error or exception condition has occurred within the library that requires operator assistance to correct, it performs the following actions:

- It adds the condition detected to the list of outstanding operator intervention required conditions maintained by the Library Manager.
- If not previously in the intervention required state, it places the library in that state and notifies all attached hosts of the state change.

The Library Manager notifies all attached hosts of an intervention-required condition in a 3494 regardless of which logical library (VTS or non-VTS) generated the intervention.

The Library Manager keeps track of the outstanding intervention required conditions. The operator can display the associated intervention required messages in the Operator Intervention window on the Library Manager console (Figure 6-3) and on the host consoles and can subsequently indicate which conditions have been resolved. To display the messages and to determine the cause of an intervention required condition and the steps required for its resolution, select **Commands**, then **Operator Intervention** from the pull-down menu.

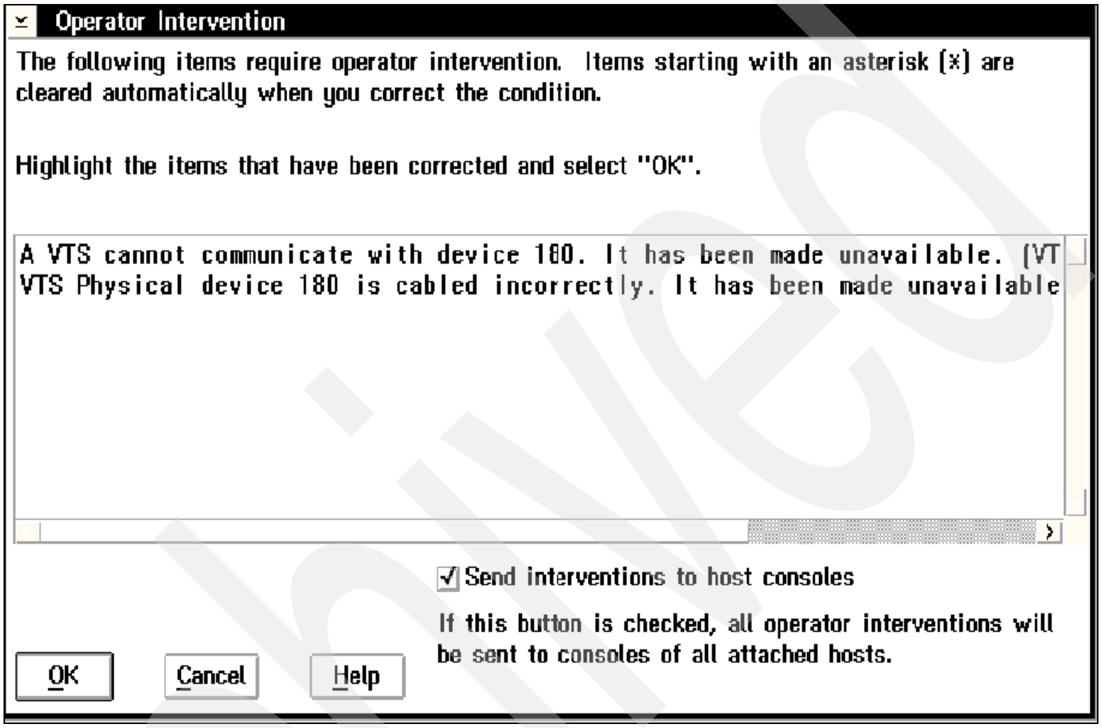

Figure 6-3 Library Manager Operator Intervention window

The notification that the library is in the intervention required state will generate the following host console message (Example 6-1) for libraries managed under the DFSMS managed tape environment:

Example 6-1 Intervention required console message sent to host

```
CBR3762E Library <library-name> intervention required
```

Some messages are retained on the console until all intervention required conditions have been cleared. When both logical libraries of a logically partitioned 3494 or 3584/3953 are connected to the host, the message is issued twice on the host console.

See the *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280 or the *IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0472 for a summary of intervention-required messages relating to VTS in general and the Advanced Function feature specifically.

6.3.1 Intervention-required messages

Additional intervention required conditions with associated messages are generated when a VTS is included in a 3494 or 3584/3953 system. These are of two types:

- Conditions detected by the Library Manager during normal (non-failure) VTS operations that require operator action. An example is when the total effective storage space of the stacked volumes in the library falls below a threshold set by the system administrator.
- Errors and failures detected within the VTS, including its TVC, managed 3590 volumes and 3590 tape drives.

Messaging enhancements

If the option, **Send Interventions to Host Console**, is enabled on the Operator Intervention screen (as shown in Figure 6-3 on page 217), a broadcast message of the operator intervention text is sent to all attached hosts. For libraries managed under the z/OS DFSMS system managed tape environment, this results in the following message (Example 6-2) being displayed at the host console:

Example 6-2 Message from the library

```
CBR3750I Message from Library <library-name>: message
```

This function was introduced at LM 521.11, LM 522.01 and LM 523.06 and allows only 70 characters of data. Therefore, <message text> will consist only of the first 70 characters of the library message text. At LM 524 the message field is increased.

Prerequisites

With the Library Manager code for the advanced functions of VTS you get the support for the extended messaging function as well. The advanced function option is supported on zSeries systems with the appropriate software support as described in 3.8, “Software requirements” on page 93.

If you want to exploit the new messaging functionality but you do not want to change the current software level running your zSeries system, install LM521.08 on your Library Manager. This code level contains the messaging improvements only.

See the *IBM TotalStorage 3494 Tape Library Operator’s Guide* - GA32-0280 or the*IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0472 for a summary of intervention-required messages relating to VTS in general and the Advanced Function feature specifically.

6.3.2 Service Information Messages (SIMs)

When the z/OS maintenance listed in 3.8, “Software requirements” on page 93 is on the system, the VTS generates SIMs to report its error. Message IEA480E is generated for supported hardware problems and lists the severity of the problem. The possible types are ACUTE, SERIOUS, MODERATE, or SERVICE. ACUTE is the most severe and SERVICE is the least severe.

The only way the VTS can communicate to the zSeries host is by way of the 3490 virtual devices defined to the host. Every SIM, regardless of what the actual hardware problem is related to (tape drive, internal VTS hardware, accessor problems, etc.), surfaces as a 3490 alert. The refcode in the SIM will have the details relating to error on the failing component. Contact your IBM service representative when you receive one of these messages. Example 6-3 shows possible console messages. See also “Statistical Analysis and Reporting System (SARS)” on page 268.

Example 6-3 SIM console messages

```
IEA480E 90A9,TCU,3490,acute alert,ser0113-A0151,MC=F1,ES=11,ID=07,
REF=E000-5FCA-0206,UM=0000-0000-80
```

```
IEA480E 0610,TCU,3490,SERIOUS ALERT SER=0113-71621,MC=F1,ES=11,
REF=0000-0001-028C,UM=0000-0000-00
IEA480E 0610,TCU,3490,SERIOUS ALERT SER=0113-71621,MC=F1,ES=11,
REF=0000-0001-01D2,UM=0000-0000-00
```

6.3.3 Simple network management protocol (SNMP) traps

In this section we discuss the Library Manager and how it supports SNMP. If you need more information about SNMP, consult *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280 or the *IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0472.

If you decide to take advantage of the Library Manager's ability to send SNMP traps, you will need to have connectivity to a LAN through either the Token Ring Adapter (FC5219) or the Ethernet Adapter (FC5220). FC5219 is not supported on the 3953 Library Manager.

Along with these features you need to have an SNMP monitoring stations capable of processing the traps sent by the Library Manager. The Library Manager will send the traps to a maximum of five such stations. The SNMP monitoring stations and the various software components needed to capture the traps are standard and are available from IBM and other vendors. Selection of the SNMP monitoring stations hardware and software is not covered in this book. Figure 6-4 shows one possible way of configuring a 3494 and a 3584/3953 with two SNMP monitoring stations.

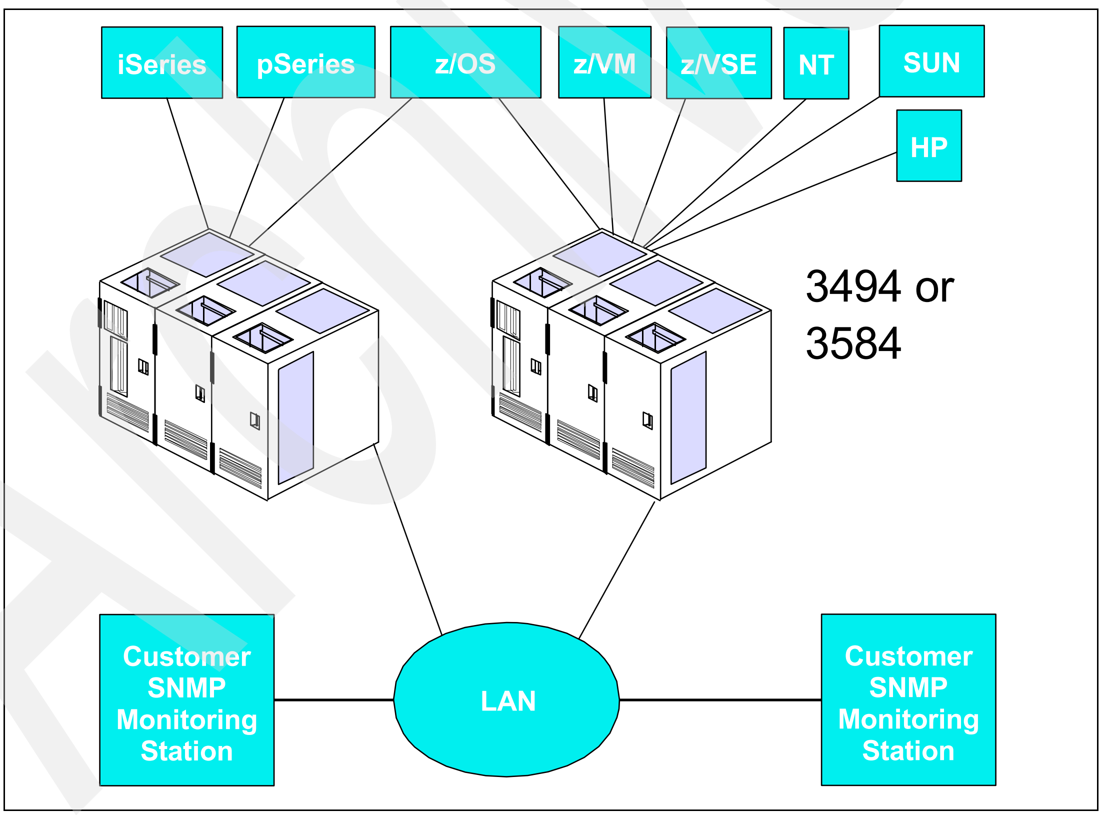

Figure 6-4 SNMP configuration diagram

The Library Manager in either 3494 or 3584/3953 offers limited SNMP support and does not contain any SNMP Management Information Base support. When certain events occur on the Library Manager, an SNMP trap message is sent out. The following events are currently defined:

OPINT: Operator Interventions. These events inform the SNMP monitoring stations of the state of the 3494. They can signify problems and also signal a need for a service call. The OPINT messages are listed in the OPINT message table in the *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280 and *IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0472. These traps should be selected, as they are very useful for monitoring the 3494, VTS, or 3953 Library Manager.

UNSOL: Unsolicited Attention Messages. These events offer additional information but are not as useful as OPINT events. They can signify drive availability and volume movement, among other activities. To send the messages, the 3494 must be online and at least one host must be active. The UNSOL messages are listed in the UNSOL message table in the *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280 or the *IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0472.

SERVC: Service Request Messages. These messages are not currently supported.

CHCK1: Library Manager Check1 Conditions. These events are posted when the 3494 has a problem that requires re initialization.

TESTM: Test SNMP Trap Messages. These events are used to test the ability of the SNMP monitoring stations to receive SNMP traps. They are used to set up the SNMP environment.

Selecting SNMP traps

The first thing you do to set up SNMP is to select the types of traps you would like the Library Manager to generate. To set the types, select **Commands** from the main menu and then **SNMP Options** to display the SNMP Trap Types window (Figure 6-5). Select the traps that you would like generated and click **OK** when finished.

Figure 6-5 SNMP Select Trap Types window

Configuring SNMP trap destinations

After the trap types have been selected, you have to specify where the Library Manager will send the traps. Select **SNMP Options** from the **Commands** pull-down menu, then select **Change SNMP Trap Destinations**. You will now see the window shown in Figure 6-6.

> **Note:** If the 3494 has HA1 Frames, the standby Library Manager does not need to be configured for SNMP. The active Library Manager will configure the standby Library Manager when it becomes active. The standby Library Manager will not send SNMP traps while it is in standby mode. This also applies to a 3953 with dual Library Managers installed.

The SNMP Trap Destinations window (Figure 6-6) is displayed by the HRMCNFIG external process. Once it is displayed, select the **trap destination** radio button. Selecting this button causes the Network Mask field to be grayed out because it will not be used. Now for each SNMP monitoring stations, select **UDP** in the Protocol field and enter the SNMP monitoring stations name and address. After typing in all this information, select **Add**. After all you have added SNMP monitoring stations, click **OK**.

> **Note:** The HRMCNFIG process is external to the Library Manager. SNMP support will not be enabled until the process completes. When you see that Change SNMP trap destinations is selectable (not grayed out), the process is complete and you can use the SNMP features.

You highlight the SNMP trap destination in the list box when you need to delete it. Then select the **Delete** push button and click **OK**. The SNMP trap destination will be deleted as soon as **Change SNMP trap destinations** is selectable (not grayed out).

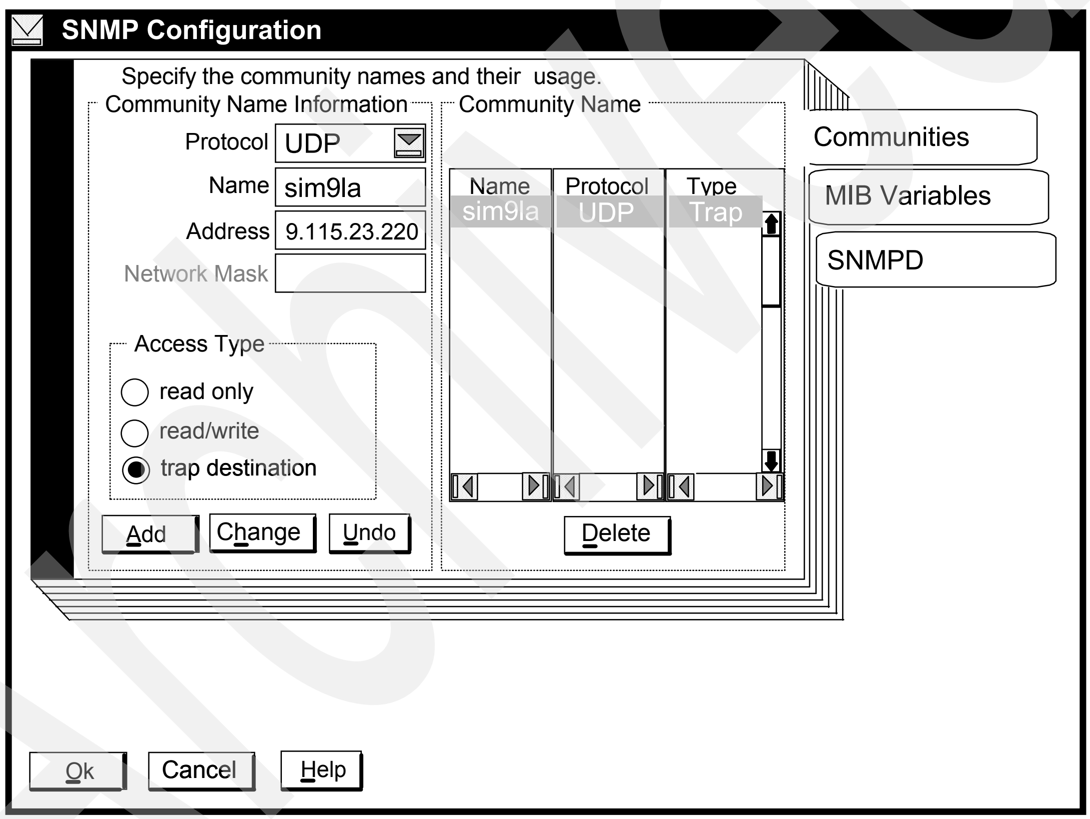

Figure 6-6 SNMP trap destination window

Starting SNMP

After you have selected the trap types and configured the destinations, you have to enable the Library Manager SNMP support. Select **Start SNMP** from the SNMP Options window. To ensure that the SNMP daemon has started, press the Control and Escape (Ctrl+Esc) keys simultaneously to bring up a list of the windows. If the SNMP daemon has started, SNMPD will be listed.

Stopping SNMP

To disable SNMP support, select **Stop SNMP** in the SNMP Options window. The Library Manager will not generate SNMP traps until you start up SNMP again.

Sending TESTM trap messages

You can verify that the Library Manager is set up correctly by sending an SNMP TESTM trap. This will send a test message to every SNMP monitoring stations that has been configured. Select **Send TESTM Trap** from the SNMP Options window. You are presented with a dialog box where you can type in a text string. After you type in the test text string and select the **Send** button, the message will be sent to all SNMP monitoring stations.

The most common problems encountered with SNMP messages are in the network area. You can confirm that the network is working by sending a test message.

6.4 Tape management

In this section we discuss the management of the physical tapes used in the library. We illustrate how to insert and eject the physical cartridges assigned to a VTS in a Library. Proper tape handling procedures should always be followed whether the tapes are in located in a 3494 or 3584. Mishandling or dropping 3590 or 3592 cartridges can seriously damage the tape to the point where the tape will not load, have permanent read/write errors or damage the tape drive. See the following documents for detailed information.

- *IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0472
- *IBM TotalStorage 3584 Tape Library Operator Guide,* GA32-0468
- *IBM TotalStorage 3494 Tape Library Operator Guide,* GA32-0449
- *IBM TotalStorage Enterprise Tape System 3592 Operator Guide,* GA32-0465
- *IBM TotalStorage Enterprise Tape System 3590 Operator Guide,* GA32-0330

6.4.1 Defining volser ranges for stacked volumes

Stacked volumes must be defined to the Library Manager before you can insert them in the 3494 or the 3584. Library Manager volser ranges are used for this purpose.

Inserting cartridges in the 3494 Tape Library or 3953/3584 Tape Library

A range definition associates a range of volsers and tape media with a logical library; when a cartridge is inserted in the Library, the Library Manager checks the range definitions to see to which logical library partition the volume is assigned. If a volume does not fall in any existing volser range, it is assigned to the non-VTS logical library, if such exists.

We recommend that you fully define your volser ranges during VTS installation even if you do not initially insert all the volumes in the library. No penalty is associated with defining long volser ranges, but such definition enables you to quickly insert new cartridges when needed, without having to first define new volser ranges. It also ensures that you cannot inadvertently insert logical volumes that conflict with your planned ranges of stacked volumes. If needed, you can later add more ranges, up to a total of 50. Existing volser ranges can be displayed via the TotalStorage Specialist or the Library Manager but only added via the Library Manager console.

Should your VTS have the Advanced Policy Management (APM) also installed (FC4001 - 4004), you can also assign your stacked volumes to different storage pools. With this feature installed, the *Volser Ranges* panel is available at both the Library Manager and 3494 Web Specialist (see Figure 6-9 on page 224).

To define volser ranges, select **System management** from the **Commands** pull-down menu, then select **Volser range for media types** to display the Volser Ranges window (Figure 6-7, or Figure 6-8, depending on LIC level). By using this window, you can:

- Add, modify, and delete volser ranges.
- Check whether a volser belongs to a specified range.
- Query the number of volumes in the library that belong to a selected range.

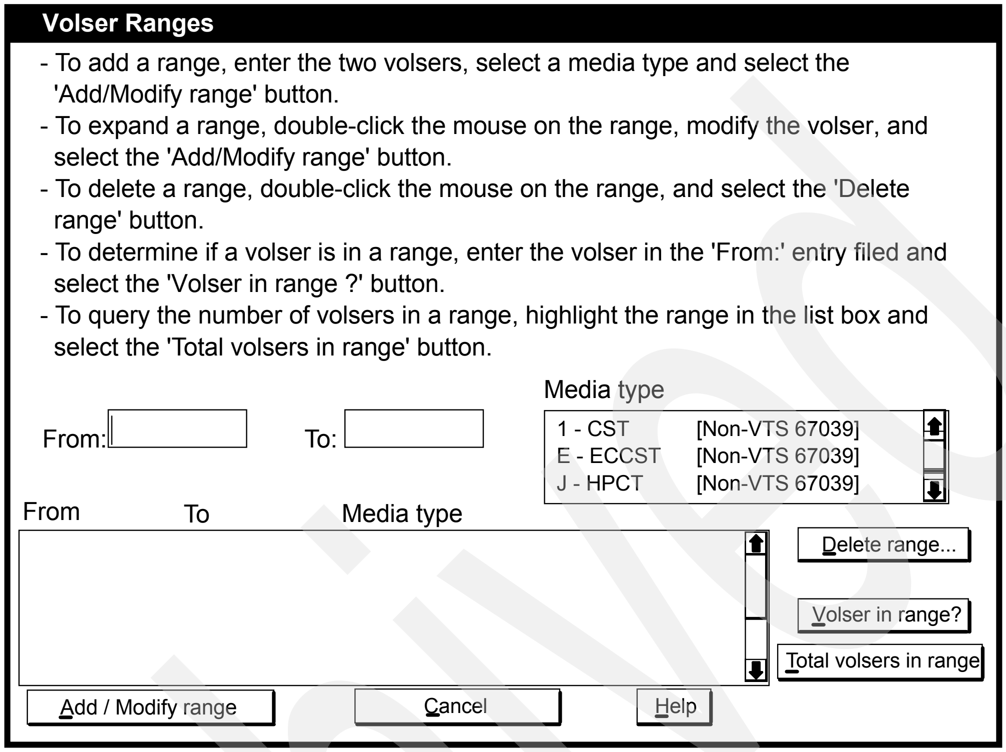

Figure 6-7 Library Manager Volser Ranges window without APM

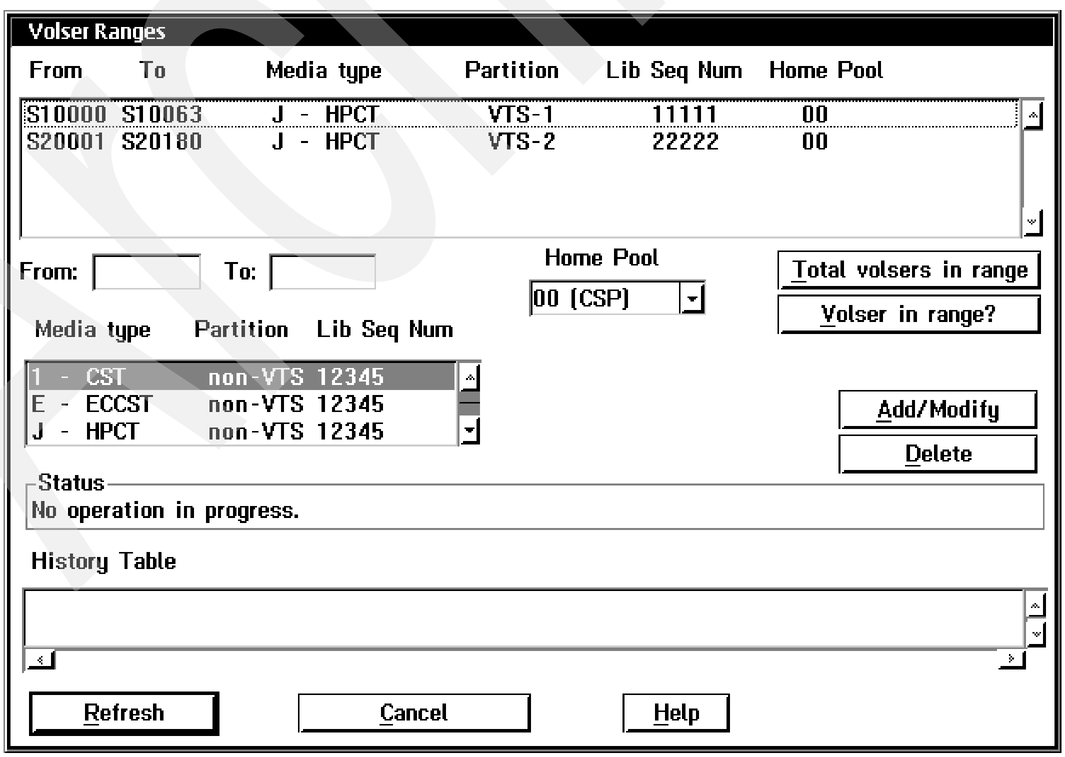

Figure 6-8 Library Manager Volser Ranges window with APM

Alternatively, you can use the Library Manager Specialist to Modify Volser Ranges (Figure 6-9).

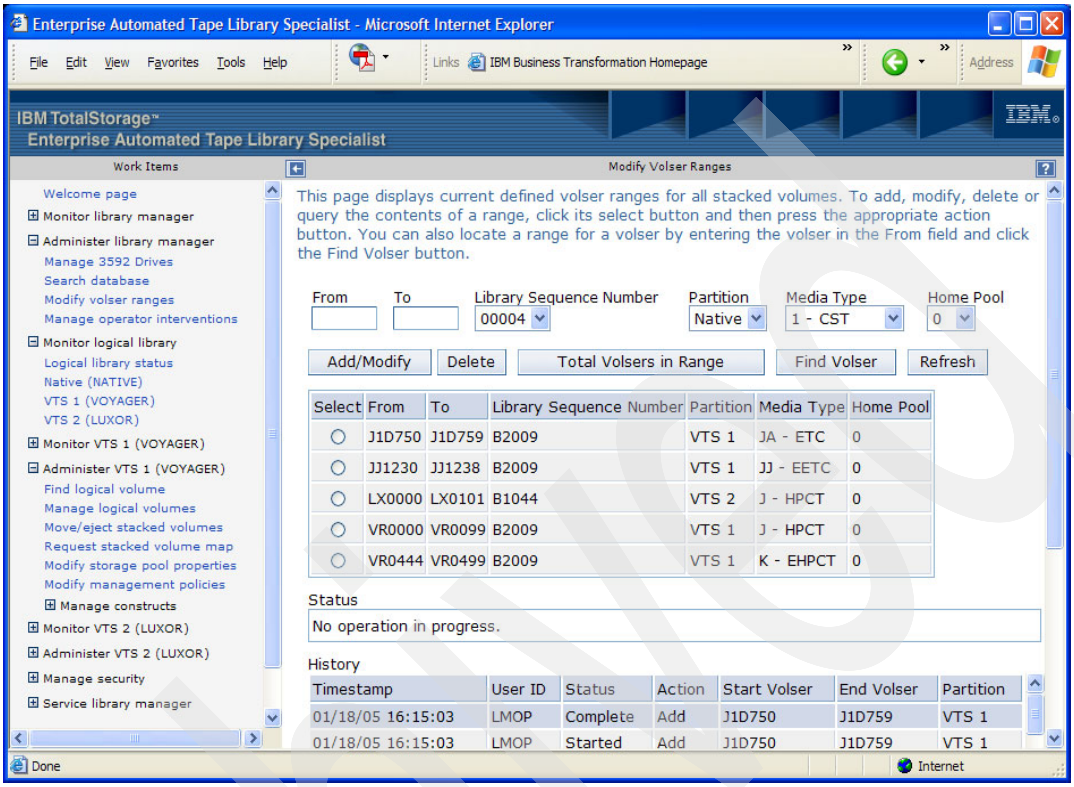

Figure 6-9 APM Modify volser range at the ETL Specialist

You may define J, K, JJ, and JA cartridges in the same VTS, provided that both the VTS hardware and microcode are at the right levels or features. See 3.6.5, "Extended Length Media Support" on page 87 for details.

To add a new range:

1. Enter the first and last volsers of the range in the From and To fields.
2. Highlight the desired media type of the range in the Media type list box.
3. If your VTS has APM Select a **Home Pool** for the volumes.
4. Select **Add / Modify range**.

Volser range specification: Volser entry fields must contain six alphanumeric characters. The two volsers must be entered in the same format. Corresponding characters in each volser must be either alphabetic or numeric. For example, AAA998 and AAB004 are of the same form, whereas AA9998 and AAB004 are not. The range includes all volsers that fall between the two volsers (inclusive) in a collating sequence where alphabetic characters are incremented alphabetically and numeric characters are incremented numerically. For example, volser range ABC000 through ABD999 would include 2000 volsers (ABC000 - ABC999 and ABD000 - ABD999).

Media type: The Media type list box contains the following entries:

```
1 - CST (non-VTS)
E - ECCST (non-VTS)
J - HPCT (non-VTS xxxxx)
K - EHPCT (non-VTS xxxxx)
```

```
JA - ETC (non-VTS xxxxx)
JJ - EETC (non-VTS xxxxx)
JR - EEWTC (non-VTS xxxxx)
JW - EWTC (non-VTS xxxxx)
1 - CST (VTS 1 yyyyy)
E - ECCST (VTS 1 yyyyy)
J - HPCT (VTS 1 yyyyy)
K - EHPCT (VTS 1 yyyyy)
JA - ETC (VTS 1 yyyyy)
JJ - EETC (VTS 1 yyyyy)
1 - CST (VTS 2 zzzzz)
E - ECCST (VTS 2 zzzzz)
J - HPCT (VTS 2 zzzzz)
K - EHPCT (VTS 2 zzzzz)
JA - ETC (VTS 2 zzzzz)
JJ - EETC (VTS 2 zzzzz)
? Unknown. This is an actual media type that is no longer used. The option is available
for backward compatibility.

1 - CST - Cartridge System Tape
E - ECCST - Enhanced Capacity Cartridge System Tape
J - HPCT - High Performance Cartridge Tape
K - EHPCT - Extended High Performance Cartridge Tape
JJ - ETC - Enterprise Cartridge in non-VTS logical library
JA - EETC - Enterprise Economy Tape Cartridge
JR - EEWTC - Enterprise Economy WORM Tape Cartridge
JW - EWTC - Enterprise WORM Tape Cartridge
```

Here, xxxxx, yyyyy and zzzzz are the library sequence numbers. Only those options that correspond to an existing logical library in the IBM 3494 are available.

> **Note:** JJ, JA, JR and JW media are for 3592-J1A tape drives only. JR and JW are only for native attached 3592 tape drives and are not supported for VTS. JJ, JA, JR, and JW media will be the only media in the list if this is on a 3584/3953 Tape Library system with VTS.

The entry you select defines the media type of the volser range and at the same time associates the range with a logical library.

When a volser range is added or modified, the Library Manager automatically combines overlapping ranges with the same media type and makes checks to ensure that the range does not conflict with existing volumes or volser ranges that are of different media type. If a conflict exists, the Library Manager rejects the add or change request. Note that from the Library Manager's point of view, VTS and non-VTS cartridges have a different media type as *HPCT for Non-VTS* and *HPCT for VTS 1*.

To check whether a range contains volumes, highlight the range in the volser range list box and select **Total volsers in range**.

You can change the media type of a volser range definition by deleting the range and then redefining it, provided that no range conflicts exist.

Volser range definitions are required for stacked volumes but are optional for native library volumes. The volser ranges help the Library Manager to determine a physical volume's media type when it is inserted into the library. The Library Manager determines a volume's media type at insert time by using the following rules:

1. The media type returned by the vision system is used as a first choice.

2. If the media type returned is for a TotalStorage 3590 or 3592 type cartridge and there is more than one logical library in the physical tape library, the volser ranges are used to determine to which logical library the volume is assigned.
3. The volser ranges are used to determine a volume's media type if it cannot be determined by the vision system. If the volume being inserted appears within one of the ranges, the range's associated media type is used. The search of the ranges is an inclusive search.
4. The system uses the default media type defined during the teaching process by the IBM System Service Representative (SSR) to determine the media type if the volser does not fall into one of the ranges.
5. The volume is ejected if there is no default media type.
6. Rules 4 and 5 do not apply to stacked volumes because they always fall into a range.

Note: There is no default media type in the 3584/3953.

6.4.2 Deleting volser ranges for stacked volumes

The Library Manager lets you delete a volser range definition at any time regardless of whether the range contains volumes or not. The range definition is only used at cartridge insert time to define a volume's media type and to assign the volume to a logical library.

We recommend that you *do not* delete stacked volser ranges that contain volumes. Deleting a volser range does not have any immediate effect on volumes that are in the range; the volumes remain in the same logical library as before. However, if a complete reinventory of the library is required, all cartridges in the library are reassigned to logical libraries according to the volser range definitions that are in effect during the reinventory.

If you have deleted a stacked volser range definition, all stacked volumes in that range are assigned to the non-VTS logical library and become unavailable to the VTS subsystem. There is also a potential for data loss, as the volumes could be assigned to a scratch category and subsequently used to satisfy scratch mounts on native TotalStorage 3590 and 3592 drives. As a safeguard against this, we recommend that you use your tape management system to deny any use of stacked volumes on the host, as described in 5.2, “DFSMSrmm and other tape management systems” on page 186.

6.4.3 Tape labels

Each stacked volume, like any other physical volume in a 3494 and 3584, must have a unique external machine and operator-readable bar code label. The label identifies the volume to the Library Manager. Currently (as at publication) labels are supplied by two manufacturers:

- Data Products (EDP) Tri-Optic
- Wright Line Tri-Code
- Dataware
- NetC

Note: Complete information can be found at:

http://www.storage.ibm.com/media/tapecartridges/index.html

These companies supply labels that conform to the common standard of IBM ATL. The bar code used is the Automation Identification Manufacturers Uniform Symbol Description version 3 (Code 39). It contains the same characters as the human readable area.

The only exception to this label requirement is when you use the **Insert Unlabeled Cartridges** function on the Library Manager **Commands** pull-down menu. However, we do not recommend that you insert volumes for the VTS by using this facility. If the external label on a volume becomes damaged, you can use this facility to reinsert the volume until its external label can be replaced. The Library Manager will eject volumes that you insert using the Insert Unlabeled Tape function during any subsequent reinventory complete system or inventory new storage of the library. Figure 6-10 shows the labels and the media of 3590 cartridges, while Figure 6-11 shows the labels and indicators for 3592 cartridges.

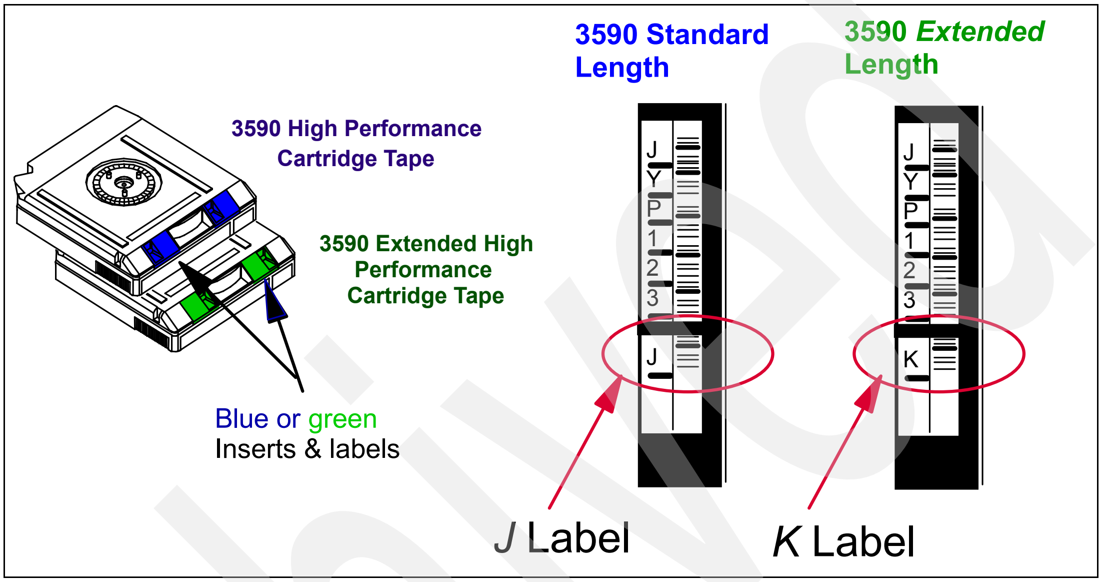

Figure 6-10 Tape labels of 3590 J and K cartridges

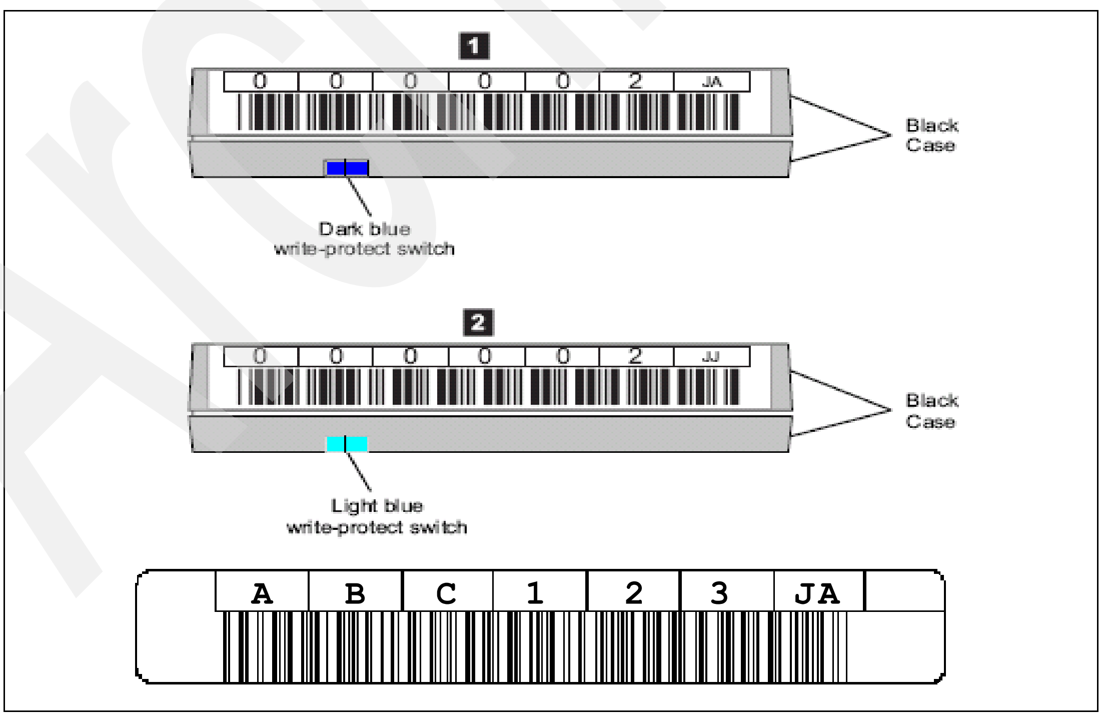

Figure 6-11 The 3592 JA and JJ tape cartridges and labels.

Volsers of all volumes, including native, stacked, and logical, must be unique within a 3494 or 3584. You have to use distinct volser ranges for the three volume types. Figure 6-12 shows an example of possible volser ranges for the different volume types. In this diagram the asterisk represents the last 5 digits of the volser. The cleaner and CE group are selected when the 3494 is first taught. These last two groups are not present on the 3953 Library Manager since the 3584 handles the cleaning and service tapes separately.

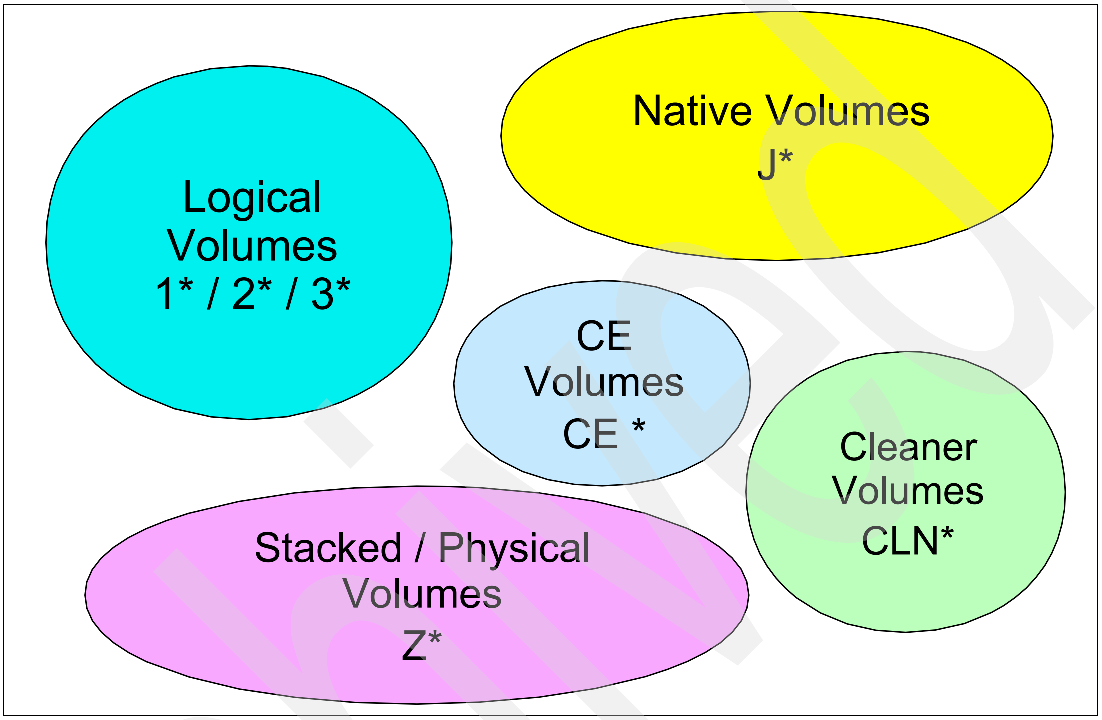

Figure 6-12 Unique volumes within library and within TCDBplex

Use the first characters of the volsers to indicate the volume's type. It will help the operator to define volser ranges and insert logical volumes in the library and to distinguish native volumes from stacked volumes. To further help the operator, you can use differently colored labels for stacked and native volumes.

The media type of each volume in the library must be identified to the Library Manager. Typically, TotalStorage 3590 cartridges are delivered with an external seventh-character media-type label, in which case the Library Manager uses that as the volume's media type. If not, the Library Manager uses the media type of the volser range in which the volume belongs. Because each stacked volume belongs to a volser range, it is not necessary for volumes to have a media-type label. However, we recommend that you use cartridges that have an external media type label.

The zSeries attached 3584 Tape Library requires that all media (either native 3592 tape drives, VTS stacked volumes or Open Systems attached) have labels with 8 characters, the last two characters being the media type. Any cartridge with a blank or unreadable label will be ejected or left in the convenience I/O station and not reported to 3953 or host. There is an option to set the whole 3584 Tape Library into 6 digit volser mode. This is not allowed when attached to a zSeries host.

Attention: It is important to ensure that the zSeries attached 3584 Tape Library is set in 8 character volser mode. This is the default setting, but you must still ensure that the 3584 is set to 8 character volser mode to work correctly.

Always handle 3590 and 3592 cartridges carefully to prevent dirt from getting on the tape. Do not try to pull the tape out of the cartridge, and never touch the tape. Keep the cartridge cover clean. Serious damage can occur if the 3590 or 3592 tape cartridges are dropped. Refer to the *IBM TotalStorage 3590 Tape Subsystem Operator Guide*, GA32-0330, and the *TotalStorage Enterprise Tape System 3592 Operator Guide*, GA32-0465 for details on how to handle 3590 and 3592 cartridges.

6.4.4 Inserting stacked volumes in the VTS

The following command has been enhanced (with the PTFs enabling the Advanced Function feature) to list the numbers of scratch stacked cartridges:

```
D SMS,LIBRARY(libname),DETAIL
```

See Example 4-5 on page 125 for a sample output of the DISPLAY SMS command.

You can use this number or the Free Storage Threshold Library Manager definition (see 4.3.7, “Define VTS management policies” on page 137) to find out if additional scratch stacked cartridges are needed.

When a physical volume is added to the tape library inventory, the Library Manager checks the volser to ensure that it is readable and is not already in the inventory. The convenience I/O station is in either *Import Mode* or *Insert Mode* depending on the capabilities of the VTS subsystems in the library and the configuration of the library.

The convenience I/O station is in *Insert Mode* if any of the following statements are true:

- There are no VTS subsystems in the library.
- There are no VTS subsystems in the library with the Advanced Function feature enabled.

The convenience I/O station is in Import Mode if the following is true:

- There is at least one VTS subsystems in the library with the Advanced Function feature enabled.

Note: Import mode can be disabled if required.

In the Import Mode, any cartridge which is found in the convenience I/O station is moved into the library and assigned to the Unassigned category (FF16). An operator must then use a Library Manager panel (Figure B-25 on page 402) to move an exported stacked volume into the Import category (FF14), or move a 3590/3592 cartridge into the Insert category for native use, or VTS stacked volume use, or to eject the cartridge.

Volumes assigned to the Insert category can then be handled using the panel in Figure 6-13.

You insert stacked volumes in the VTS library in the same way you insert native cartridges, by using either the convenience input/output station, the high capacity input/output facility, or empty cells in the library.

Inserting cartridges in a 3584/3953 Tape Library system

You have two options when inserting stacked volumes into the 3584 for the attached VTS:

- You can insert the cartridges into any open cell location in the 3584. After the door is closed the 3584 performs an automatic inventory of that frame. The 3584 notifies the 3953 Library Manager and it updates the database accordingly. The volser of these inserted volumes must be part of the range already defined for use by that VTS. Any volumes not already in an assigned volser range must be processed by the 3584 Web Specialist.

- The second option is to use the convenience I/O station on the 3584. Insert notification must first be enabled in the 3584. The 3584 Cartridge Assignment Policy (CAP) is used to define which volumes are assigned to which logical library. (See the CAP Assignment screen shown in Figure 4-19 on page 136.)

VTS stacked cartridge processing

When you insert a cartridge in the 3494 or 3584, the Library Manager checks the volser range definitions and determines whether the volume should be assigned to a VTS logical library. If so, the Library Manager places the volume in the VTS insert category (FF01) and notifies the corresponding VTS controller, which takes ownership of the volume and changes it to the appropriate stacked volume scratch category (FF03). The volume is dedicated for use in that VTS only; the Library Manager does not use it to satisfy mount requests for drives in the other logical libraries. Because the Library Manager does not place a stacked volume in the normal insert category (FF00) or notify the host that a cartridge was inserted, the host does not gain any knowledge of stacked volumes.

A Library Manager function is provided to manage the insert process. The Manage Insert Volumes panel (Figure 6-13) allows for actions to be taken for physical volumes that are in the insert category (FF00). The purpose of this panel is to allow unclaimed physical volumes to be ejected.

They could be unclaimed, because the library is a VTS-only library (no native 3590 or 3592 drives) or the cartridge is a CST or ECCST cartridge and the Library Manager volser ranges did not cover some volumes. Note that in a 3494 without VTS, there is no way to eject an unclaimed volume from the Library Manager:

Commands —> **Systems Management** —> **Manage Insert Volumes**

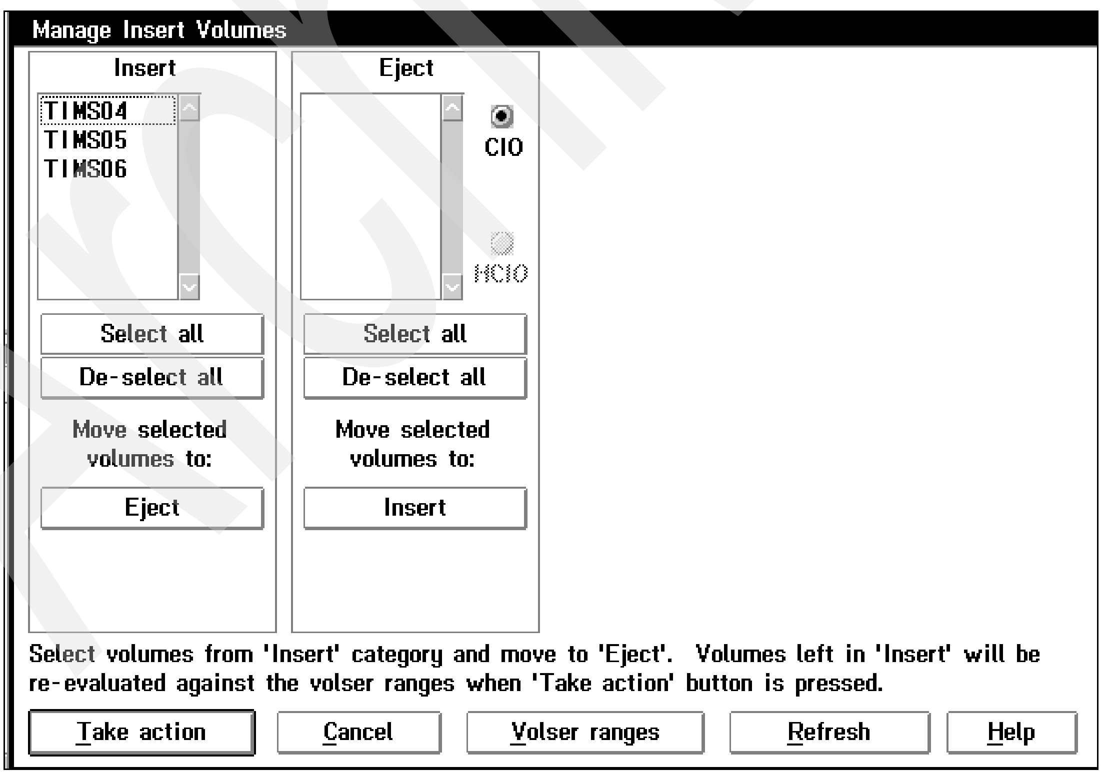

Figure 6-13 Manage Insert Volumes window

The Library Manager makes checks to ensure that the volume you insert is not already in the 3494 or 3584 as a physical (stacked or native) or logical volume and rejects the insert as a duplicate volume if it is.

You do not need to initialize stacked volumes. The VTS automatically checks a stacked volume for an internal volume label when the volume is used for the first time and labels it with an IBM standard label if there is no internal volume label. If there is an internal volume label and it does not match the volser expected, the Library Manager ejects the stacked volume to the convenience I/O station and notifies all attached hosts of an Intervention-required with this message:

```
A stacked volume has a label mismatch and has been ejected to the
convenience I/O station. Internal: xxxxxx, External: yyyyyy
```

Note that the VTS does not check the internal label at cartridge insert time, but when the VTS tries to mount the stacked volume for the first time.

We recommend that you have at least 50 extra stacked volumes inserted because of the way reclaim processing works. If you have more than 50 scratch cartridges, reclaim will be enabled when it will be non-invasive to VTS activity and according to the inhibit reclaim schedule. See “Reclamation and reconciliation” on page 137 for more information about reclaim processing. When the number of empty stacked volumes falls below the number required for GB of data specified by the Free Storage Threshold on the Library Manager VTS Management Policies window (see 4.3.7, “Define VTS management policies” on page 137), the Library Manager raises an intervention required condition indicating that:

```
Free storage threshold has been crossed.
```

To resolve the condition, insert more stacked volumes into the library. A subsequent reclamation process may also increase the number of empty stacked volumes but does not automatically clear the intervention required condition; the operator has to clear it manually. The condition does not trigger VTS space reclamation to start.

In a logically partitioned library, physical shelf space is not dedicated to the logical libraries; stacked volumes can reside anywhere on the library shelves among native volumes. Because stacked volumes are mounted only on the tape drives assigned to the VTS, you can minimize accessor movement by having stacked volumes in cells close to the VTS drives. Having cartridges next to the assigned tape drives where they can be mounted will also help the tape operator if he or she ever has to operate the library in Manual mode. To achieve this, insert stacked volumes manually in empty cells in the library on the wall with the tape drives or in adjacent frames.

Note: If your VTS is Import/Export capable, use care; an Exported Stacked Volume, once removed, will become a VTS Scratch Stacked Volume if reinserted using this process.

If your 3494 Tape Library operates in fixed home cell mode, the cartridges always stay in the cells where you initially inserted them. If your 3494 operates in floating home cell mode, the cartridges will be placed where they optimize performance. They may gradually move around the library but generally tend to stay close to the VTS drives. The 3584 Tape Library does not operate in floating cell mode. The tapes are returned to the same cell they were last inventoried in.

If you insert cartridges in empty cells of the 3494 Tape Library by opening the doors, you have to update the Library Manager inventory by selecting **Perform Inventory Update** from the **Commands** pull-down menu. If you have **Enable Inventory** Update selected in the **Commands** pull-down menu, frames are automatically inventoried when the doors are closed and the Library Manager is placed in Auto mode.

See 6.9, “Inventory update” on page 259 on how to perform inventory update.

Note: With the 3584 Tape Library, if a door is opened, an inventory of that frame will always take place before the accessor becomes ready. A new inventory of that frame’s cartridges is then sent to the 3953 Library Manager to update its database.

6.4.5 Ejecting VTS stacked volumes

Stacked volumes are dedicated to a specific VTS subsystem and must remain in the 3494 that contains the VTS. You cannot move data out of a VTS by ejecting stacked volumes from the library; data on a stacked volume is usable only within the VTS where it was created. However, you may need to eject stacked cartridges from a VTS for media or maintenance reasons. The VTS enables you to do that. Ejection of stacked volumes can be initiated at either the LM or the 3494 Specialist if you have Advanced Policy Management installed.

Notes:

1. The VTS provides a disaster recovery function that enables data on stacked volumes to be recovered on another VTS subsystem. See 6.13, “Disaster recovery” on page 274 for details.
2. The APM or Advanced Function feature with the Import/Export component allows you to move logical volumes in and out of a VTS. The operations described here relate to ejecting stacked volumes, not exported stacked volumes, as the stacked volumes holding logical volumes to be exported are called.

To eject a stacked volume, use the Library Manager **Commands** pull-down menu and select **System management** followed by **Eject a VTS stacked volume**. Specify the volser of the stacked volume on the **Eject A Stacked Volume** window (see Figure 6-14 for the LM panel and Figure 6-15 for the 3494 Specialist).

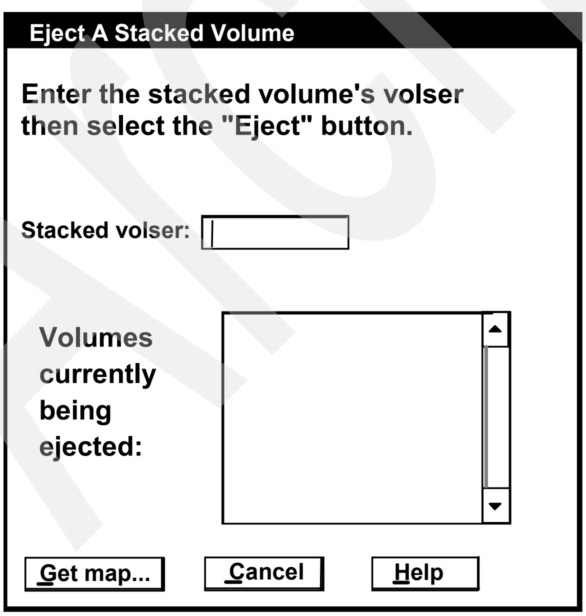

Figure 6-14 Library Manager Eject a Stacked Volume window

When a stacked volume eject is initiated, the eject process of the VTS gets a list of all logical volumes on the stacked volume. Once it has the list, it recalls the active logical volumes into the TVC. After that, the logical volumes are placed on the copy queue. When all active data has been copied, the VTS requests the Library Manager to eject the volume.

You will want to perform this ejection process during a low usage time for the VTS because each logical volume occupies space in the TVC that could have held other logical volumes. Although the wait should be minimal, there could also be a delay for host requests looking for logical volumes on the stacked volume being ejected. Be aware that this can be a time-consuming process, as up to 300 GB of data may have to be copied to other stacked volumes.

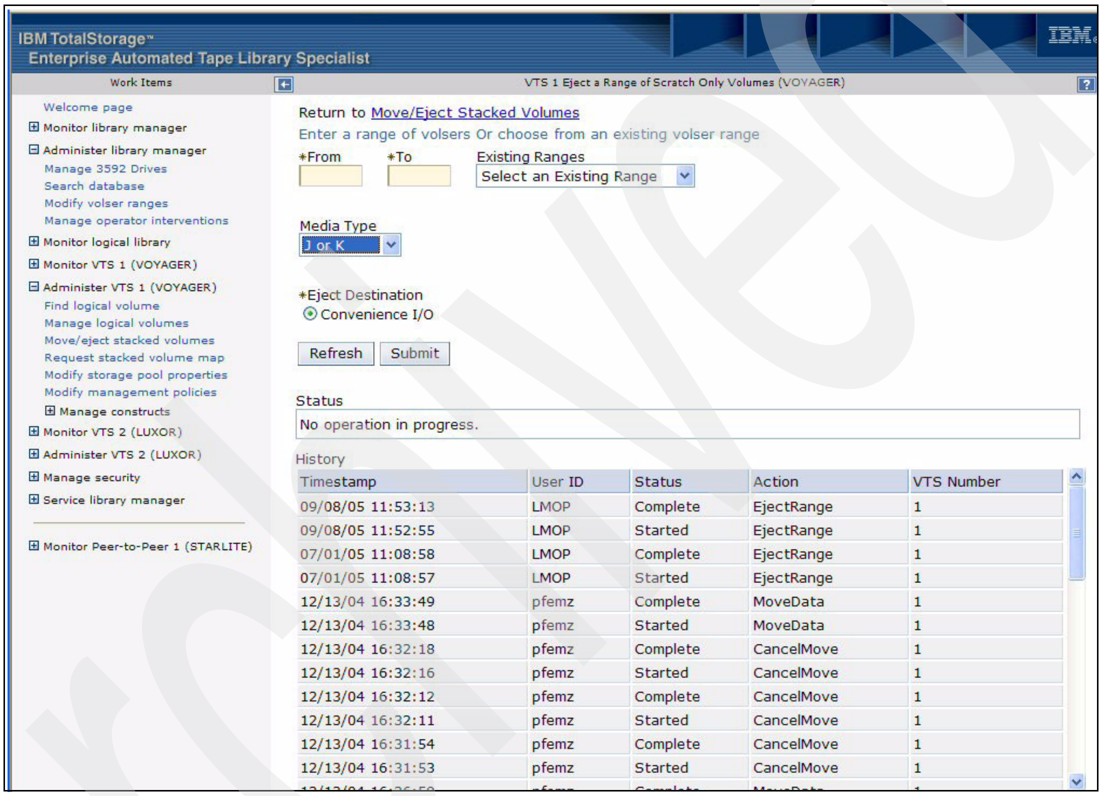

Figure 6-15 3494 Specialist stacked volume ejection panel

If you need to eject a volume for use outside the VTS, select an empty stacked volume; the ejection of an empty volume is a quick operation. To find out which stacked volumes are empty, use the Library Manager Search Database for Volumes window to display volumes in VTS stacked volume scratch categories FF02 and FF03. See 6.8.1, “Search database for volumes” on page 250.

> **Attention:** With VTS LIC level 2.26.xx.x and above, pooling is enabled. All stacked volumes are assigned a category of FF04 for both **private and scratch** volumes. See 6.4.6, “Move and eject stacked volumes” on page 234.

You cannot eject a stacked volume from the tape library by using host system commands. For example, in a DFSMS-managed tape environment, IDCAMS command **`LIBRARY EJECT,volser`** against a stacked volume fails with the following message because the tape configuration database (TCDB) does not contain a record for stacked volumes:

```
CBR1064I Command rejected. Volume serial number <volser> undefined
```

Do not manually remove a stacked volume from the library. Doing so will cause jobs to fail that try to access logical volumes resident on the stacked volume.

When the Library Manager cannot mount a stacked volume to satisfy a mount request from the VTS, it raises an intervention required condition indicating one of the following reasons as shown in Example 6-4:

Example 6-4 Logical volume mount failure - intervention required message

```
Mount of logical volume xxxxxx on physical volume yyyyyy failed
Mount of logical volume xxxxxx failed because physical volume yyyyyy is misplaced
Mount of logical volume xxxxxx failed because physical volume yyyyyy is inaccessible
Mount of logical volume xxxxxx failed because physical volume yyyyyy is not in the library
```

The VTS subsystem indicates to the host that the mount was completed. When the host issues the first I/O operation to the virtual device, the subsystem will fail the command and indicate an equipment check. The host will then fail the job and generate console messages. The only command that will actually do anything is the Rewind/Unload command.

Example 6-5 shows the messages that are generated as a result of a logical volume mount failure for volume V00001 on device 1FD5.

Example 6-5 Logical volume mount failure - console message

```
IEC501A M 1FD5,V00001,SL,NOCOMP,GENVTSP8,STEP4,VTS42.TEST74

IOS000I 1FD5,AD,EQC,**,0600,,**,,GENVTSP8613
 10488035000000206011(B310000102004113)0000(500007FF)CE3FD69000050500

IEC147I 613-54,IFG0194K,GENVTSP8,STEP4,SYSUT2,1FD5,OAM274 VTS42.TEST74
```

The IOS000I message reports that an equipment check was received for device 1FD5 and contains the sense data returned by the subsystem. When a logical volume recall fails, the volume is not available in the TVC. Sense byte 14 contains X'02', which indicates this condition. The IEC147I message reports that the failure was detected as the tape volume was opened. The abend and return code (613-54) indicates that the volume serial number was unreadable, had no label, or the volume was uninitialized, however in a VTS it is also the indication that the recall failed.

When a logical volume mount is failed in this way, the operator will need to review the intervention required conditions at the library and based on the intervention reported, take action prior to resubmitting the job. Use the **Find logical volume's home** function of the Library Manager to find out which stacked volume the logical volume resides on, then use the **Search database for volumes** function to determine the status of the stacked volume. See 6.8, “Querying the Library Manager database” on page 250.

When the VTS determines that a stacked volume is damaged, it notifies the operator of the problem and does not write any more data on the stacked volume. VTS determines that a stacked volume is damaged if the volume gets a permanent I/O error on two different drives. See “Read-only status volume recovery” on page 267 for more details.

6.4.6 Move and eject stacked volumes

Beginning with microcode release 2.26.xx.x for the VTS and Library Manager microcode level 526.xx, storage pools were introduced. With the addition of FC4001 (Advanced Policy Management) for the VTS, you could take full advantage of these new features. With these levels of microcode came changes to the way the Library Manager presented information and added more functionality to manage the storage of the library.

The Move/Eject Stacked Volumes window (Figure 6-16) allows you to move and eject stacked volumes. The following actions can be performed:

- Move a range of volumes to another pool.
- Move a range of scratch only volumes to another pool.
- Move a quantity of scratch only volumes to another pool.
- Move logical volumes between pools.
- Eject a range of volumes.
- Eject a range of scratch only volumes.
- Eject a quantity of scratch only volumes.
- Cancel active move requests.
- Cancel active eject requests.

To initiate the move or eject process, perform the following actions:

1. Select the **Commands - System Management** pull-down menu
2. Click **Manage Constructs and Pools**.
3. Click **Move/Eject Stacked Volumes.** The window shown in Figure 6-16 will be displayed.

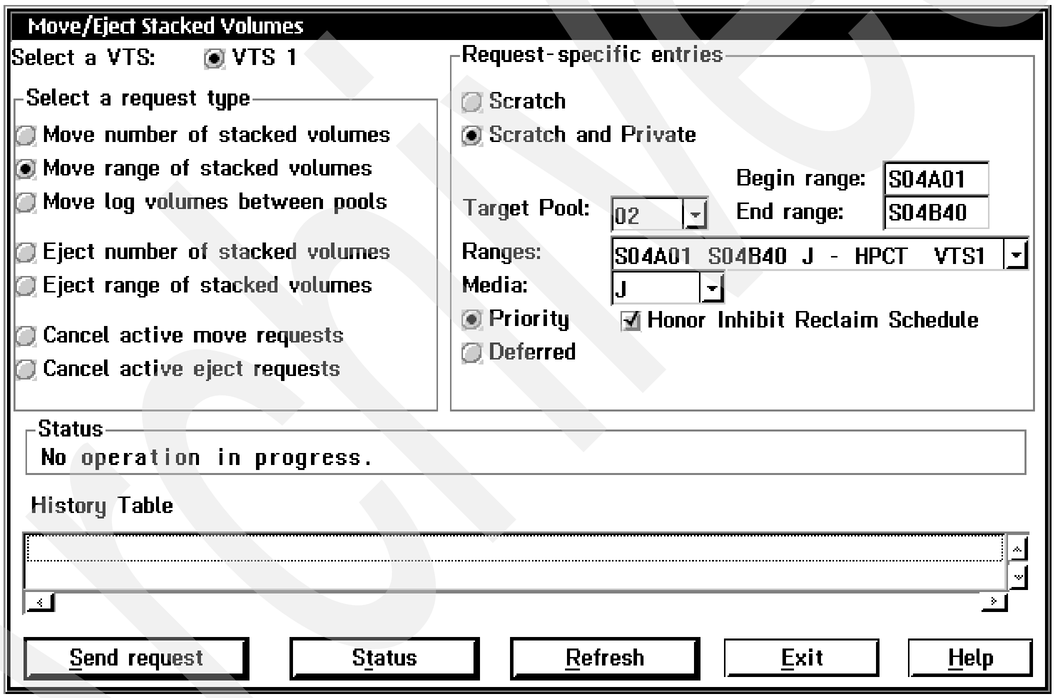

Figure 6-16 Move Stacked Volumes window

Move stacked volumes

You can move stacked volumes in various ways, which we describe in the following sections.

Moving a range of volumes to another pool

Follow this procedure to move a range of volumes to another pool:

1. Select **Move range of stacked volumes**.
2. In the Request-specific entries, do the following steps:
   a. Select **Scratch and Private**.
   b. Define the volser range in the **Begin range** and **End range** fields or select a range using the **Ranges** pull-down.
   c. Select the pool where the volumes will be moved to in the **Target Pool** pull-down.

d. Select the media type for the request in the **Media** pull-down.

J/K: Indicates that either J or K media types can be moved/ejected.
J: Indicates the default is the High Performance Cartridge Tape (HPCT).
K: Indicates the default is the Extended High Performance Cartridge Tape (EHPCT).
JA/JJ: Indicates that either JA or JJ media types can be moved/ejected.
JA: Indicates the default is the Enterprise Economy Tape Cartridge (EETC).
JJ: Indicates the default is the Enterprise Tape Cartridge (ETC).

e. Select either **Priority** or **Deferred**. Priority move requests are actively processed by the VTS. A deferred move request occurs as part of the normal reclaim function.

f. If **Priority** was selected, checking the **Honor Inhibit Reclaim Schedule** prevents requests from being performed during the times inhibited by the reclaim schedule

3. Select the **Send Request** push button.

Moving a range of scratch only volumes to another pool

Follow this procedure to move a range of scratch only volumes to another pool:

1. Select **Move range of stacked volumes**.
2. In the Request-specific entries, do the following steps:
   a. Select **Scratch**.
   b. Define the volser range in the **Begin range** and **End range** fields or select a range using the **Ranges** pull-down.
   c. Select the pool where the volumes will be moved to in the **Target Pool** pull-down.
   d. Select the media type for the request in the **Media** pull-down:

J/K: Indicates that either J or K media types can be moved/ejected.
J: Indicates the default is the High Performance Cartridge Tape (HPCT).
K: Indicates the default is the Extended High Performance Cartridge Tape (EHPCT).
JA/JJ: Indicates that either JA or JJ media types can be moved/ejected.
JA: Indicates the default is the Enterprise Economy Tape Cartridge (EETC).
JJ: Indicates the default is the Enterprise Tape Cartridge (ETC).

3. Select the **Send Request** push button.

Moving a quantity of scratch only volumes to another pool

Follow this procedure To move a quantity of scratch only volumes to another pool:

1. Select **Move number of stacked volumes**.
2. In the Request-specific entries section, do the following steps:
   a. Specify the number of volumes to move in the **Number of vols:** field.
   b. Select the pool where the volumes will be taken from in the **Source Pool** pull-down.
   c. Select the pool where the volumes will be moved to in the **Target Pool** pull-down.
   d. Select the media type for the request in the **Media** pull-down:

J/K: Indicates that either J or K media types can be moved/ejected.
J: Indicates the default is the High Performance Cartridge Tape (HPCT).
K: Indicates the default is the Extended High Performance Cartridge Tape (EHPCT).
JA/JJ: Indicates that either JA or JJ media types can be moved/ejected.
JA: Indicates the default is the Enterprise Economy Tape Cartridge (EETC).
JJ: Indicates the default is the Enterprise Tape Cartridge (ETC).

3. Select the **Send Request** push button.

Moving logical volumes between pools

Follow this procedure to move logical volumes between pools:

1. Select **Move logical volumes between pools**.
2. In the Request-specific entries section, do the following steps:
   a. Define the volser range in the Begin range and End range fields or select a range using the Ranges pull-down.
   b. Select the pool where the data will be moved to in the Target Pool pull-down.
   c. Select the media type for the request in the Media pull-down:

      J/K: Indicates that either J or K media types can be moved/ejected.
      J: Indicates the default is the High Performance Cartridge Tape (HPCT).
      K: Indicates the default is the Extended High Performance Cartridge Tape (EHPCT).
      JA/JJ: Indicates that either JA or JJ media types can be moved/ejected.
      JA: Indicates the default is the Enterprise Economy Tape Cartridge (EETC).
      JJ: Indicates the default is the Enterprise Tape Cartridge (ETC).

   d. Select either **Priority** or **Deferred**. Priority move requests are actively processed by the VTS. A deferred move request occurs as part of the normal reclaim function.
   e. If **Priority** was selected, checking the **Honor Inhibit Reclaim Schedule** prevents requests from being performed during the times inhibited by the reclaim schedule.
3. Select the **Send Request** push button.

Note: When this option is used, the logical volume construct information is *not* changed to reflect what has been moved. If getting the construct information is important, then use the host means of moving logical volumes between pools (mount an LVOL, change constructs, then dismount).

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, without FC4001 installed, only the eject and cancel functions are available.

Eject stacked volumes

Figure 6-17 shows the selection on the Library Manager console to eject stacked volumes.

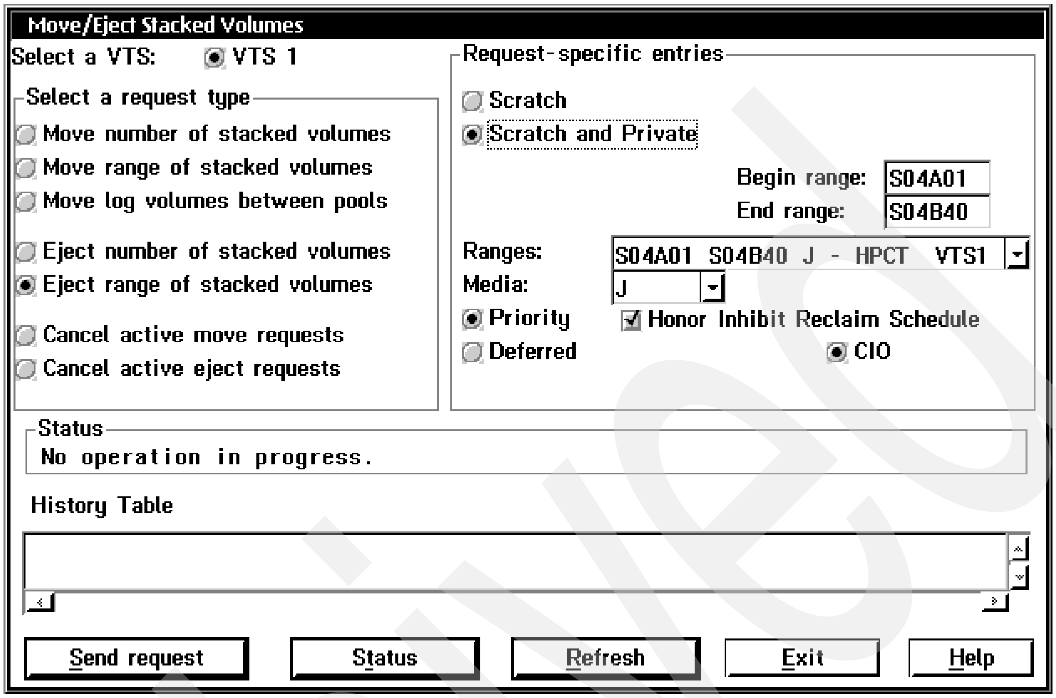

Figure 6-17 Eject Stacked Volumes window

Ejecting a range of volumes

Follow this procedure to eject a range of volumes:

1. Select **Eject range of stacked volumes** (Figure 6-17).
2. In the Request-specific entries, do the following steps:
   a. Select **Scratch and Private**.
   b. Define the volser range in the **Begin range** and **End range** fields or select a range using the **Ranges** pull-down.

   > **Note:** If you wish to only eject only one volume, select a range starting and ending with the same volser.

   c. Select the media type for the request in the **Media** pull-down.

   J/K: Indicates that either J or K media types can be moved/ejected.
   J: Indicates the default is the High Performance Cartridge Tape (HPCT).
   K: Indicates the default is the Extended High Performance Cartridge Tape (EHPCT).
   JA/JJ: Indicates that either JA or JJ media types can be moved/ejected.
   JA: Indicates the default is the Enterprise Economy Tape Cartridge (EETC).
   JJ: Indicates the default is the Enterprise Tape Cartridge (ETC).

   d. Select either **Priority** or **Deferred**. A priority eject request is forced to the top of the queue so it is performed immediately. A deferred eject request occurs as part of the normal reclaim function.

   e. If **Priority** was selected, checking the **Honor Inhibit Reclaim Schedule** prevents requests from being performed during the times inhibited by the reclaim schedule.
3. Select the **Send Request** push button.

Ejecting a range of scratch only volumes

Follow this procedure to eject a range of scratch only volumes:

1. Select **Eject range of stacked volumes**.
2. In the Request-specific entries, do the following steps:
   a. Select **Scratch**.
   b. Define the volser range in the **Begin range** and **End range** fields or select a range using the **Ranges** pull-down.
   c. Select the media type for the request in the **Media** pull-down.

      J/K: Indicates that either J or K media types can be moved/ejected.
      J: Indicates the default is the High Performance Cartridge Tape (HPCT).
      K: Indicates the default is the Extended High Performance Cartridge Tape (EHPCT).
      JA/JJ: Indicates that either JA or JJ media types can be moved/ejected.
      JA: Indicates the default is the Enterprise Economy Tape Cartridge (EETC).
      JJ: Indicates the default is the Enterprise Tape Cartridge (ETC).

   d. Select where the cartridges will be ejected to by clicking the radio button next to **CIO** Convenience I/O or **HCIO** High Capacity I/O Facility (if installed).
3. Select the **Send Request** push button.

> **Note:** The 3953/3584 does not support HCIO.

Ejecting a quantity of scratch only volumes

Follow this procedure to eject a quantity of scratch only volumes:

1. .Select Eject number of stacked volumes.
2. In the Request-specific entries, do the following steps:
   a. Specify the number of volumes to eject in the **Number of vols:** field.
   b. Select the pool where the volumes will be ejected from in the **Source Pool** pull-down.
   c. Select the media type for the request in the **Media** pull-down.

      J/K: Indicates that either J or K media types can be moved/ejected.
      J: Indicates the default is the High Performance Cartridge Tape (HPCT).
      K: Indicates the default is the Extended High Performance Cartridge Tape (EHPCT).
      JA/JJ: Indicates that either JA or JJ media types can be moved/ejected.
      JA: Indicates the default is the Enterprise Economy Tape Cartridge (EETC).
      JJ: Indicates the default is the Enterprise Tape Cartridge (ETC).

   d. Select where the cartridges will be ejected to by clicking the radio button next to **CIO** Convenience I/O or **HCIO** High Capacity I/O Facility (if installed).
3. Select the **Send Request** push button.

> **Note:** The 3953/3584 does not support HCIO.

Canceling active move/eject requests

Follow this procedure to cancel active move/eject requests:

1. Select **Cancel Active Move Requests** or **Cancel Active Eject Requests**.
2. Select an option in the Target Pool drop-down list.
   – To cancel actions for a specific pool, select the pool.
   – To cancel actions for all pools, select 'All Pools'
3. Specify the cancellation of either **Priority**, **Deferred**, or **All** Requests.

Send Request: Sends the move/eject request.

Status: Opens the Move/Eject Stacked Volumes — Status window.

Refresh: Refreshes the Move/Eject Stacked Volumes window.

Exit: Closes the Move/Eject Stacked Volumes window.

Help: Provides help about the Move/Eject Stacked Volumes 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.

6.4.7 Move/Eject Stacked Volumes (Status)

The Move/Eject Stacked Volume (Status) window allows you to view your current move/eject requests (Figure 6-18).

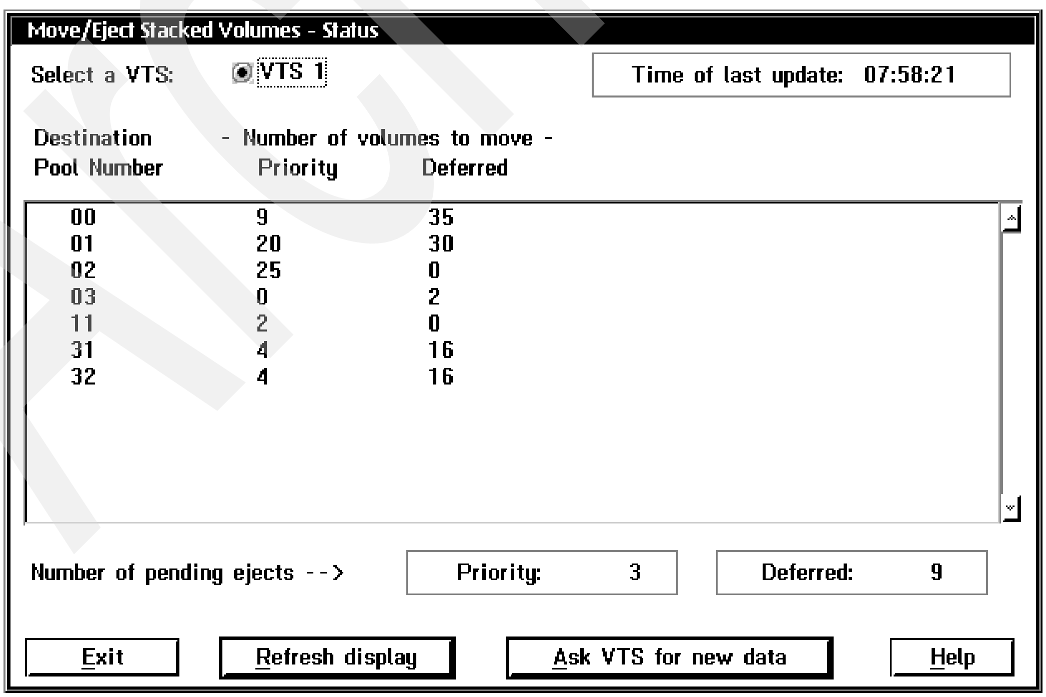

Figure 6-18 Move/Eject Stacked Volume (Status)

This window displays the following information:

Destination Pool Number

- The storage pool where the range of volumes specified in the request are being moved.

Number of volumes to move — Priority

- The number of active priority move requests. Priority move requests are actively processed by the VTS.

Number of volumes to move — Deferred

- The number of active deferred move requests. A deferred move is a request where the move occurs as part of the normal reclaim function.

Number of pending ejects:

- The number of pending ejects are separated into two categories:
  - **Priority:** The number of active priority eject requests. A priority eject request is a request that is forced to the top of the queue, so that it gets completed immediately.
  - **Deferred:** The number of active deferred eject requests. A deferred eject request is a request where the eject occurs as part of the normal reclaim function.

Exit: Closes the Move/Eject Stacked Volumes (Status) window.

Refresh display: Refreshes the data on the Move/Eject Stacked Volumes (Status) window. Check the Time of Last Update to determine if the data is current.

Ask VTS for new data: Sends a request for new data to all VTSs that are Advanced Policy Management Capable. The panel data is refreshed automatically after the VTS responds to the request.

Help: Provides help about the Move/Eject Stacked Volumes (Status) window.

6.5 Inserting logical volumes into the VTS

You insert logical volumes in a VTS by defining their volsers and their media type in the Library Manager's Insert Logical Volumes window (Figure 6-19). Physical cartridge movement is not involved. To access the window, select **System management** on the Library Manager **Commands** pull-down, then select **Insert VTS logical volumes**.

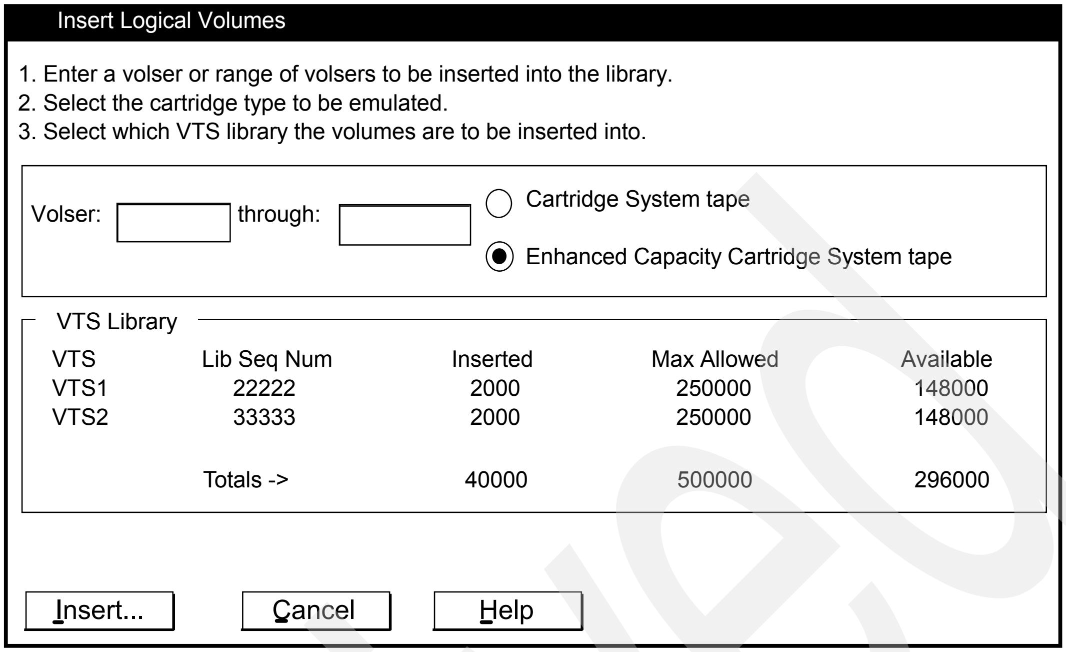

Figure 6-19 Library Manager Insert Logical Volumes window pre-LM LIC level 527

To insert logical volumes:

- Enter a volser or a range of volsers in the *Volser* and *through* fields.
- Select the cartridge type to be emulated.
- Select into which VTS library the volumes are to be inserted.
- Select the **Insert** button.
- Once insert is complete, select the **Cancel** button to exit this menu.

With APM available, you can also assign constructs at logical volume insert time (see Figure 6-20). However, if you are running your VTS in a DFSMS enabled host environment, we do not recommend that you assign any constructs. This is because the appropriate constructs names will be passed through your DFSMS routines as part of the original file allocation request (file sequence one) and then actioned at rewind / unload (EOV) time. As such, any predefined names could be overwritten.

Volser range specification

Volser entry fields must contain six alphanumeric characters. The two volsers must be entered in the same format. Corresponding characters in each volser must both be either alphabetic or numeric. For example, AAA998 and AAB004 are of the same form, whereas AA9998 and AAB004 are not. The range includes all volsers that fall between the two volsers (inclusive) in a collating sequence where alphabetic characters are incremented alphabetically and numeric characters are incremented numerically. For example, volser range ABC000 through ABD999 would result in an insert of 2,000 logical volumes (ABC000 - ABC999 and ABD000 - ABD999).

Up to a total of 500,000 logical volumes can be inserted in a Library Manager for each of the two VTS installed with appropriate features installed. If the insert would cause more than 1,000,000 logical volumes in a library (3494 or 3953/3584) or more than 500,000 in a single VTS, the insert is aborted.

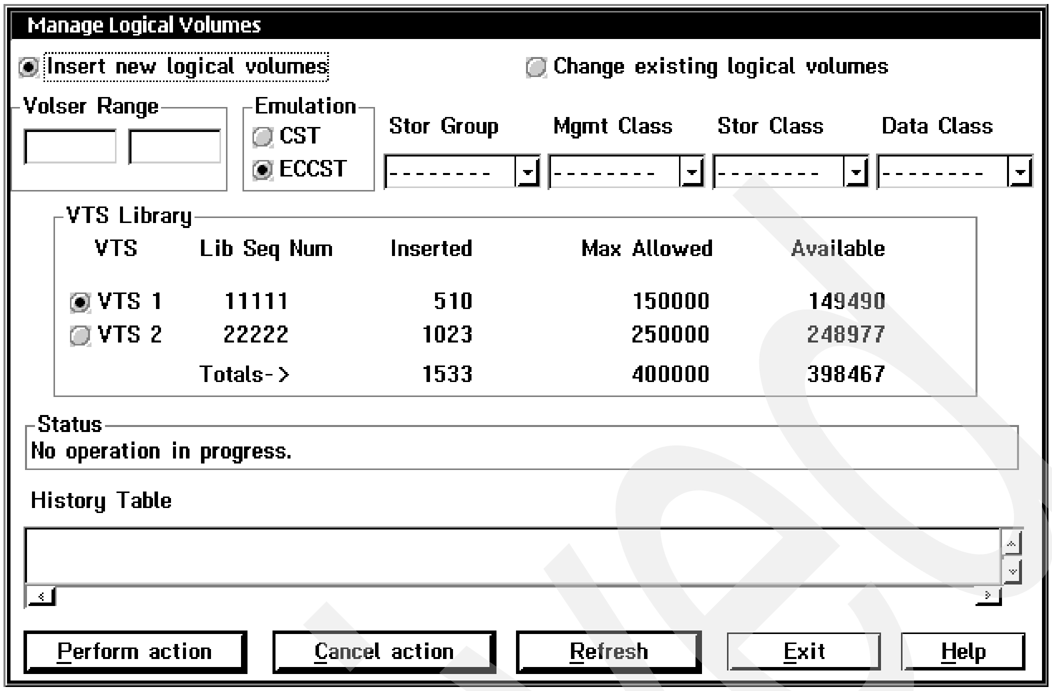

Figure 6-20 Library Manager Insert Logical Volumes window post LM LIC level 527

Volsers must be unique within a physical library. The Library Manager makes checks to ensure that the logical volumes you define do not conflict with existing volumes in the library or with any of the physical volser ranges defined for the library. If the logical volser range overlaps with a physical volser range, the logical volume insert is refused; no logical volumes are inserted. If there is no range conflict, the logical volumes are inserted, except for those whose volser is a duplicate of an existing volume in the library.

Unlike physical volser range definitions, a logical range definition cannot be displayed after the insert has completed. However, there is no limit on how many times you can insert new logical volume ranges.

You do not have to initialize logical volumes. The VTS automatically initializes virtual volumes with an IBM standard label when they are mounted for the first time.

When you insert logical volumes in a 3494 or 3953, the Library Manager assigns them to the insert category (FF00) just as it does with non-VTS volumes. The newly added logical volumes then go through the normal cartridge insert processing; the Library Manager notifies attached hosts of volumes in the insert category, and the host subsequently assigns them to another Library Manager category.

Important: Be aware that insert processing is a time-consuming process that is potentially performance impacting. You can roughly estimate that it takes one second per volume to complete the insert process. Insert processing could take many hours if you insert large numbers of logical volumes at a time. The recommended maximum number of logical volumes you should insert at any one time is 10,000. Wait until the first group of has completed the insert process before adding more. See 4.3.4, “Define (insert) logical volumes” on page 130 for additional information.

We recommend that you initially insert only the number of logical volumes that are required to handle the workload you put in the VTS in the near future, and gradually add new ones as more scratch volumes are required. Inserting volumes in small increments is fast and does not cause unnecessary overhead in the form of logical volumes that the system has to manage.

Be aware that if your VTS LIC is below 2.26.xx.x, the VTS does not invalidate a logical volume when the host returns it to scratch. Data on the logical volume is preserved and continues to take up stacked volume space. From the point of view of VTS, a scratch logical volume is an active volume even though from the application's point of view it no longer contains valid and usable data.

The Library Manager selects volumes from a category to satisfy nonspecific mounts in a first-in, first-out order. When a volume returns to the scratch category, it goes at the end of the selection list, so the Library Manager selects all volumes in the category before re-selecting that volume. Eventually, all logical volumes in the VTS get used and contain some data. Therefore, if you define more logical volumes than your operations require, you may end up wasting space.

To illustrate this point, we give you an example of the extra physical cartridges you would require. Here are the assumptions that were made for this example:

1. Using a VTS with 3590 B1A drives
2. Average logical volume size of 400 MB
3. Number of free stacked volumes is 50
4. Stacked volume utilization of 50%
5. Average compression ratio of 2.5 to 1

Using the above assumptions, we run through the calculations in "Sizing the number of stacked volumes" on page 81, with 20,000 logical volumes, ending up with 690 stacked volumes needed. Using the same assumptions with 150,000 logical volumes, we end up with 4850 stacked volumes needed.

> **Note:** From a VTS perspective, it doesn't matter if these 150,000 logical volumes have a SCRATCH or PRIVATE status in the host tape management system. Once used, they will all occupy space on stacked cartridges. The implementation of Expired Volume Management will greatly reduce this constraint.

As you can see, there is a significant difference in the number of stacked volumes needed.

If you find you have defined too many logical volumes, the best solution is to follow the procedure in 6.7, "Ejecting logical volumes from the VTS" on page 248 to delete them. If your operational procedures do not let you easily eject all of the extra logical volumes, there is another solution. You can issue jobs to write over the unneeded logical volumes with a single record. This will invalidate the data on the previous logical volumes and take a minimal amount of space on the stacked volume.

To change the media type of logical volumes, you need to eject the volumes first, then re-insert them as new. Data on the volumes will be lost.

In a DFSMS managed tape environment, we recommend that you set the Entry Default Use Attribute for your VTS library to PRIVATE to make sure that logical volumes you insert into the library are protected. The tape management system should handle the process of switching the volumes to a scratch category as required.

Use the ISMF Tape Library application to set the Entry Default Use Attribute on the Tape Library Define panel (Figure 6-21).

```
 TAPE LIBRARY DEFINE                   Page 1 of 2

 Command ===>

 SCDS Name    : SMS.SCDS
 Library Name : VTS1

  DESCRIPTION ===> 'Virtual Tape Server 1'

  LIBRARY ID                 ===> _____     (00000 to FFFFF)
  CONSOLE NAME               ===> ________
  ENTRY DEFAULT DATA CLASS   ===> ________
  ENTRY DEFAULT USE ATTRIBUTE ===> P        (P=PRIVATE  S=SCRATCH)
  EJECT DEFAULT              ===> P         (P=PURGE  K=KEEP)

Media Type:          Scratch Threshold     Media1  . . . . 0
    Media2 . . . . 400            Media3  . . . . 0         (0 to 999999)
    Media4 . . . . 0              Media5 . . . . 0         (0 to 999999)
    Media6 . . . . 0              Media7 . . . . 0         (0 to 999999)
    Media8 . . . . 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 6-21 ISMF panel: Tape Library Define

When the number of scratch volumes in the library falls below the media type's scratch threshold specified on the ISMF Tape Library Define panel, OAM issues the following message:

```
CBR3660A Enter {MEDIA1|MEDIA2} scratch volumes into library <library-name>
```

To resolve the condition, insert new logical volumes, or run tape expiration processing to return expired logical volumes to scratch status.

To display the number of scratch volumes in the library and the corresponding thresholds, use the command:

```
DISPLAY SMS,LIBRARY(library-name),DETAIL
```

6.6 Expired volume management

Logical volumes in a Virtual Tape Server (VTS) are stored on physical stacked volumes after the host closes the logical volume. After a period of time, the data on a logical volume that is not intended for long term archive, is eventually expired through the host tape management system and the logical volume is placed in scratch status. Previous to Library Manager LIC 526, returning a logical volume to scratch did not change the status of the data associated with the volume from the VTS's point of view until the logical volume was re-written or otherwise modified. This resulted in the VTS continuing to manage the data as if it were still active, even though it had really been scratched by the user.

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

A VTS connected to a Library Manager running code 526 or higher has the ability to associate a limited “grace period” with the logical volumes scratch category. This option provides you additional flexibility in managing the data in a VTS. The deletion of expired logical volume data eliminates the need for the VTS to manage logical volume data that has already been expired at the host.

Attention: Once the data associated with a logical volume has been deleted, it cannot be recovered without assistance from IBM development.

A customer can estimate the amount of expired but active data the VTS is managing by multiplying the following information:

- The # of scratch logical volumes (from the TMS or Display library details through z/OS)
- Average logical volume size (from the SMF94 data, but it is the compressed value)

For example, if the VTS has 10,000 scratch logical volumes and the average volume size is 100 MB, then the VTS is managing 10,000 X 100 = 1,000 GB of expired data. Assuming that the VTS uses 3590E drives that stores approximately 20 GB on a cartridge, that represents an additional 50 cartridges.

6.6.1 Fast-Ready category

An option is provided on the **Define Fast-Ready** definition menu (Figure 6-22). This option will delete the data associated with logical volumes that have been returned to scratch status. The option makes use of the existing Fast-Ready category concept. Scratch categories are identified as Fast-Ready categories. This indicates to the VTS that the data on volumes in categories identified as Fast-Ready is expired. An additional parameter has been added to the Library Manager panel that defines the fast-ready categories to enable the deletion of expired data.

The new parameter, Expire Time, specifies the amount of time in hours, days or weeks, the data will continued to be managed by the VTS after a logical volume is returned to scratch before the data associated with the logical volume will be deleted. A minimum of 24 and a maximum of 32,767 hours (approximately 194 weeks) can be specified. Specifying a value of zero means that the data associated with the volume is to be managed as it was prior to the addition of this option, meaning that it is never deleted. In essence, specifying a value (other than zero) provides a “grace period” from when the logical volume is returned to scratch until its associated data is eligible for deletion. A separate expire time can be set for each category defined as fast-ready.

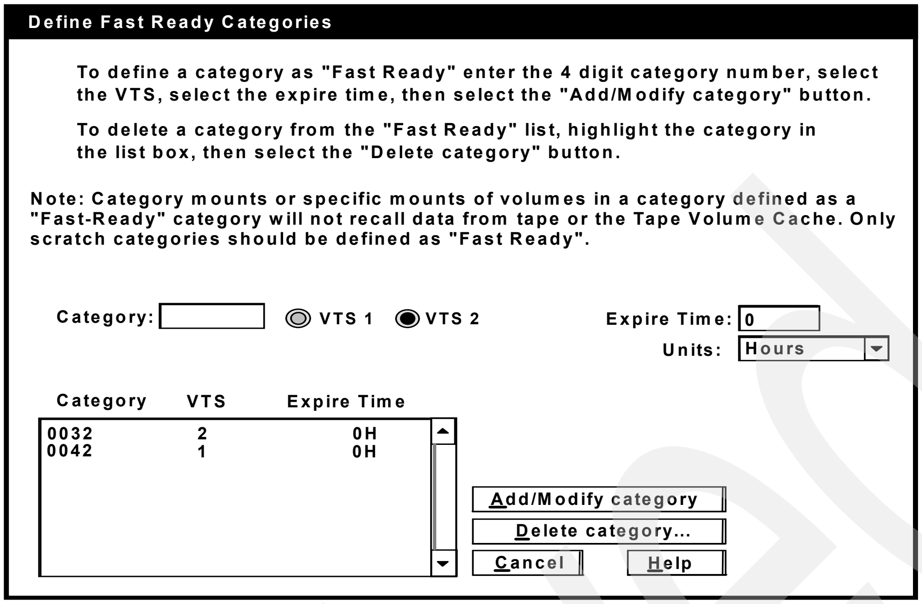

Figure 6-22 Define Fast Ready Categories window

Establishing the expire time for a volume occurs as a result of specific events or actions. The possible events or actions and their effect on the expire time of a volume are as follows:

- **A volume is mounted:** The data associated with a logical volume will not be deleted, even if it is eligible, if the volume is mounted. Its expire time is set to zero, meaning it will not be deleted. It will be re-evaluated for deletion when its category is subsequently assigned.
- **A volume's category is changed:** Whenever a volume is assigned to a category, including assignment to the same category it is currently in, it is re-evaluated for deletion.
- **To a category with a non-zero expire-time:** If the category has a non-zero expire time, the volume's data will be eligible for deletion after the specified time period, even if its previous category had a different non-zero expire time.
- **To a category with an expire-time of zero:** If the volume's previous category had a non-zero expire time or even if the volume was already eligible for deletion (but has not yet been selected to be deleted) and the category it is assigned to has an expire time of zero, the volume's data is no longer eligible for deletion. Its expire time is set to zero.
- **A category's expire time is changed:** If a user changes the expire time value through the fast-ready definition panel on the Library Manager, the volumes assigned to the category are re-evaluated for deletion.
- **Non-zero to zero:** If the expire time is changed from a non-zero value to zero, volumes assigned to the category that currently have a non-zero expire time are reset to an expire time of zero. If a volume was already eligible for deletion, but had not been selected for deletion, the volume's data is no longer eligible for deletion.
- **Zero to non-zero:** Volumes currently assigned to the category continue to have an expire time of zero. Volumes subsequently assigned to the category will have the specified non-zero expire time.
- **Non-zero to non-zero:** Volumes maintain their current expire time. Volumes subsequently assigned to the category will have the updated non-zero expire time.

Once a volume's expire time has been reached, it is eligible for deletion. Not all data eligible for deletion will occur in the hour it is first eligible. Once an hour on the half hour, the Library Manager selects up to 500 eligible volumes for data deletion. The volumes are selected based on the time that they became eligible, with the oldest ones being selected first. Up to 500 eligible volumes for the first VTS in the library are selected first, followed by up to 500 eligible volumes in the second VTS, if it exists.

When the Library Manager database is backed up, the expire-time for all volumes will be set to zero in the backup file. The expire times in the database itself will not be changed.

The status of a volume's expire time can be viewed via the query database functions provided through the Library Manager.

6.6.2 Scratch volume recovery

The advantage for this method of managing the data is that if a user determines that a volume was mistakenly returned to scratch, they only have to return the volume to private status to recover from the mistake, as long as they have not re-used the volume or the "grace period" has not expired. The method to recover depends on the tape management system used. For example, for DFSMSrmm, the following command will return the volume to private status and increase its retention period, including communicating the change to the VTS and Library Manager (see the DFSMSrmm Guide and Reference for complete details of the command):

```
RMM CHANGEVOLUME yyyyyy STATUS(USER) RETPD(days) OWNER(userid), where yyyyyy is the
volser
```

6.7 Ejecting logical volumes from the VTS

Logical volumes are not physical entities that can be individually removed from the library. They reside on stacked volumes with many other logical volumes. If you issue an EJECT for a logical volume, all data on that volume will be lost.

Attention: There is no way to recover the data on the logical volume once the EJECT command is processed.

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.

Note: This fact has proven to be cumbersome for volumes that happen to be in the error category (000E). An easy way to eject such volumes is to use ISMF panels to set these volumes to the PRIVATE status. The volume status is propagated to DFSMSrmm. You can use DFSMSrmm to subsequently assign the volume to the SCRATCH status and eject it.

Ejecting large numbers of logical volumes can have a performance impact on the host and the library.

Tapes that are in INSERT status can be ejected by the resetting of the return code via the CBRUXENT exit. This exit is usually provided by your tape management system vendor.

Once the tape is in SCRATCH status, follow the procedure for EJECT processing based on whether your environment is system-managed tape or BTLS. You will also need to follow the procedure specified by your tape management system vendor. For DFSMSrmm, issue the **RMM CHANGEVOLUME volser EJECT** command. If your tape management system vendor does not specify how to do this, you can use one of the following commands:

```
z/OS command LIBRARY EJECT,volser
IDCAMS command LIBRARY EJECT,volser (for BTLS)
ISMF EJECT line operator for the tape volume
```

The eject process fails if the tape is in another status or category. For libraries managed under DFSMS system managed tape, the system command **LIBRARY EJECT,volser** issued to a logical volume in PRIVATE status fails with this message:

```
CBR3726I Function incompatible error code 6 from library <library-name>
         for volume <volser>.
```

Note: In a DFSMS system-managed tape environment, if you try to eject a logical volume and get this error, OAM notifies the tape management system. This is done through the OAM eject exit CBRUXEJC before the eject request is sent to the tape library. The Library Manager will eventually fail the eject, but the tape management system has already marked the volume as ejected. Prior to APAR OW54054 there was no notification back that the eject has failed.

Failed Eject Notification was added to OAM with APAR OW54054 and currently in all supported releases of DFSMS. Any tape management system supporting this notification can use this function.

If your tape management system is DFSMSrmm, you can use the following commands to clean up the RMM CDS for failed logical volume ejects and to re synchronize the TCDB and RMM CDS:

```
RMM SEARCHVOLUME VOL(*) OWN(*) LIM(*) INTRANSIT(Y) LOCATION(vts) -
    CLIST('RMM CHANGEVOLUME ',' LOC(vts)')

EXEC EXEC.RMM
```

The first RMM command asks for a list of volumes that RMM thinks it has ejected and writes a record for each in a sequential data set called *prefix.EXEC.RMM.CLIST*. The CLIST then checks that the volume is really still resident in the VTS library and, if so, it corrects the RMM CDS.

Attention: Limiting the number of outstanding ejects to a couple thousand total per system will limit exposure to performance problems.

There are considerations to be aware of when ejecting large numbers of logical volumes. APAR OW42068 introduced Peer-To-Peer toleration support. This in effect treats all libraries as if they had the same eject processing restrictions as a Peer-to-Peer. This support is currently in all supported releases of DFSMS. There is a 1000 logical volume eject limit that the Peer-to-Peer VTS can handle at any one time. 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. Additionally, there is a restriction on the number of eject requests on the device services' queue. All of these conditions can have an impact on the host's performance.

So the recommended limit for the number of outstanding 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 zSeries hosts to list the outstanding and the active requests.

```
F OAM,QUERY,WAITING,SUM,ALL
F OAM,QUERY,ACTIVE,SUM,ALL
```

6.8 Querying the Library Manager database

The Library Manager enables you to search the Library Manager database and view selected volumes in the 3494 according to specified search criteria. You can search for both logical and physical volumes.

Use the **Database** pull-down menu (Figure 6-23) on the action bar to select one of three search functions.

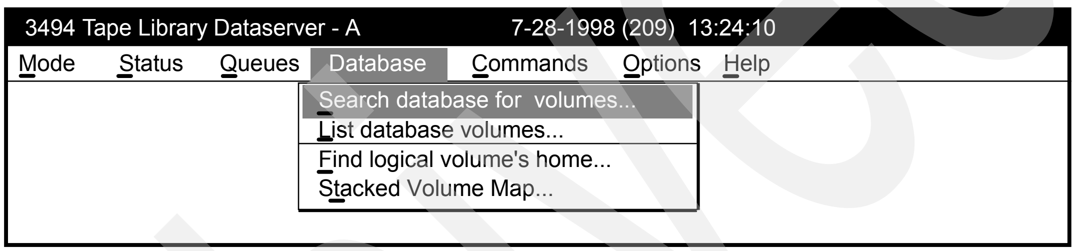

Figure 6-23 Library Manager Database pull-down menu

6.8.1 Search database for volumes

Select the **Search database for volumes** function from the **Database** pull down menu to view selected volume information on the Library Manager console. Specify your search criteria in the Search Database for Volumes window.

Figure 6-24 shows the respective ETL Specialist Welcome panel to query the Library Manager database. Select **Search database** from the **Administer Library Manager** work items.

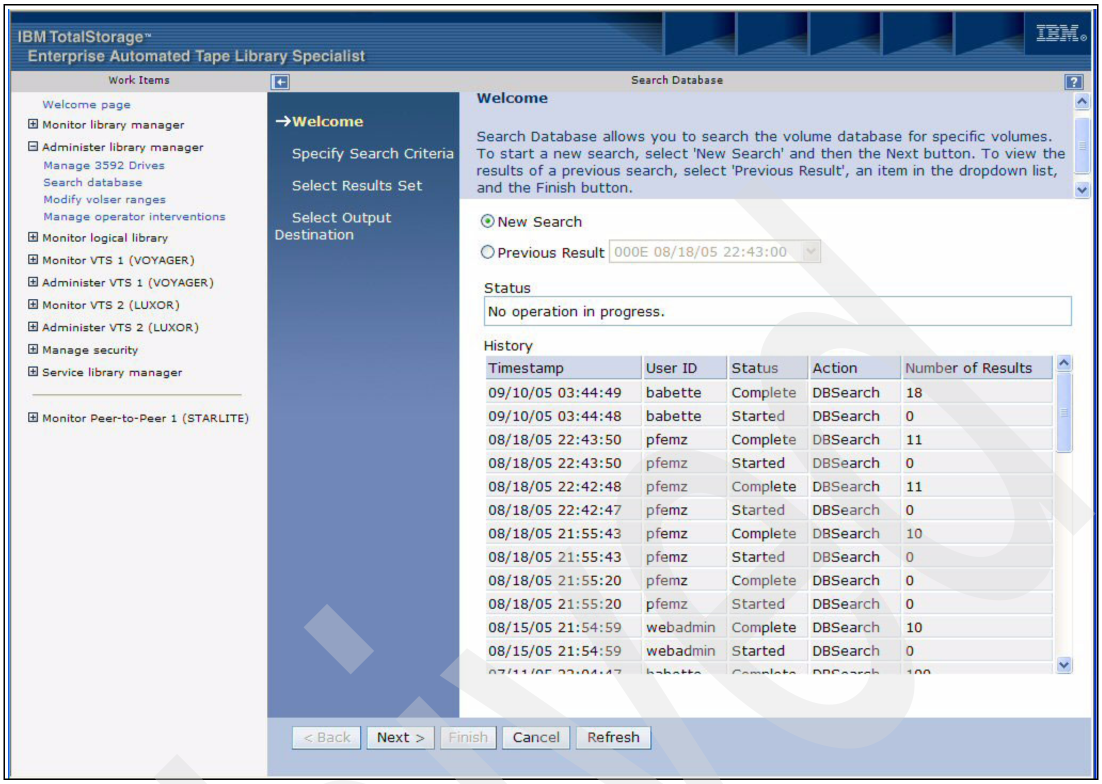

Figure 6-24 ETL Specialist Search Database for Volumes window

When clicking **Next**, the panel shown in Figure 6-25 is displayed, which allows you make your selections for the search criteria. You can specify the following search criteria:

- Volser
- Category
- Media type
- Device
- Partition
- Expire time
- Home rack
- Home row
- Home column
- Home pool
- Current rack
- Current row
- Current column
- Current pool
- APM construct names
  - Storage group
  - Management class
  - Storage class
  - Data class
- Volume flags
  - Misplaced
  - Unreadable
  - Mounted
  - Inaccessible
  - Manual mode

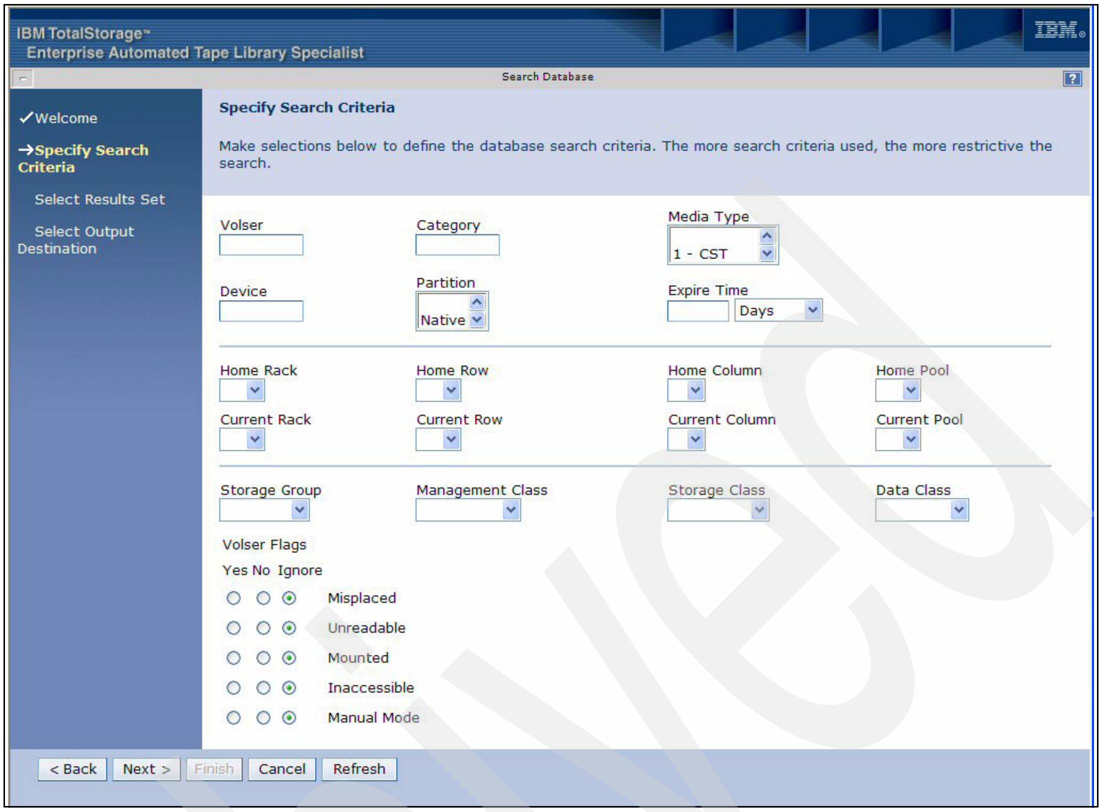

Figure 6-25 Search Database window: Select Search Criteria

After specifying your search criteria and clicking **Next**, the window shown in Figure 6-26 is displayed. In this panel you can select which values you want to have displayed.

You can select the following values to be displayed for the volumes that meet the search criteria specified before:

- Volser
- Category
- Media type
- Device
- Partition
- Expire time
- Home rack
- Home row
- Home column
- Home pool
- APM construct names
  - Storage group
  - Management class
  - Storage class
  - Data class
- Volume flags
- Category order
- Mount date
- Number of mounts

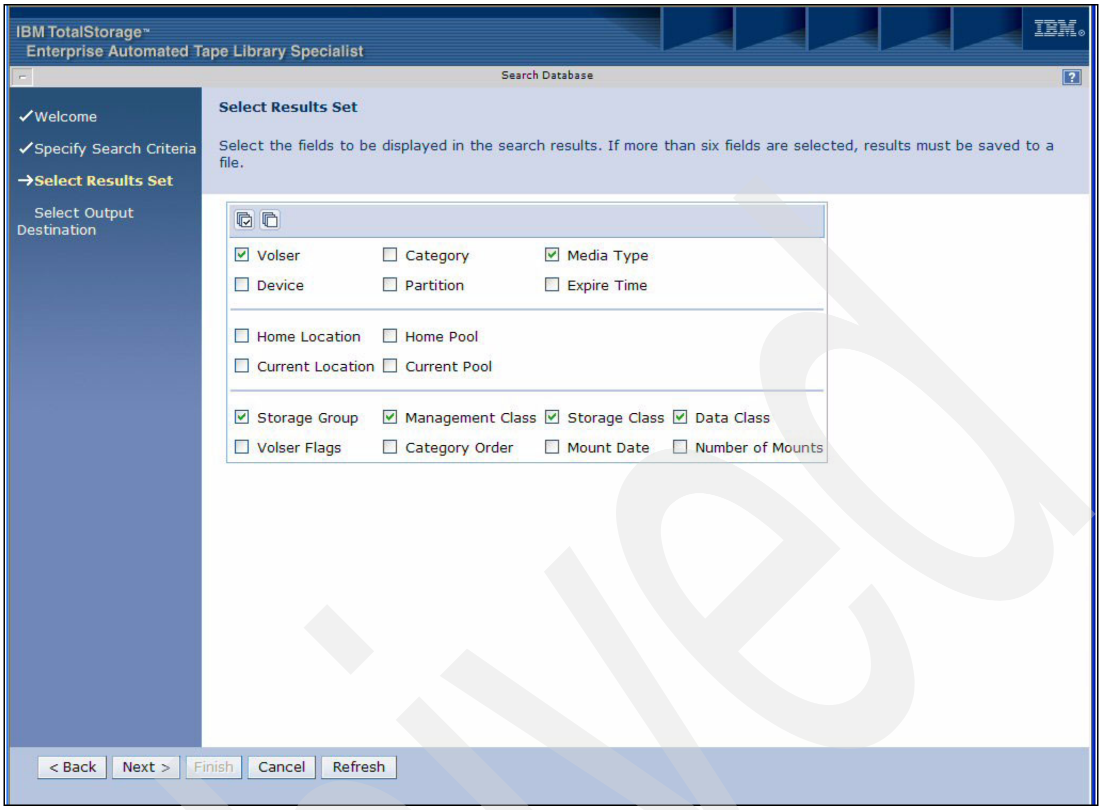

Figure 6-26 Search Database window: Select Results Set

Depending on the number of columns selected in the window shown in Figure 6-26, you can select whether you want to have your output displayed on the screen or written to a file (see Figure 6-27).

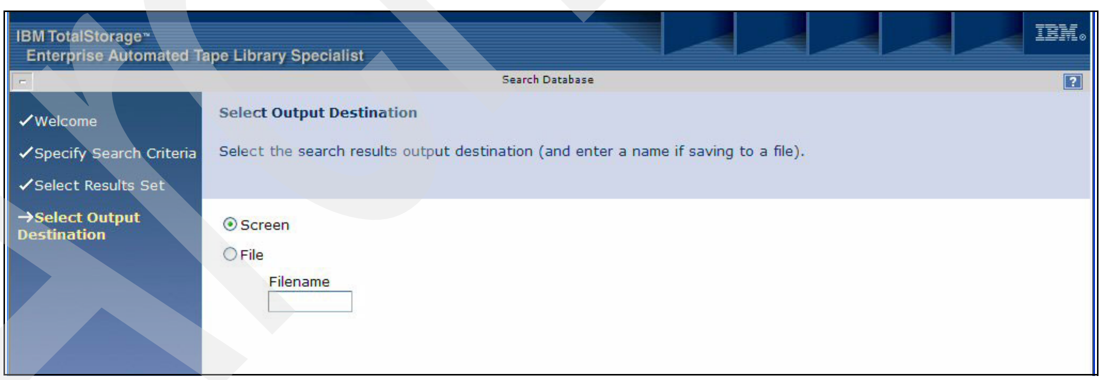

Figure 6-27 Search Database window: Select Output Destination

If you select between one and six columns for the Results Set, you can select between Screen and File on the panel shown in Figure 6-27. If you select seven or more columns, the output will be written to a file.

Figure 6-28 shows the sample output of a search selecting all volsers.

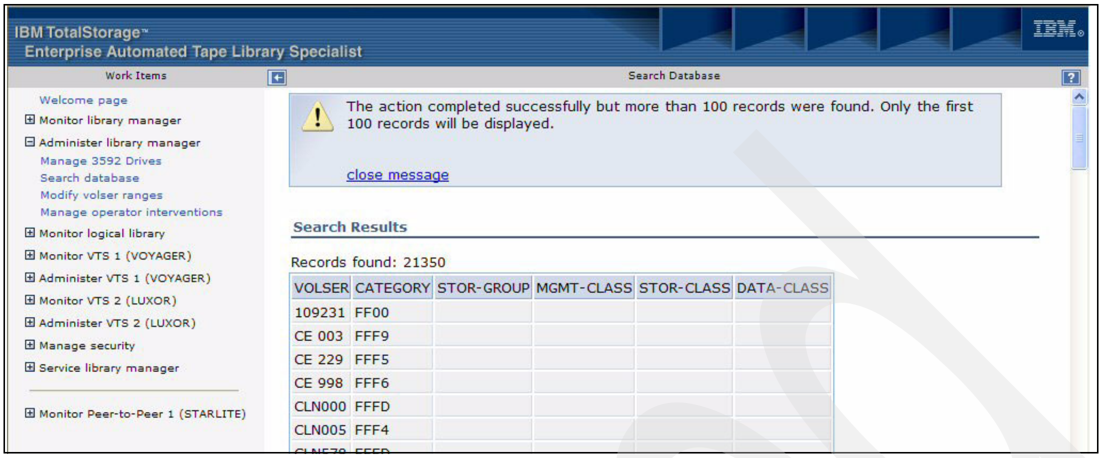

Figure 6-28 Search Database window: Search Results

As you can see from the categories listed in the example shown in Figure 6-28, these are not private VTS volumes and they do not have construct names assigned.

6.8.2 Find logical volume's home

Select **Find logical volume's home** to find out on which stacked volume a logical volume resides. Specify the logical volume's serial number on the resulting Find A Logical Volume's Home window. Refer to Figure 6-29 or Figure 6-30 as appropriate for your situation. With a VTS a LIC level 527 or greater, this search can be performed at either the LM or the 3494 Specialist.

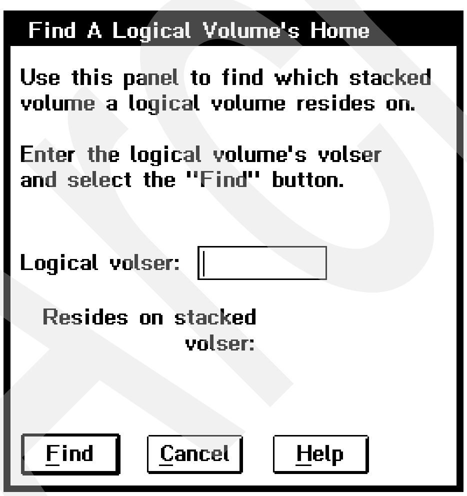

Figure 6-29 Find a Logical Volume's Home window pre-527 LIC level

If the logical volume has been written on a stacked volume, the volser of the stacked volume that holds the active copy of the logical volume is displayed in the window. The search cannot succeed if the volume has not yet been copied on any stacked volume.

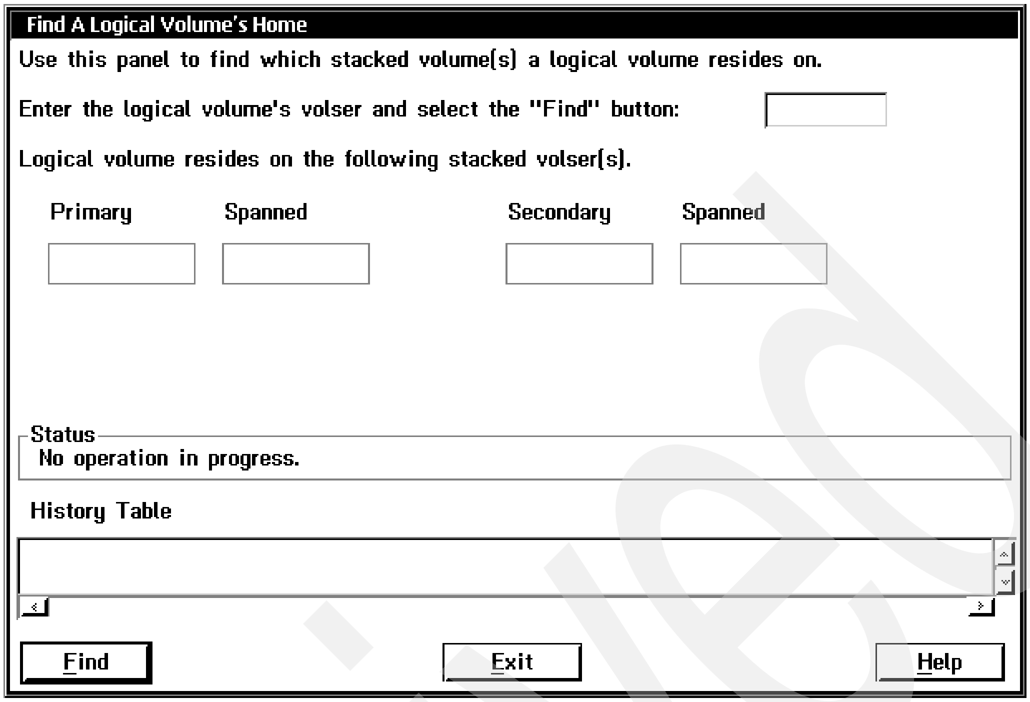

Figure 6-30 Find a Logical Volume's Home window post-527 LIC level

The logical volume can potentially be stored on up to four stacked volumes as shown in Figure 6-30. Here are the definitions for any of the volume types a logical volume could exist on.

Primary: The primary physical volume that contains the specified logical volume.

Spanned: If the logical volume does not fit entirely on the primary stacked volume, the remainder of the data is written to this spanned primary stacked volume. The data that makes up the requested logical volume has been split between the primary stacked volume and the primary spanned volume.

Secondary: If a secondary copy of a logical volume has been made (as defined by the volume's management class), this is the physical volume that contains the secondary copy of the logical volume. A logical volume can only have a secondary copy if Advanced Policy Management is installed and enabled and the volume's management class indicates that a secondary copy should be made.

Spanned: If the logical volume does not fit entirely on the secondary stacked volume, the remainder of the data is written to this spanned secondary stacked volume. The data that makes up the requested logical volume has been split between the secondary stacked volume and the secondary spanned volume.

If a logical volume is in use at the time you issue the query, the window indicates that the volume is open.

The corresponding 3494 Specialist panel is shown in Figure 6-31.

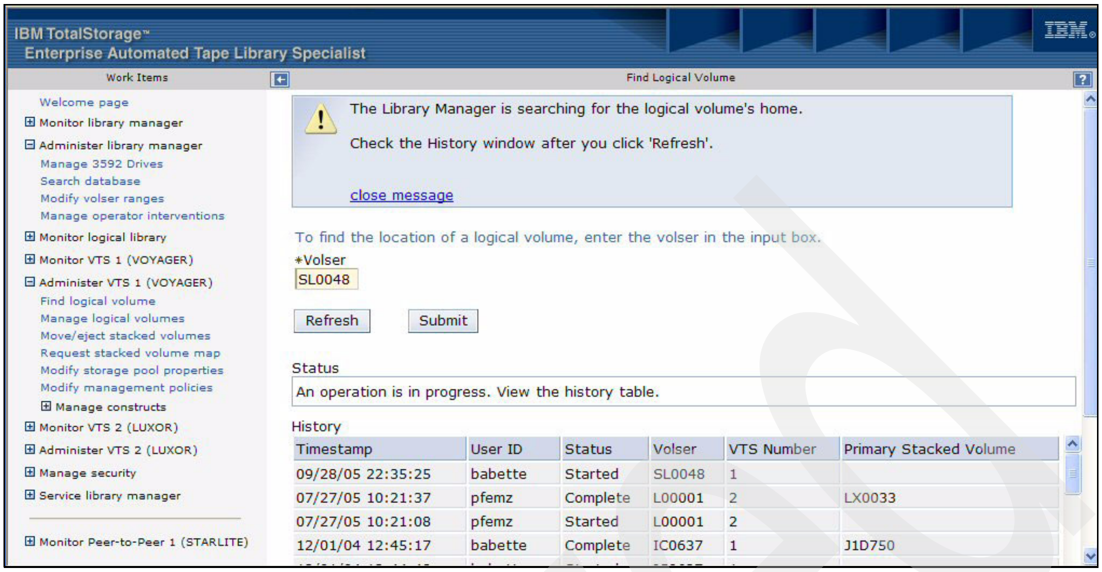

Figure 6-31 ETL Specialist Find Logical Volume Panel

After the request has completed, you can see the details by clicking the respective line of the History table, and the panel shown in Figure 6-32 is displayed.

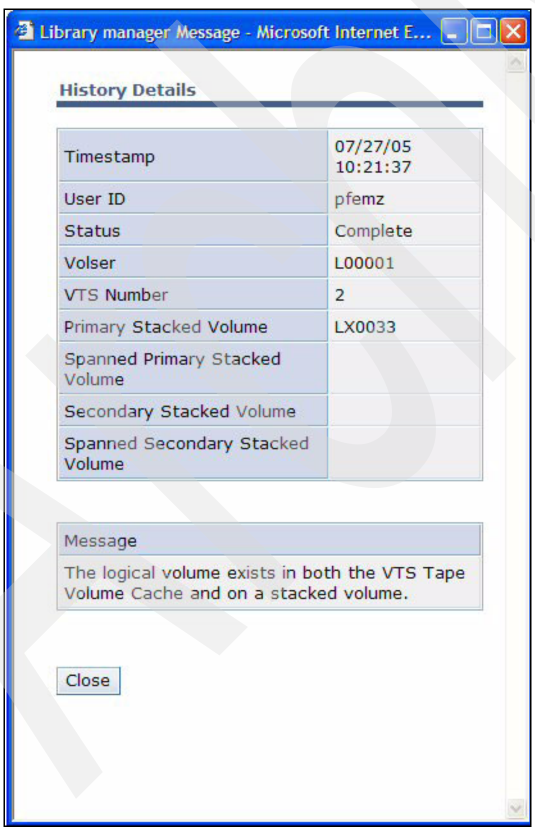

Figure 6-32 History Table Details window

6.8.3 List logical volumes on a stacked volume

There may be certain situations where you would like to find out which logical volumes are on a stacked volume. If you have a stacked volume that is marked Read-Only, you may want to find out which logical volumes are on the stacked volume before issuing the eject. To perform this function, you can use the LM panels of the 3494 or 3953 Library Managers. However, the easiest way is to use the ETL Specialist panels, as they provide a file which you can store directly on the workstation or PC you are using to create the Stacked Volume Map. Select **Request Stacked Volume Map** from the **Administer VTSx** work items to see the panel shown Figure 6-33.

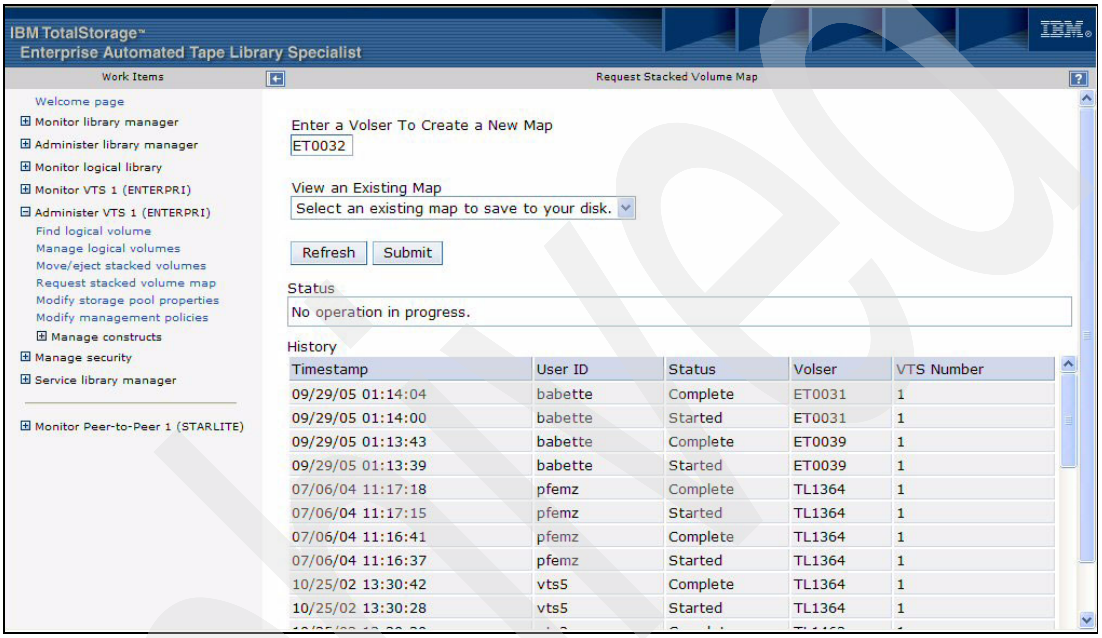

Figure 6-33 ETL Specialist window: Request Stacked Volume Map

The file that is created will have a file named STxxxxxx.MAP, where xxxxxx is the volser of the stacked volume. The file will contain some header information and a list of all of the logical volumes on the stacked volume.

Some of the volumes may be in the form L1-nnn where nnn is a number. The L1-nnn volumes are log files that are created by the VTS. They are used by service and are listed so service can find them easily.

If a logical volume spans a second stacked volume, the text 'SPAN' will be in the record. See the *IBM TotalStorage 3494 Tape Library Operator Guide*, GA32-0280 for more information.

Figure 6-34 shows the Stacked Volume Map created for stacked Volume ET0032.

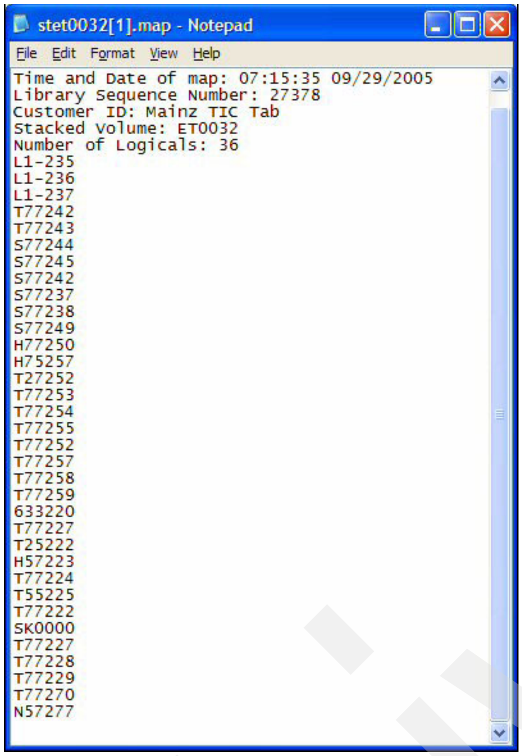

Figure 6-34 Stacked Volume Map

6.8.4 List Database Volumes function

The **List Database Volumes** function enables you to create a flat file containing a customized list of selected database columns. Specify your selection criteria in the List Database Volumes window (Figure 6-35).

> You can also obtain a list of volumes directly on your workstation or PC through the ETL Specialist:
>
> 1. Select **Search Database** from the **Administer VTS** work items.
> 2. Define your search criteria.
> 3. Define your results set.
> 4. Define your output destination.
>
> See 6.8.1, “Search database for volumes” on page 250 for more details and a description of the panels.

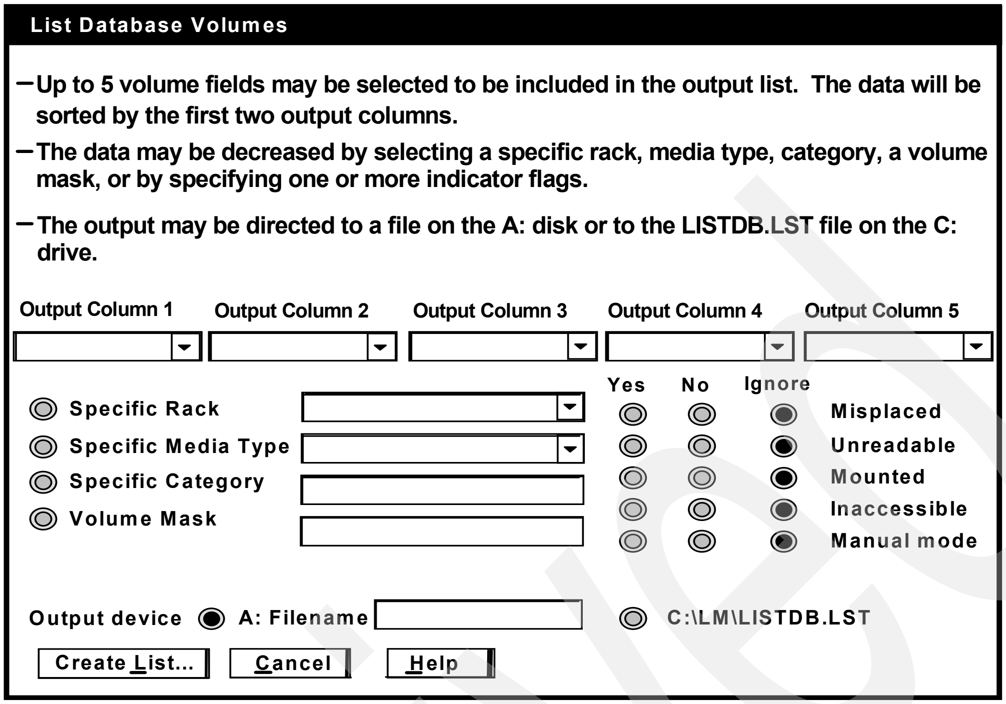

Figure 6-35 List Database Volumes window

Use the Output Column fields to specify which columns to include in the output file. Use the various search criteria fields to specify which volumes to include in the output file.

The **List Database Volumes** function lets you use the same volume selection criteria as the **Search Database** for Volumes function. In addition, it enables you to find out which volume is in a specific rack in the 3494 or 3584.

You can write the flat file to a selectable file name on a diskette in the Library Manager floppy disk drive, or to the C:\LM\LISTDB.LST file on the Library Manager hard disk.

Restriction: The Library Manager must be offline for the List Database Volumes operation.

6.9 Inventory update

On the 3494 Tape Library

The Inventory Update facility allows for the following operations, which you can select from the Commands pull-down:

- **Disable inventory update:** This disables inventory update, allowing doors to be opened and closed without performing inventory update.
- **Enable inventory update:** This enables inventory update to take place whenever the tape library system is returned to Auto/Online mode after an enclosure door is opened and closed.
- **Perform inventory update (Full):** This performs an immediate inventory update. This option is available only if the tape library system is in Auto/Online mode.
- **Perform inventory update (Partial):** This performs inventory on just the frames on which the doors have been opened.

During inventory update processing, the actuator checks all cartridge cells and drive feed slots for cartridges, reads all bar code labels and updates the Library Manager database as follows:

- Cartridges not previously in the Library Manager inventory are set to insert category (FF00)
- Cartridges that are in the Library Manager database but not found during the inventory are placed in the Manually ejected category (FFFA)
- Cartridge cell locations are updated if necessary

Logical volumes are not affected by inventory update.

> **Note:** All other library activities (mounts, demounts,...) are held during inventory update processing. Requests from the hosts are queued until the library is back in Auto mode.

When the **Enable inventory update** option is selected, a partial inventory update is automatically performed on those library frames whose front door has been opened while the library was in Pause or Manual mode. This process checks all of the cartridge storage cells in the frames that had doors opened, and depending on the selection made during the teach process by the IBM service representative, may also check any adjacent frames to the frames that had doors opened. When the **Disable inventory update** option is selected, you must manually start the Inventory update process.

Refer to the *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280, for more information about the Inventory update facility.

Inventory commands on 3953 for 3584 Tape Library

Since the 3953 does not directly control the accessor's in the 3584 Tape Library, there must be a means of updating the Library Manager's database with any changes in the inventory of physical cartridges in the 3584. One method the 3584 uses is to inventory any frames that have had the door opened during a pause of the library. But there are times when this is method alone is not enough. Performing the Request Inventory Upload command will satisfy this requirement.

The Request Inventory Upload window (Figure 4-36) allows you to manually update the physical volume inventory from the 3584 Tape Library. This procedure is necessary after insertion of a volume into the 3584 Tape Library that does not match any defined 3584 Tape Library Cartridge Assignment Policies. If you do not respond to the insert notification at the tape library operator panel, the inserted volume will not be assigned to a logical library.

> **Attention:** Before you use the Request Inventory Update window, a system administrator must first designate the correct logical library for the inserted volume using the Web Specialist of the 3584 Tape Library.

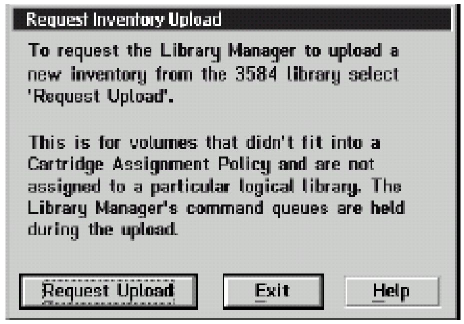

Figure 6-36 Request Inventory Upload

Refer to the *IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0473, for more information about inventory functions of the 3584.

6.10 Operating the VTS from a host

One of the original design objectives of the VTS was that it would not require host software changes. In consequence, no new host commands or new host messages are associated with the VTS. The VTS advanced function features (FC4000-4004) require host software changes and introduces new messages to facilitate the Import/Export processing. Refer to Appendix B, “VTS Import / Export” on page 369 for details. At the host, you use the same commands and procedures as you use with the 3494 or 3953 Library Manager today.

Error and status conditions noted at the library that relate to a VTS are handled using existing software facilities. For a Library managed under DFSMS, intervention required conditions are reported on the host console with this message, with no indication of severity:

```
CBR3762E Library <library-name> intervention required
```

> **Note:** If **Send Interventions to Host Console** is enabled, some new, condition-specific, host console messages are generated. See “Messaging enhancements” on page 218 for further explanation.

When your Library is logically partitioned, you operate the logical libraries as you would operate separate IBM 3494 tape libraries. For example, you can vary the logical libraries online and offline to the host independent of each other.

Note that the physical tape drives attached to the VTS are not known to the host operating system and are not directly accessible by host programs.

Refer to the following product manuals for guidance on how to operate a 3494 on the different host platforms:

- *z/OS Object Access Method Planning, Installation and Storage Guide for Tape Libraries*, SC26-0427
- *Basic Tape Library Support User's Guide and Reference*, SC26-7016
- *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141

6.10.1 Data security pre-Release 7.4

In some systems, execution of the `data security erase` channel command causes random data patterns to be written on a volume to its end-of-tape. Instead, with VTS, the end-of-data (EOD) mark is repositioned at the logical block location at which the command was issued. Any data beyond that logical block location is no longer accessible.

When the virtual volume is copied to a stacked volume, only the data up to the EOD mark is written. Any data beyond that point is ignored; it is no longer accessible by the host.

The VTS does not take any action or provide any function to erase data on invalidated copies of a logical volume. Although the previous copies of a logical volume are not readily accessible, data on them remains on physical volumes. These copies are not written over until the corresponding stacked volume is reclaimed and the space occupied by the data is subsequently overwritten by new logical volumes.

6.10.2 Secure data erasure

With the introduction of microcode Release 7.4 for the VTS, the data security erasure process has changed. Physical volume erasure is done on a pool basis and is controlled by an additional reclamation policy for each pool. When a logical volume is selected for erasure the entire physical stacked volume this logical is stored on will be over written with random pattern of data. All the logical volumes on this tape that are not selected for erasure will be copied to another stacked volume.

Once the data erasure has been started for a volume, it cannot be stopped or interrupted. This can affect the performance of the VTS when large stacked cartridges, 300 GB JJ tapes for example, are flagged for Data Security Erasure. A VTS drive will be in use until all the logical volumes are removed from the select stacked cartridge and the random pattern has been written to entire length of the tape. See Figure 4-22 on page 143 for setting this policy.

> **Attention:** This process once started, will not be interrupted for recalls or reclaims. Depending on the number of backend drives, mount times for recalls of logical volumes may increase.

6.10.3 Stand-alone support

Stand-alone support is required, for example, for restoring from a virtual volume when SMS is not active.

In the 3494 Tape Library

The Library Manager provides a Stand-Alone Device facility that allows you to set a tape device (virtual or native) in a 3494 in stand-alone mode. To the host, the device then appears as if it were not in a tape library. This mode allows you to use devices in the 3494 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.

In the 3584/3953

The 3953 Library Manager does not have control over the accessor and convenience input/output station as in the 3494. This makes the option to Mount From Input Station not available in the 3584/3953 system. The rest of the facility is the same as the 3494 Library Manager. See Figure 6-37.

Stand-Alone Device facility

The Stand-Alone Device facility supports virtual devices with the exception that the Mount from Input Station feature is not supported with virtual devices. Using the Stand-Alone Device facility, you can:

- IPL from a virtual tape device
- Use virtual tape volumes with stand-alone programs that do not support mounting of tapes on native attached tape drives resident in the library.

To set a virtual device in stand-alone mode and to request the Library Manager to mount a virtual volume on it, select **Stand-alone device** from the **Commands** pull-down menu, then select **Setup stand-alone device.** Provide the following input on the Setup Stand-alone Device window (Figure 6-37), which shows the IBM 3953 Library Manager panel.

1. Enter a virtual device address in the Enter device field.
2. Enter the volser of the virtual volume in the Volser field.
3. Select **Mount a single volume** to mount the specified volume into the specified virtual drive.
4. Select **Do not change volume category**.

When you select **OK**, the Library Manager requests the VTS to mount the virtual volume and make the virtual device ready. If necessary, the VTS then recalls the virtual volume from the stacked volume into the TVC.

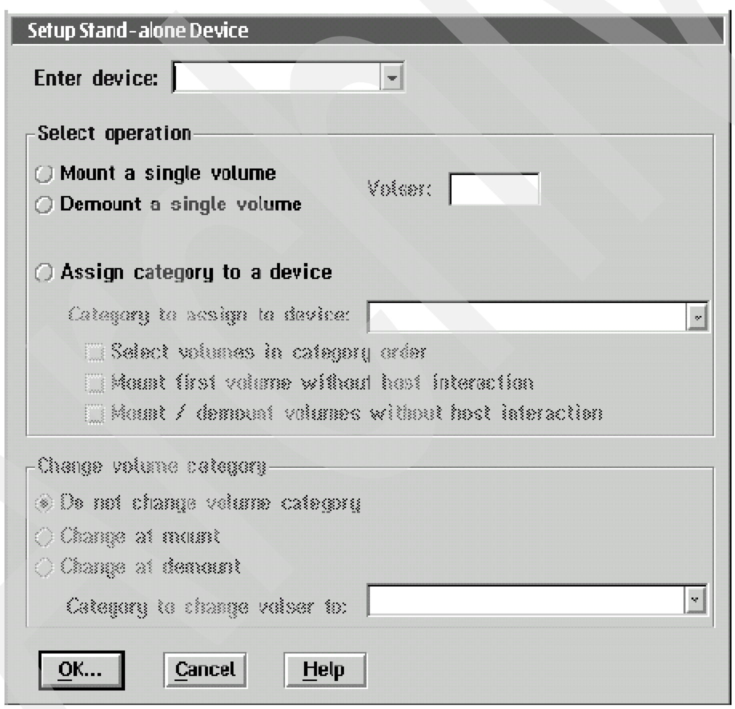

Figure 6-37 Stand-alone device support in the 3584/3953

Note: Vary off the device that is being used in stand-alone mode from all hosts except the host that is being used in this special mode. This will prevent unwanted interaction from other hosts connected to the VTS.

You cannot set the physical TotalStorage 3590 drives or 3592 drives associated with a VTS in stand-alone mode.

To unload and demount a virtual volume from a drive that is in stand-alone mode, select the **Demount a single volume** and **Do not change volume category** push buttons on the Setup Stand-alone Device window, then select **OK**.

To take a virtual device out of stand-alone mode, use the Reset Stand-alone Device window, which you can access by selecting **Stand-alone device** from the **Commands** pull-down menu, then selecting **Reset stand-alone device**.

See “Stand-alone support” on page 262 for an example of how to IPL Stand-Alone Services from a virtual device and use it to restore a dump data set from virtual tapes.

6.11 Error recovery

In this section we document the error conditions that can occur in a VTS and, where appropriate, the actions you should take to recover from them.

Because the VTS has been designed to operate with as little manual intervention as possible, many of the error conditions can be resolved with little or no action on your part.

6.11.1 IBM VTS resilience

The VTS, the IBM TotalStorage Enterprise Automated Tape Library 3494 and the IBM TotalStorage 3584 Tape Library include a number of features that provide resilience or backup in the event of a failure:

- Holding the TVC on RAID fault-tolerant disk arrays
- Providing two data paths to the IBM 3590 and 3592 VTS tape drives
- Dual communications paths between VTSs and Library Managers
- Providing up to sixteen ESCON ports
- Providing up to eight SCSI ports
- Providing up to eight FICON ports
- Holding four to eighteen IBM 3590 and 3592 tape drives
- Offering a second Library Manager hard disk drive or second Library Manager
- Making available an optional dual gripper (with 3494, standard on 3584)
- Providing the ability to create dual copies of logical volumes
- Offering an optional second accessor and service bay with the IBM TotalStorage High Availability Tape Frames (HA1 Frames)

There are, however, some components that are not fully protected. In the event of a failure of one of the following unprotected components, the VTS will not be available for use:

6.11.2 VTS controller

The VTS controller consists of CPU, memory, adapters, internal disks, and power supplies. In the Model B18 VTS, adapters and internal disks are located inside the controller, therefore, if an adapter or internal disk requires servicing, the entire VTS’s operations will be affected. If the VTS controller fails, the data stored on the VTS cannot be processed.

The Model B10 VTS and Model B20 VTS have their adapters and internal disks in a separate *I/O Drawer* and *Boot/Database Drawer*. Adapter cards are in the I/O Drawer separately located outside of the controller. If there is an adapter card failure in the VTS, it will not impact the availability of the VTS. The affected resources may be degraded, but you will have complete access to the data stored in the VTS. Servicing of failed adapters, power supplies, and disk drives are in most cases a concurrent operation. Quiesced resources during the repair operation will be transparent to the operation of the VTS.

AIX boot images are in the disks of the Boot/Database Drawer and are mirrored for redundancy. Replacement of a failed boot disk is not concurrent. Replacement of one of the other four mirrored disks used for the internal VTS operation are concurrently replaceable to normal VTS operation. With the release of LIC 7.2 for the VTS, a background routine runs on a monthly basis to detect latent HDISK errors. If a latent error is detected on a scan, an SIM and an RSF call home are generated. This will allow for proactive repairs of degrading HDISKs.

6.11.3 Power distribution

There is only one power control compartment (PCC) fed by a single line cord on each 3494-L1x Frame, 3494-Dxx Frames and early manufactured Model B18 VTS. If this fails, the associated frame loses complete power. If the Early Model VTS B18 fails in this way, all the data on the VTS is not available for processing. If the VTS' associated drive frame loses power, the VTS will remain powered up, but if the failed frame contains all of the VTS drives (as in a B10 or B18 with 6 drives or a B10 or B20 with only 3592 drives) the VTS will halt operation because it has no back-end tape drives to copy the logical volumes to. All subsequent VTS models have two primary control compartments and two power cords. This ensures continued availability of the VTS even in the case of the failure of one power source. The 3953 and the 3584 also have dual power source to eliminate this as a single point of failure.

Note: Be aware that the Unit Emergency Switch (EPO) of a stand-alone VTS is separate from the switch on the Library Manager. Setting the VTS EPO to the “OFF” position will not affect the power on the Library Manager. But setting the 3494 Tape Library EPO to the “OFF” position will affect the power on the VTS. Also be aware that removal of power in this manner can result in loss of customer data. There is no EPO switch on the 3584 or 3953.

6.12 Recovery scenarios

In this section we discuss some of the potential recovery scenarios you might be required to perform. Figure 6-38 shows some of the possible areas where your involvement might be required to perform these recovery actions. Most of the errors which require operator attention are notified through Intervention Required message in the Library Manager and hosts. Refer to “Messaging enhancements” on page 218 for a description on how to activate notification of intervention conditions on the host.

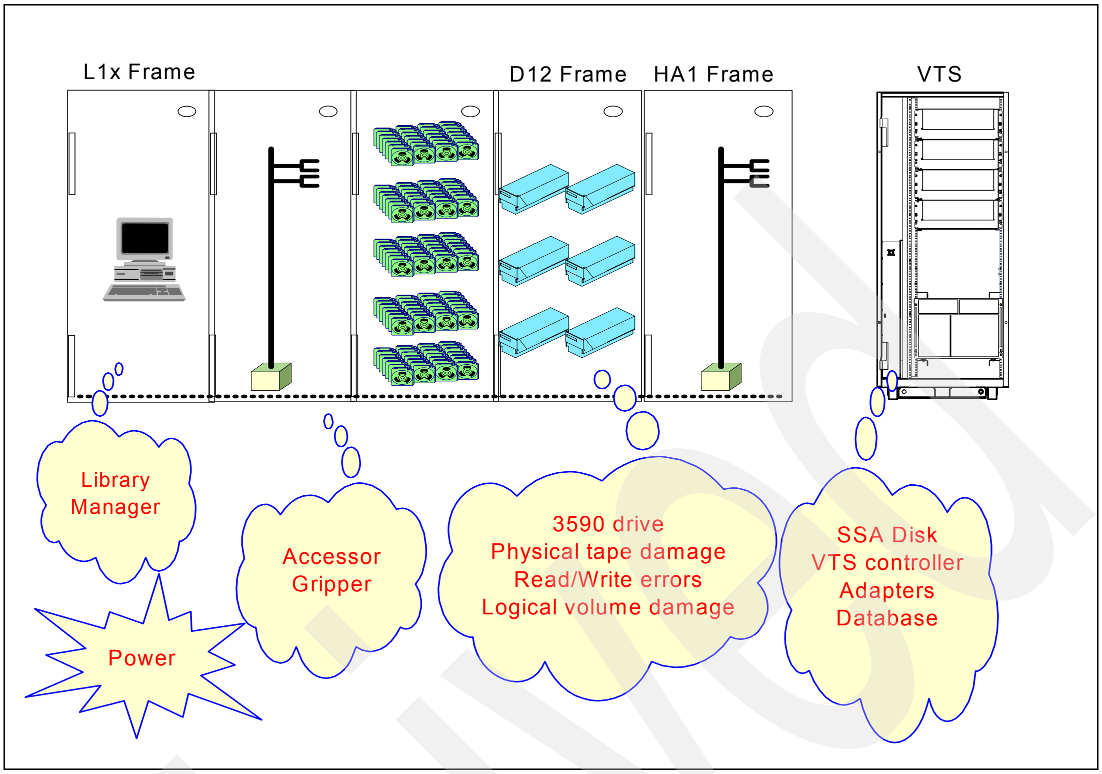

Figure 6-38 Failure components

Permanent write error

When the VTS attempts to copy a virtual volume from the TVC to a stacked volume and a permanent write error occurs, the stacked volume is swapped to another tape drive and the copy operation for the entire virtual volume is attempted a second time. If the second attempt fails, it is likely that the tape media caused the problem. The stacked volume is placed in read-only status, the copy of the virtual volume which cannot be completed will not be included on the stacked volume and that stacked volume is unloaded. A new stacked volume is mounted and the copy operation of the virtual volume from the TVC is performed from the beginning.

Permanent read error

When the VTS attempts to recall a logical volume from a stacked volume and a permanent read error occurs, the stacked volume is swapped to another tape drive and the recall is retried. If the attempt to recall fails on the second drive, the VTS indicates on the first I/O operation issued by the host that a drive equipment check has occurred. The host job which requested the logical volume which cannot be made available will be terminated. The VTS places the stacked volume in read-only status. When the host application attempts to read beyond the failure point or uses a tape positioning command that attempts to position the tape beyond the failure point, the application receives notification of a permanent error, as with a real IBM 3490E. The backup copy can be used to recover at this stage, thereby not causing the host job to fail. The dual copy feature helps with this.

Read-only status

When permanent read or write data errors occur on stacked tape volumes and are placed in read-only status, the following Intervention Required message is posted to all hosts (see Example 6-6).

Example 6-6 Read-only status

```
CBR3750I MESSAGE FROM LIBRARY lib: Stacked volume xxxxxx is in Read-Only status with a
reason code of yyyyy. (VTS z)
```

This informational message comes directly from the Library Manager. OAM is only the messenger of the message. Reason codes are explained in Table 6-1.

Table 6-1 Read-only reason codes

Reason code	Description
1039	Volume made read-only because it was a filling volume after a disaster recovery
1069	Volume made read-only because a database backup is missing
1087	Volume made read-only by VTS, could be due to I/O error

The VTS performs a recovery of read-only status stacked volumes as an independent process. Customer data on logical volumes which reside on the stacked volume continue to be accessible, but no new data will be written to that stacked volume. The Intervention Required notification will be sent to hosts for the first volume to be placed in read-only status within a 24 hour period. Any subsequent volumes that may have been placed in read-only status are visible at the Library Manager console.

> **Note:** A customer initiated eject operation is also available from the Library Manager to force the VTS to copy the logical volumes from the stacked volume to a new stacked volume. See "Eject stacked volumes" on page 238 for details.

Read-only status volume recovery

Automated read-only recovery is set as the default operation mode and will manage volumes in read-only status in a transparent manner. This is an option in the service panels of the VTS. The VTS stacked volumes will be placed in read-only status if they experience one of the following conditions:

- Permanent media error during copying, recall, or reclamation
- Error when writing the from the Tape Volume Cache to the stacked volume
- Excessive temporary media errors
- VTS drive conversion from 3590-B1A to 3590-E1A tape drives
- VTS drive conversion from 3590-B1A to 3590-E1A or H1A tape drives

Once an hour, the VTS checks for the presence of stacked volumes in read-only status. For each stacked volume found in read-only status, the VTS recalls the logical volumes into the tape volume cache and makes them eligible to be copied to another stacked volume, therefore removing the logical volumes from the damaged stacked volume one by one. Access to logical volumes on read-only status stacked volumes for recall is not impacted, except for the potential of re-occurrence of a permanent error if recalled again.

The VTS starts the read-only recovery as soon as it can. There is only one read-only recovery in process at any given time. It also cannot occur while an **`eject volume`** command is underway, or while reclamation is running.

When all stacked volumes of the read-only tape have been read into the cache and the tape ejects, then the read-only recovery was successful and the following Intervention Required message is sent to the host (see Example 6-7).

Example 6-7 Read-only status volume recovery

```
CBR3750I MESSAGE FROM LIBRARY lib: A Read-Only stacked volume xxxxxx has been ejected to
the Convenience Input/Output Station.(VTS z)
```

If an unrecoverable read error occurs on a logical volume in this process, the Intervention Required message is posted to the attached hosts, as you can see in Example 6-8.

Example 6-8 Read-only status volume recovery

```
CBR3750I MESSAGE FROM LIBRARY lib: Logical volume xxxxxx was not fully recovered from
damaged stacked volume yyyyyy. (VTS z)
```

If damaged logical volumes must be recovered from the ejected stacked volume, you should call the SSR to open a PMR and initiate the recovery action at the IBM support center.

With the dual copy feature (as part of APM), use of the logical volume backup copy is initiated when a recall failure occurs on the logical volume primary copy or for stacked volumes flagged as read-only. If a recall error occurs:

1. The VTS queries the Logical Volume Database for the location of the backup volume.
2. Flags in the Logical Volume Database are used to confirm integrity of the primary and backup files.
3. At this point the recall is tried again on the primary logical volume, which now contains a reference in the database to the backup location of the logical volume. If the recall is successful, the VTS again reschedules the copy of the logical volume.

If the recall fails again on the mount, the files are swapped again and the mount is failed.

Any failing stacked volumes are marked as read-only.

6.12.1 Statistical Analysis and Reporting System (SARS)

This is a method for determining whether read and write errors are caused by the tape media or by the drive hardware. The SARS microcode runs in the drive and keeps information about how the tape media is performing (read/write errors, and so on) over its life. This media information (Volume SARS or VSARS) is saved in the Volume Control Region (VCR) portion of the tape and travels along with the tape from drive to drive.

The VCR contains Device Block ID Map, Media Statistics and Format Identification. Information about how the drive hardware (Hardware SARS or HSARS) is performing is written in the drive memory and is saved into the flash memory every 8 hours. The SARS microcode runs when a cartridge is unloaded from the drive. At this time the media and hardware performance is updated and various algorithms are run to check if a media or hardware problem exists.

Volume SARS

The volume information saved on the media includes various types of errors experienced by the tape (read/write, temporary/permanent, and so on), the physical location of the error on the tape, the serial number of the drive where the error occurred and a judgment about the quality of the drive responsible for the error. The media algorithms (currently implemented in the code) are quite complicated and require the media to experience errors on multiple different drives before any media action is required.

The number of different drives varies depending upon the type of error, the physical location of the error, and the quality of the drives involved in the error. When VSARS determines that there is a problem with the media, it generates a Media Information Message (MIM) stating that the tape should be re-written or that the tape should be copied to a new tape and then discarded. If VSARS determines that a tape is bad, it can logically write-protect the tape.

Hardware SARS

The hardware information kept for the drive includes various types of errors experienced by the drive, Error Correction Code (ECC) rates and a SARS volume identifier for the media involved in the error. HSARS is intended to find hardware that is gradually going bad over a period of time causing read/write errors. It will not point to hardware that suddenly fails. HSARS is based on various software filters that charge up or down based on the number and type of errors and ECC rates that occur in a given number of GB processed. A quality assessment of the drive hardware is generated based on the value of the filters. When HSARS determines that hardware needs to be replaced, a Service Information Message (SIM) is generated that points to the hardware needing to be replaced.

SARS impact on read-only recovery

The VTS has a number of conditions which turned a tape to read-only status (permanent write error, database does not fit on the tape, and so on). Now, there is another source of read-only status tapes: Those which had a high number of temporary errors, or other occurrence which causes SARS statistics to be high. Activating VSARS in a VTS environment will result in a temporary increase in the number of read-only volumes, but these read-only volumes will not prevent access to any data. The number of permanent errors will decrease as marginal cartridges are removed.

Drive swapping

Both z/OS DFSMS and BTLS support automatic internal swap inside the tape library for non-VTS drives. You can see DDR host messages for any non-VTS drives, in case of a media error, but not for VTS virtual IBM 3490E drives or VTS physical drives. Drive swap is an internal VTS process and does not involve the dynamic device reconfiguration (DDR) z/OS host processing.

Refer to the *Recovery and Reconfiguration Guide*, SA22-7623, for more information about z/OS DDR functions.

6.12.2 Hardware conditions

In this section we describe potential hardware failure scenarios.

IBM 3590 or 3592 tape drive failure

When the VTS determines that one of its tape drives is not operating correctly and requires service (likely that the drive has excessive read or write errors), the drive is marked off-line and an SSR must be called. The following Intervention Required message (see Example 6-9) is displayed on the Library Manager Console.

Example 6-9 3590 tape drive failure

```
CBR3750I MESSAGE FROM LIBRARY lib: Device xxx made unavailable by a VTS. (VTS z)
```

Operation of the VTS continues with a reduced number of drives until the repair action on the drive is complete. To recover, the SSR repairs the failed tape drive and makes it available for the VTS to use once again.

SSA adapter failure

The TVC is unavailable for use and therefore the VTS and virtual volumes stored on it are unavailable until the repair action is completed. To recover, the SSR repairs the failed SSA adapter. Once repair is completed, processing on the VTS can continue.

Note: In Model B10 VTS and Model B20 VTS, adapter card failure in VTS does not impact VTS availability. Service of a failed adapter can be deferred to non-peak VTS operating time.

Channel card failure

If an ESCON or SCSI channel adapter fails, it is marked unavailable and the VTS continues operation with one less channel adapter until the repair action is completed. To recover, the SSR repairs the failed channel adapter. Once repair is completed, processing on the VTS can continue.

SSA disk drive failure

The VTS uses a RAID fault-tolerant implementation stored on SSA disk drives as a tape volume cache. If a disk drive fails, the array automatically reconfigure itself, removing the failed disk drive from the array, adding the spare disk drive to the array and rebuilding the data onto the new disk drive. Reconfiguration of the array and the repair of the failed disk drive are concurrent. To recover, the SSR repairs the failed SSA disk drive.

VTS controller failure

If the controller fails, the VTS and virtual volumes stored on it are unavailable until the repair action is completed. To recover, the SSR repairs the VTS controller.

Power failure

User data is protected in the event of a power failure, as it is stored on the TVC. Any host jobs reading or writing to virtual tapes will fail as they would with a real IBM 3490E and will need to be restarted once the VTS is available again. When power is restored, the VTS recovers access to the TVC, using information available from the VTS database and logs. Once power is restored, restart the IBM Magstar 3494 Tape Library. The Library Manager restarts the VTS.

VTS database corruption

To recover from corruption of the VTS database you have to perform a so-called disaster recovery. This is a VTS process further described in 6.13, “Disaster recovery” on page 274.

The VTS controller maintains a database of information about the location and status of logical volumes on the stacked volumes it manages. When a stacked volume has been filled with logical volumes, a backup of the entire database is placed at the end of the filled stacked volume. The database contains a time and date stamp that identifies when the backup was performed.

When the database copy operation is complete, a message is sent to the attached hosts (Example 6-10).

Example 6-10 VTS database backup to physical tape

```
CBR3750I MESSAGE FROM LIBRARY lib: VTS Database Backup written to Physical Tape  xxxxxx.
```

The disaster recovery process causes the VTS controller to load the stacked volumes, locate the latest version of the database and restore it. Any logical volumes written after the last backup of the VTS database are lost. When the restoration is complete, a message is displayed on the Library Manager console informing you of the date and time when the VTS database was restored.

VTS software failure

If a problem develops with the VTS software, the VTS issues an Intervention Required message to the Library Manager and attempts to recover. In the worst case, this can involve a reboot of the VTS controller. If the problem persists, you need to contact your IBM service representative. The Intervention Required message (Example 6-11) is sent to the Library Manager.

Example 6-11 VTS software failure

```
CBR3750I MESSAGE FROM LIBRARY lib: Virtual Tape System z has a CHECK-1 (xxxx) failure
```

The VTS internal recovery procedures handle this situation and restart the VTS.

Library Manager failure

If the Library Manager fails in a 3494 that contains a VTS, the situation is the same as with any other 3494. Current tape tasks complete, but no new tasks can start until the Library Manager switchover is complete or the problem has been solved. To recover, fix the problem that caused the Library Manager to fail. If the HA1 Frames are not installed, the 3494 and the VTS are unavailable until the repair action is complete and the Library Manager is restored. If HA1 Frames are installed, the second Library Manager takes over. The library will be taken away from the hosts and will be unavailable until take over is complete.

As with the 3494, the 3953 can also be equipped with dual Library Managers. This will enable continued operation if the active Library Manager fails by switching over to the stand-by Library Manager. During the switchover current tape tasks complete on the VTS, but no new tasks can start until the switchover is complete.

For detailed scenarios on how to handle failure scenarios with HA1 Frames, refer to *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632.

Note: On the 3584/3953 Tape Library the HA1 and the second Library Manager are separate features. Ordering the HA1 feature for the 3584 does not give you a second Library Manager but does give you the second accessor.

Accessor failure and manual mode

If the HA1 Frames are not installed, failure of the accessor results in the library being unable to automatically mount physical volumes in response to host requests for data. When a host requests a logical volume for recall, the request is converted into a message to mount the stacked volume where it is located. Fast-Ready scratch mounts do not require a physical mount until the data is to be copied to a stacked tape. The message is displayed on the Library Manager Console for manual mode operation. For detailed information about running the 3494 in manual mode, see the *3494 Tape Library Operator's Guide*, GA32-0280. For details on running the 3584/3953 Tape Library system in manual mode see *IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0472.

If HA1 Frames are installed, the second accessor takes over, then you can call your IBM service representative to repair the failed accessor.

Gripper failure

If a gripper fails, it can be assumed to be an accessor failure if a dual gripper is not installed. If a gripper fails on a dual gripper accessor, the library operations will continue with one gripper. While the gripper is being repaired, the accessor is not available, so the library should be run in manual mode until the repair is complete. On the 3584, the accessor can only come with two grippers.

If the HA1 Frames are installed, the second accessor will be used until the gripper is repaired. For detailed information about running the IBM Magstar 3494 Tape Library in manual mode, see the *3494 Tape Library Operator's Guide*, GA32-0280. For details on running the 3584/3953 Tape Library system in manual mode see *IBM TotalStorage 3953 Library Manager L05 Operator Guide*, GA32-0472.

Out of stacked volumes

If the Tape Library runs out of stacked volumes, then copying to the 3590 or 3592 tape drives will fail and an Intervention Required message is sent to the Library Manager. All further logical mount requests are delayed by the Library Manager until more stacked volumes are added to the Tape Library. To recover, insert more stacked volumes. Copy processing can then continue.

Broken leader block

If a stacked volume has a broken leader block, it causes a load failure on the 3590 drive. The drive is made unavailable, the volume is marked as inaccessible and an Intervention Required message is generated indicating that a Load/Unload failure occurred. All subsequent accesses to logical volumes on the stacked volume fail because the stacked volume is unavailable.

The 3592 does not have a leader block. It has a metal pin that is grabbed by the feeding mechanism in the 3592 tape drive to load the tape on to the take-up spool inside the drive. If this pin gets dislodge or damaged, then follow the *IBM TotalStorage Enterprise Tape System 3592 Operators Guide* - GA32-0465.

> **Important:** Repairing 3590 or 3592 tape should only be done for data recovery. Once the data has been moved to a new volume, you should replace the repaired cartridge.

To recover, follow the procedure for load failures to remove the cartridge from the drive, repair the leader block and reinsert the stacked volume. Eject the volume as detailed in "Eject stacked volumes" on page 238, to force the VTS to copy the logical volumes from the stacked volume to a new stacked volume. Once the stacked volume has been copied and ejected successfully, destroy it.

Broken tape

If a 3590 or 3592 tape cartridge is physically damaged and unusable (the tape is crushed, or the media is physically broken, for example), the VTS cannot recover the contents. This is the same for any tape drive media cartridges today. You can generate a list of logical volumes that are on the stacked volume. See 6.8.3, "List logical volumes on a stacked volume" on page 257 for details. Please consult with your SSR, as there may be further avenues of recovery available from within IBM.

Logical mount failure

When a mount request is received for a logical volume that is not resident in the tape volume cache, the VTS subsystem must recall the volume from a physical tape. A logical mount failure could result from a number of problems. The most likely causes of a logical mount failure might be the accidental removal of a stacked volume, or perhaps a stacked volume that is stuck in a drive. To recover, solve the problem with the stacked cartridge, and make it accessible to the VTS functions. For a stuck volume, contact your IBM service representative.

If the subsystem cannot successfully recall the logical volume, it will still indicate to the host that the mount completed. When the host issues the first I/O operation to the virtual device, the subsystem will fail the command and indicate an equipment check. The host will then fail the job and generate console messages. For example, the following messages (Example 6-12) are generated as a result of a logical volume mount failure for volume **JYP001** on device **1FD5**.

Example 6-12 Mount request

```
IEC501A M 1FD5,JYP001,SL,NOCOMP,GENVTSP8,STEP4,VTS42.TEST74

IOS000I 1FD5,AD,EQC,**,0600,,**,,GENVTSP8613
 10488035000000206011(B310000102004113)0000(500007FF)CE3FD69000050500

IEC147I 613-54,IFG0194K,GENVTSP8,STEP4,SYSUT2,1FD5,OAM274 VTS42.TEST74
```

The IOS000I message reports that an equipment check was received for device 1FD5 and contains the sense data returned by the subsystem. When a logical volume recall fails, the volume is not available in the Tape Volume Cache. Sense byte 14 contains X'02' which indicates this condition.

The IEC147I message reports that the failure was detected as the tape volume was opened. The abend and return code (613-54) indicates that the volume serial number was unreadable, had no label, or the volume was uninitialized, however in a VTS it is also the indication that the recall failed. When a logical volume mount is failed in this way, the operator will need to view the Intervention Required conditions and based on the intervention reported, take action prior to resubmitting the job.

Note: For a logical volume where a dual copy has been created, the above error can still occur if the recall of the dual copied logical volume fails.

Orphaned logical volume

This occurs when the VTS database has a reference to a logical volume but no reference to its physical location. This could result from hardware or internal software errors. When it does occur, any data on the logical volume is lost. When this error occurs, contact your SSR.

Internal-external label mismatch

If a label mismatch occurs, the stacked volume is ejected to the convenience Input/Output station and the Intervention Required condition is posted at the Library Manager console (see Example 6-13).

Example 6-13 Label mismatch

```
CBR3750I MESSAGE FROM LIBRARY lib: A stacked volume has a label mismatch and has been
ejected to the Convenience Input/Output Station.
Internal: xxxxxx, External: yyyyyy
```

The host is notified that Intervention Required conditions exist. Investigate the reason for the mismatch. If possible, relabel the volume to use it again.

Failure during reclaim

If there is a failure during the reclamation process, the process is managed by the VTS controller microcode. No user action is needed; recovery is managed internally.

Excessive temporary errors on stacked volume

When a stacked volume is determined to have an excessive number of temporary data errors, to reduce the possibility of a permanent data error, the stacked volume is placed in read-only status.

6.12.3 FICON support enhancements

Previously, channel errors on tape devices resulted in the job abending, requiring the job to be resubmitted. This was largely due to the fact that the Tape Command sets do not contain any positioning information to aid in recovery. The host access methods will now keep track of the block ID of the device as each chain is sent to the device.

If a command chain does not complete successfully due to a channel error, the host will reposition the tape device back to the block before the chain was issued and continue with the job with the chain that failed. To the user, the job is not abended or interrupted.

The details of the recovery are as follows:

- The host keeps track of the block ID before the start of a command chain.
- As each command chain completes successfully, the block ID is incremented by the number of blocks in the chain.
- If the current chain incurs a channel error, the host initiates the following recovery:
  - A SENSE command is sent to the device to determine its current position.
  - A REWind command is issued to the device to bring it back to the beginning of tape.
  - A LOCate command is issued to position the device back to the block ID it was in before the chain began.
  - The current command chain is issued again to the device.
- If the command chain is successful, the job continues as though there was no error.

Note: The channel recovery was added to the host software as part of the FICON support. However, you do not require FICON in order to obtain the benefits.

6.13 Disaster recovery

In the event that a VTS is not usable, because of interruption of utility or communication services to the site or significant physical damage to the site or the VTS itself, access to the data managed by the VTS is restored through automated processes designed into the product. The recovery process assumes that the only elements available for recovery are the stacked volumes themselves. It further assumes that only a subset of the volumes are undamaged after the event. If the physical cartridges have been destroyed or irreparably damaged, recovery is not possible, as with other cartridge types. It is important that you integrate the VTS recovery procedure into your current disaster recovery procedures.

The disaster recovery process is a joint exercise that requires your involvement as well as that of your SSR.

Note: Data created on a VTS with Extended High Performance Option (EHPO) feature cannot be successfully recovered on a Model B18 VTS without EHPO feature. Data produced on a VTS with LIC 2.26.xx.x and above cannot be recovered to a back-level VTS.

6.13.1 VTS database

The microcode in the VTS maintains a database indicating where the current accessible version of a logical volume is located on the stacked volume. This database is copied to 3590 or 3592 tape cartridge automatically as a backup when the stacked volume cartridge is filled. This backup of the database contains a complete logical to physical volume data map for the stacked volume with header marking times and date stamps. Figure 6-39 depicts the VTS database and its backup location.

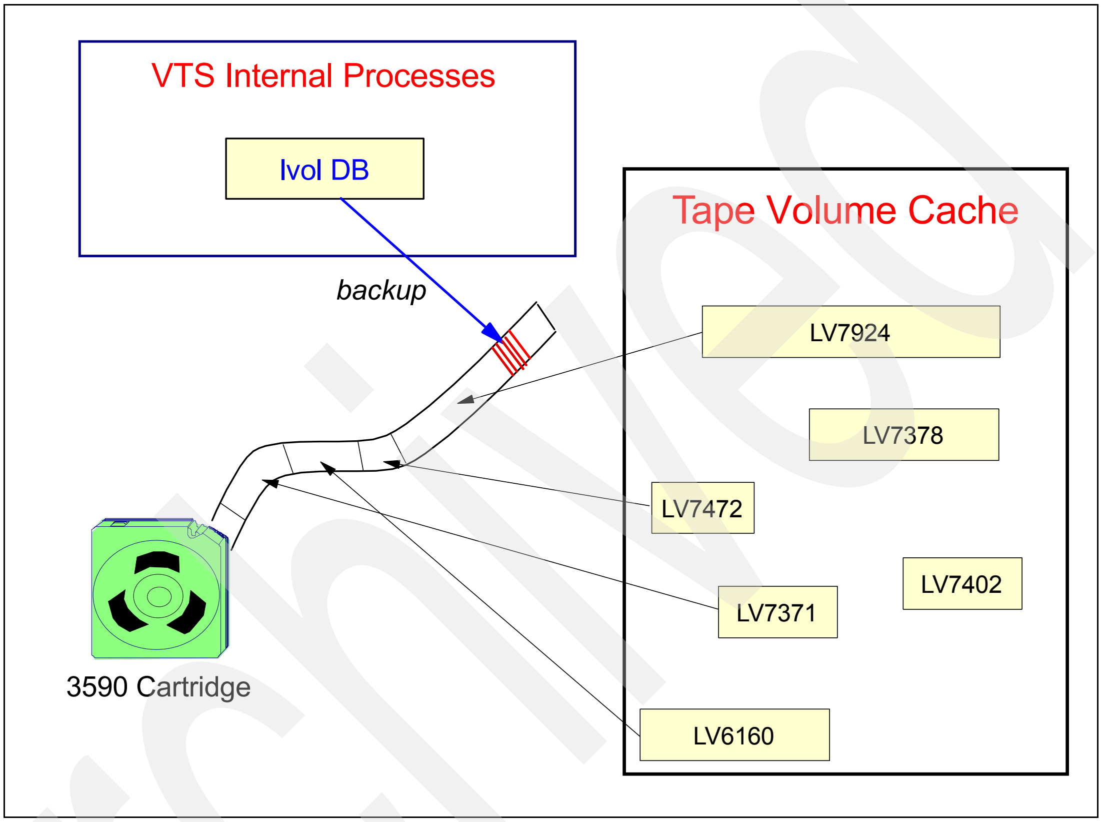

Figure 6-39 Logical volume database and its backup in 3590 tape cartridge

The backup is placed at the end of the stacked volume (physically at the beginning of tape) before cartridge is unloaded. If this database became unavailable, any logical volumes written after the last backup will not be recoverable. Disaster recovery procedures relies on a restored database from the latest backup on the cartridge. In the figure, virtual volumes LV7371, LV6160, LV7472 and LV7924 are copied to a stacked volume but LV7402 and LV7378 are not copied yet. All logical volumes copied to tape are updated in the VTS database but not backed up to tape yet, none of the six volumes in figure would be recovered if the VTS database became inaccessible.

In the sections that follow, we document the process for recovering the contents of a VTS when any situation occurs preventing access to the data in VTS.

6.13.2 Actions to take at the damaged site

When the site or VTS itself is damaged so that it is no longer available, take the following actions to resume processing with the VTS:

- If possible, have the Fast Migration procedure performed at the damaged site:

  If the automated library (3494 or 3584) and VTS are still functional, ask the SSR to perform a Fast Migration procedure. This procedure ensures that all data has been moved from the TVC to a stacked volume, and that two complete database backups are written two backup tapes.

- Remove the stacked volumes from the 3494 or 3584:

  Remove the undamaged stacked volumes from the library. Ensure that all stacked volumes are removed. If you leave (or lose) any physical cartridges, the logical volumes they contain are no longer available for processing and cannot be added at a later date.

- Take the removed physical cartridges to the recovery site.

6.13.3 Actions to take at the recovery site

Take the following actions at the recovery site:

- Restore the tape management system catalog:

  You must restore the tape management system catalog before invoking the disaster recovery operation at the Library Manager. Restoring the catalog allows correct synchronization between the tape management system and the Library Manager. The tape management system determines whether a logical volume is scratch or private. If you do not restore the tape management system catalog before selecting the disaster recovery operation, logical volumes will get the Default Use Attribute assigned, which is defined during ISMF library definition. For the VTS, the use attribute is usually defined to be scratch. Therefore, without the tape management catalog restored, all logical volumes will be in scratch status.

- Set up the DFSMS tape catalog:

  The 3494 Tape Library or 3953/3584 Tape Library with the VTS must be defined before starting disaster recovery at the library. You must either define a new TCDB or, if you have restored the one from the disaster site, using IDCAMS IMPORT RECONNECT, you must alter the library definition to replace the library ID with the ID that identifies the recovery library. If you do not replace the ID, DFSMS will not accept the logical volumes because it thinks they are already in another library. You can use the ISMF Tape Library Application Selection panel, specifying the **Alter** option.

- Have the SSR begin the stacked volume recovery process:

  As indicated below, some of the steps in this process are dependent on whether or not it was possible to perform the Fast Migration procedure at the damaged site. Here is an overview of the steps in the process:

  – You will need to provide the SSR with the volser ranges of the physical volumes that you are going to put into the recovery VTS.

  – When the SSR tells you to, place all of the physical stacked volumes into the library.

    As the number of volumes may be large, open the library frame doors and manually place the volumes into any library cell, except for the high capacity input/output cells (if configured). There are no specific cells for the location of the stacked volumes. Once the volumes are in the library, the SSR can continue with the recovery process.

– If a Fast Migration procedure was performed at the damaged site, then the Fast Restore procedure must be used to restore the current database that was backed up on 2 tapes at the damaged site.
– Otherwise, if the Fast Migration procedure was not performed at the damaged site, the disaster recovery process must be used.

The disaster recovery process is an option on the Library Manager service panel. The process scans each of the stacked volumes in the subsystem to find the volume with the most recently written database; this database is used to restore the subsystem. If the most recent volume cannot be used to recover, the next most recent is used.

- Reinsert the logical volumes:

  When the SSR has indicated that the disaster recovery process is complete, there is still one more action that you must take: reinsert the logical volume serial number ranges into the Library Manager, with the library online to the attached hosts. As the Library Manager processes the volumes you are adding, it notifies the attached hosts and the Library Manager inventory, the TCDB, and the tape management system are synchronized.

Once the hosts and the library have completed synchronization, the logical volumes managed by the VTS are available for host access.

6.13.4 Disaster recovery processing time

The time taken to perform the disaster recovery process depends heavily on the number of stacked volumes, drives and logical volumes you have. Several steps are involved in performing disaster recovery. The steps and their estimated times are described below in Table 6-2. Times for steps 1 through 3 are sequential. Times for steps 4 and 5 can overlap.

Table 6-2 Disaster recovery processing time in 3494 Tape Library

Recovery steps	Processing time
1. Reinventory Library	5 min per library frame[a]
2. Find most recent database on physical volumes[b]	90 sec per cartridge[c]
3. Reload database and shutdown/restart the VTS	20 min
4. Add logical volumes to the Library Manager inventory	4 min per 1000 logical volumes
5. Upload the Library Manager inventory and re-synchronize with the tape management system and operating system	1 sec per logical volume

a. Time is for one accessor to inventory one frame. If the 3494 had the HA1 feature installed, then each accessor could participate in the inventory at the same time.
b. This step only applies when the Fast Migration and Fast Restore Procedures could not be used. In those procedures the stacked volume with the most recent database is already known.
c. All 3590 VTS drives are used for this operation. Total time is (number of cartridges * 90 seconds) / number of drives. Time will be less if 3592 taped drives are installed in the 3494.

Note: The above noted times and the following example are for a 3494 Tape Library. If a 3584/3953 Tape Library is involved the times will be different. Where the 3494 takes about 5 min to inventory a frame, the 3584 takes only 60 seconds. Tape mount times for the 3584 accessor and specific 3592 tape drive functions (load, rewind, etc.) are also faster.

With the following configuration, we can estimate the disaster recovery time:

Three-frame library (L1x Frame, D12 Frame, Model B18 VTS)
Eight hundred 3590 physical volumes
Six 3590 drive configuration
25,000 logical volumes

This is the time required for the above steps:

Reinventory Library: 3 * 5 min = 15 min
Find most recent database on physical volumes: (800 * 90 sec)/6 = 200 min
Reload database and shut down/restart the VTS controller: 20 min
Add logical volumes to the Library Manager inventory: 25 * 4 min = 100 min
Upload the Library Manager inventory and re-synchronize with the tape management system and operating system: 25,000 * 1 sec = 417 min

Total disaster recovery time:

Steps 1-3 = 235 min
Steps 4-5 = 417 min (step 5 overlaps 4)
Approximately 11 hours

6.13.5 Disaster recovery implications

At the completion of the recovery process, the VTS and the Library Manager contain database and inventory records as well as status information for the logical volumes as found in the most recent database backup on the undamaged stacked volumes. Depending on the following conditions, some customer data and logical volumes may not have been recovered:

- The stacked volumes containing the latest database backups may have been destroyed.

  Any changes to the location or status of logical volumes since the last found database backup are lost. Some logical volume records may be lost. It is possible, however, that the database will contain the location of the previous use of logical volumes and that data is accessible to the host even though it is not the latest one.

- One or more stacked volumes found in the restored database may have been destroyed during the event.

  Although the VTS database and the Library Manager inventory may have a record of the logical volumes that resided on the missing stacked volumes, the data itself is lost.

- A virtual volume may not have been closed at the time of the event.

  The VTS database has no record of the virtual volume's location on a stacked volume and the data is lost.

- Logical volumes may have been written to a stacked volume between the time the last database backup was made and the disaster.

  The VTS database has no record of the new logical volume locations on the stacked volume and the data is lost. It is possible, however, that the database may contain the location of the previous use of the logical volume and that data is accessible to the host.

7

VTS performance and monitoring

In this chapter we describe the underlying elements that determine and influence the performance of the VTS. We also describe what actions to take, when necessary, to improve VTS performance.

This chapter includes:

- An overview of the shared tasks which are running in the VTS controller
- A description of a VTS monitoring and performance evaluation methodology
- A walk-through of a VTS capacity planning case study
- Changes to VTSSTATS reporting with the introduction of Advanced Policy Management

We discuss the VTS shared resources so that you can understand the impact that contention for these resources has on the performance of the VTS.

The monitoring section can help you to understand the important values recorded in the SMF94 type record. By using these records, you will be able to perform simple health checks as well as in-depth investigations on performance issues that may arise with the VTS. It can also help you recognize the symptoms that indicate the VTS is at or near its maximum capacity. The information provided can help you to evaluate the options available to improve the throughput or performance of the VTS.

The capacity planning case study illustrates some guidelines and techniques for the management of virtual and stacked volumes associated with the VTS.

7.1 Introduction

The virtual tape subsystem requires that we take a new view of tape performance, compared with older non-virtual tape systems. This is due to the number of new elements brought in to the overall solution, including Tape Volume Cache (TVC), a VTS controller running AIX (and storage management software), virtual tape drives, virtual and logical volumes, and stacked volumes. This architecture can provide significant benefits to the tape processing environment, but a new approach must be taken in order to effectively size, monitor, and manage the performance of the VTS.

7.2 VTS model performance attributes

Since its introduction, significant improvements have been made in VTS capacity and performance. At the time of publication, there are basically four configurations of the VTS. They are: Model B10, B16, B18 or B20 VTS. The first Virtual Tape Server released by IBM was the Model B16 VTS; this was followed by the Model B18 VTS, the Model B10 VTS, and the Model B20 VTS.

Models B16 and B18 have been removed from marketing and cannot be purchased as a new configuration.

7.2.1 Model B10 VTS and Model B20 VTS

The Model B10 VTS and Model B20 VTS feature an advanced RISC controller and include the capability of the existing VTS with advanced function, peer-to-peer copy, and extensive server attachment. EHPO and the Performance Accelerator feature are now provided as standard functions that ensure superior performance over all previous VTS models:

- The Model B20 VTS provides:
  - 256 virtual tape addresses
  - Up to eight FICON channels
  - Up to 16 ESCON channels
  - Up to eight SCSI bus connections.
- The Model B10 VTS provides:
  - 64 virtual tape addresses
  - Up to four FICON channels
  - Up to four ESCON channels
  - Up to eight SCSI bus connections.
- Use of eight Virtual Tape Controllers in a PtP VTS configuration to address up to 256 virtual tape drives

These changes combine to provide significant throughput increases.

In this section, we address primarily the performance attributes for the VTS. We provide guidelines for monitoring your system. Later in this chapter, we introduce the IBM TotalStorage Enterprise Tape Library (ETL) Specialist, the Bulk Volume Retrieval System (BVIR), and TotalStorage ETL Expert; these tools are used for displaying and monitoring the 3494, the VTS, and its volumes.

Note: More actual and detailed performance measurements and results are published in the Performance White Papers, which you may access via the IBM TechDocs Home Page:

`http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs`

Then search for *VTS performance.*

7.2.2 VTS Performance Accelerator

The Performance Accelerator is a combination of hardware and software improvements featuring a new file system, additional processors for the VTS controller and Licensed Internal Code VTS which has been available since September 1999. For the Model B18 VTS, this functionality was provided as a feature upgrade. With the Performance Accelerator, the performance characteristics of Model B18 VTS are dramatically improved. The Model B10 VTS and Model B20 VTS have this function as standard.

Performance Accelerator introduction

For the Performance Accelerator implementation, channel enhancement with data compression was required as a prerequisite. For the Model B18 VTS, data compression which was provided by Extended High Performance Option (FC3400), a feature which had to be ordered separately. For the Model B10 VTS and Model B20 VTS, data compression function is provided as a standard feature.

Performance Accelerator feature

The Performance Accelerator feature is a combination of hardware and software improvements, which provide in conjunction with the new Extended Performance ESCON Channels a balanced subsystem. The essential elements of the Performance Accelerator feature are:

- A new file system (GPFS)
- Additional processing hardware and storage
- Improved Licensed Internal Code

One principal element of the Performance Accelerator feature is the increase of processing power from two to four processors. Another is the buffering of the incoming data to the processor, either from an host or tape, in logically contiguous blocks of 256 KB, that are written in parallel across the disks of the tape volume cache (Full Stride Write). These enhancements are accomplished by the new file system GPFS (General Parallel File System) and result in an increase of the MB/s throughput of the TVC.

The Performance Accelerator feature and the Extended Performance ESCON Channels for VTS mode B18 provided additional performance options for customers to either allow additional storage access during a period of peak activity or reduce the duration of peak activity for jobs requiring higher transfer rates. Both are standard for the VTS model B10 and B20.

A comparison of bandwidth numbers of various VTS models, based on initial modeling and measurements and assuming a 3:1 compression ratio, is shown in Figure 7-1 and described in more detail in the actual Performance White Paper.

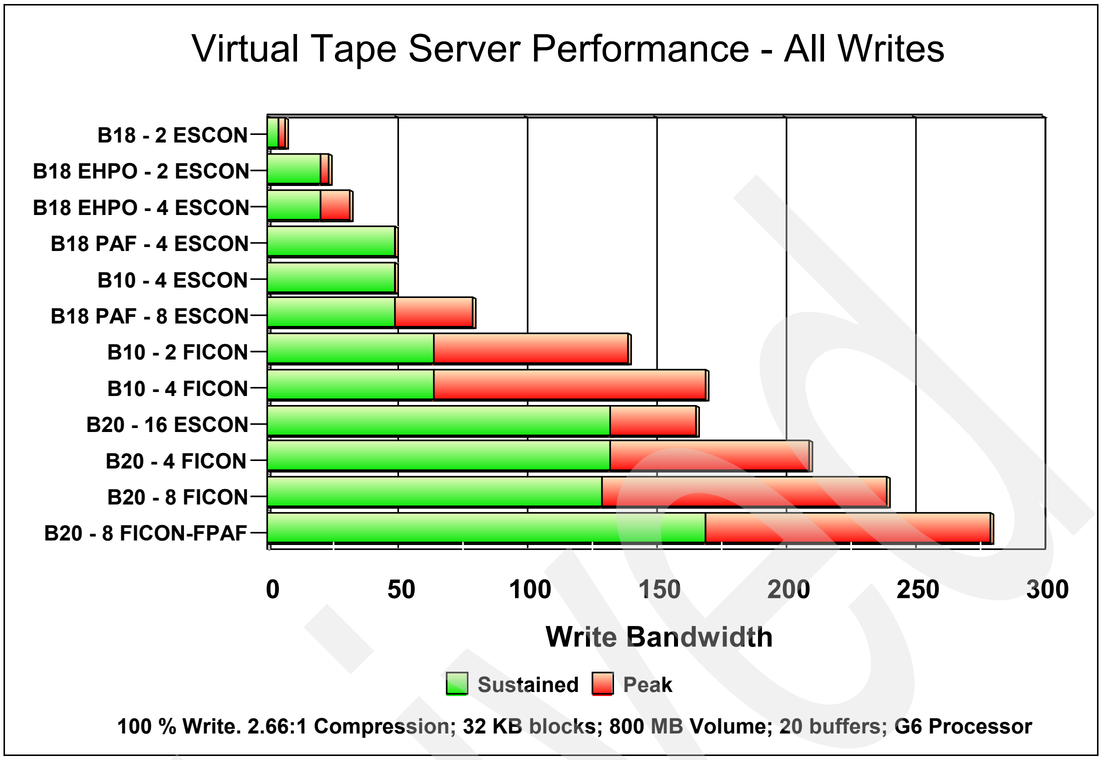

Figure 7-1 Performance comparison

In general, performance depends on such factors as total system configuration, VTS tape volume cache capacity, the number of tape drives available to the VTS, the number of channels, read/write ratio as well as data characteristics like blocksize and mount pattern.

All these performance factors and numbers for different configurations are discussed in 7.4, "Monitoring and evaluating VTS performance" on page 290.

General Parallel File System overview

IBM's **General Parallel File System (GPFS**) provides global access to data from any of the hosts within a cluster or within a pSeries processor. It was initially released on the RS/6000 SP in 1998 using a software simulation of a Storage Area Network (SAN) called the IBM Virtual Shared Disk (VSD). The product has been running on cluster configurations of up to 512 nodes supporting demanding I/O loads and fault tolerance requirements since its introduction.

GPFS supports very large file systems and stripes data across multiple disks for higher performance. It offers many of the standard AIX file system interfaces by allowing most applications to execute without modification or recompilation. These capabilities are available while allowing high speed access to the same data from all nodes of the SP system.

The parallel nature of GPFS guarantees that the entire file system is available to all nodes within a defined scope and that file system services can be safely applied to the same file system on multiple nodes simultaneously. It also means that multiple records can be safely written to the same file simultaneously by being striped across several disks. The GPFS improves the performance of both parallel and sequential programs. Sequential programs are automatically parallelized by GPFS. For more information about the GPFS, refer to *GPFS: A Parallel File System*, SG24-5165, and *GPFS on AIX Clusters; High Performance File System Administration Simplified*, SG24-6035.

Improved Licensed Internal Code

With the Performance Accelerator feature (FC5236 for Model B18 VTS, standard function for Model B10 VTS and Model B20 VTS) installed, the management of tape volume cache content is modified to take full advantage of all available processing resources. This has been achieved by the improved efficiency of the High Performance Streaming File System and the configuration of the cache. The primary changes of the VTS Licensed Internal Code apply to the throttling function, which controls the VTS throughput in a different way when the Performance Accelerator feature is installed.

A further explanation of throttling is contained within 7.4, “Monitoring and evaluating VTS performance” on page 290.

7.3 Shared resources

In the process of writing scratch volumes or copying and recalling virtual volumes on physical stacked volumes, components are shared by tasks running on the VTS controller. Some of these tasks represent customer work (such as scratch mounts) and other tasks are associated with the internal operations of the VTS (such as reclamation). All these tasks must share the same resources, especially the VTS processor, the TVC and physical tape drives. Contention may occur for these resources when heavy demands are placed on the VTS subsystem. To manage the use of shared resources, the VTS uses various resource management algorithms, which can have a significant impact on the level of performance achieved for a specific workload.

In this section we discuss the effects on performance of the following shared resources:

- RISC/6000 processor cycles
- Tape Volume Cache (TVC) management
- Physical drives
- Physical stacked volumes

7.3.1 Processor cycles

All tasks running in the VTS controller require a share of processor cycles. These tasks include the emulation of virtual drives, copy tasks, and recall tasks. For example, if there were eight active virtual drives, two copy tasks, and one recall task active at one point, then there would be a total of eleven concurrent tasks. The processor cycles would be shared by these tasks through a time-slicing multiprocessing algorithm.

Figure 7-2 shows examples of functions that use processor cycles.

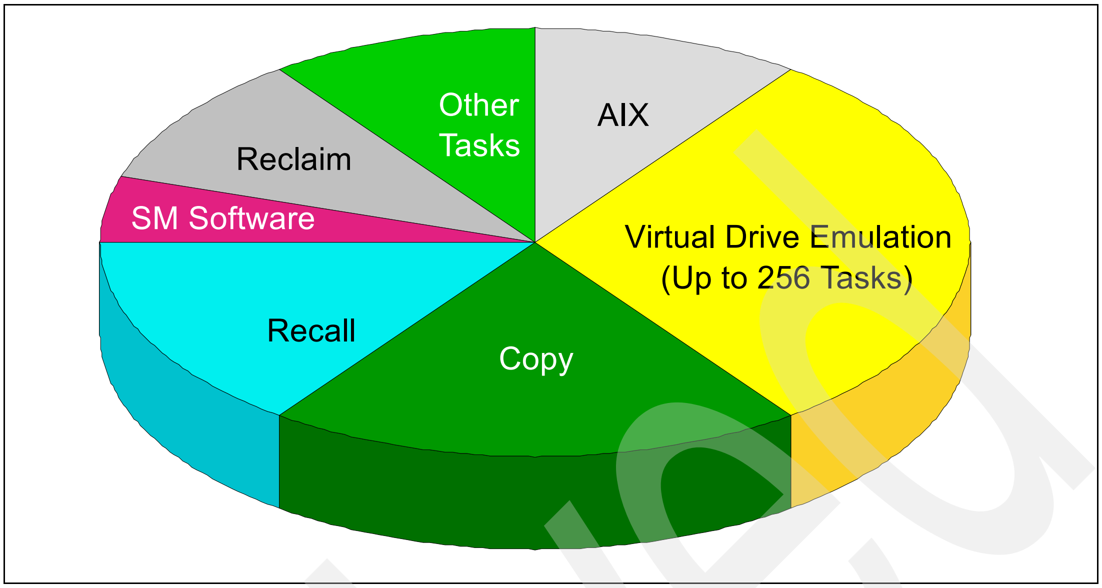

Figure 7-2 Processor cycles

The VTS monitors the utilization of the processor and uses processor utilization or idle time in many of its internal task management algorithms. High processor utilization can adversely effect channel throughput and logical volume mount times.

The VTS does not report processor utilization as part of its performance statistics, therefore you cannot monitor this metric.

7.3.2 Tape volume cache management

All virtual volumes are written and read by the host into and out of the TVC. The process of copying virtual volumes to stacked volumes and the recall of the virtual volumes from stacked volumes is transparent to the host.

The primary objectives of TVC management are to ensure that sufficient freespace is available for new or old virtual volumes and to maximize the number of read cache hits. The primary mechanisms which are used by the VTS to manage the cache are the internal allocation of physical drives, logical volume fragmenting, and throttling.

> **Note:** The software management algorithms used by the VTS for TVC management, including throttling values, are internal processes that cannot be controlled by the customer.

Also see “Cache Management exploitation” on page 97.

Figure 7-3 shows the TVC processes.

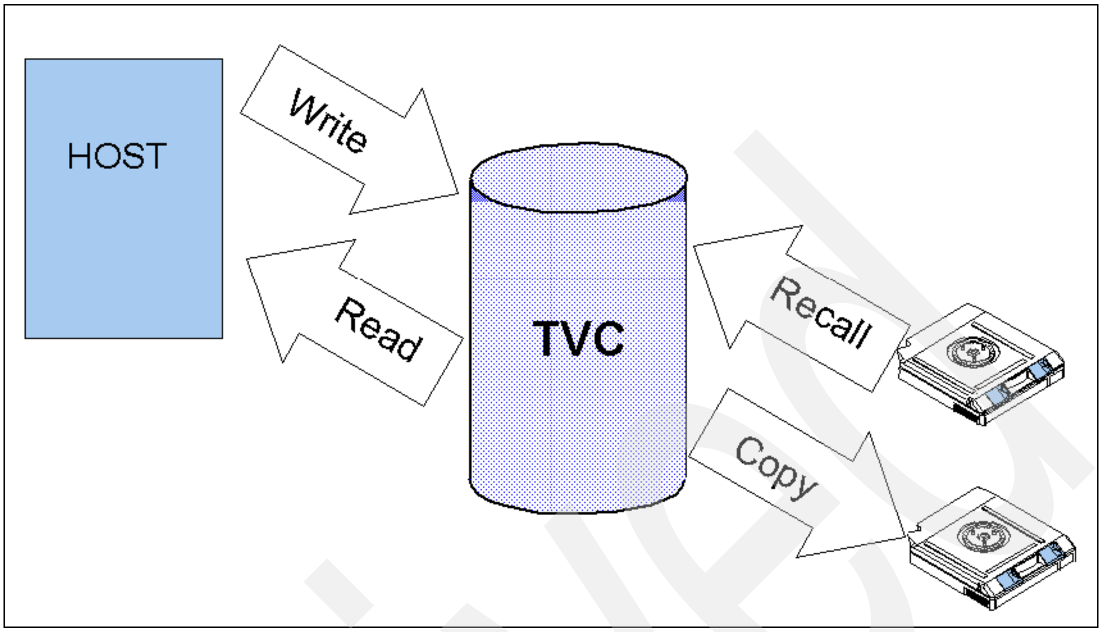

Figure 7-3 Tape volume cache processes

Physical drive allocation

Write operations for new volumes, appending to old volumes and read operations requiring the recall of a volume into TVC, all require freespace in the TVC for storage of data. When freespace becomes too low, or the space occupied by volumes which are closed and ready to be copied to stacked volumes (copy queue) becomes too high, the VTS will, if possible, increase the number of copy tasks allowed. In addition, if freespace becomes too low, the VTS reduces the number of recall tasks allowed, therefore making more physical drives available for copying. The VTS will always reserve at least one drive for the copy function and one drive for recalls.

Logical volume fragmenting

Copied volumes are eligible for fragmenting, which is the process of reducing the data portion of the virtual volume in cache down to a 'stub'. A stub, 4 KB in size, contains the information recorded on the beginning of the volume (such as the header). Fragmenting generally works on a first in, first out basis.

The process of fragmenting reduces used space within the TVC and therefore “creates” the space required for new volumes to be written and old volumes to be recalled. The VTS will fragment logical volumes in the cache only when cache freespace becomes low. Before a logical volume can be fragmented, it must have been copied to a stacked volume.

Figure 7-4 shows a typical example of the TVC contents.

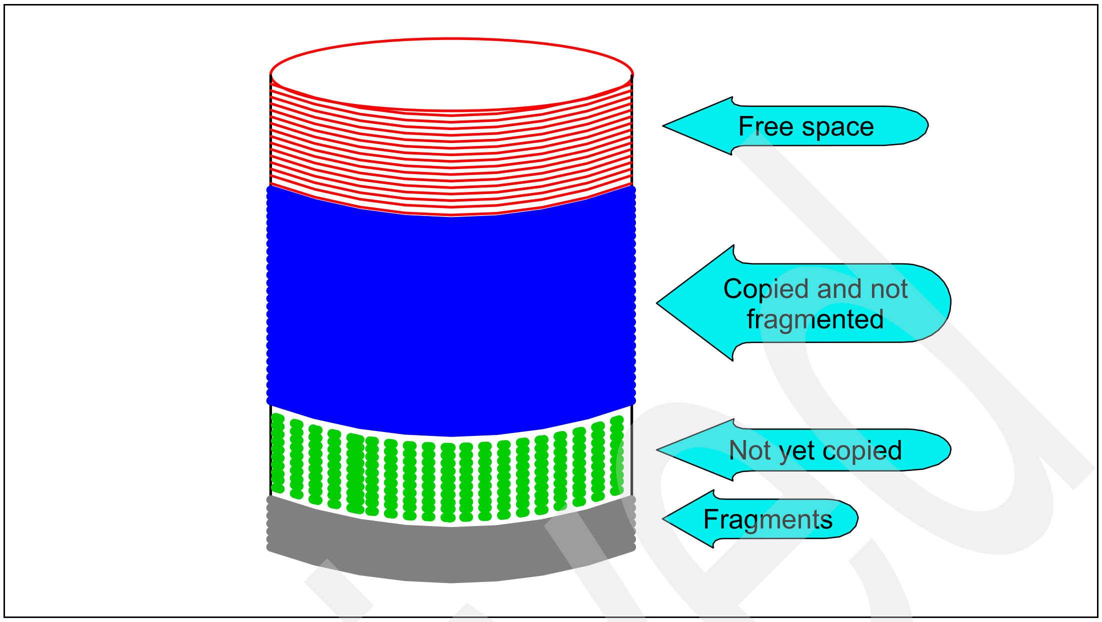

Figure 7-4 Tape volume cache contents

Host write throttling

In order to make sure that there are sufficient processing cycles available to perform high priority tasks, the VTS may introduce a delay in response to host write operations. This is called throttling. High levels of throttling can cause a reduction in host write activity, which may result in elongated or erratic job run times.

The VTS calculates throttling values for each of the following at regular intervals:

- Amount of TVC freespace
- Amount of TVC non-premigrated data
- Number of TVC logical volumes
- Number of active recalls (for early VTSs only)
- Size of the PtP copy queue

The actual amount of throttling that the VTS applies at any given time depends on the above factors and may vary depending on the VTS model and/or LIC level. The *recall* throttles no longer apply if your VTS is equipped with the Performance Accelerator Feature (standard on Model B10 VTS and Model B20 VTS).

Freespace throttling: Once the low freespace threshold is crossed, the amount of freespace throttling is inversely proportional to the amount of TVC freespace remaining. As long as the amount of freespace is above the threshold, freespace throttling will not occur.

Figure 7-5 shows VTS actions when cache freespace is low.

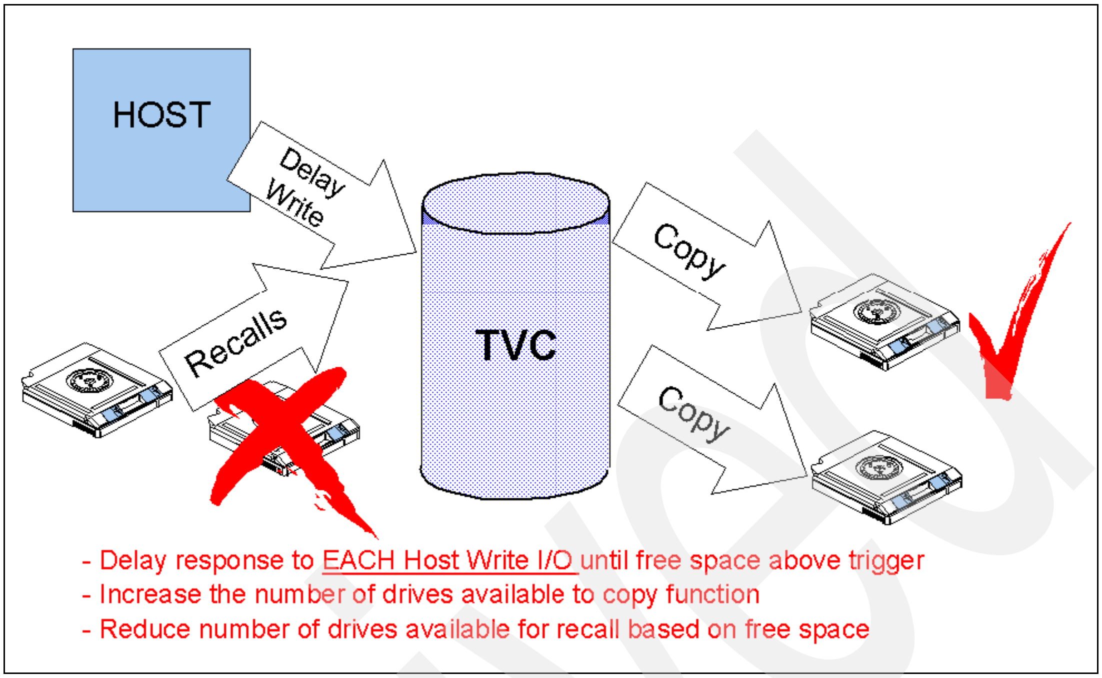

Figure 7-5 Cache freespace low actions

Copy queue throttling: The VTS introduces host write throttling if either the number of copy tasks or the amount of data waiting to be copied exceeds certain thresholds.

Figure 7-6 shows actions when the copy queue value is high.

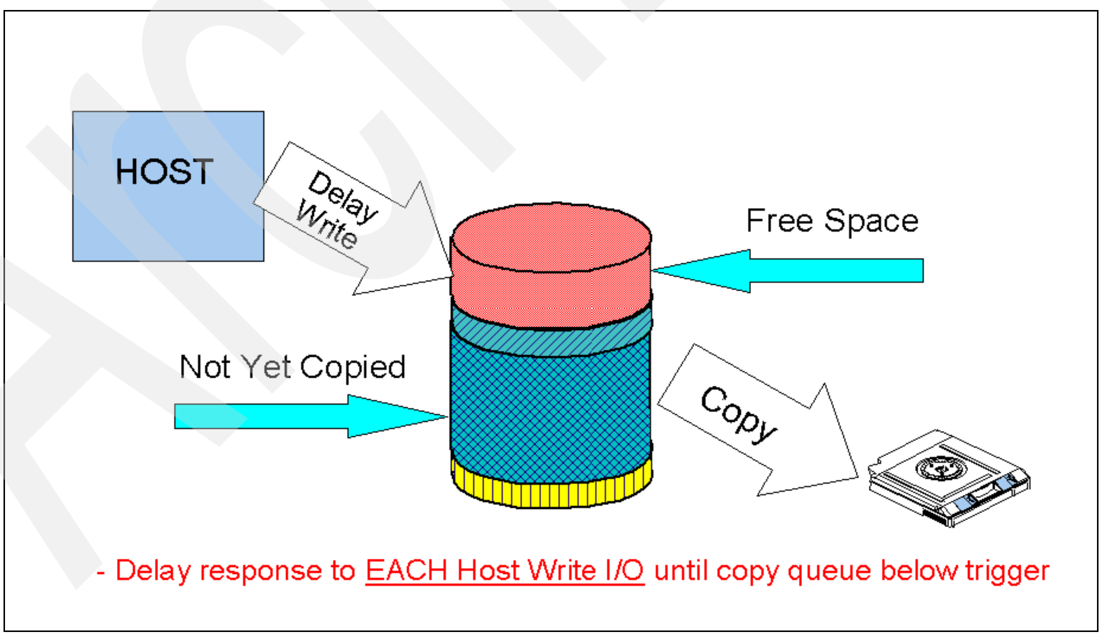

Figure 7-6 Copy queue high actions

Pre-migrate throttling: This type of write throttling will occur when the pre-migrate threshold is reached. This threshold again varies depending on configuration and activation of 2000 or 4000 MB logical volume support. If 2000 or 4000 MB logical volume support is not activated, the threshold is roughly between 60% and 80% of the total TVC space. If 2000 or 4000 MB logical volume support is activated, then the threshold is reduced to roughly between 40% to 50% of the total TVC space. Hitting this threshold means that there are a lot of volumes in the TVC that have not been destaged to physical tape, and the 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.

7.3.3 Physical drive availability

The physical drives are managed by the VTS internal management software and cannot be accessed from any attached host. These drives are used by the VTS for the mounts required for copying virtual volumes to stacked volumes, recalling virtual volumes into the cache and reclaiming stacked volume space.

The availability of VTS physical drives for certain functions can significantly effect VTS performance. The VTS will manage the internal allocation of these drives as required for various functions, but it will always reserve at least one physical drive for recall and one drive for copy.

Tape Volume Cache Management algorithms will also influence the allocation of physical drives. For example:

Cache freespace low: The VTS will increase the number of drives available to the copy function and reduce the number of drives available for recalls.

Copy queue high: The VTS will reduce the number of drives available for recall down to a minimum of one drive in order to make drives available for the copy function.

The number of drives available for recall or copy will also be reduced during reclamation or Import/Export processing.

If the number of drives available for copy is restricted, this can lead to limiting the number of virtual volumes in the cache which are eligible to be migrated. This can lead to freespace or copy queue size throttling being applied.

If the number of drives for recall is restricted, this can lead to elongated virtual mount times for logical volumes being recalled.

Recall performance is highly dependent on both the placement of the recalled logical volumes on stacked volumes and the order in which the logical volumes are recalled. To minimize the effects of volume pooling on sustained write performance, volumes are pre-migrated using a different distribution algorithm (when the LIC levels reach 527 LM and 2.26 VTS).

This algorithm chains several volumes together on the same stacked volume for the same pool. This process can be seen in “Changes for copying data to tape” on page 98. This can change recall performance, sometimes making it better, possibly sometimes making it worse. Other than differences in performance due to differences in distribution over the stacked volumes, recall performance should be the unchanged for LIC levels 527 LM and 2.26 VTS.

7.3.4 IBM 3592 and IBM 3590 physical tape drives

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

VTS B10

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.

VTS B20 base

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 benefit for other workloads such as single-volume recalls, reclaim, and import/export.

VTS B20 with FPAF

The FICON Performance Accelerator Feature - FPAF (feature code 5250) was introduced with the availability of FICON for 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. Furthermore, 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`

7.3.5 Physical stacked volumes

There are two instances where the physical stacked volumes can impact VTS performance:

- When multiple concurrent recalls occur for logical volumes that reside on the same stacked volume, these recalls will be processed serially. Although the VTS will not demount the stacked volume until all of the recalls currently queued for that volume are processed, the second and subsequent recalls from the same stacked volume will experience elongated mount times.

- When the minimum amount of scratch stacked volumes falls below a certain level:
  - For a VTS whose LIC is prior to LM 527 and VTS 2.26, this value is 10. At this point, the VTS will automatically raise the priority of the reclamation process and will assign two physical drives and reclaim until it reaches 15 available scratch stacked volumes. If this happens, the total number of drives available for copy or recall activity is reduced.
  - For a VTS whose LIC is at LM 527 and VTS 2.26 or higher, this figure is set at less than two empty physical volumes. At this point, the VTS will automatically raise the priority of the reclamation process until it has two or more empty physical volumes available.

7.4 Monitoring and evaluating VTS performance

With the same link described in 3.5, "Tape analysis and sizing the VTS" on page 72, 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. You also have access 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 flatfiles called DAILY and HOURLY may be used as input to a graphic tool that generates multiple charts.

To have SMF94 records written, you must have the statistics enabled on the Library Manager console. Select **Menu** —> **Service** —> **Utilities** —> **Enable/Disable Functions** and finally **Statistics Messaging**.

7.4.1 Tools download

All the VTS monitoring/evaluating tools that are discussed in this chapter 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 partners tools can be located at the following Internet site:

  http://testcase-yellow.boulder.ibm.com

- Customer tools can be located at the following Internet site:

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

VTSSTATS

Every hour the 3494 Library Manager sends VTS statistics records to all attached z/OS and z/OS hosts. These statistics are recorded as SMF type 94 records. See Appendix F, "SMF type 94 record layout" on page 461 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.

VTSSTATS and all other mainframe based tape analysis tools are in the file IBMTOOLS.EXE. The IBMTOOLS.TXT file gives you complete download and installation instructions. The spreadsheet programs are available in either Lotus123 format (VTSGR123.EXE) or in Excel® format (VTSGRCXCL.EXE). Each version has its corresponding TXT file with installation instructions.

VTSSTATS can be used to monitor VTS Drive and Tape Volume Cache (TVC) activity and for doing trend analysis to see where the bottlenecks are and whether an upgrade, such as, for example, adding additional physical tape drives or additional TVC may improve the overall performance of the VTS. VTSSTATS is not a projection tool, but it provides the basis for an overall health check of the VTS and performs the following tasks:

- Analyzes and reports on SMF 94 data.
- Summarizes VTS activity on an hourly and daily basis:
  - 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.
- 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
- Reports Tape Volume Cache Activity:
  - MBs read and written to TVC
  - MBs read and written to physical tape
  - Virtual volume size
  - Logical volumes managed

7.4.2 Critical performance values

To be able to make the best use of the information provided in the reports and files, you will need to 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 VTS(s), you should note the following information regarding your configuration and setup:

- Model (Standard, EHPO or PAF)
- Cache size
- Number of physical drivesD12 Frame
- Number of virtual volumes defined to the Library Manager
- Number of stacked volumes physically available in the ATL
- Number and type of channels attached to the VTS

You can find out the number of stacked and logical volumes in use by the VTS by interrogating the Library Manager console and searching for volumes in your specific category or range. See *IBM TotalStorage 3494 Tape Library Operator's Guide,* GA32-0280 for details on the VTS Library Manager console.

Use the following list of questions 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.

   VTSSTATS daily reports summarize the various SMF type 94 fields by day.

   The field in question is called SMF94VRA (Average time in seconds used to complete a mount request on a virtual drive in the last hour).

   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?**

   Figure 7-7 shows the VTSSTATS Hourly Virtual Drive Activity Report.

HOURLY VIRTUAL DRIVE ACTIVITY (CHANGE)

```
1HOURVIRT PTPVTS=EXAMP1 - HOURLY VIRTUAL DRIVE 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

      VIRTUAL  VIRTUAL DRIVES    VIRTUAL
      DEVICES  CONCUR MOUNTED.    MOUNT     TOT  -----FAST READY-----  ------MOUNTS--------  ----MOUNTS NOT------  %MNTS  --VVOL RESIDENCY
      CONFIG   maybe allocated* TIME (SEC)  VMTS       MOUNTS           SATISFIED IN CACHE   SATISFIED IN CACHE   NOT    ------(SEC)-----
 HOUR           MAX  MIN  AVG    MAX   AVG          #   MAX  MIN  AVG    #   MAX  MIN  AVG    #   MAX  MIN  AVG  SAT     MAX   MIN   AVG
smf94==>(vdc)  (vdx)(vdn)(vda)  (vrx)            (vfr)  ---SECONDS---  (vmh) ---SECONDS---  (vms)  ---SECONDS---        (vvx)  (vvn)  (vva)
 ptpnote:        *5   *5   *5     *5    *5          *5       *5          *5        *5          *5       *5        *5      *5    *5    *5

 00       256   165   17   86    601    16   997   715   54    2   11   269   77    2   15    13  601  214  354    1   6724     1   273
 01       256   161   26   72    724    22  1036   629   34    2    9   376  105    2   15    31  724  188  396    2   3767     5   190
 02       256    53    9   26    576    20   392   262   33    2    5   116   98    2   10    14  576  151  389    3   5804     1   256
 03       256    75    8   34    595     8   644   522   46    2    6   119   99    1    9     3  595  341  406    0   8692     4   154
 04       256    65   11   31    660    31   468   406  104    2    6    34   77    2   10    28  660  267  432    5   1690     3   240
 05       256    60    9   36    544    11   377   290   45    2    5    82  147    2   17     5  544  243  319    1   3357     1   317
 06       256    58   11   30    392    12   394   317   38    2    5    69  104    1   12     8  392  233  336    2   3786     3   206
 07       256    75    8   36   1039    40   276   139   11    1    3   114   11    2    5    23 1039   83  442    8   4008     3   248
 08       256    46   20   33   1119    82   152    99   20    2    3    28   95    1    4    25 1119  289  484   16   4911     1   313
 09       256    54   18   34    735    75   176    98   50    2    4    51  101    1    6    27  735  214  468   15  16869     3   486
 10       256    72   23   49   1387   172   210   122   47    2    4    40   80    1   20    48 1387  314  729   22  11034     2   347
 11       256    84   42   63   1091   139   243   164  565    2    8    33   88    1   15    46 1091    5  699   18  12502     5   657
 12       256    84   24   53    751    56   330   247   15    2    3    50  103    2   29    33  751  164  500   10  18058     3   924
 13       256    49   21   35   1070    59   187   148   40    2    5    18  462    1   14    21 1070  164  480   11  10221     2   350
 14       256    67   23   43    839   169   189   111    8    2    3    24   93    2   17    54  839  284  579   28  15819     1   556
 15       256    60   29   38    866    74   179    94    6    2    3    55  103    1   10    30  866  229  418   16  18273     4   236
 16       256    94   44   75   1458    92   215    96   18    2    3    88  131    1   14    31 1458  218  592   14  22717     6   558
 17       256    91   65   76   1137   272   102    35    8    2    3    29   83    1   14    38 1137  375  718   37  26494     5  1796
 18       256   103   52   77    842    64   276   149   32    2    4    99  101    1   11    28  842  393  580   10  30523     4  1311
 19       256    77   44   59   1007    53   276   206   32    2    4    52   91    1   11    18 1007  388  739    6  16502     3   207
 20       256    80   43   61   1125    26   409   346   24    2    3    48   94    1   28    15 1125  237  558    3  22517     5   195
 21       256    95   45   66    826    40   367   256   34    1    4    87  122    1   12    24  826  305  535    6   3586     4   191
 22       256    81   46   63    687    40   254   166   39    2    4    72  108    1   16    16  687  388  525    6  26822     4   303
 23       256    82   17   58   1330    56   256   165   25    2    4    69  120    1   19    22 1330  345  570    8  31860     4  3814
                                             ----- -----                  -----                -----                ---
DATE                                          8405  5782                   2022                  601                  7
 vdxnote: Since new mounts take priority over dismounts, this field could show a value higher than the number of concurrently ALLOCATED
          drives from the host view during an hour when there were many mount requests in a short period of time.
peer-to-peer notes: *1=Vts1/Vts2, *2=max(V1,V2), *3=min(V1,V2), *4=V1+V2, *5=Complib
```

Figure 7-7 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, this 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 296. 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 294.

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 freespace 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.

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. Failure to establish proper Fast Ready category definitions obviously will increase physical drive and accessor activity and will have a negative effect on performance, especially on virtual mount times. Refer to 4.3.5, "Define Fast Ready categories" on page 133 for more information about setting up the fast-ready categories.

– **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. The number depends on the effective size of the TVC and the time between references of logical volumes. When we say "effective size", we are referring to the size of the TVC, regardless of whether or not you have the EHPO feature, which provides compression before data is written in the TVC. With the EHPO feature installed, you can fit more logical volumes in the TVC, therefore, allowing all of the logical volumes to stay in the TVC for a longer period of time.

– **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%?**

See Figure 7-7 on page 293 (Virtual Drive Activity Report).

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.

If you are constantly exceeding this guideline, and your overall average mount time is greater than 30 seconds, you may want to investigate why you are having so many recalls. If you cannot prevent or reduce these recalls by rescheduling or removing workload, then a larger cache capacity may improve your cache hit ratio.

4. **How much data is my VTS managing and how much more data can it store before it becomes full?**

   Also, much more information is provided:

   – 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 the appropriate section.

   Figure 7-8 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 7-8 VTS Hourly TVC Activity Report

A description of the appropriate SMF fields is provided 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. Reconciliation is the process during which the VTS determines the most recent copy of a virtual volume, it deletes all other instances of that volume from the active volume list. Therefore, once all the volumes have been used at least once, the 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 the logical volumes are being used until the next reconciliation. Reconciliation is run every 18 hours unless the system is very busy. After 24 hours, it will run regardless of how busy the system is. After reconciliation has identified any volumes eligible for reclamation, then reclamation will start unless it is inhibited by schedule or lack of drives. Virtual volumes that exist in the TVC but have not been copied to tape are not included. See “Reclamation and reconciliation” on page 137 for a detailed description.

– **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. The number of empty stacked volumes is multiplied by the estimated amount of data that could be stored on one of these volumes. The Model B18 VTS with EHPO, Model B10 VTS and Model B20 VTS provide compression into the TVC; therefore, further compression when writing to the stacked volume is unlikely, and the capacity of a stacked volume will be approximately 10 GB of TVC data for the 3590-B1A tape drives and 20 GB of TVC for the 3590-E1A tape drives and 30GB of TVC for the 3590-H1A tape drives. With this configuration, the value presented has been interpreted into host data capacity by multiplying the number of scratch physical volumes times 10 (or 20 for 3590-E1A or 30 for 3590-H1A), multiplied by the historic compression ratio. Because the compression ratio can fluctuate, the value presented can also fluctuate.

– **# 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?**

   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 7-9 shows the VTSSTATS Hourly Physical Drive Activity.

HOURLY PHYSICAL DRIVE ACTIVITY (CHANGE)

```
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---  --PHYSICAL DRIVES--    --SECONDS TO MOUNT--   ACCSSOR_MNTS          TOTAL  ----------------------  VIRTUAL
        ---DRIVES----  --CONCURRENT MOUNTS    ---PHYSICAL DRIVES--   (VTS+NATIVE)           MNTS  --PHYSICAL MOUNTS FOR--  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 7-9   VTS Hourly Physical Drive Activity

A description of the associated SMF records are shown here:

– **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. Refer to “Reclamation and reconciliation” on page 137 for a detailed overview.

Other possible reasons for high reclaim mount numbers include these:

- There was a previous time interval with high cache usage rate prevented reclamation.
- There was a recent change of reclaim threshold on the Library Manager.
- There was 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.

6. **Are you experiencing elongated or erratic job run times?**

   Elongated or erratic job run times can 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 that 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.

7.4.3 APM reports

The following reports are available as part of the Advanced Policy Management:

- *Hourly preference level virtual mounts* shows the mounts in PG0.
- *Hourly preference level cache statistics* shows for how long data is kept in cache for PG1
- *Hourly cache IART mounts statistics* shows the statistics for logical volumes which use the host IART value.
- *Hourly pool active data and reclamation* shows how many cartridges are available and how many can be reclaimed.
- *Hourly volume poolset statistics* shows the media usage for all pools.
- *Hourly volume pool media statistics* shows the number of media used per pool.

Figure 7-10 shows a sample report for the Hourly volume pool media statistics for Pool 1.

by POOL

```
HOURPOLM GROUP: ...... -HOUR VOLUME POOL MEDIA STATISTICS  VTSNUM=70277                              VTSSTATS(05145)
______________________EXAMPLE_____________________                                        Run Date: 17:14 25MAY2005
Report Date 01APR2005/00:42(FRI)                                                            Data timestamp: 00:42:42
Logical Day Start hour: 00                                                                           LM ref hour: 01

====>POOL NUM=  1 (3590H)    RECLAMATION POOL NUM=  1

     LOG VV    MB OF   MB       STATC STATC BRROW BRROW STATC STATC BRROW BRROW  AVG AGE   MAX AGE   AVG AGE   MAX AGE
     ON PHYS   POOL    WRTTN    SCRTH PRVAT SCRTH PRVAT SCRTH PRVAT SCRTH PRVAT  RES DATA  RES DATA  FULL VOL  FULL VOL
 HR  IN POOL   DATA    TO POOL  MDIA0 MDIA0 MDIA0 MDIA0 MDIA1 MDIA1 MDIA1 MDIA1    DAYS      DAYS      DAYS     DAYS
 94S2=>alvip   adivp   dwtplh   sssvc spsvc bssvc bpsvc sssvc spsvc bssvc bpsvc    aaord     maord    aaofpsv   maofpsv

 00  195805  85383978  171671     0     0     1  1167     0     0     1  1781       60       511        59       513
 01  196235  85447896  153175     0     0     1  1168     0     0     1  1781       60       511        59       513
 02  196487  85550890  178457     0     0     0  1169     0     0     2  1781       60       511        59       513
 03  196675  85650222  141874     0     0     0  1171     0     0     2  1782       60       511        59       513
 04  194139  84772083   89798     0     0     0  1171     0     0     2  1783       59       511        59       514
 05  192219  84084811  119084     0     0     0  1171     0     0     2  1783       59       511        59       514
 06  192727  84192495  147286     0     0     0  1173     0     0     1  1785       59       511        59       514
 07  193059  84293334  146372     0     0     1  1174     0     0     1  1786       59       511        59       514
 08  193621  84417037  168295     0     0     0  1176     0     0     2  1787       59       511        59       514
 09  193977  84576065  155741     0     0     1  1177     0     0     1  1790       59       511        59       514
 10  194159  84650626   96833     0     0     0  1179     0     0     2  1790       60       511        59       514
 11  194369  84739638  134753     0     0     0  1179     0     0     2  1790       60       511        59       514
 12  194677  84875771  171510     0     0     0  1179     0     0     2  1794       60       511        59       514
 13  194973  84981206  134282     0     0     0  1178     0     0     2  1796       60       511        59       514
 14  193773  84593637  114954     0     0     0  1178     0     0     3  1798       60       511        59       514
 15  193354  84460657  142411     0     0     0  1178     0     0     3  1800       60       511        59       514
 16  193580  84541621  134490     0     0     0  1179     0     0     3  1801       60       511        59       514
 17  193752  84643938  123293     0     0     0  1179     0     0     3  1802       60       512        59       514
 18  193971  84715652  153932     0     0     0  1179     0     0     3  1803       60       512        59       514
 19  194189  84800694   85306     0     0     1  1172     0     0     3  1804       60       512        59       514
 20  194437  84864790   88543     0     0     0  1171     0     0     3  1801       60       512        59       514
 21  194677  84928679   86921     0     0     0  1171     0     0     4  1800       60       512        59       514
 22  194936  85029797  119086     0     0     0  1171     0     0     2  1803       60       512        59       514
 23  195087  85121792  100860     0     0     0  1171     0     0     3  1805       60       512        59       514
```

Figure 7-10   Hourly Volume Pool Media Statistics report

7.4.4 Hardware capacity

If the workload is appropriate or, in any case unchangeable, and the VTS is indicating signs of stress, 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 or throttling is occurring, consider installing additional tape drives.

Also, depending on your current system configuration, consider the installation of the FICON Performance Accelerator Feature - FPAF.

To help with the potential configuring of any VTS upgrade use the BatchMagic tool.

7.5 Using Library Manager windows to monitor the VTS

In this section we show and describe a few useful windows which can be used to display information on performance and capacity of the VTS subsystem and how they relate to the SMF94 type record. We do not cover all windows available on the Library Manager here. See *IBM TotalStorage 3494 Tape Library Operator Guide*, GA32-0280 for more detailed information on Library Manager windows.

Note that both the IBM 3494 Library Manager and the IBM 3953 Library Manager provide basically the same information and panels. We therefore only show the 3494 Library Manager panels in the following sections.

You can access the windows from the **Status** pull-down menu (Figure 7-11) on the Library Manager action bar.

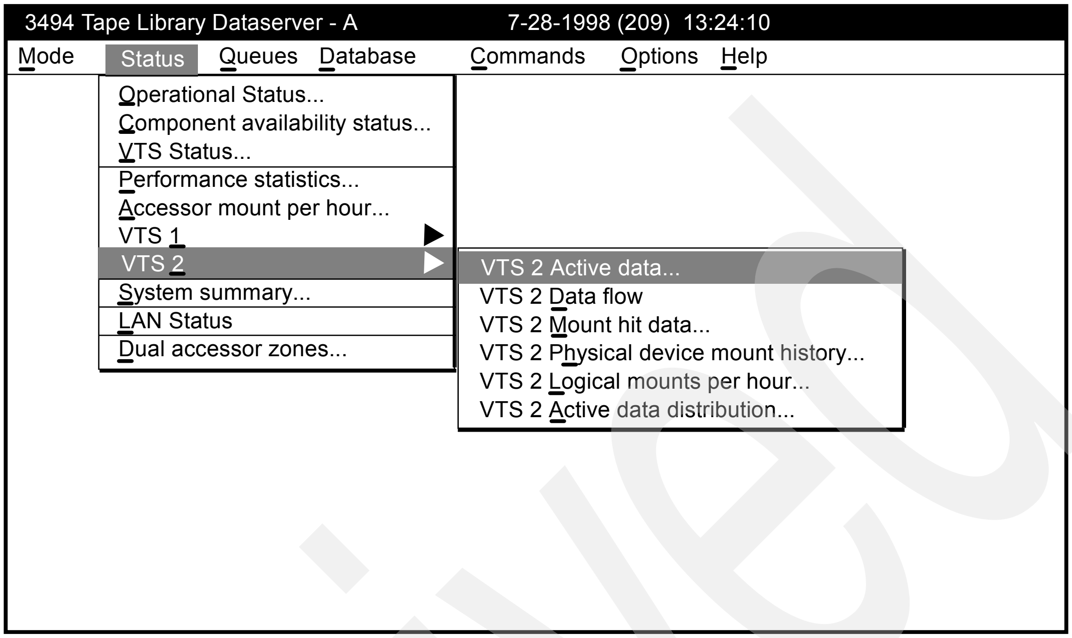

Figure 7-11 Status pull-down menu

7.5.1 Performance statistics

The Performance Statistics window (Figure 7-12) displays statistics related to physical activity in the library. This includes a brief mount summary for the drives attached to the VTS. Actions involving VTS logical volumes are not part of these statistics.

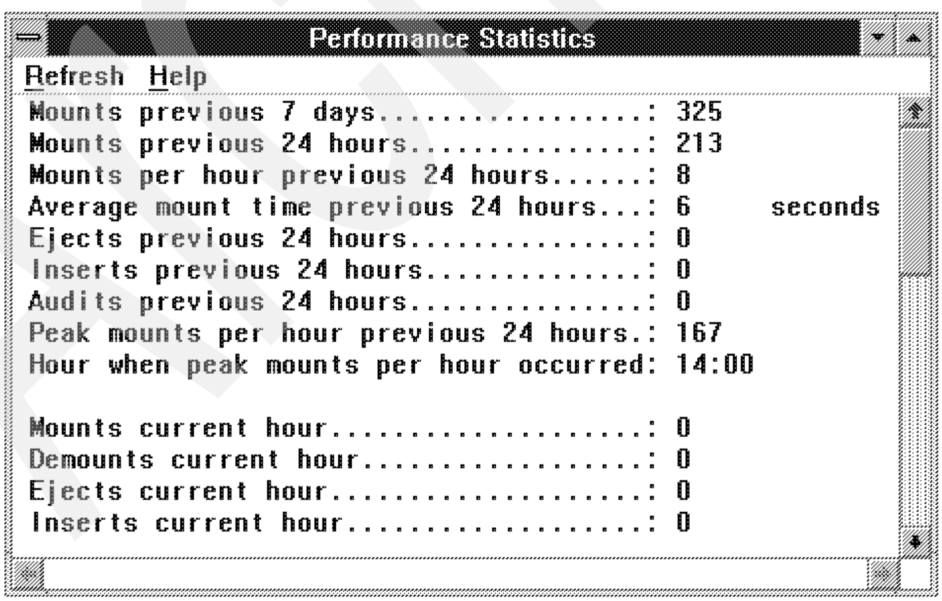

Figure 7-12 Performance Statistics window

7.5.2 VTS active data

The Active Data window (Figure 7-13) displays a graph showing the amount of active data, amount of free storage, maximum active data and a free storage alarm level for the stacked volumes in a VTS. The data stored on stacked volumes is from the TVC and may be compressed when the VTS has the EHPO feature; therefore, the graphs do not represent actual host data bytes, but they represent the compressed volume sizes as stored in the TVC. A separate graph is available for each VTS in the library. Data is displayed for the previous 29 days as of midnight and for the current day on an daily snapshot.

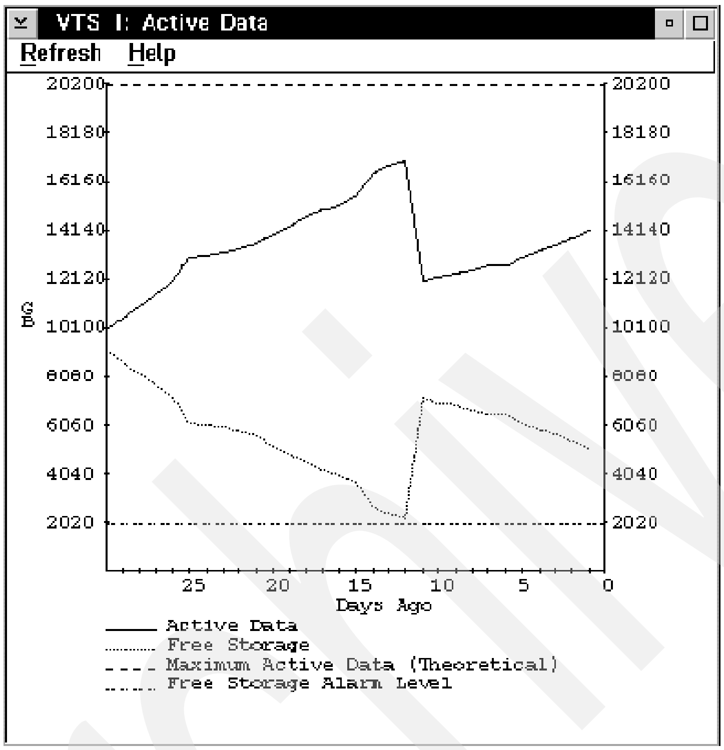

Figure 7-13 VTS Active Data window

Active data (SMF94VBA capacity used in MB) is the total size, as stored in the TVC, of all active logical volumes that have been copied to stacked volumes. Virtual volumes that exist in the TVC but that have not been copied to tape are not included. Invalidated copies of logical volumes (duplicate volumes which have not been deleted from the active volume list during the reconciliation process) are included in this field.

> **Note:** Invalidated copies of logical volumes reduce the amount of active data you can store on the stacked volumes. See 4.3.5, “Define Fast Ready categories” on page 133 for a detailed description.

Maximum active data is the total capacity of all stacked volumes that have been inserted in the VTS. It uses the actual compression ratio achieved for all data copied from the TVC to full stacked volumes to calculate the Maximum Active Data. Because data received from the host attachments has previously been compressed into the TVC, compression to the stacked volume from the TVC will be approximately 1 to 1.

Free storage (SMF94VEC capacity available) is the total capacity of all empty stacked volumes in the library calculated by using a compression ratio as described for Maximum Active Data above. Partially filled stacked volumes are not included in this calculation. Free Storage gives an indication of how much data from the TVC can be added to stacked volumes currently in the VTS. You can display, via the Library Manager, how many scratch stacked volumes you have currently available by searching on category FF03. See *IBM TotalStorage 3494 Tape Library Operator Guide*, GA32-0280.

Free storage alarm level is a threshold to warn you when to add more stacked volumes to the VTS library. If the number of empty stacked volumes available is less than the number of stacked volumes required to store the amount of TVC data specified by the Free Storage Threshold (GB), the Library Manager signals an intervention required condition to notify the operator to add more stacked volumes. See "Reclamation and reconciliation" on page 137 for information on the recommended threshold settings. You can modify the alarm threshold on the Library Manager's VTS Management Policies window (see *IBM TotalStorage 3494 Tape Library Operator Guide*).

7.5.3 Virtual Tape Server data flow

The Data Flow window (Figure 7-14) displays a graph showing the amount of data written to and read from the channel **(SMF94VBW and SMF94VBR)**. Data is displayed for the previous 24 hours and is unique to the particular VTS that was selected. The numbers are based on the amount transferred at the channel interface and reflect the (uncompressed) amount of host data. The current hour's data is designated by a diamond-shaped marker.

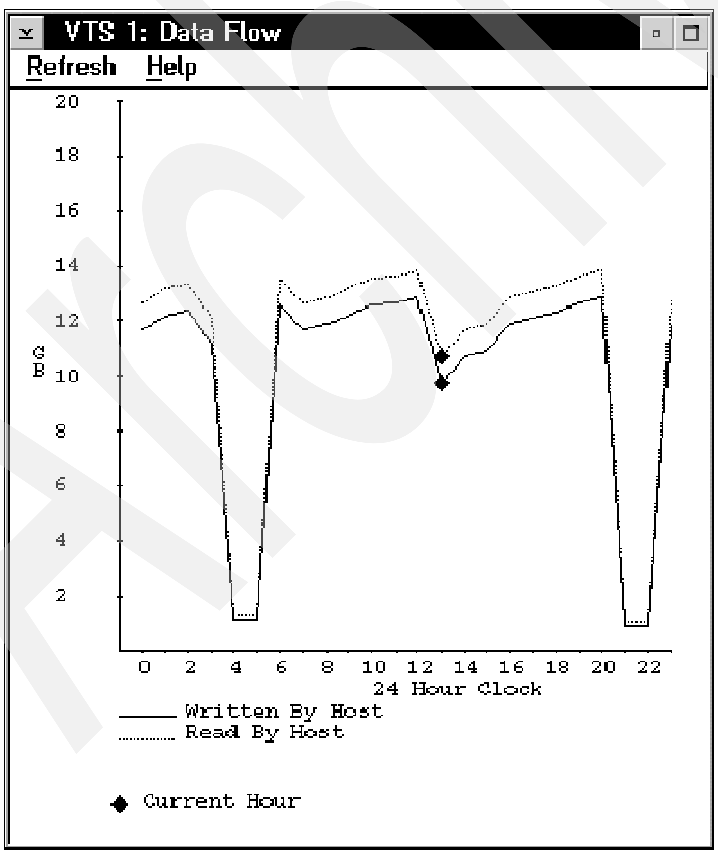

Figure 7-14 VTS Data Flow window

The Written By Host and Read By Host curves show the amount of data successfully transferred through the host channels to and from virtual volumes during the hour. The bytes are credited to the hour in which the virtual volume is closed.

7.5.4 Virtual Tape Server mount hit data

The Mount Hit Data window (Figure 7-15) shows the distribution in percentage of three types of logical mounts:

- Fast-ready hits (SMF94VFR)
- Cache hits (SMF94VMH)
- Cache miss, mounts requiring a recall (SMF94VMS)

Data is displayed for the previous 24 hours. The current hour's data is designated by a diamond-shaped marker.

The graph displays three lines, one for each type of mount, indicating its percentage of the total number of mounts for an hour. To display the total number of logical mounts per hour, use the Logical Mounts Per Hour window (Figure 7-15).

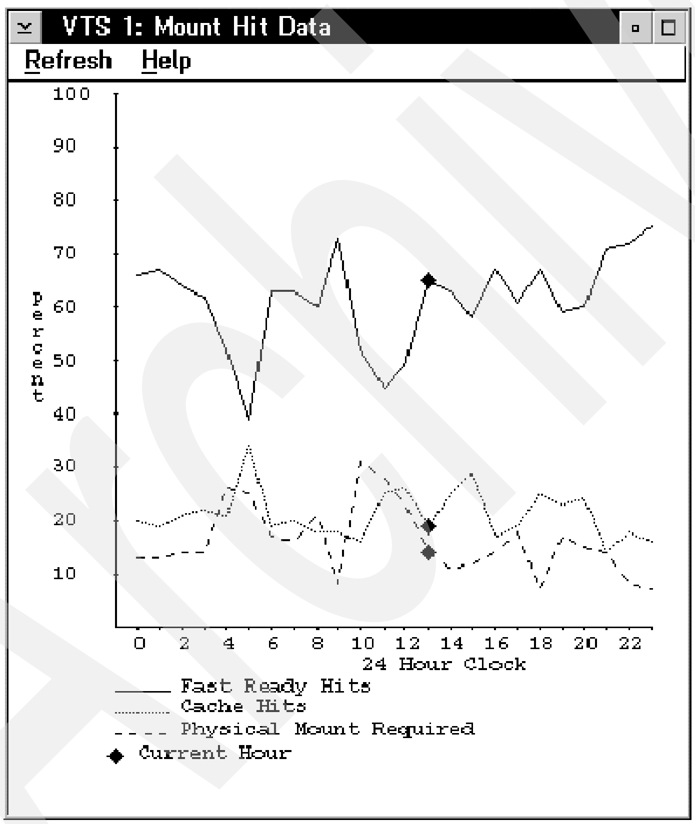

Figure 7-15 VTS Mount Hit Data window

Fast Ready hits (SMF94VFR) means the number of times the host requested a mount from a specific volume category, and the category was designated as “fast-ready”. Typically, these are scratch mounts. Because this type of mount does not require any recall of data from tape, it is the fastest. To benefit from Fast Ready hits, assign the Fast Ready attribute to your scratch volume categories. If this value is 0, you should check your Fast Ready categories settings described in 4.3.5, “Define Fast Ready categories” on page 133.

Cache hits (SMF94VMH) shows the number of virtual mounts, where the volume to be mounted still resides in the TVC. This type of mount does not require any recall of data from tape.

Physical mount required (SMF94VMS) is the number of times a logical volume had to be recalled from a stacked volume into the TVC to satisfy a virtual mount request. While a stacked volume is mounted, the VTS can recall more than one logical volume from it. Therefore, the graph does not represent the number of times a stacked volume actually had to be mounted in response to a virtual mount request. This is the slowest type of virtual volume mount.

7.5.5 VTS physical device mount history

The Physical Device Mount History window (Figure 7-16) displays a graph showing the average, maximum and minimum number of physical tape drives used at one time to mount stacked volumes. Data is displayed for the previous 24 hours. The current hour's data is designated by a diamond-shaped marker.

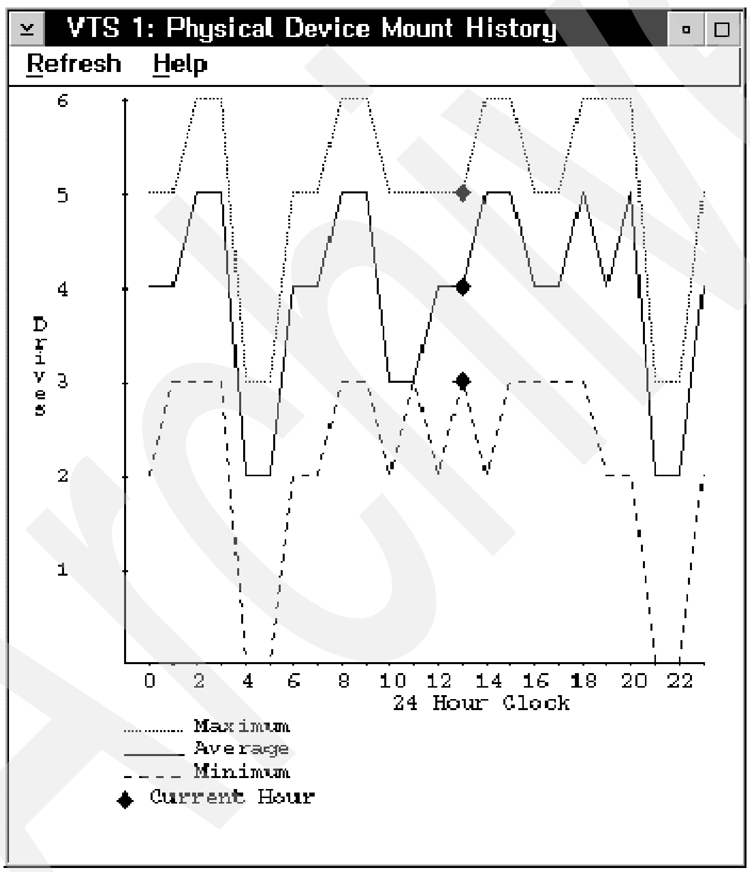

Figure 7-16 VTS Physical Device Mount History window

Average (SMF94VTV) is the average number of physical drives that were used concurrently to mount stacked volumes during the hour. The sampling rate for calculating the average is 10 seconds.

Maximum (SMF94VTX) is the maximum number of physical drives used concurrently to mount stacked volumes during the hour.

Minimum (SMF94VTN) is the minimum number of physical drives used concurrently to mount stacked volumes during the hour.

7.5.6 Virtual Tape Server logical mounts per hour

The Logical Mounts Per Hour window (Figure 7-17) shows the number of logical mounts per hour. The figure is the sum of Fast Ready mounts, cache hit mounts and cache misses (Recall). Data is displayed for the previous 24 hours. The current hour's data is designated by a diamond-shaped marker.

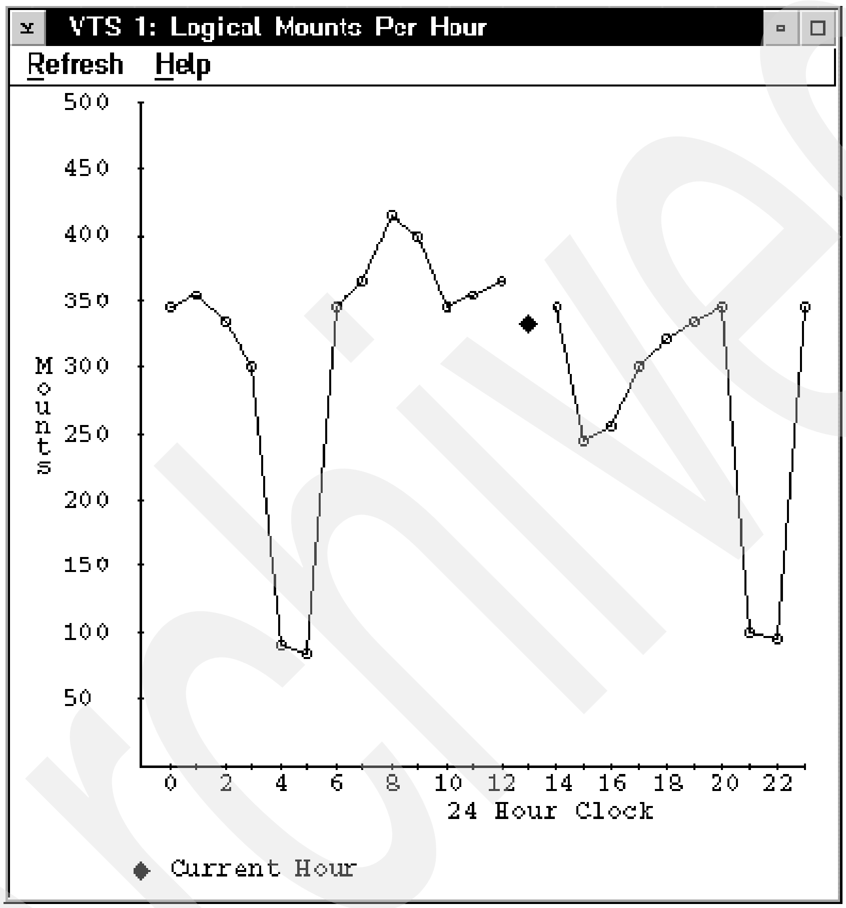

Figure 7-17 VTS Logical Mounts Per Hour window

To display the percentage distribution of logical mounts, use the **Mount Hit Data** window (Figure 7-15 on page 303).

7.5.7 Active data distribution

The Active Data Distribution window (Figure 7-18) displays a graph showing the distribution of the data on the stacked volumes. Each bar goes to the height that signifies the number of tapes in that group. There is a bar for each 5% increment in active data. You use this window to identify the number of tapes that will be freed up if you were to change the reclaim threshold. See 2.1.8, “Reconciliation and reclamation” on page 17 for more details.

At the bottom of the window you can identify the number of active volumes, that is the number of stacked volumes that have active data on them. This does not include stacked volumes that are in the process of being filled with new data.

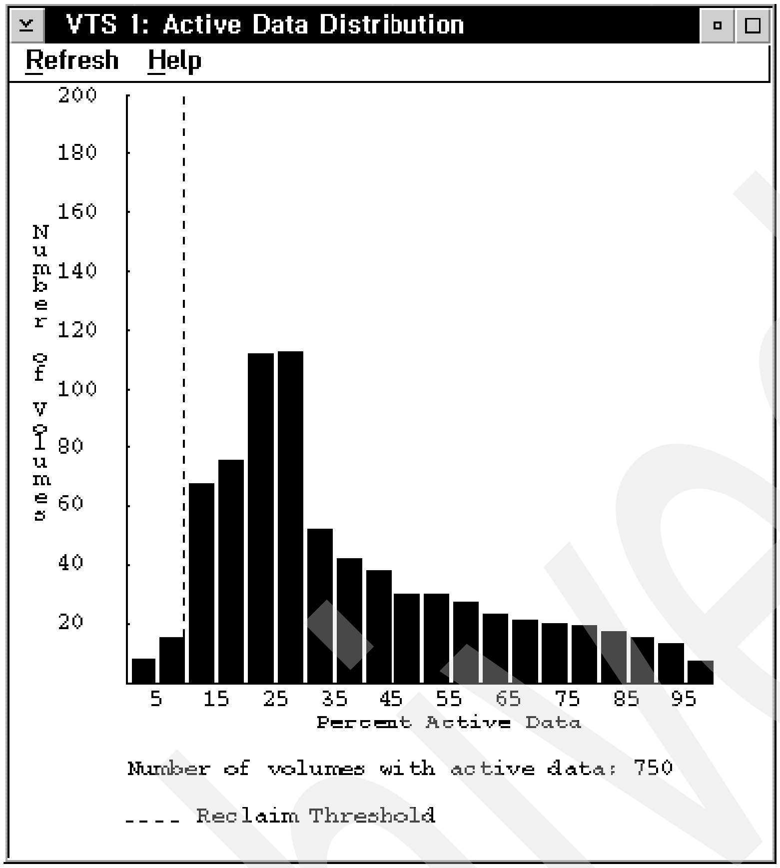

Figure 7-18 VTS Active Data Distribution window

7.5.8 Active data distribution per pool

The Active Data Distribution per Pool window (Figure 7-19) is available with Advanced Policy Management. This window displays a graph showing the distribution of the data on the stacked volumes at the **pool level**. Each bar goes to the height that signifies the number of tapes in that group. There is a bar for each 5% increment in active data. You use this window to identify the number of tapes that will be freed within a pool up if you were to change the reclaim threshold.

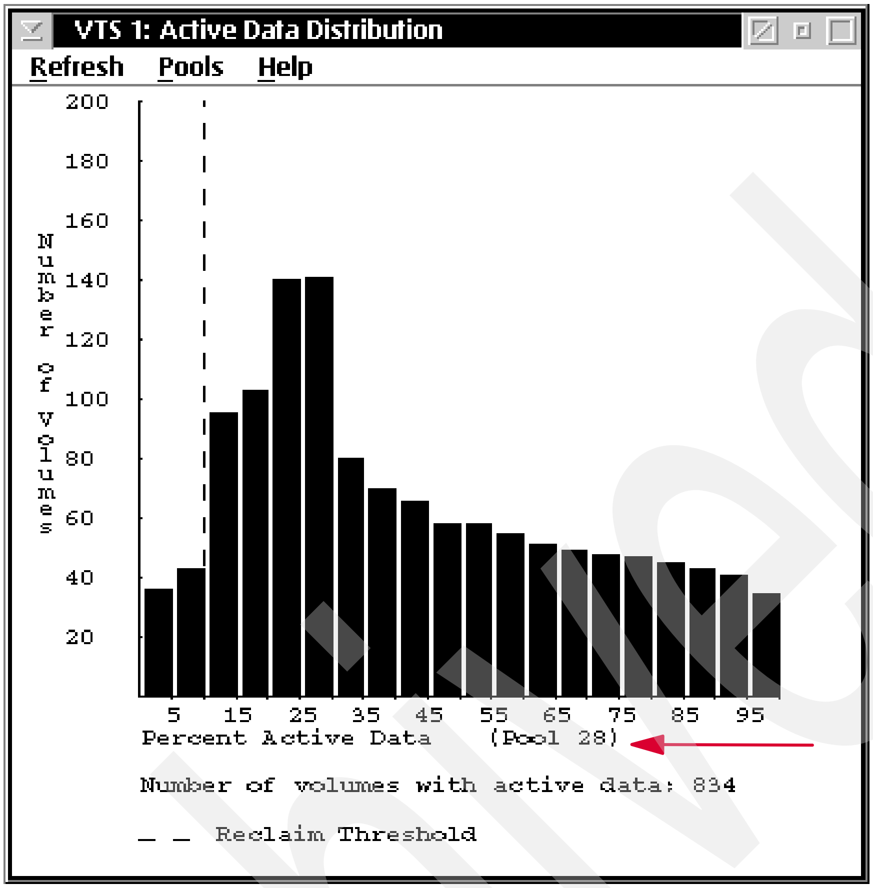

Figure 7-19 VTS Active Data Distribution per Pool window

7.5.9 VTS specific logs

The VTS keeps track of certain events that occur in a wraparound log. This log contains only a limited number of events or activities. When diagnosing certain software situations, you may want to look at these logs. If there is a situation that arises where the logs are needed, contact your IBM service representative.

7.5.10 Other considerations

To display the number of storage slots in the tape library, issue the following command, as shown in Figure 7-20:

```
DISPLAY SMS,LIBRARY(library-name),DETAIL
```

The number, represented under Total Slots, reflects the physical storage slots in the entire 3494, and is not a count of the virtual volume capability of the library.

```
DISPLAY SMS,LIBRARY(LIBVTS1),DETAIL

CBR1110I OAM library status:
TAPE     LIB  DEVICE    TOT ONL AVL TOTAL EMPTY SCRTCH  ON OP
LIBRARY  TYP  TYPE      DRV DRV DRV SLOTS SLOTS  VOLS
LIBVTS1  VL   3494-L10  128 100  50  1443   800   146   Y  Y
----------------------------------------------------------------------
MEDIA      SCRATCH       SCRATCH     SCRATCH**
TYPE         COUNT     THRESHOLD    CATEGORY**
MEDIA1          16            25        0011**
MEDIA2         130            25        0012**
----------------------------------------------------------------------
OPERATIONAL STATE:AUTOMATED
ERROR CATEGORY SCRATCH COUNT:          12
SCRATCH STACKED VOLUME COUNT:         243
PRIVATE STACKED VOLUME COUNT:         400
----------------------------------------------------------------------
Library supports import/export.**
Library supports outboard policy management.**
Host initiated import in process
```

Figure 7-20 Library Status Display window

Note: In Figure 7-20, ** denotes new fields introduced with Advanced Policy Management.

An SMF type 21 record (Error Statistics by Volume) is written when a tape volume is demounted. This record contains the statistics for the volume while it was mounted.

For a virtual volume, the fields in the SMF type 21 record that represent activity at the physical tape drive contain zeros. These fields include:

- Number of temporary errors
- Number of permanent errors
- Number of erase gaps
- Number of cleaner actions

The number of bytes read and written is reported for a virtual volume on the SMF type 21 record.

Mount messages for virtual volumes are reported in the same way as for native volumes in a 3494. In z/OS, mount messages are written in the Syslog and Joblog but are not displayed on the system console.

7.6 VTS statistical data

The 3494 accumulates statistics on the activity of the tape library that results from responding to all hosts attached to the library. At the end of an hour, the Library Manager calculates composite statistics about the activity of all devices in the library during the hour and sends the statistics to all z/OS attached hosts.

The library statistical record presents current information on the library's status and summarizes activity during the last hour.

The library statistical record has a length of 256 bytes; only the first 128 bytes contain values if a VTS is not installed in the 3494. When a VTS is installed in a 3494, the bytes above 128 contain virtual mount statistics and operational information about the VTS. The existing statistics in bytes 0 through 127 continue to reflect all physical movement of cartridges within the 3494, including the movement of VTS stacked volumes. When a VTS is not installed, bytes above 128 contain X'00'.

7.6.1 Collecting the statistical records

In SMS managed tape and BTLS environments, the z/OS tape error-recovery program (module IGE0001E) reads the statistical record sent by the Library Manager and writes an SMF type 94 record (IBM Tape Library Statistics) that includes the activity of all devices in the tape library. The SMF type 94 record has been enhanced to support the new statistics related to the VTS. The maintenance that is needed to get this support is listed in 3.8.1, "Software requirements for ES/3090, ES/9000, S/390, or zSeries" on page 93. Appendix F, "SMF type 94 record layout" on page 461 shows the format of the SMF type 94 record. The new statistics for the VTS are generated by the IECSMF94 macro.

If you have Performance Reporter for MVS (5695-101), support is provided for the new VTS related fields in the SMF type 94 record. The maintenance required to support these new fields is listed in Section 3.8.1, "Software requirements for ES/3090, ES/9000, S/390, or zSeries" on page 93.

In a non-z/OS environment, the VTS statistical records are not collected by any host. This applies to VM hosts and SCSI hosts connected to a VTS. However, the VTS operating system collects the statistical data, so that they can be browsed in the Library Manager's log files and from the TotalStorage Library Specialist.

In order to collect the statistical data when no z/OS system is connected to the Library Manager, you need to extract them directly from the Library Manager. See 7.6.4, "Monitoring the VTS with VTSLOGRP" on page 310.

7.6.2 Analyzing SMF type 94 and LOGREC MDR records

The SMF type 94 record is sent to each attached z/OS host on an hourly basis. Because the VTS does outboard data management, it is a little different from other tape devices. The SMF type 94 records can be used to analyze the VTS for capacity and performance planning. You can use any of the standard tools to print out the SMF type 94 records. We give some examples in 7.4, "Monitoring and evaluating VTS performance" on page 290 on how to use these records to monitor the workload currently running within the VTS.

VTSSTATS reports data blocksize and compression ratios. This data can be used to evaluate VTS performance. A methodology for this evaluation is given in 7.4, "Monitoring and evaluating VTS performance" on page 290.

7.6.3 Monitoring

RMF may be used to monitor disconnect time for the VTS virtual drives. Disconnect time is reported by RMF and not by the SMF type 94 records from the VTS. See page 298 for more information on RMF processing.

7.6.4 Monitoring the VTS with VTSLOGRP

The Library Manager of a VTS builds statistical data every hour. This data contains information for all activities relevant to the VTS. On the z/OS host these records are recorded in the SMF type 94 record and can be processed by your normal SMF reporting tools.

In a non-z/OS environment, the VTS statistical records are not collected by any host. This applies to VM hosts and SCSI hosts connected to a VTS. The VTS operating system collects the statistical data to its own Library Managers log files.

In order to collect the statistical data when non-z/OS system is connected to the Library Manager, you need to refer them directly from the Library Manager or Tape Library Specialist.

VTSLOGRP was developed as a reporting package which can be used to analyze the data from the Library Manager log files for performance and capacity planning.

VTSLOGRP can be run on multiple platforms, including the z/OS platform running z/OS. Currently, the supported platforms are OS/2®, VM, z/OS and Windows. For Windows you must have the OBJECT REXX package installed.

The main reports are broken down into the same categories as VTSSTATS, which is covered in detail in section 7.4, “Monitoring and evaluating VTS performance” on page 290. The reports have the following headings:

- **Physical:** Reports on physical drive activity.
- **Cache:** Reports on transient volume cache activity.
- **Virtual:** Reports on virtual volume activity, count, time and category.

There are a number of reports which have already been built which you can use. One of the most useful reports is the *Hit Parade Report* (Figure 7-21), which focuses on key fields of the statistical record.

```
VTSLOGRP V1.L3                              VTS LIBRARY STATISTICS
                         DAILY HIT PARADE                 FOR VTS 60499

  CACHE MISS COMPARED TO VIRTUAL MOUNTS & STAGE

             CACHE    TOTAL FAST RDY CACHE HT PHYSICAL
DATE          MISS VIRT.MNT      MNT      MNT STAGE MT

27APR1999       54       482       299       129       59
28APR1999       44       398       250       104       55
29APR1999       42       472       316       114       55
30APR1999       28       453       306       119       39
01MAY1999       24       340       248        68        24
02MAY1999         2       217       206         9         9
03MAY1999       75       491       311       105       84


% CACHE MISS COMPARED TO VIRTUAL MOUNTS & STAGE
```

Figure 7-21 Daily hit parade - cache miss

This report can be used to analyze your cache activity for the relevant day. Key things to look for in this report are:

- Total amount of cache misses
- Total virtual mounts, this is the value of cache hits, cache misses, and Fast Ready mounts
- Total Fast Ready mounts, this is the number of mounts for scratch processing
- Total physical stage mounts, this number shows you how many times a stacked volumes was mounted to stage a virtual volume back into cache

See 7.4, “Monitoring and evaluating VTS performance” on page 290 for more detailed descriptions of these values.

Other reports and files are produced which can be uploaded to spreadsheet packages, such as Lotus® 123 or Microsoft Excel. See 7.6.2, “Analyzing SMF type 94 and LOGREC MDR records” on page 309 for information on downloading this product from the intranet site.

7.7 Using the Tape Library Specialist to monitor the VTS

In this section we discuss the IBM TotalStorage Enterprise Tape Library Specialist (Specialist) and show example screens with monitor and statistical data.

7.7.1 Specialist displays

The Specialists are a family of tools used for displaying and monitoring IBM tape library and DASD product status. These tools do not provide reports, but can be used for online queries about the status of the Enterprise Tape Library VTS, its components and the distributed libraries. See Figure 7-22.

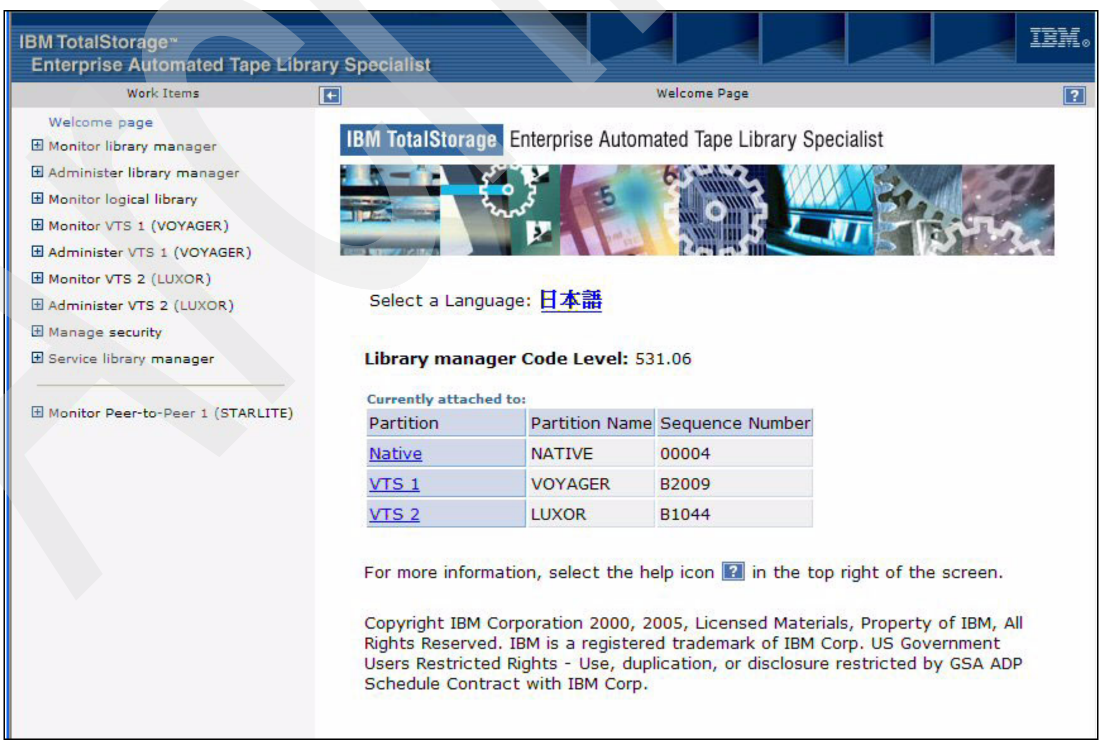

Figure 7-22 IBM TotalStorage Tape Library Specialist: home page

The IBM TotalStorage 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).

Library Manager information includes:

- System summary
- Operational status and operator interventions
- Component availability
- Performance statistics
- Command queue
- LAN hosts 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

The real time statistics include workload and throttling statistics; the same information that has been added in the SMF94 record.

7.7.2 Specialist prerequisites

The Library Manager must be connected to your system's LAN using FC5219 (Token-Ring Adapter) or FC5220 (Ethernet Adapter). During the installation process, the IBM System Service Representative (SSR) will set up TCP/IP on the Library Manager to use your assigned TCP/IP host name and TCP/IP address (and router information, if necessary). You can help the installation process if you obtain the following information before the installation starts:

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

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

You are required to have a commonly-used browser to view the information provided by the Specialist. Microsoft Internet Explorer version 5.0 or Netscape Navigator Version 4.7 with Javascript and Java™ enabled provide compatible capability. A text-based Web browser is not supported.

The Specialist resides within the Library Manager and is part of the microcode. Once attached to the customer network, it is available to the customer via a Web browser. There is no special “setup” required on the browser, however, the initialization phase will take a little longer when you start the connection for the first time ever on any client. This is expected due to the Java code download, it should only occur once. The general topography of the IBM TotalStorage ETL monitoring products is depicted in Figure 7-23.

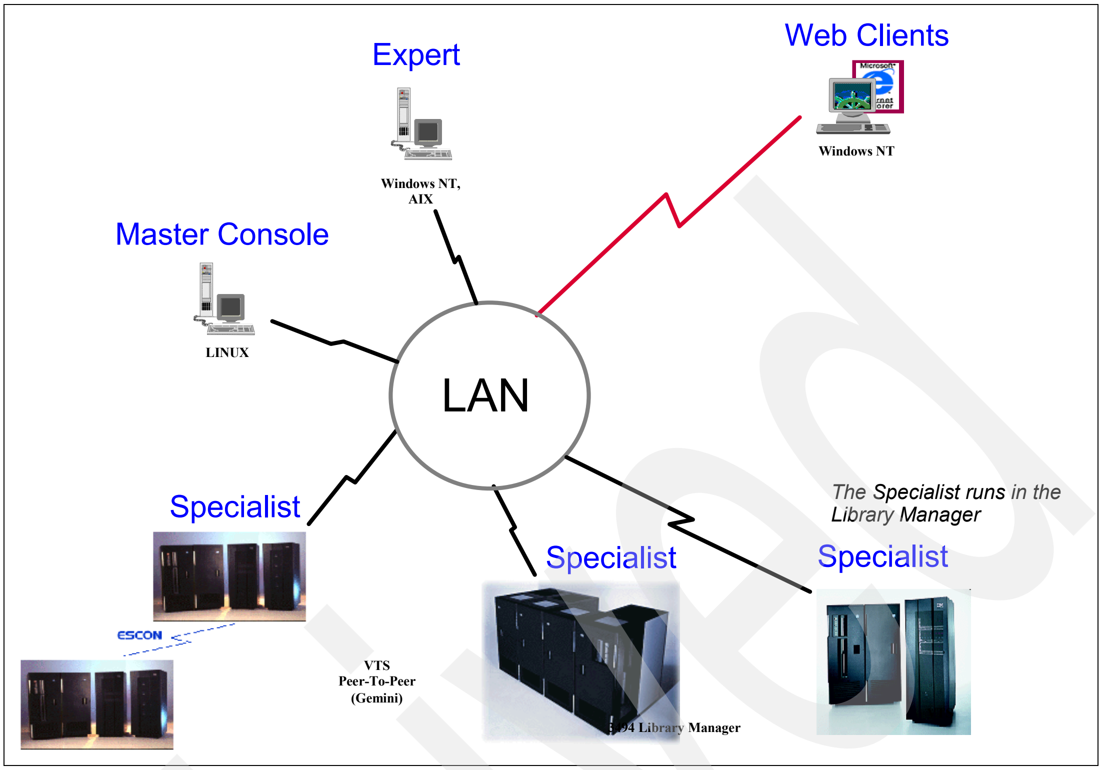

Figure 7-23 Enterprise Tape Library Specialist connection

Some of the panels can be used for online queries about the status of the Enterprise Tape Library VTS, its components, and the distributed libraries (see Figure 7-24). Others, which require Administration access, can be used to change your existing environment.

7.7.3 VTS status screen

The VTS Status screen shown in Figure 7-24 gives you an overview on the VTS and its capabilities. For the VTS in this example, you can see that:

- It is part of a PtP VTS.
- It has the Advanced Policy Management feature installed.
- Up to 300.000 logical volumes can be inserted.
- Both 3590 and 3592 tape drives are installed.
- Larger logical volumes are supported, but CE has not activated 2000 or 4000 MB support.
- It supports Secure Data Erase.

Figure 7-24 VTS Status panel

On the left hand side of the panel shown in Figure 7-24, you can see the work items to monitor the VTS:

Active data: Shows the amount of active data plus the maximum amount of data and the free storage threshold over the past 30 days.

Active data distribution: Gives an overall view of the amount of active data distributed over the stacked (physical) volumes. All pools can be viewed or pools can be viewed individually.

Data flow: Displays the amount of data written to and read from the channel. Data is displayed for the previous 24 hours. There is a text-only table below the graph.

Logical mounts per hour: Displays the number of logical mounts per hour, which includes the sum of fast ready mounts, cache hit mounts, and physical mounts (recalls). Data is displayed for the previous 24 hours.Mount hit data

Mount hit data: Displays the distribution in percentage of three types of logical mounts: fast ready hits, cache hits, and physical mounts required. Data is displayed for the previous 24 hours.

Physical device mount history: Displays the maximum, average, and minimum numbers of physical tape drives used at one time to mount stacked volumes. Data is displayed for the previous 24 hours.

Real time statistics: Displays a table that summarizes the VTS real time statistics. From here you have a direct link to the Backstore Media Count window which displays the media statistics for the common scratch pool and all used pools. See Figure 7-25 for an example.

Move/eject status: Displays two tables which show the status of in progress eject and move stacked volume requests.

Volser ranges: Summarizes the current defined Volser ranges for the VTS stacked volumes.

Management policies: Display the current VTS management policy settings for inhibit reclaim and the free storage threshold. These values can be changed at the Library Manager console, but not from the Specialist panel.

Category attributes: Displays the Library Manager categories that have the Fast Ready attribute set. The Fast Ready attribute can only be changed from the LM console.

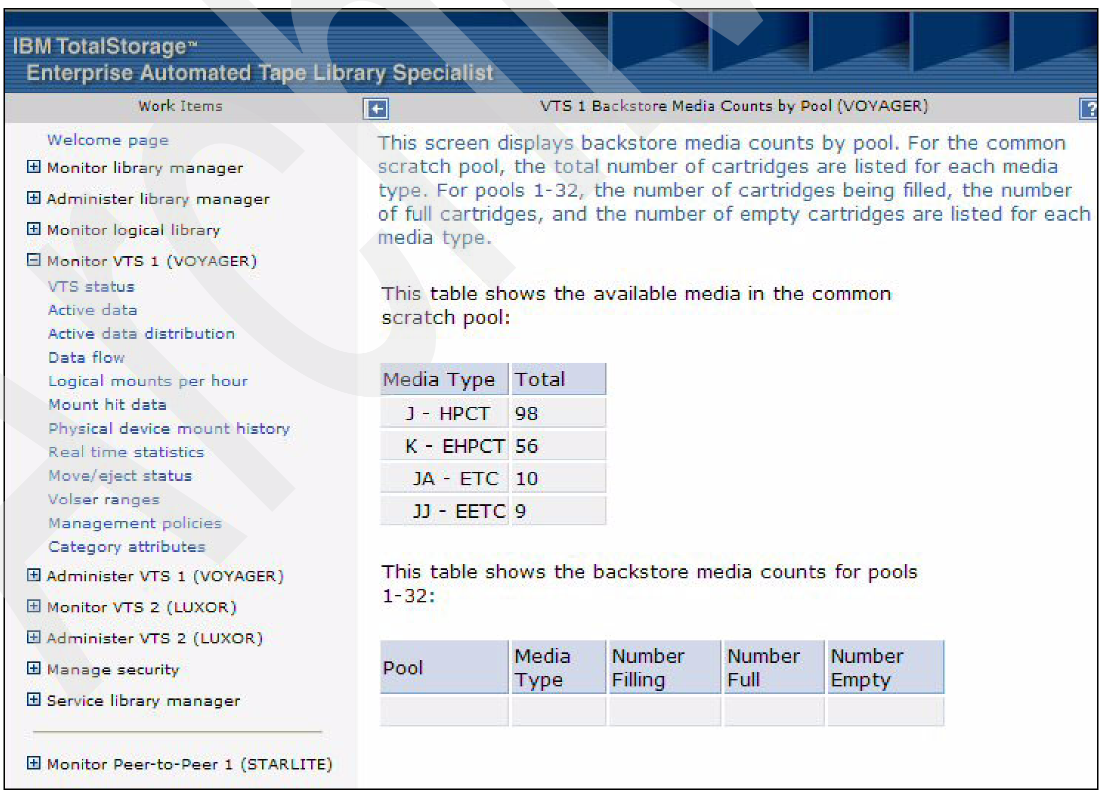

Figure 7-25 Backstore Media Counts

As you can see in Figure 7-25, there are currently no pools active. All cartridges are in the common scratch pool, and the VTS has both 3590 and 3592 drives and cartridges.

7.8 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

7.8.1 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 SSR 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, i.e. 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, i.e. 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.

Figure 7-26 shows the process flow of BVIR.

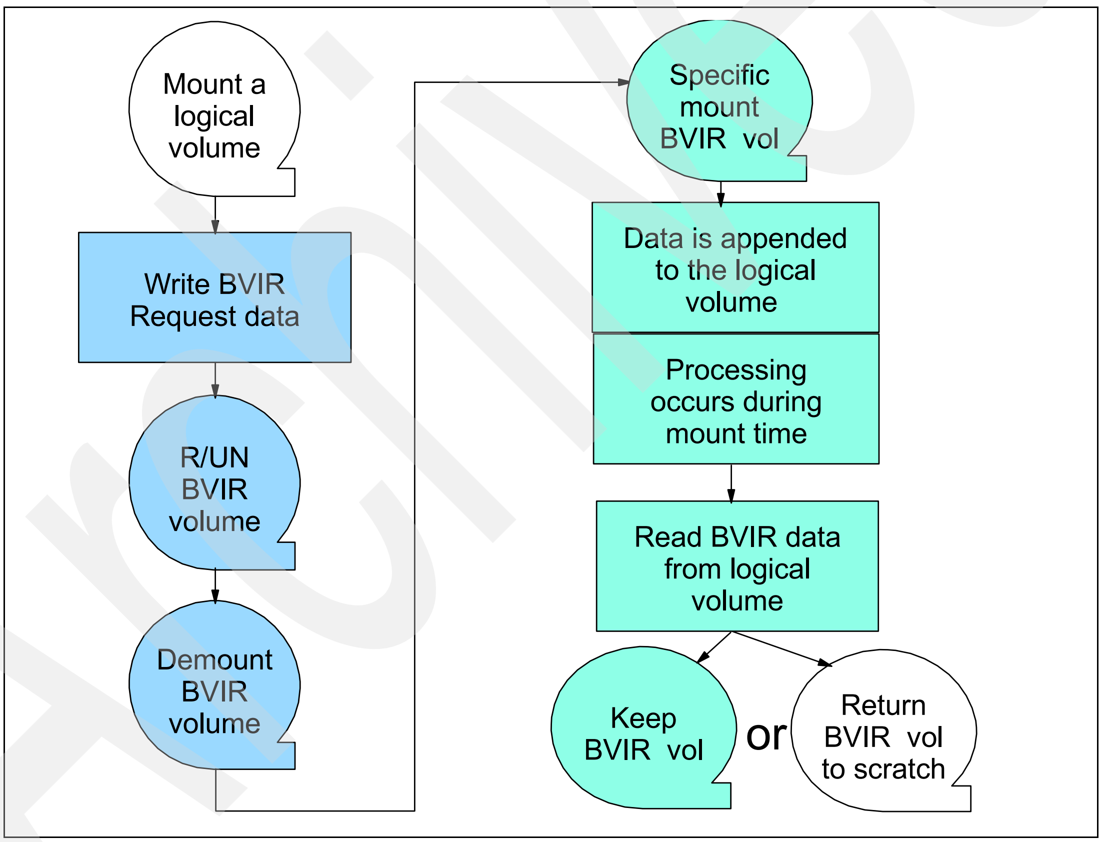

Figure 7-26 BVIR Process flow

The building of the response information does require a small amount of the resources of the VTS. We commended that you do not use the BVIR function to “poll” for a specific set of information, and only issue one request 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.

7.8.2 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.

7.8.3 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 that you use normal scratch allocation methods for each request — that is, DISP=(NEW,CATLG). In this way, any of the available scratch logical volumes in the VTS can be used. Likewise, we recommend that when the response volume's data is no longer needed, you return the logical volume 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 inTable 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 Note: 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.		

7.8.4 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 the date 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** 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, 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 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 onto another physical volume and if so, which part it is.

Size 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
```

7.8.5 Sample JCL

The following are fragments of JCL that show how to use a standard IBM utility, IEBGENER, to request and process the requested data for the BVIR function.

The 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 JES2 (this will not work for JES3 because it will not demount/remount the volume between steps). This example 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

7.9 IBM Enterprise Tape Library Expert

The IBM TotalStorage ETL Expert is a program whose purpose is to help you manage your Enterprise Storage Server® (ESS) and Enterprise Tape Library (ETL) using a Web browser interface. The primary panel is shown in Figure 7-27.

Figure 7-27 TotalStorage ETL Expert: home page

TotalStorage ETL Expert helps you do the following tasks:

- Performance management
- Asset management
- Capacity management

Performance management for ETL

In your role as a system administrator, you may need to analyze performance information when you, an application owner or an end user, feels that something has gone wrong with applications performance. In this case, you will need to figure out which parts of your installation may be degrading performance — for example, application programs, data base management systems, host processors, or I/O subsystems.

The Enterprise Tape Library (ETL) is no longer a collection of physical tape drives. This is particularly true with Virtual Tape Server (VTS) or Peer-to-peer VTS. These are a complex of intelligent controller, disk cache, and physical tape drives in an integrated library. Their purpose is to emulate 3490E tape drives at the front end and use reliable, high performance tape drives along with high capacity tape media at the back end. Therefore you need different performance indexes from what you used to have for native tape drives. Just taking the physical specification of tape drives into account will not help monitor/analyze performance.

The IBM TotalStorage ETL Expert can gather performance information for all kinds of ETL, whether they be a 3494 native library, VTS or peer-to-peer VTS or indeed, any combination of the three. This is especially useful for VTS. Since TotalStorage ETL Expert shows you the VTS's activity along with the back end native drive's activity, you can correlate this data to determine how your VTS is working and help in problem determination and capacity planning.

Like ESS, these performance statistics have been available for z/OS as system management facility (SMF) records. An ETL generates performance statistics every hour (SMF Type 94) and z/OS can store this data. However, you need to have a post edit program to create a report from these SMF records. While the ETL has been available to Open Systems and open-z/OS mixed environments, ETL statistics had previously not been available since SMF records are only available for z/OS.

TotalStorage ETL Expert allows you to manage ETL's performance no matter what kind of host platform is used. Furthermore, it can also get more up-to-date status/statistics than SMF records.

Asset management for ETL

An IBM ETL can contain up to three libraries. From a host system's viewpoint, all of these can be seen as a "tape library", but the physical configuration is very different for each. For example, a VTS is comprised of VTS controller, Tape Volume Cache (TVC) and its back end native tape library. You may want to review which VTS uses which native tape library. Or you may want to review how many physical tape drives in a native library are allocated for a VTS's use. The TotalStorage ETL Expert shows you the relationship between libraries and the configuration is made plain and simple.

Capacity management for VTS/ETL

The TotalStorage Expert can show the capacity information for VTS — for example, how much data is active in back-end tape cartridges, and how many empty tape cartridges are available for stacked logical volumes. This information is important not only from a capacity planning perspective, but also from the VTS's performance viewpoint, since the tape cartridge utilization is tightly aligned with the VTS's back-end operation.

Recommended hardware environment for AIX version

This is the recommended hardware:

- pSeries processor.
- 512 MB of RAM.
- 441 MB of disk space for TotalStorage ETL Expert and included prerequisites.
- Additional disk space for TotalStorage ETL Expert history data (depends on number of nodes managed and customer-defined data retention policy). A minimum of 260 MB is recommended.
- CD-ROM drive (for installation).

Recommended hardware environment for Windows NT version

This is the recommended hardware:

- PC, workstation, or server with Pentium® II class or better processor.
- 256 MB RAM.
- 395 MB of disk space for TotalStorage ETL Expert and included prerequisites (total space available on the hard drives may need to be increased if ETL Expert and the prerequisites are split among more than one logical drive).
- Up to 25 MB of disk space on the system drive, usually C:, for new and updated drivers, DLLs, and so on — this number depends on products already installed and the latest level of service pack applied to Windows NT 4.0.
- Additional disk space for TotalStorage ETL Expert history data (depends on number of nodes managed and customer-defined data retention policy). A minimum of 260 MB is recommended.
- Up to 50 MB of temporary disk space on the system drive at installation time for use by the setup utility.
- 512 KB cache.
- CD-ROM drive (for installation).
- VGA display or better.

When the TotalStorage ETL Expert installs, in addition to the Expert code, the installation process installs DB2 UDB, which provides the database storage capability for the asset, capacity, and performance information; the IBM WebSphere® Application Server, which provides the capability for remote viewing and management; and the Sage product from Viador Corp, which provides capability for graphing and for customizing queries. The combined memory requirement for these applications is such that we recommend, in Windows NT environments, that the Expert be run in a server of 256 MB memory in production environments.

8

Upgrade scenarios

In this chapter we discuss a variety of migration scenarios and give you the information you need to migrate your current VTS model (B18, B10, or B20) to a new VTS, to upgrade existing VTS models having 3590 tape drives to 3592 tape drives, as well as attachment of the 3494 VTS to 3584 Tape Library via the IBM TotalStorage 3953 Tape System.

We describe the enhancements which are available with the Model B10 VTS, and Model B20 VTS. We discuss the general considerations and migration scenarios in detail with description of the installation tasks that must be performed to implement a new IBM VTS in your installation when you already have a VTS.

> **Note:** We use the term *library* to refer to the 3494 Tape Library and the 3584/3953 Library system. In most cases what applies to the 3494 also applies to the 3584/3953. When specific differences occur, we explain what each difference is and what the specifics are for each library.
>
> We also use the term *tape drive* to refer to both the 3590 models and 3592 models of tape drives. Where there are differences, they are also specified.

For upgrading your existing 3494 with a VTS, refer to Chapter 3, “Preinstallation planning and sizing” on page 57.

8.1 VTS upgrades

Depending on your current VTS model and features installed, the processing power of VTS shows dramatic variations. Here we describe and compare the components in each of the VTS models. Some enhancements can be implemented to your currently running VTS with minimal changes, but some enhancements require VTS model conversion.

8.1.1 Upgrades available

The Model B18 VTS, Model B10 VTS, and Model B20 VTS are stand-alone frames that can be installed up to 14 m for the SCSI 3590 and up to 25 m for the fibre 3592 attached tape drives installed in the 3494. Installation outside the 3494 allows for concurrent maintenance of the VTS while the 3494 itself can continue to operate with native tape drives or with another VTS. This applies to the 3584/3953 configuration as well. Although in this configuration the Library Manager is also located outside of the automated library. This allows the 3584 to operate autonomously of the VTS or Library Manager when Open Systems hosts are directly connected to the 3584.

The VTS contains the VTS controller and the Tape Volume Cache. Attachment to the hosts is through the VTS controller, using ESCON, FICON, or SCSI. Fibre Channel hosts also can be connected to VTS through IBM 2108 SAN Data Gateway with or without a IBM 2109 Fibre Channel Switch.

VTS controller

The VTS controller in the Model B18 VTS, Model B10 VTS, or Model B20 VTS communicates with the Library Manager through a LAN connection. Table 8-1 compares VTS controllers for the models shown.

Table 8-1 VTS controller comparison

Model B18 VTS	Model B10 VTS	Model B20 VTS
RS/6000 Model H50	RS/6000 Model H80	RS/6000 Model H80
Two way processor[a] 1 GB of memory	Two way processor 2 GB of memory	Four way processor[b] 4 GB of memory
LAN connection to Library Manager	LAN connection to Library Manager	LAN connection to Library Manager

a. With FC5236, additional processor and memory are provided.
b. With FC5250, a six way processor is installed.

FC5250 - FICON performance accelerator feature

Feature code 5250 provides a new six way processor card for the Model B20 VTS. This feature is supported from VTS LIC 2.27.600.0 and above. This upgrade improves the performance of the data transfer to the back-end VTS tape drives and to the host, when FC3415 and/or 3416 are installed.

Tape Volume Cache

The VTS Tape Volume Cache is available in many sizes from 72 GB to 5,184 GB (compressed). The actual usable size varies depending on the compression ratio that you can achieve in your environment. With the ESCON High Performance Option (EHPO), data is compressed at the channel interface before it is written to the TVC. Without the EHPO feature these sizes are respectively 72 GB, 144 GB, 216 GB, or 288 GB in Model B18 VTS. EHPO is standard function with Model B10 VTS and Model B20 VTS.

Table 8-2 summarizes the tape volume cache sizes configurable to each VTS models.

Table 8-2 Tape volume cache capacity (assume 3:1 compression)

Features	Base B18	B18 VTS with FC3200 or 3400	B18 VTS with performance accelerator** (FC5236)	B10 VTS	B20 VTS
(1)3702	72 GB	216 GB	N/A	N/A	N/A
(2)3702	144 GB	432 GB	432 GB	N/A	N/A
(3)3702	216 GB	648 GB	N/A	N/A	N/A
(4)3702	288 GB	864 GB	864 GB	N/A	N/A
(1)3703	72 GB	216 GB	N/A	N/A	N/A
(1)3704	144 GB	432 GB	648 GB	648 GB	N/A
(1)3705	288 GB	864 GB	1296 GB	1296 GB	N/A
(2)3705	N/A	N/A	2592 GB	N/A	2592 GB
(4)3705	N/A	N/A	5184 GB	N/A	5184 GB

**Disk storage configurations with capacity greater then 864GB require installation of FC5236 (performance accelerator feature)

Host attachment

The Model B18 VTS comes with two standard ESCON channels. The Extended Performance ESCON Channels in conjunction with ESCON High Performance Option (EHPO) provide compression of the tape data in the TVC. The data is compressed by the Extended Performance ESCON Channels. EHPO is a standard function for the Model B10 VTS and Model B20 VTS.

A 3494 B10 or B20 VTS can attach to zSeries host systems through FICON channels. As part of the installation plan, you need to analyze performance and distance requirements carefully. The following considerations may be used as a starting point:

- Short wavelength attachments provide for a direct link distance of up to 300 meters between the VTS and the host using 50 micron fiber cables and up to 125 meters using 62.5 micron fiber.
- Long wavelength attachments provide for a direct link distance of up to 10 km between the VTS and the host using 9 micron fiber cables.
- Short and long wavelength attachments provide for up to 100 km distance between the VTS and the host using appropriate switches and Dense Wave Length Division Multiplexors (DWDM).

For ESCON attachments, support is not provided through more than one dynamic switch. For FICON attachments, support is provided through two dynamic switches. For availability or performance reasons, you can install additional ESCON, FICON or SCSI channels. The number of channels installed does not impact any other configuration option. Table 8-3 summarizes host attachments configurable to each VTS models.

Table 8-3 Host attachment comparison

Model B18 VTS	Model B10 VTS	Model B20 VTS
0/ 2 / 4 / 8 ESCON	0 / 2 / 4 ESCON 0 / 2 / 4 FICON 0 / 4 / 8 SCSI	0 / 4 / 8 / 16 ESCON 0 / 4 / 8 FICON 0 / 4 SCSI
0 / 2 / 4 SCSI		

> **Note:** In addition to having a single channel type on a VTS Model B10 or VTS Model B20, you can also mix ESCON and FICON or ESCON and SCSI channels. FICON and SCSI are not supported. If the VTS is attached to a 3584 through a 3953, then the SCSI channel feature is not supported. FICON is only available on VTS models B10 and B20. ESCON and SCSI channels cannot be intermixed on the Model B18 VTS.

Library configuration

A single 3494 can contain up to two VTSs: any two combinations of Model B18 VTS, Model B10 VTS and Model B20 VTS are possible to configure in a single 3494. Each IBM VTS in the library has its own library sequence number, which is also called *library id* and is defined as a separate logical library.

With VTS LIC 2.32.740.x and above, the 3494 VTS Model B10 and B20 can now be connected to a 3584 Tape Library with the addition of the 3953 Library Manager Model L05. The 3953 provides the support for 2 VTSs and support for native zSeries attached 3592-J70 controllers. For further information, see the redbook, *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning, Implementing, and Monitoring,* SG24-6789.

Communication from all model VTS to the Library Manager uses a LAN connection only. In a mixed configuration of a 3494 Tape Library with both VTS attached and native attached control units, it is possible to have a LAN only configured Library Manager. The requirement for the RS-422 connections can be removed if no Open System attached 3590 or 3592 tape drives are included in your configuration or Opens Systems attached 3590 and 3592 tape drives are configured to use the service LAN via the EASH (Ethernet Attached Serial Hub). This feature can eliminate the need for the RS-422 connection in the Library Manager. In addition, both the 3590 A60 and the 3592 J70 can be configured to use this internal service LAN to communicate with the Library Manager.

8.1.2 Physical tape drives

The configuration options for 3590Magstar tape drives in the D12 Frame remain the same when attaching to the Model B18 VTS, Model B10 VTS, or Model B20 VTS.

- The Model B18 VTS and Model B20 VTS support from four to twelve IBM 3590 tape drives. A maximum of six 3590 tape drives can be installed into a single D12 frame. A second adjacent D12 frame must be configured if the VTS is to have more than six 3590 tape drives.
- The Model B10 VTS can support from four to six 3590 tapes drives in a D12 frame or from four to twelve Fibre attached 3592 tape drives within a D22 Frame.
- The B20 VTS can also support from four to twelve Fibre attached 3592 drives within a D22 frame in a *Homogeneous* configuration. The VTS Model B20 can also support a mixture of four to six SCSI attached 3590 drives within one D12 Frame and from four to twelve Fibre attached 3592 tape drives in a D22 Frame. This is a *Heterogeneous* configuration and allows for a maximum of 18 tape drives within two 3494 Dxx Frames. These two frames are no longer required to be adjacent.

Table 8-4 VTS tape drive configuration

	B18	B10	B20 Homogeneous	B20 Heterogeneous
3590 B1A, E1A, H1A	4 to 12 adjacent D12 frame only	4 to 6 single D12 frame	4 to 12 adjacent D12 frame	4 to 6 IBM 3590 single D12 frame *plus* 4 to 12 IBM 3592 single D22 frame
3592 J1A, E05	not supported	4 to 12 single D22 Frame	4 to 12 single D22 frame	

The VTS homogeneous or heterogeneous 3592-J1A configuration is supported with the introduction of VTS R7.3 and above. With VTS R7.4+5, the IBM 3592 Model E05 tape drives are also supported in the VTS. The 3592-E05 runs in J1A Emulation mode and must be installed in the same frames as the Model J1A tape drives. Existing 3592-J1A tape drives cannot be replaced with E05 models, but you can add 3592-E05 tape drives to a new or an existing VTS configuration. Adding the 3592-E05 to an existing configuration is nonconcurrent.

Important: The 3592 tape drive models are not supported on the VTS Model B18.

Restriction: The 3590 tape drive models are not supported when the VTS models B10 or B20 are attached to an IBM 3584 Tape Library. If migration or upgrade of your VTS incudes moving the VTS to another tape library, make sure that the tape drives in the source and in the target library are compatible.

8.1.3 Virtual tape drives

With appropriate TVC capacity and features, the Model B18 VTS can have a maximum of 128 virtual drives. With the Model B20 VTS, this can be increased up to 256. A sufficient number of virtual tape drives provides flexible assignment of virtual tape drives to specific systems and makes management of tape drives simple. Table 8-5 summarizes the number of virtual tape drives supported on each of the VTS models.

Table 8-5 Virtual tape drives comparison

Model B18 VTS	Model B10 VTS	Model B20 VTS
32/64/128 Virtual tape drives (2/4/8 control units)	64 Virtual tape drives (4 control units)	128/256 Virtual tape drives (8/16 control units)

Feature FC5265 provides a way of incrementing the number of virtual devices defined in the VTS. This feature is for the VTS Model B20 only. This feature replaces FC5264. When the B20 has FC4010 installed, 2 FC5265s support 128 devices with 4 VTCs. When FC5265 is not installed on a B20 with FC4010, the PtP supports 128 devices with 8 VTCs.

Where the installation of additional FC5264 was disruptive, the installation of FC5265 is concurrent to the VTS and Library Manager after the first FC5265 is installed. You can have zero, two, or four of these feature codes installed. The default number of devices supported on the VTS Model B20 is 128. This is the case if FC5265 is not installed.

8.1.4 Virtual volumes

Prior to VTS R7.4, you could define up to 250,000 logical volumes for a single VTS Model B18, B10 or B20. If you install two VTSs in a single 3494, you will be able to define up to 500,000 logical volumes in the configuration. With VTS R7.4 and above, the maximum number of logical volumes have been increased to 500,000 for each VTS. This has doubled the maximum number of definable virtual volumes in the Library Manager to 1,000,000.

To increase the logical volumes beyond 250,000, the logical volume expansion feature, FC4036, must be installed. There is a 50,000 logical volume increase for each FC4036 installed. The first FC4036 installation is disruptive to the VTS, but installing an additional FC4036 is concurrent. There is a maximum of five FC4036 expansion features allowed.

Restriction: FC4036 is supported only on VTS Model B20 with 3590-E1A, -H1A and 3592 tapes drives. FC5245, Dual Path Attachment, is a prerequisite for this feature.

Virtual volumes that were created in a Model B18 VTS which has been upgraded to a Model B10 VTS, or Model B20 VTS can be read and appended to in the new VTS. Compression is performed on a record level as the appended data is sent to the TVC. Therefore a single virtual volume can contain compressed and uncompressed records and allows you to use the DISP=MOD to process a virtual volume.

Because of the compression, the data stored on a virtual volume is now comparable to the data stored on a physical 3490 cartridge. With the use of 800 MB ECCST, up to 2.4 GB of data can be stored into one virtual volume. With VTS LIC release 2.32.740.xx, the logical volume size has also increased. Logical volumes of 1000 MB and 2000 MB are supported on the Model B10 and B20. The VTS Model B20 can also support 4000 MB logical volumes.

Restriction: For 4000 MB logical volume sizes, you require a Model B20 with maximum cache size (4 FC3705) and 3592 tape drives.

Advanced Policy Management allows you to take full advantage of all the enhanced features of the VTS. It provides the control you require to manage your VTS no matter what the configuration. Some VTS configurations and features require APM installed before you can exploit these features. This feature provides policy management for tape volume cache residency, physical volume pooling, dual volume copy, copy mode control, and the import/export of logical volumes. The feature code ordered depends on the size of the TVC you have installed in the VTS. FC4001 for up to 250 GB, FC4002 for up to 500 GB, FC4003 for up to 1000 GB, and FC4004 for 2000 GB.

8.1.5 Hardware configuration upgrades

Table 8-6 shows sample model conversion scenarios from B18 to the B20 VTS.

Table 8-6 Model conversion examples

From	Configuration changes	To
3494 B18 VTS FC2711 FC3400 FC3412 x 2 FC3702 x 2 FC4000 FC5234 FC9109	Specify B18 to B20 model conversion Remove FC3400 Remove FC5234 Specify FC3702 to FC3705 feature conversion Specify FC3702 to FC3705 feature conversion Add FC3412 x 2 Add FC2713 Add FC5235 Add FC9021 to L1x Frame Add FC9021 to HA1 Frames (if present) Add FC9011 to D12 Frame	**3494 B20 VTS** FC2711 FC2713 FC3705 x 2 FC3412 x 4 FC4000 FC5235 FC9109
B18 VTS FC2711 FC3400 FC3412 x 2 FC3418 x 2 FC3705 x 2 FC4000 FC5234 FC5236 FC9109	Specify B18 to B20 model conversion Remove FC3400 Remove FC5234 Remove FC5236 Convert FC3418 x 2 to FC3412 x 2 Add FC2713 Add FC5235 Add FC9021 to L1x Frame Add FC9021 to HA1 Frames (if present) Add FC9011 to D12 Frame	**B20 VTS** FC2711 FC2713 FC3705 x 2 FC3412 x 4 FC4000 FC5235 FC9109
B18 VTS FC2711 FC3400 FC3412 x 2 FC3418 x 2 FC3705 x 2 FC4000 FC5234 FC5236 FC5237 FC5264 FC9109	Specify B18 to B20 model conversion Remove FC3400 Remove FC5234 Remove FC5236 Remove FC5237 and D12 Frame FC9012 Add FC3412 x 2 Add FC3418 x 2 Add FC2713 Add FC5235 x 2 Add FC5264	**B20 VTS** FC2711 FC2713 FC3705 x 2 FC3412 x 4 FC3418 x 2 FC4000 FC5235 FC5264 FC9109
B10 VTS FC3705 FC3412 x 2 FC2711 FC2713 FC4000 FC5235 FC9109	Specify B10 to B20 model conversion Add FC3412 x 2 Add FC3705	**B20 VTS** FC2711 FC2713 FC3412 x 4 FC3705 x 2 FC4000 FC5235 FC9109

From	Configuration changes	To
Notes FC2711: Remote Support Switch FC2713: Master Console For Service FC3400: Extended High Performance Option FC3412: Extended Performance ESCON Channels FC3418: Additional ESCON Channels FC3702: Disk Storage Capacity for Model B18 VTS FC3705: 288/432 GB Disk Storage FC4000: Advanced Function (import/Export) FC5234: 18 m SCSI Drive Cables FC5235: 20 m SCSI Drive Cables FC5236: Performance Accelerator FC5237: Additional SCSI Drive Bus FC5264: 64 Additional Virtual Drives FC9109: Attach to S/390, zSeries FC9011: Additional Drives Support FC9021: Virtual Drive Enhancement		

3494 and VTS accessibility during upgrades

Table 8-7 lists the general tasks required during upgrades to the VTS. This applies to both a 3494 and 3584/3953 with a VTS.

Table 8-7 Upgrade tasks

Task	VTS accessibility	Library accessibility
Update HCD	Yes	Yes
Update JES3 definitions	Yes	Yes
IPL / activate IODF	n/a	n/a
Vary virtual drives offline	No	Yes
Migrate using the VTS Upgrade/Conversion Utility Diskette TVC	No	Yes
Copy VTS DB and file fragments	No	Yes
Vary the VTS and LM offline	No	No[a]
Update LIC of the Library Manager	No	No
Update LIC of the VTS	No	No
Install Accelerator Feature	No	Yes
Install additional tape drives	No	No
Convert tape drives to new models	No	Yes
Install new Channels (ESCON,FICON or SCSI)	No	Yes
Install additional TVC	No	Yes

a. If the VTS being upgraded was attached to a 3584 with Open Systems hosts attached, the Library would still be accessible to these hosts. Any native attached zSeries 3592 tape drives would be inaccessible with the 3953 Library Manager offline.

8.1.6 Software implementation

A Migration upgrade of a VTS does not necessarily require much effort in the way of software implementation. However, you should consider whether the following topics apply to your particular VTS:

Additional virtual devices: If your upgraded VTS provides more virtual devices than previous VTS does, you need to update your HCD definitions to define the additional devices to the subsystem. These changes can be done in advance of the hardware migration. In addition, when you are using the LIBRARY-ID and LIBPORT-ID parameters in the HCD definition as described in 4.2.5, “HCD support for library and port IDs” on page 121, you need to define additional LIBPORT-IDs. If you are using JES3, you need to update the JES3 definitions to add the new devices. This change can be done in advance and be activated during a planned IPL.

Additional ESCON or FICON Channels: If you add additional ESCON or FICON channels, or if your upgraded VTS has more ESCON or FICON channels than your previous VTS, you need to update the HCD definitions. You should do this in advance of the hardware migration.

Compression: To benefit from the compression of the data in the TVC, make sure that the SMS data classes that you assign to tape data sets written to the VTS have defined COMPRESSION=YES. If you define COMPRESSION=NO, you switch off compression when data is written in the TVC. This wastes space in the TVC and may decrease the overall VTS performance.

Library ID and library name: One physical 3494 can consist of up to three logical libraries: one logical library per VTS and one library for all native tape drives in the 3494. Each of these logical libraries has its own library sequence number or library ID, which the IBM service representative defines during hardware installation. Usually, the last five characters of the machine serial number are used. When installing a zSeries attached 3584, the 3953 requires the same consideration. The 3953 Library Manager supports the same configuration as a 3494 Library Manager, a maximum of 2 VTS libraries and a native library. Therefore a library ID is required for each of these three logical library partitions.

The library ID is used when defining the system-managed library in the ISMF Library Define panel. When you convert from one model of VTS to another, the library ID remains the same. Therefore no software changes are required for migration in place as in migration scenario 1.

When you add a VTS to an existing library or when you install an additional 3494 or 3584/3953 including a VTS, the new VTS logical libraries will require a new library ID. You need to define a new logical library through ISMF. You also need to update your SMS constructs.

Import/Export: If you want to implement the Import/Export function, software updates are required. You must install PTFs that provide the support for the new functions. To implement Import/Export, additional definitions are required in your tape management system. We describe planning and implementation information for the Import/Export function in Appendix B, “VTS Import / Export” on page 369.

Open Systems Attachment: If you plan to attach Open Systems to the VTS, refer to 4.8, “VTS Open System attachment” on page 167.

8.1.7 Volume and data migration

There are situations where moving data that has been created on logical volumes in your current VTS is required. In some cases this might be the only option. The logical volumes physically reside on tape in the automated library connected to the VTS. In some of the migration scenarios, after the installation of the new VTS model, you can immediately continue your normal processing. The circumstances are dependent on what upgrade is being undertaken and which scenario you choose.

Depending on options installed on the new VTS, logical volume capacity has be increased. All volumes are compressed in the TVC, so the capacity of the volumes effectively increases from 400 MB for MEDIA1 or 800 MB for MEDIA2 to up to 1.2 GB or 2.4 GB depending on the compression ratio achieved. New larger logical volumes were introduced with VTS LIC 2.32.740.0, 1000 MB, 2000 MB and 4000 MB. When using these larger logical volumes compressed capacity could be as high as 12 GB.

Important: There are no new media types with the introduction of the new larger logical volume sizes of 1000 MB, 2000 MB, and 4000 MB. There are still only MEDIA1 and MEDIA2 type logical volumes. To use the new larger volume sizes, the correct data class must be specified within DFSMS and the corresponding data class name must be configured correctly in the Library Manager. To use this feature, the VTS also requires the Advanced Policy Management feature to be installed.

8.1.8 VTS LIC level considerations

Data migrations involve transferring VTS data, LM data, and physical tapes from one (source) library to another (target) library.

Currently, the only restrictions on data migrations are that the code on the target library has to be equal to or greater than the source library and that the source VTS cannot have GPFS when the target VTS has JFS. Introduced with the Denali (LIC 2.26.xx.x) level of VTS microcode and its pooling database structure, migrating from non-pooling to a pooling environment requires an VTS database conversion. Also, copying from a pooling library to another pooling library may not always be straightforward anymore. The source library could have pooling enabled (FC4001 installed) and the target library could have pooling disabled (an unsupported scenario) or vice versa. The following sections will give an overview of each of the scenarios described in Table 8-9.

The following migration restrictions apply to the 527 LM and 2.26 VTS LIC levels:

1. Migrating from LIC 2.26.x or above to 2.23.31.88 or lower
2. Migrating from a source VTS with FC4001 installed to a target without FC4001 installed
3. Migrating from a VTS that has E1A drives to a VTS with B1A drives or H1A to E1A or B1A
4. A general restriction is that the target VTS must have a code level equal to or greater than the source VTS.

New code has been added to enforce these restrictions.

Additional considerations for levels above 527

In this section we list the prerequisites for Library Manager, VTS, and tape drives.

Library Manager prerequisites

If the source or target VTS has FC4001 installed, both Library Managers must be at (529.10A or higher). **The target LM code must be at the same as the source LM code.**

VTS prerequisites

VTS general prerequisites: FC4001

- If the source VTS has FC4001 installed and the target VTS does not have FC4001 installed, FC4001 must be installed on the target VTS before continuing with this procedure.

VTS B18 ONLY prerequisites: FC5236

- If the target VTS does not have FC5236, the source VTS must not have FC5236.
- FC5236 is standard on models B10 and B20, and therefore will not be displayed in the machine VPD.

VTS B20 ONLY prerequisites: FC4036

- If the source VTS has FC4036 installed and the target VTS does not have FC4036 installed, FC4036 must be installed on the target VTS before continuing with this procedure.

Tape drive prerequisites

Refer to Table 8-8 for the detailed prerequisites for tape drives.

Table 8-8 Tape Drive prerequisites.

Source VTS tape drives	Target VTS tape drives
3590-B1A	3590-B1A or E1A, H1A with or without 3592
3590-E1A	3590- E1A, H1A with or without 3592
3590-H1A	3590-H1A with or without 3592
3592	3592 with or without 3590 drives
3592 and 3590-B1A	3592 and 3590-B1A, E1A or H1A
3592 and 3590-E1A	3592 and 3590-E1A or H1A
3592 and 3590-H1A	3592 and 3590-H1A

Data migration scenarios

Table 8-8 shows some possible migration scenarios. Only the VTS LIC levels are described.

Table 8-9 Migration scenarios

Number	Source VTS	Source FC4001	Target VTS	Target VTS FC4001
1	pre 2.26VTS	NA	pre 2.26 VTS	NA
2	pre 2.26 VTS	NA	2.26 or greater	No
3	pre 2.26 VTS	NA	2.26 or greater	Yes
4	2.26 or greater	No	2.26 or greater	No
5	2.26 or greater	No	2.26 or greater	Yes
6[1]	2.26 or greater	Yes	2.26 or greater	No
7	2.26 or greater	Yes	2.26 or greater	Yes
8[1]	2.26 or greater	No	pre 2.26 VTS	NA
9[1]	2.26 or greater	Yes	pre 2.26 VTS	NA

Number	Source VTS	Source FC4001	Target VTS	Target VTS FC4001
Notes: NA = Not Applicable [1] Not allowed in current configurations.				

Scenario 1: Pre-2.26 VTS to pre-2.26 VTS

Migrate process: No changes (Already exists today)
Restore process: No changes (Already exists today)

Scenario 2: Pre-2.26 VTS to 2.26+ VTS/No Pooling

Migrate process: No changes (Already exists today)
Restore process: Will build a new token DB, if one does not exist and will build the new ADSM DB structure to support FC4001 to save time by avoiding an extra code activate

Scenario 3: Pre-2.26 VTS to 2.26+ VTS/Pooling

Migrate process: No changes (Already exists today)
Restore process: Will build a new token DB, if one does not exist, and will build the new ADSM DB structure to support FC4001 to save time by avoiding an extra code activate

Scenario 4: 2.26+ VTS/No Pooling to 2.26+ VTS/No Pooling

Migrate process: Uses a new online migration process
Restore process: Verifies that both VTSs are Denali and continues current restore (exists today)

Scenario 5: 2.26+ VTS/No Pooling to 2.26+ VTS/Pooling

Migrate process Uses a new online migration process
Restore process Verifies that both VTSs are Denali and continues current restore (exists today)
Prerequisite: Enable Pooling at the target LM before going online

Scenario 6: 2.26+ VTS/Pooling to 2.26 VTS/No Pooling

Migrate process: Uses a new online migration process
Restore process: Verifies that both VTSs are Denali and continues current restore (exists today)
Prerequisites: Enable Pooling on target LM as a prerequisite and enable Pooling on target VTS via FC4001 as a prerequisite

Scenario 7: 2.26+ VTS/Pooling to 2.26+ VTS/Pooling

Migrate process: Uses a new online migration process
Restore process: Verifies that both VTSs are Denali and continues current restore (exists today)
Target VTS Pooling state:Remains Pooling
Target LM Pooling state:Remains Pooling
Recommended Installation Instruction modification:None

Scenario 8: 2.26+ VTS/No Pooling to Pre 2.26 VTS

Migrate process: Uses a new online migration process
Restore process: Verifies that target VTS is pre-Denali when source was Denali
Prerequisites Update LM and VTS to Denali as a prerequisite

8.1.9 3494 Frame upgrades

Different considerations apply depending on the migration scenario that you select. To convert an existing VTS into a new model VTS, an MES must be ordered for the frame model conversion. To install the new VTS into an existing library in addition to the current VTS, you need to order an MES including the additional frames to be installed, and may include some additional features on the L1x Frame. If you are installing a completely new library including a VTS, see Appendix G, "VTS feature codes" on page 499 for details.

The Library Manager LIC level to support the Model B18 VTS must be 523 or higher. The Library Manager LIC for supporting the Model B10 VTS or Model B20 VTS is 527.08 or higher. If the new VTS is to be installed in an existing 3494, that does not already have the Enhanced Library Manager feature (FC5045) or the IBM 3494-HA1 frames installed, the new Library Manager hardware is automatically provided as part of the MES. FC5045 is the minimum requirement and provides the Library Manager with a 400 MHZ processor and 256MB of memory. There are FC5046 (PCI Library Manager), or FC5047 (LAN PCI Library Manager) also available.

Additional Storage Unit Frame (S10 Frame): This frame contains up to 400 cartridge storage cells only. No tape drive can be installed into this frame. It is physically shorter than the Dxx frames and this makes it impossible to use this frame for anything but storage.

Drive Unit for VTS Frame (D12 Frame): This frame contains up to 290 cartridge storage cells and up to six Magstar tape drives. It cannot be used like a D12 Frame frame for Magstar tape drives attached through SCSI to its hosts; it can be attached to a Model B10, B18, or B20 VTS frame only. All the 3590 Tape drives associated with the D12 frame must be of the same type (either all 3590-B1A, 3590-E1A or 3590 H1A). There can be from zero to six 3590 tape drives in a D12 Frame. When no tape drives are installed in a D12 Frame then there is a maximum storage capacity of 400 cartridges.

SCSI Drive Unit Frame (D12 Frame): The SCSI drive unit frame contains additional cartridge storage (up to 290 cartridges) and provides the necessary hardware for installation of three or four IBM TotalStorage 3590 tape drives that will be attached to host system SCSI adapters. You need to order feature code 5503 on the L1x Frame for this conversion. If you need to install five or six Magstar tape drives, you need to order RPQ 8B3167 in addition.

ESCON Drive Unit Frame (D14 Frame): The ESCON drive unit frame contains additional cartridge storage (up to 345 cartridges) and provides the necessary hardware for installation of IBM TotalStorage 3590 tape drives and one 3590 Controller that will be attached to host system ESCON adapters.

3592 Drive Unit Frame (D22 Frame): The Library Frame D22 for the 3494 supports the 3592 tape drive. The Base configuration of no tape drive units and 400 tape cartridge storage cells can be expanded to a total of twelve 3592 tape drives and 250 tape storage cells. The D22 Frame may be attached with FC9010 to an IBM 3494 VTS. When so attached, the D22 Frame with four through twelve 3592 tape drives, can be integrated with a 3494 VTS to provide an ESCON or FICON attached VTS system.

8.1.10 IBM 3584/3953 Tape Library configuration

When installing a VTS onto an IBM 3584 Tape Library with the IBM 3953 Library Manager, you also require an IBM 3953 Model F05 Tape Frame along with the new VTS. In the 3584 you also need support for the 3592 tape drives. See the *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning, Implementation and Monitoring*, SG24-6789 for more concise information.

IBM 3953-F05 Frame

This frame's configuration can contain up to the following:

- Two 3953 Library Managers Model L05
- One 3592 Tape Controller Model J70 (used for native attached 3592 to a zSeries host)
- IBM TotalStorage Master Console (MC)
- Ethernet Switches for communications between the 3953 Library Manager, 3592 Tape Controller, VTS, and the IBM TotalStorage Master Console
- Two Fibre Channel Switches for connection of 3592 tape drives to the IBM 3592 Tape Controllers Model J70
- Four Fibre Channel Switches for VTS and 3592 communications (two per VTS)
- Power for all components including support for dual line cords
- Monitor and keyboard/mouse including a monitor/keyboard/mouse switch to allow the single monitor/keyboard/mouse to be switched between two 3953 Library Managers in a rack and the Master Console, if present in the rack

IBM 3584-L22 (Base Frame)

The 3584 L22 Base Frame is designed for 3592 tape drive models J1A and E05, and for 3592 data cartridges. The 3584 L22 Base Frame has 58 to 260 cartridge slots and support for up to 12 tape drives with an incremental reduction of storage slots for more than four drives or with the additional I/O station installed.This frame has a smaller footprint than the 3584 L32 frame. The 3584-L22 frame is designed with an optimized gripper for use with LTO or 3592 tape cartridges. Up to 12 logical libraries (one per tape drive) can be configured for each frame.

Each 3584-L22 frame has a standard 16-slot cartridge I/O station for importing or exporting 3592 tape cartridges from the library without requiring a re-inventory. An additional 16-slot cartridge I/O is optionally available for either LTO or 3592 data cartridges. Libraries containing a mixture of LTO and 3592 drive technologies must have one LTO I/O station and one 3592 I/O station. For bulk loading of tape cartridges, the library door can be opened. Each time the library door is closed, a bar code reader mounted on the auto changer scans the cartridge labels enabling a re-inventory of the cartridges in the library frame in typically less than 60 seconds. A door lock is included to restrict physical access to cartridges in the library.

IBM 3584-D22 Frame

The 3584 D22 Expansion Frame is designed for 3592 Model Tape Drives and 3592 data cartridges. Up to 15 Model D22 Expansion Frames may be added to the 3584 Model L22, L32, or L52 Base Frame to increase 3592 cartridge storage or drive capacity. Each Model D22 supports up to 400 3592 cartridge slots and up to 12 3592 Tape Drives, with an incremental reduction of storage slots for each set of four tape drives installed. Each frame can have up to 12 logical libraries or 12 control paths (one per tape drive). This frame can optionally be configured as service bay B.

8.2 Planning for VTS model upgrades

In this section we describe general migration considerations and various upgrade paths from your current VTS configuration to a VTS configuration that meets your requirements. Figure 8-1 shows possible upgrade paths between VTS models.

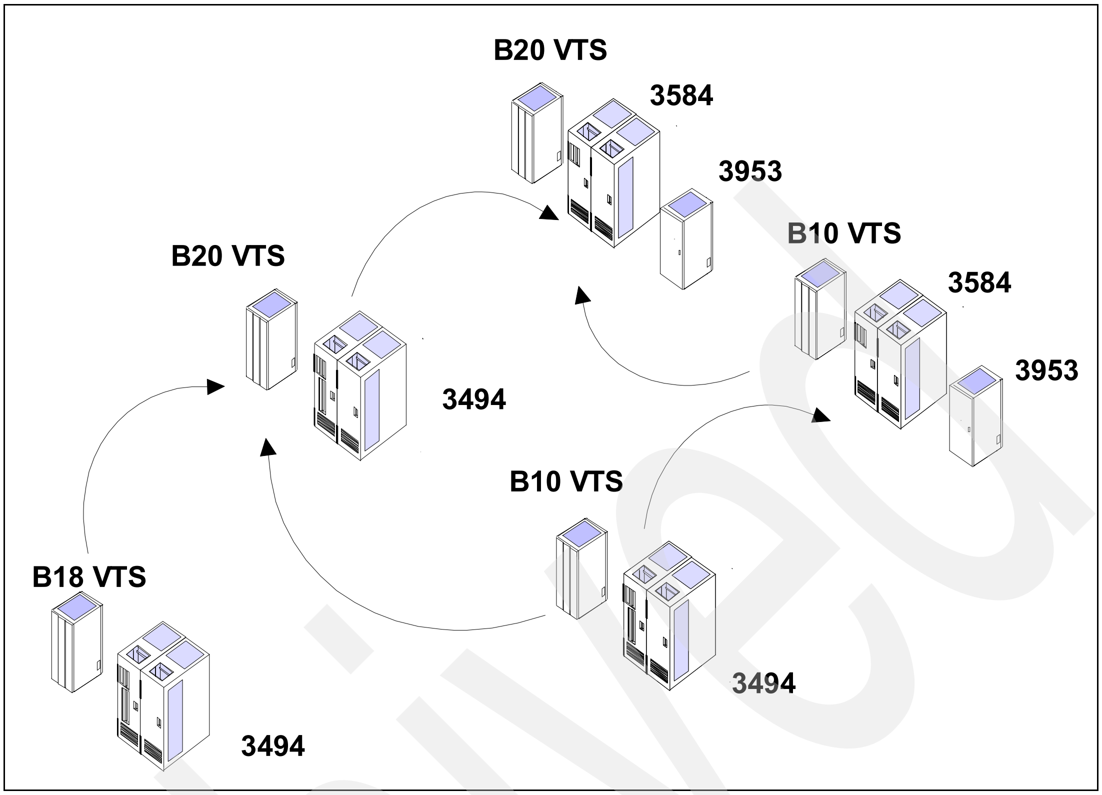

Figure 8-1 VTS upgrade paths

Migrating an existing VTS to a new VTS

> **Attention:** These migration scenarios can apply to the 3584/3953 Tape Library. But there may be different prerequisites and limitations applying specifically to this upgrade path that do not apply to a 3494 Tape Library. See *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning, Implementing, and Monitoring,* SG24-6789,and ask your local IBM representative for approved migration scenarios when migrating to the 3584/3953 Tape Library configuration.

The scenarios for migration of an existing VTS to a new VTS are as follows:

- You can add a new VTS to an existing library containing the current VTS. This will also require an additional drive frame or an upgrade to an existing frame to house the required tape drives. During this upgrade, both VTSs are used in parallel, where new logical volumes are written to the new VTS and logical volumes existing in the current VTS are processed from there. After data migration is complete, the current VTS is either kept and redeployed for other work or it is removed from the library and its associated drive frames MES'ed to remove the tape drives.

  This scenario is referred to as *parallel operation within one library*. For details refer to "Scenario 1: Parallel operation within one library" on page 344.

- You can install an additional 3494 with a new VTS. Both libraries are used in parallel until migration is complete.

  We call this scenario *parallel operation with two libraries*. For details, refer to "Scenario 2: Parallel operation with two libraries" on page 349.

Table 8-10 lists the advantages and drawbacks of the two migration scenarios.

Table 8-10 Migration scenario comparison

	Advantages	**Drawbacks**
Scenario 1 Parallel operation within one library	No transition hardware to be installed Excellent testing possible Good fallback Preferable solution if target configuration with two VTS	Customer involvement Data movement Additional software definitions needed (HCD, DFSMS)
Scenario 2 Parallel operation with two libraries	Limited outage time Extended testing possible Excellent fallback possible	Transition hardware to be installed/removed Customer involvement Data movement Additional software definitions needed (HCD, DFSMS)

- Scenario 1 requires slightly more outage time, and requires about the same changes to your software definitions. You might want to choose Scenario 1 if you plan for two VTS in a single library.
- Scenario 2 requires the installation of an additional library. You might want to choose this scenario if your installation does not allow any downtime or if your target environment is to be expanded while keeping the original VTS for specific workloads.

8.3 Scenario 1: Parallel operation within one library

Scenario 1 explains migration within one library using parallel operation. Figure 8-2 shows the source and target configuration when you migrate to an new VTS using a parallel operation within one 3494 Tape Library.

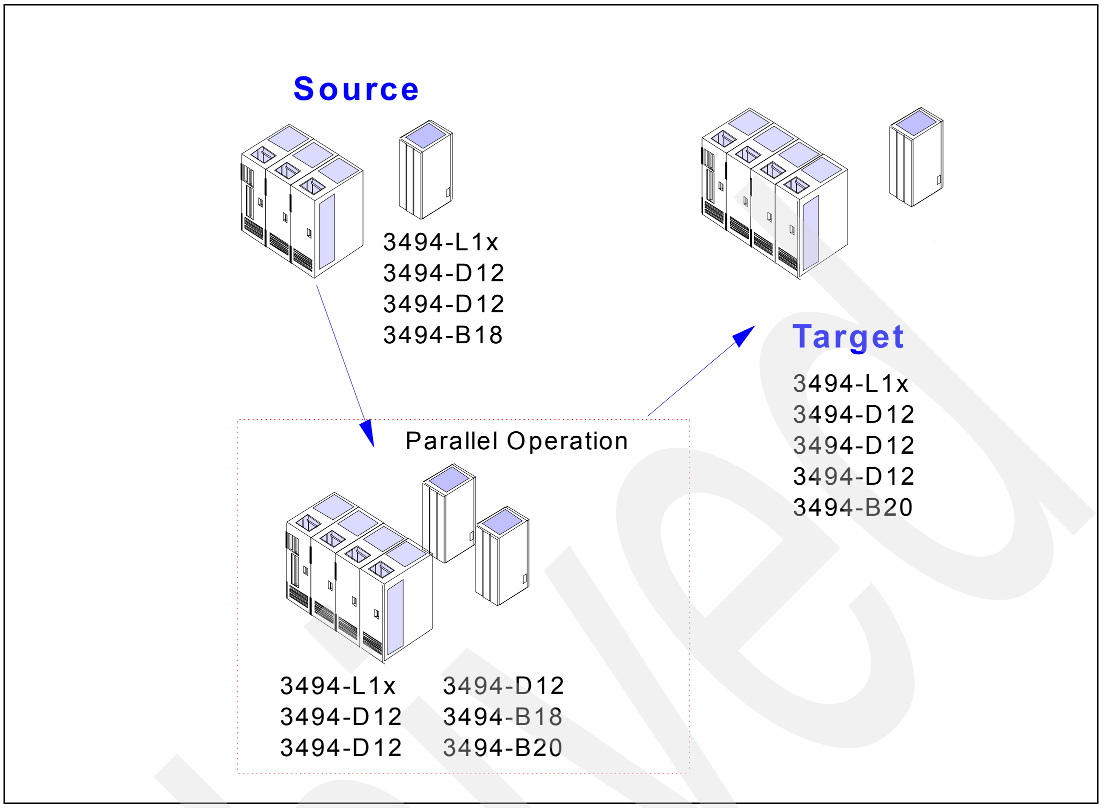

Figure 8-2 Migration scenario 1: parallel operation within one library

In migration Scenario 1, you add an additional new VTS (Model B20 in the figure) to the source 3494 Tape Library. This new VTS will become VTS2 in the 3494 configuration. Because you create a new logical library within the same physical tape library, the software implementation steps are the same as if you were installing a complete new library. After installation of the new VTS, you start moving your data from the source VTS into the new VTS. This can be achieved by either copying the data from the source VTS to the new VTS, writing new data to the new VTS and waiting for the data in the source VTS to expire, or using a combination of both.

After you have completely migrated your data to the new VTS, you may choose what to do with the source VTS. Because this migration scenario does not include a model conversion from a source VTS to an new VTS. Therefore we recommend this migration scenario if you are planning to add a second VTS to the same 3494 which you already have installed. You may upgrade the original VTS to an new VTS any time later, if your throughput and processing requirements require it.

You may want to use the source VTS for other workload, for example, for the attachment to a test LPAR, or you may want to move workloads to the source VTS that do not benefit from the EHPO compression, for example, data sets that are already host compressed.

8.3.1 Preparing for migration

In preparing for the new VTS hardware installation, you should plan for hardware and software changes.

Hardware installation

Perform the following tasks before the hardware installation:

1. Decide on the configuration of the new VTS regarding:
   - The number of 3590 or 3592 tape drives.
   - The number and type of host channel attachments.
   - The size of the TVC.
   - Plan for a sufficient number of storage cells to store the new stacked cartridges required to support the new logical library.
2. Prepare and complete the physical installation planning, including the floor plan.
3. Decide on the volser range for logical volumes for use in the new VTS.
4. Order additional cartridges and labels.

During migration, additional stacked cartridges may be required to hold the logical volumes of the additional VTS. You may want to use a volser range adjacent to the one you are using for the source VTS. If the new VTS is to use 3592 tape drives, then new media as well as a new volser range is required.

For additional information about hardware planning and preparation, see Chapter 3, “Preinstallation planning and sizing” on page 57.

Software-related considerations

To the host, the implementation of a second VTS appears as if a new library were being added, because every VTS represents a logical library. Therefore, there is no difference in the software implementation when you install a completely new library (3494 or 3584/3953) including the VTS, or when you add a VTS to an existing library. We summarize the steps required for implementation in a zSeries environment and refer you to the appropriate sections in Chapter 4, “Implementation” on page 107 for details. For implementation in a VM/ESA or VSE/ESA environment, refer to “DFSMS/VM” on page 161.

1. **Define the hardware configuration:**

   See 4.2, “Hardware I/O configuration definition” on page 111 for details of HCD definitions for a new library and for PARMLIB updates such as setting the MIH value for the new virtual tape drives.

2. **If applicable, update the JES3 definitions:**

   See “JES3 sample initialization deck definition” on page 156 for a sample JES3 initialization deck definition.

3. **Identify the data to be moved to the new VTS:**

   If you are planning to move all workload to the new VTS, you need to change the ACS routines to direct the complete VTS workload to the new VTS. You may, however, want to leave part of the workload in the source VTS, for example, if tape data is already compressed on the host. Identify the tape workload that you want to move from the source VTS and from native tape drives to the new VTS.

4. **Define a data migration plan:**

   You can either copy your tape data to the new VTS, or you can write new data to the new VTS and use the source VTS to read existing data, or you can use a mixture of both strategies. For details on the methods to move data, see:

   - Section 5.7, “Moving data into and out of VTS” on page 196
   - Section 5.7.1, “Moving data into the VTS: phased method” on page 197
   - Section 5.7.2, “Moving data into the VTS: quick method” on page 198
   - Section 5.7.5, “Combining methods to move data into the VTS” on page 203

Develop a data migration plan depending on the method you have selected and on the time you can spend on moving the data.

5. **Prepare the SMS definitions:**
   - Define the new library through ISMF.
   - Define new data classes, if required. Make sure that the compression parameter in the data class definition is set to YES for data to be written to the new VTS. You can also use existing data class definitions that have been used for the source VTS.
   - Define a new storage class, if required. Because the storage class in a system-managed tape environment is used only to determine whether the tape volume is to be system-managed or not, you can assign any storage class you want. If you use the storage class name for further filtering in the ACS routines, you may want to define a new one.
   - Define or update one or more storage group related to the new library. Because the storage group is used to direct allocations to a specific SMS-managed tape library, you should define at least one storage group that is connected only to the new VTS library.
   - Prepare the update of the ACS routines. You have to update the ACS routines to direct tape allocations to the new VTS library. Depending on the test plan and data migration plan, you have to update the ACS routines more than once. Be sure that the new ACS routines are activated only after the new VTS is online, otherwise tape allocations to this new library will be unsuccessful.
6. **Update your tape management system:**

   Define the logical volser range to your tape management system as scratch volumes. Although you do not have to predefine scratch volumes with DFSMSrmm, we recommend predefining the volsers. If the volsers are predefined, INSERT processing on the host will be performed much faster after the logical volser range has been defined on the Library Manager.
7. **Define a test plan:**

   Before moving workload to the new VTS, you may want to test it in your environment. Set up a test plan including typical processing in your environment. If you are planning to move additional workload from native tape drives to the VTS, you should also include this in your test plan.

Plan for having these steps completed before the hardware installation of the new VTS. We recommend that you perform all updates in advance. If you prefer, you can also perform the updates after the hardware has been installed.

8.3.2 Migration steps

In this section we summarize the migration steps for hardware installation and software implementation of the new logical library.

Upgrade the Library Manager

Depending on the model of your existing 3494 Tape Library, you may be required to upgrade the Library Manager.

Install the new VTS subsystem

For 3590 drives, an additional D12 Frame has to be added to your existing 3494, or an existing D12 Frame can be used for attachment of the new VTS. During the installation of the D12 Frame, the 3494 and the source VTS are unavailable for the hosts. If you intend to use 3592 drives for the VTS, a D22 frame must be installed or an existing Dxx frame can be converted to a D22, or an Open Systems attached D22 Frame can be changed to VTS attached.

During installation of the new VTS drive frame, the IBM service representative performs these activities:

- If converting an existing D12 Frame, install the SCSI extender in the D12 Frame.
- If a new D12 Frame is added to your configuration, the SCSI extender will be preinstalled.
- If you are using 3592 drives, install a new D22 Frame or upgrade an existing Dxx Frame
- Install the new VTS and cable it to the SCSI extender or install the fibre cables to D22 Frame.
- Install the appropriate tape drives (3590 or 3592)
- Install the attachment concentrator and cables from the new VTS to the L1x Frame.
- Teach as the new configuration.

The IBM service representative performs a TEACH CURRENT operation and a subsequent inventory and tests the new VTS.

Some of the above activities are concurrent to the use of the other VTS and any native tape drives (if installed) in the library.

Setup of the New VTS subsystem

After completing the hardware installation, you have to set up the new VTS. On the Library Manager console, you need to:

- Define an additional volser range for stacked volumes.
- Define the volser range for logical volumes.
- Set the VTS category attributes.
- Insert the volser range for logical volumes.
- Define the VTS management policy.

See 4.3, “VTS definition from the Library Manager” on page 127 for details on these implementation tasks. On completion of the Library Manager definitions, you should:

- Insert additional stacked TotalStorage 3590 or 3592 cartridges.
- Reinventory the 3494.
- Vary online the new VTS logical library and the logical drives
- Activate the ACS routines that allow you to direct test allocations to the new VTS logical library.
- Start testing the new VTS.

Migrate the tape workload

On completion of your testing of the new VTS, start migrating the target workload into the new VTS. To do so, you just need to change the ACS routines:

- For a phased migration, use existing filter lists and change the storage group assignment in the storage group ACS routine per group of data sets or per application at the pace you want to migrate your tape data to the new VTS.

- For a quick migration, change all storage group assignments for VTS candidates to point to the new VTS. Use the tools described in 5.7.3, “Products to simplify the task” on page 201 to copy logical volumes residing in the source VTS.
- For a partial migration to balance the tape workload between both VTSs, redirect tape allocations as described before for the data sets identified in your migration plan and copy tape data as required.

After you have completed your migration, you may have obsolete logical volumes in the source VTS. To free up space on the stacked TotalStorage 3590 cartridges, consider ejecting logical volumes from the source VTS. For details on ejecting logical volumes, see 6.7, “Ejecting logical volumes from the VTS” on page 248. Eventually, you may want to reassign stacked volume ranges to the new VTS as data is transferred away from the source VTS.

8.3.3 Migration summary

Table 8-11 summarizes the installation tasks of migration scenario 1. Note that these tasks may vary depending on the actual configuration.

Table 8-11 Migration tasks with scenario 1

Task	VTS accessibility	3494 accessibility
Upgrade LM	No (Yes with HA1 Frames)	No (Yes with HA1 Frames)
Add D12 or D22 Frame	No	No
Convert Dxx Frame to support Drives	No	No
Install new VTS[a]	Yes	Yes

a. The last part of the VTS installation is a Teach of the Library Manager. This is not concurrent to the 3494 Tape Library nor the current VTS.

8.4 Scenario 2: Parallel operation with two libraries

Scenario 2 explains the migration process using parallel operation with two tape libraries. When you install an additional library (3494 or 3584/3953) including the VTS, you do not experience downtime during the hardware installation and migration path. As outlined in Figure 8-3, the target configuration can include completely removing the 3494 with the source VTS or keeping it for other workload. Or you may also want to upgrade the source VTS to a new VTS after you have completed the migration.

Figure 8-3 shows the source and target configuration when you migrate to a Model B20 VTS using a parallel operation with two libraries.

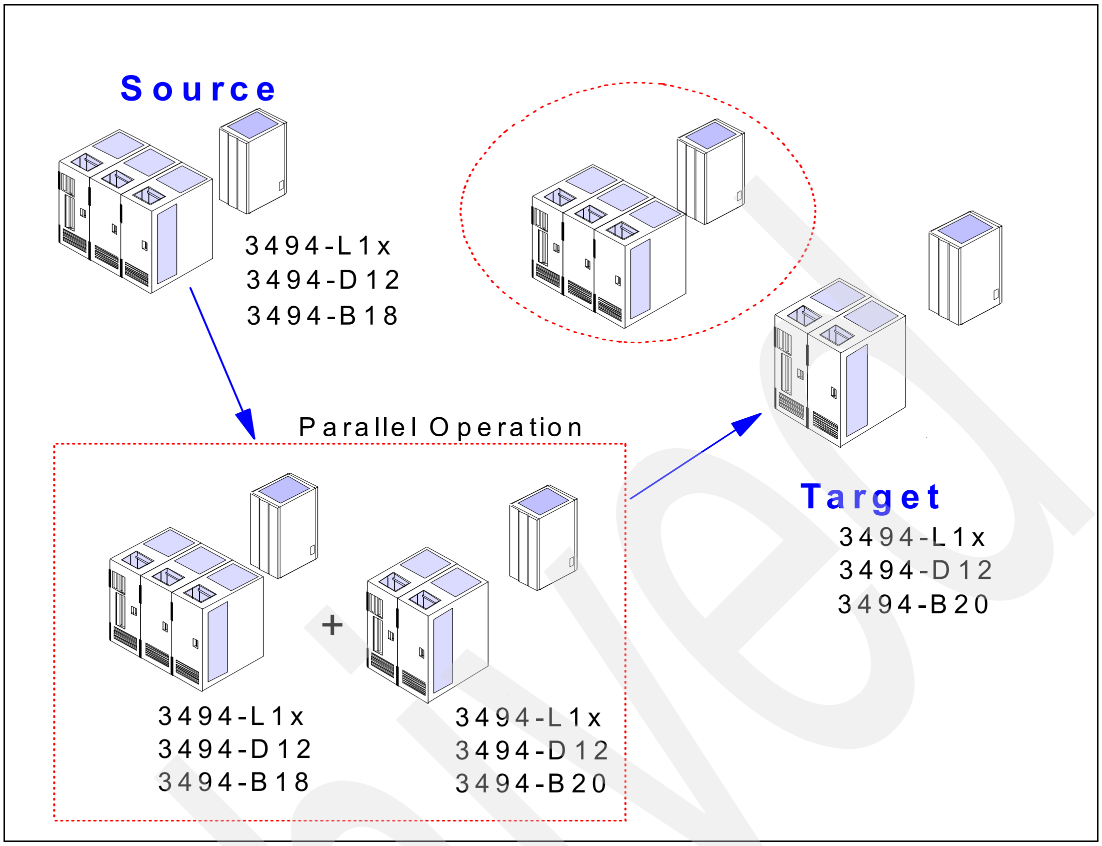

Figure 8-3 Migration scenario 2: parallel operation with two libraries

Depending on your target configuration, you can choose between two different ways of migrating your workload to the target VTS:

- **Host controlled data migration:**

  This migration method is comparable to the migration method described in "Scenario 1: Parallel operation within one library" on page 344 and is not disruptive to your tape processing. It requires that the target library contain additional stacked cartridges, because you cannot use the physical cartridges in the old VTS at the same time as the new VTS uses them. Host controlled data migration can be used independently of the physical configuration of your source and target libraries and independently of whether you plan to keep both VTS or plan to de-install it.

  The method of moving your data under the control of the hosts from one VTS to another VTS is the same for two VTSs in the same library and for two VTSs in two different libraries. Therefore we do not repeat the data migration steps in this section. See "Software-related considerations" on page 346 and "Setup of the New VTS subsystem" on page 348 for details.

> **Important:** Migration through host software is the only supported migration path from a VTS on a 3494 to a new VTS on a 3584/3953 at this time.

- **Physical cartridge movement:**

  This migration method is disruptive to your tape processing and is largely performed by the CE. You must remove the stacked TotalStorage 3590 cartridges from the 3494 containing the old VTS and placed in the 3494 containing the new VTS. The CE creates a backup of the source Library Manager's database and restores it on the Library Manager of the target library. The CE also backs up the old VTS databases and file fragments in the TVC are restores them on the target VTS. This procedure is not generally applicable because the physical configuration of source and target 3494 may vary. You need to check with your IBM representative to determine whether you can use this method.

For details see “Migration steps for physical cartridge movement” on page 351.

8.4.1 Migration steps for physical cartridge movement

Figure 8-4 shows the source and target configuration when you migrate to a new VTS using physical cartridge movement.

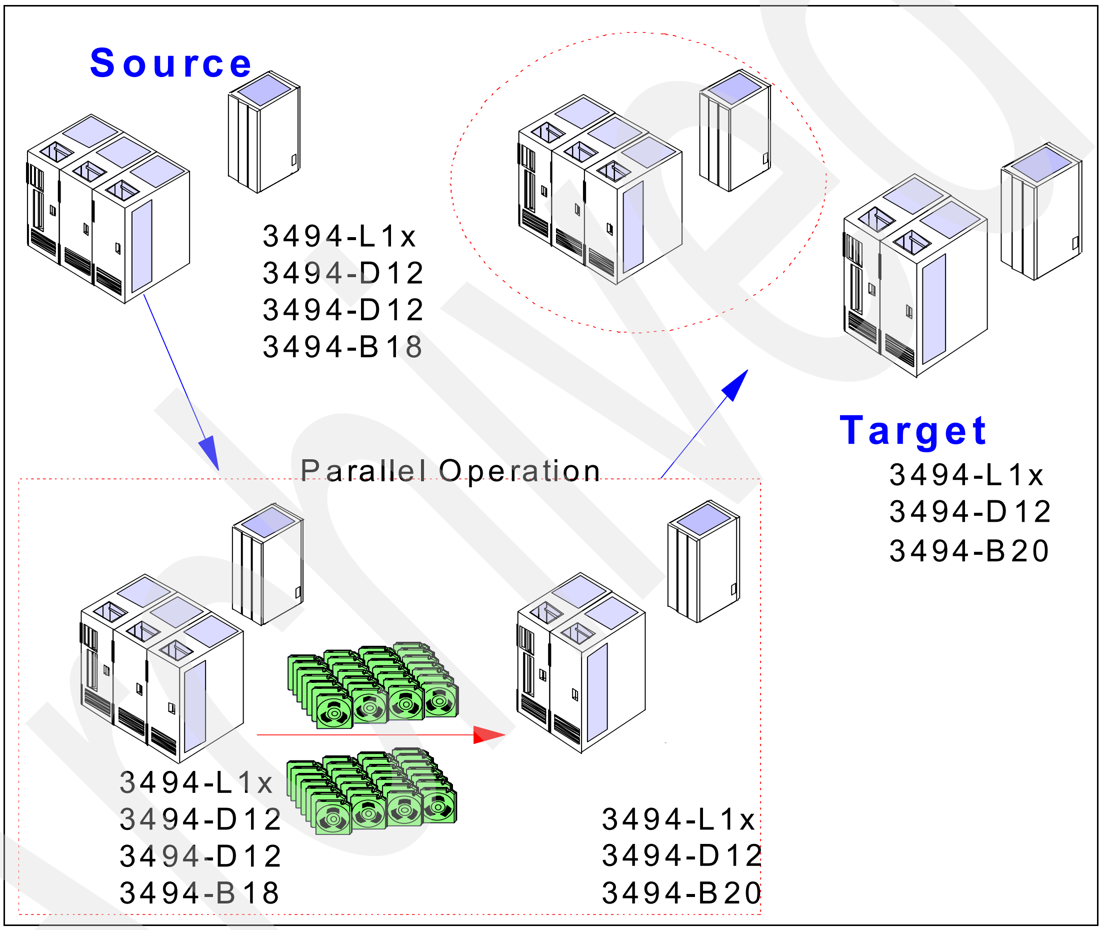

Figure 8-4 Physical cartridge movement

- **Install an additional tape library:**
  - There must be at least one L1x Frame, one Dx2 frame and one new VTS frame.
  - If the new tape library is a 3584/3953 configuration, define the logical library. See *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning and Implementation,* SG24-6789 for details.
  - Define the HCD definitions for the new library.
  - Define the new library through ISMF.
  - Define or update SMS constructs for the new library, especially storage groups.
  - Prepare ACS routine changes.
  - If applicable, update JES3 definitions.
  - Vary the new Model B20 VTS online.

Some of the subsequent tasks are done by the IBM service representative using specialized tools.

- **Move the content of the source VTS to the target VTS library:**
  - Vary source VTS offline to all hosts.
  - Initialize concurrent Migrate using the VTS Upgrade/Conversion Utility Diskette to drain the TVC.
  - Vary source VTS offline to the Library Manager.
  - Copy all TVC resident logical volumes to stacked cartridges using the Migrate VTS Upgrade/Conversion Utility Diskette function.
  - Copy TVC fragments and VTS database to service cartridge (3590 or 3592).
  - Copy Library Manager database to diskette.
  - Move and load stacked cartridges into the target VTS.
  - This might involve physical transportation of data if the new VTS is in a different location. The duration and difficulty of this step should not be underestimated.
  - Update new Library Manager database from originating source Library Manager.
  - Restore VTS database and TVC fragments.
  - Reinventory the library for physical cartridges, keeping the logical volumes.
  - Vary target VTS online.
  - Update TCDB and tape management system CDS regarding location of the logical volumes if needed.
- Redirect the workload to the target VTS library, using nondisruptive adjustments in the DFSMS ACS routines and constructs.
- Dispose of the source VTS library or use it for other workload or upgrade to a new VTS library.

8.4.2 Migration summary

Table 8-12 summarizes the installation tasks of migration scenario 2 when the physical cartridges are moved from the source to the new library. Note that these tasks may vary depending on the actual configuration.

Table 8-12 Migration tasks with scenario 2 and moving physical cartridges

Task	VTS accessibility	Library accessibility
Upgrade old VTS[a]	No	Yes
Upgrade LIC of source LM	No	No
Install new library	Yes	Yes
Install new VTS	Yes	Yes
Update host definitions[b]	NA	NA

a. Upgrading the source VTS and Library Manager may be required depending on the microcode levels of the source and target libraries. Also, feature code changes on the source VTS may be required.
b. Updating HCD definitions on the host may require an IPL.

8.5 Upgrade to the IBM Enhanced 3590 and the 3592

In this section, we describe the IBM 3590 Enhanced Models and how to migrate from an IBM 3590 Base Model based VTS configuration to an IBM 3590 Enhanced Model based VTS configuration. We also describe how to migrate from IBM 3590 Enhanced Models to the 3592 and its associated media.

8.5.1 Introduction to the IBM 3590 Enhanced Models

The IBM 3590 Enhanced Models offer improvements in tape performance, capacity, and connectivity:

- Native data transfer rate improvement of over 50% to 14 MB/s.
- Up to 34 MB/s sustained data rate with maximum data compression.
- A doubling of the current native cartridge capacity to 20 GB that uses a new 256-track format.
- A new digital channel technology and predictive failure analysis to further enhance reliability.
- Investment protection with upgrades from the current IBM 3590 Base Models and backward read compatibility of current cartridges created by IBM 3590 Base Models.
- The IBM 3590 Enhanced Models use the same cartridges as the IBM 3590 Base Models. The IBM 3590 Enhanced Models read and write the higher density track formats and read the 128-track format created by IBM 3590 Base Models. The main reason to upgrade to IBM 3590 Enhanced Models is to triple the data storage capacity of the current TotalStorage 3590 data cartridges. IBM 3590 Enhanced Models can generate 384-track format cartridges which stores three times as much data as 128-track format cartridges written by IBM 3590 Base Models and a 50% increase over the 3590-E1A model capacity. Customer data growth can easily be accommodated without needing extra TotalStorage 3590 cartridges or extra cartridge slots.

8.5.2 Introduction to the IBM 3592 tape drives

The IBM 3592 tape drives offer improvements in tape performance, capacity, and connectivity:

- Native data transfer rate improvement of 2.5 times to 40 MB/s over the 3590-H1A and almost 5 times over the original 3590-B1A
- Up to 120 MB/s sustained data rate with maximum data compression
- A ten-fold increase of the current native cartridge capacity to 300 GB that utilizes the new J-type cartridge

Restriction: The 3590 type cartridges are not compatible with the 3592. New media is required in the 3494 Tape Library to support attachment of the 3590 to the VTS.

Table 8-13 summarizes the IBM 3590 tape drive characteristics.

Table 8-13 IBM 3590 tape drive characteristics

Characteristics	Specifications			
	IBM 3590-B1A	**IBM 3590-E1A**	**IBM 3590-H1A**	**IBM 3592-J1A**
Data Capacity	10 GB / 20 GB	20 GB / 40 GB	30 GB / 60 GB	JJ-type 60GB JA-type 300 GB
Device Data Rate	9 MB/s	14 MB/s	14 MB/s	40 MB/s
Max. Device Data Rate (assuming 3:1 compression ratio)	27 MB/s (Ultra SCSI), 16 MB/s (SCSI-2 F/W)	34 MB/s (Ultra SCSI)	34 MB/s (Ultra SCSI)	120 MB/s (FIBRE)
Data Transfer (FICON - 2Gb/s)	Max. 17 MB/s	Max. 17 MB/s	Max. 17 MB/s	Max 200 MB/s
Data Transfer Rate (ESCON)	Max. 17 MB/s	Max. 17 MB/s	Max. 17 MB/s	Max. 17 MB/s
Data Transfer Rate (Ultra SCSI)	Max. 40 MB/s	Max. 40 MB/s	Max. 40 MB/s	Max. 40 MB/s
Data Transfer Rate (SCSI-2 F/W)	Max. 20 MB/s	N/A	N/A	N/A

8.5.3 Planning considerations

The VTS Model B18 or VTS Model B10 must consist of either all IBM 3590 Model B1A tape drives, IBM 3590 Model E1A tape drives or IBM 3590-H1A tape drives. Mixed tape drive configurations are not supported. However, if two VTSs are installed in an IBM 3494, one VTS could use one 3590 drive type and the second one could use a second drive type. The 3592 tape drives are not supported on the VTS Model B18.

The VTS Model B20 can consist of a mixed configuration of 3590 tape drives and 3592 tape drives providing the 3590 tape drives are of all the same type and located in the same 3494-D12 frame and the 3592 drives are installed in a second 3494 drive frame Model D22.

There are no host software requirements for implementation of IBM 3590 Enhanced Models or the 3592. This is one of the clear advantages of the IBM VTS: the attached hosts only see virtual IBM 3490E drives and virtual IBM 3490E cartridges, which makes this upgrade completely transparent to the host.

Export/Import dependencies

Exported stacked volumes created on IBM 3590 Model B1A tape drives can be imported into IBM 3590 Model E1A tape drive or IBM 3590-H1A tape drive based configurations. Exported stacked volumes created on IBM 3590 Model E1A tape drives cannot be imported into IBM 3590 Model B1A tape drive based configurations. Exported stacked volumes created on IBM 3590-H1A tape drives cannot be imported into IBM3590-E1A or 3590-B1A tape drive based configurations.

IBM 3590 Model B1A to IBM 3590 Model E1A or 3590-H1A migration

The IBM 3494, which houses the VTS drives, can stay online, as long as two drives remain in a VTS without FC5236 and FC4000 or 4 drives remain in a VTS with FC5236 or FC4000 (see Table 8-14).

Table 8-14 Migration details for B1A to E1A or H1A

VTS type	VTS with 4 - 7 drives	VTS with 8 to 12 Drives
VTS with: GPFS OR Import/Export	Full Replacement	Partial Replacement (Leave 4 Drives)
VTS without: GPFS AND Import/Export	Partial Replacement (Leave 2 Drives)	Partial Replacement (Leave 2 Drives)

If an IBM 3494 Tape Library contains two VTSs, they can be upgraded independently, therefore one VTS can be configured one drive type and the second can be configured with another drive type.

IBM 3590 Enhanced Model upgrade includes following activities:

- Upgrade Library Manager LIC to 523.00 or higher
- Install VTS AIX 4.3.2 or higher
- Install VTS LIC to 2.17.10.0 or higher
- Run the drive conversion diskette. (This marks the partially filled tapes read-only).
- Conversion of IBM 3590 Model B1A tape drive to 3590 Model E1A tape drive or to 3590 Model H1A.

The total upgrade time for the installation of the IBM 3590 Enhanced Model upgrade depends on the current LM level, the VTS code level, number of installed, and the installed VTS features.

3590 (all models) to 3592 Migration B10 or B20 only

There are two scenarios to upgrade a 3494 VTS with IBM 3590 tape drives to 3592 tape drives. Concurrent migration of from 3590 to 3592 is only supported on the Model B20 VTS. There is no field upgrade of the VTS Model B10, so migration with the VTS Model B10 will require the installation of additional VTS model B10 with 3592 tape drives. Then, through host software, copy the data from VTS1 (source B10 with 3590 tape drives) to VTS2 (new B10 with 3592 tape drives). Then remove VTS1.

Migration of a VTS Model B20 can be done as described above with the VTS Model B10 or concurrently in Heterogeneous mode. If the source VTS Model B20 currently has 6 tape drives, then a 3494-D22 will have to be installed or an existing 3494 Dxx Frame will have to be converted to a D22 frame to house the 3592 drives. Also additional features on the VTS Model B20 will need to be installed and a Teach of the Library Manager will be required to pick up the change in the frame configuration. The B20 features and the new D22 frame or conversion of an existing Dxx frame will not be concurrent. The migration from the 3590 media to the 3592 media is concurrent. There are also microcode considerations with both the VTS B20 and the 3494 Library Manager. VTS LIC 2.31.730.x and 530.xx on the library Manager are both required to support the migration.

Migrating from a VTS Model B20 with more than six 3590-x1A tape drives will require a conversion of one of the D12 frames containing the VTS back-end drives to a D22 frame where the new 3592 drives will be installed. Conversely a new D22 Frame with 3592 tape drives can be installed as long as one of the D12 frames associated with the VTS is reconfigured as no longer being attached to the VTS.

Note: A Library Manager Teach function is required whenever a frame is reconfigured on a 3494 Tape Library. This is not concurrent to library operations.

8.5.4 Operational considerations

In this section we briefly summarize the operational differences between the IBM 3590 and 3592 tape drives.

3590 type drives

After the IBM 3590 Enhanced Model upgrade, the VTS will start to write only in 256-track or 384-track mode onto scratch stacked volumes and will not append to partially filled lower capacity formatted stacked volumes. All filling volumes will be flagged as read-only and the VTS will start using scratch stacked volumes for writing (filling). All of the filling volumes will get reclaimed by the read-only recovery process and eventually returned to scratch status.

The VTS will write double or triple the amount of data on each TotalStorage 3590 tape cartridge. Rewriting of the Volume Control Region (VCR) portion in 256-track or 384-track format takes 42 seconds and occurs only once per stacked volume. This effect is negligible. Currently there are no performance implications. Read and write characteristics from a host point of view are similar. Read misses from the TVC are handled at the same speed, because the mechanical characteristics like load, search and rewind didn't change.

Immediately after the IBM 3590 Enhanced Model upgrade, the LM panels and SMF94 statistics will report an increase in the amount of *Free Storage*. After reclamation, more and more lower capacity format cartridges will return to scratch and are eligible to be written in higher capacity format mode. Gradually the amount of reported free storage capacity will increase, as can be seen, on the LM panels and SMF record 94 reports.

There is no external indication on the amount or percentage of stacked volumes that are in 256-track or 384-track format.

> **Note:** To force stacked volumes to be reclaimed and re-written in the higher capacity format, you may want to consider temporarily increasing the “Reclaim Threshold Percentage” (refer to “Define VTS management policies” on page 137).

After the IBM 3590 Enhanced Models upgrade, the LM panels and SMF94 statistics will show twice the amount of *Maximum Active Data*.

Immediately after the IBM 3590 Enhanced Models upgrade the LM panels and SMF94 statistics will show the same amount of *Free Storage Threshold*. As a consequence, the alarm now goes off at half the number of scratch cartridges. We recommend that you increase accordingly the Free Storage Threshold after the IBM 3590 Enhanced Model upgrade. This is to ensure that the alarm goes off at the same amount of scratch cartridges as with the previous model of tape drives. See “Define VTS management policies” on page 137 for a detailed discussion on the Free Storage Threshold mechanism.

3592 type drives

There are no differences in operating the Model B10 VTS or Model B20 VTS with 3592 tape drives or 3590 tape drives. With the exception of the requirement to define a new storage pool for the new media.When both drive types are installed in the same 3494 there will be no mixing of media even if the Model B20 VTS is using both drive types in heterogeneous mode. Free Storage Threshold values should be calculated accordingly based on the media being used, 60 GB for the JJ, or 300 GB for the JA 3592 media.

8.5.5 Migrating from 3590 media to 3592 media

In this section we describe how to move all the logical volumes from your existing 3590 media, J or K type cartridges, to the new 3592 media (JJ or JA). A detailed explanation of this migration can be found in the "Setting Up Data Management Definitions" section of the IBM White Paper entitled IBM TotalStorage Virtual Tape Server Using 3592 In a VTS. It can be found here:

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

Options

There are two options when converting to the new media. The first option, and the only option available to the mode B10, is migrating the data through the host to a second installed VTS with the 3592 tape drives. The source VTS can be a Model B18, Model B10 or Model B20 with 3590 tape drives with the target being a Model B10 or Model B20.

The second option is only supported with the Model B20 VTS. In Heterogeneous mode you will have both the 3590-x1A and the 3592 tape drives installed in the same 3494 Tape Library in separate Dxx Frames and attached to the same VTS. Then the logical volumes can be migrated from the 3590 stacked volumes to the new 3592 stacked volumes using the features of APM. This will allow a pool of 3590 media to be moved to another pool consisting of 3592 media through the normal VTS reclamation process.

Upgrading to 3592 tape drives with second VTS

A second VTS model B10 or Model B20 can be installed in the current library or a second library. Then the logical volumes are moved from the current VTS to the new VTS through host software. See "Scenario 2: Parallel operation with two libraries" on page 349 for more information if you will be installing a second library and VTS as your solution.

Figure 8-5 shows this scenario where a second VTS is installed inside the same library.

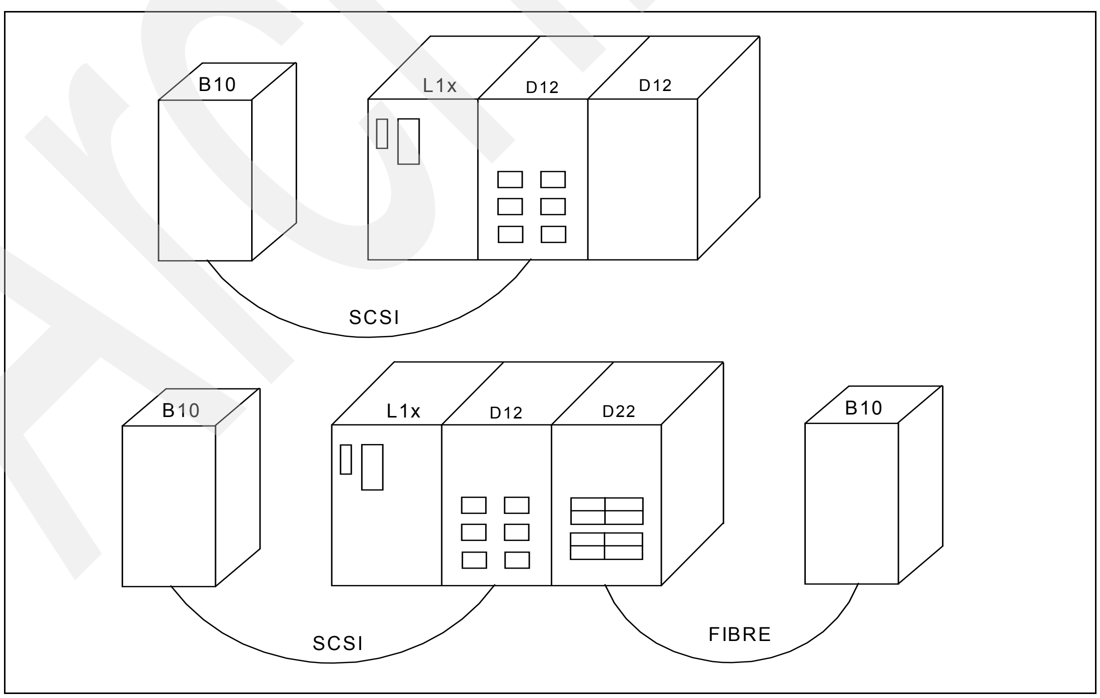

Figure 8-5 Migrating from 3590 to 3592 media with second VTS added

Software considerations

These are the considerations for software:

- Host definitions with HCD and DFSMS must be created for the new logical library.
- Logical devices (64 for B10 or up to 256 for B20) must be defined with HCD.
- APAR OA08777 is recommended
  - This provides updates to SMF 94 records for enhanced statistics reporting and additional storage pool control settings

Hardware considerations

These are the considerations for hardware:

- Provide a place for the 3592 drives to be installed:
  - Convert one D12 frame to a D22 frame.
  - Add one additional D22 frame.
  - Convert one Open Systems attached D22 frame to VTS-attached.
- All frame upgrades or installs require a Teach of the 3494 Tape Library to be usable. This is a non-concurrent activity.
- The VTS must be upgraded with FC5238 to provide fibre connection to the 3592 tape drives.
- VTS LIC 2.31.730.x or above is required.
- The 3494 Library Manager requires LIC 530.xx or above.

Migration steps

These are the steps for migration:

- Upgrade the 3494 Library Manager to 530.xx or above, if required.
  - This non-concurrent to the operation of the 3494 Tape Library and all of its components.
- Upgrade the new 3494 VTS to Release 2.31.730.x or above, if required.
- Upgrade the new VTS model with FC5238 - Fibre Card Drive Attachment. This is not concurrent to VTS operations.
- Upgrade existing D12 or D22 frame or install new D22 or D12 frame
- Install the required 3592 tape drives and cable to VTS.
- Teach the 3494 Tape Library with the new configuration. This includes the new drive frame configuration, the new VTS sequence number (VTS Lib Name) and stacked volume range. This is a non-concurrent operation to the whole of the 3494 Tape Library.
- Insert the appropriate number of 3592 media (JJ or JA type)
- Vary new logical library and its associated logical devices online.
- Use host software to copy data from old VTS logical library to the new VTS logical library.

Migrating from 3590 to 3592 media using APM

Figure 8-6 shows the migration of a B20 VTS from twelve 3590 tape drives to six 3590 and twelve 3592 tape drives. This is the maximum drive configuration which is currently available. To achieve the target configuration, one D12 frame and its six 3590 drives have to be detached from the VTS. The frame can be replaced later, after the new 3592 ape drives have been installed.

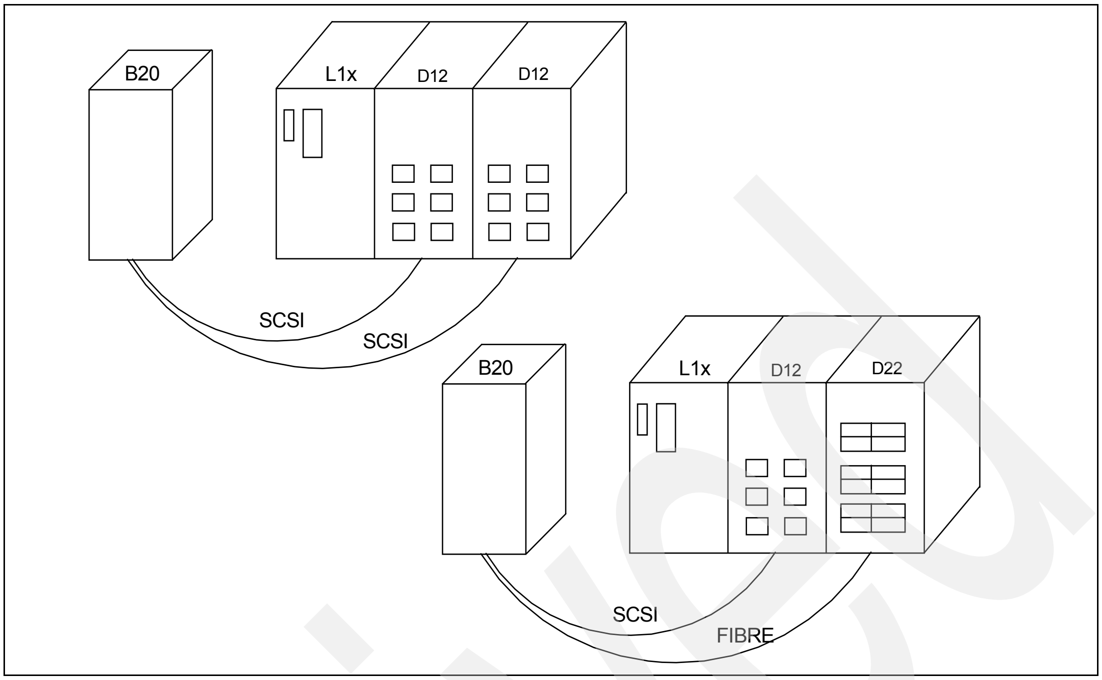

Figure 8-6 Migrating from 3590 to 3592 media in heterogeneous configuration (B20 only)

Software considerations

These are the considerations for software:

- No host software upgrade is required — z/OS V1 R3 or above.
- APAR OA08777 is recommended. This provides updates to SMF 94 records for enhanced statistics reporting and additional storage pool control settings.
- Additional storage groups may be required in heterogeneous configurations.

Hardware considerations

These are the considerations for hardware:

- There are several options in upgrading the 3494 Tape Library to support the new 3592 tape drives:
  - If the VTS has more than six 3590 tape drives installed in 2 adjacent 3494-D12 frames, you can convert one of VTS' D12 frames to a D22 frame or install a new D22 frame in the library. One of the VTS attached D12 frames must then be MES'ed to a D12-0 frame by removing the 3590 tape drives.
  - If the VTS has only six 3590 tape drives in one D12 frame, you can convert a second D12 frame in the library to a D22 frame or you can add a new D22 frame.
  - All frame upgrades or new installs require a Teach of the 3494 Tape Library to be usable. This is a non-concurrent activity.

> **Note:** In a heterogeneous configuration the D12 with 3590 tape drives and the D22 frame with the 3592 tape drives do not need to be adjacent in the 3494 Tape Library.

- The VTS requires FC5238 to provide fibre connection to the 3592 tape drives.
- The VTS requires FC4001 (APM) installed.
- The VTS requires release 2.31.730.x microcode or above.
- The 3494 Library Manager requires release 530.xx of microcode or above.

Migration steps

These are the steps for migration:

- Upgrade the 3494 Library Manager to 530.xx or above. This is non-concurrent to the operation of the 3494 Tape Library and all of its components.
- Upgrade the 3494 VTS model B20 to LIC Release 7.3 (2.30.720.40) or above. This is non-concurrent to the VTS model B20 only. The second VTS or the other native attached zSeries and Open Systems tape drives, if installed, will run as normal.
- Upgrade VTS model B20 with FC5238 - Fibre Card Drive Attachment. This upgrade is also non-concurrent to the VTS operation.
- Upgrade existing D12 or D22 frame or install new D22 or D12 frame with the required 3592 tape drives. Some of this upgrade may be done concurrent to the 3494 Tape Library's operation, but the final Teach of the Library Manager will require the library to be offline to all hosts.
- Insert new 3592 media to support the current VTS. Now the logical volumes can be migrated from the 3590 stacked volumes to the 3592 stacked volumes. This is done by using pooling features of APM.

Data migration methods

There are two methods for moving the logical volumes from 3590 media to 3592 media. Both use the same process of reclaiming data from a pool of 3590 media and writing it to 3592 media, but one method is more direct while the other is a more conservative approach. A detailed explanation of this migration can be found in the section "Setting Up Data Management Definitions" of the IBM White Paper entitled *IBM TotalStorage Virtual Tape Server Using 3592 In a VTS*. It can be found here:

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

In Appendix C, "VTS implementation step-by-step" on page 419 there are some step-by-step examples of setting the appropriate configuration settings that may assist in your understanding of these migration procedures. Also see Chapter 4, "Implementation" on page 107 for more information on where to make the changes required for these migration methods.

Method 1

Method 1 involves using the new media with only new constructs, and any new logical volumes will be written to these new pools specified to use the 3592 media. This would be used at first for testing and then for live data at a controlled pace. This method is suitable for the first time you migrate data to the new media. If you have other libraries to migrate after this, then Method 2 might be more typical.

Steps

Here are the steps you need to perform:

- Define new 3592 pools — see "Defining stacked volume pool properties" on page 143.
- Define new constructs that target new pools — see Figure 4-24 on page 146.
- Change the Host ACS routines to use the new constructs.
- Gradually direct data to the new pools.
- Set an appropriate Inhibit Reclamation Schedule — see "Inhibit Reclaim Schedule" on page 139.
- Change reclamation target of all 3590 pools to corresponding 3592 pools — see Figure 4-20 on page 140.

- Change old storage groups and management class constructs — see Figure 4.4.3 on page 147.
- Activate new reclamation policies — see Figure 4-20 on page 140.
- Let reclamation do its work.
- Issue moves to flush 3590 pools and eject 3590 tapes — see “Move stacked volumes” on page 235.
- Change empty 3590 pools to 3592 — see “Defining stacked volume pool properties” on page 143.
- Verify that no 3590 pools are specified for the VTS.
- Remove 3590 drives.
- Re-Teach the library.

Method 2

This method involves the most direct way to move all the data over to the new media. This method is more suitable for subsequent library migrations or when the process is fully understood and you are comfortable with all aspects of the migration.

> **Note:** Advanced Policy Management can be controlled by the web specialist server as well as the Library Manager panels. In the following steps the web specialist panels will be referenced, where ever possible, instead of the Library Manager panels. Both methods can be done from either interface.

Steps

Here are the steps you need to perform:

- Define new 3592 pools — see Figure 4-22 on page 143.
- Change constructs to target new pools — see Figure 4-24 on page 146.
- Change the storage group and management class constructs on the LM.
- Redefine 3590 pools — see Figure 4-22 on page 143.
- Set an appropriate Inhibit Reclamation Schedule.
- Change reclamation target of all 3590 pools to corresponding 3592 pools — see Figure 4-22 on page 143.
- Activate new reclamation policies.
- Let reclamation do its work.
- Issue Moves to flush 3590 pools and eject 3590 tapes.
- Change empty 3590 pools to 3592 — see Figure 4-22 on page 143.
- Verify that no 3590 pools are specified for the VTS.
- Remove 3590 drives.
- Re-Teach the library.

A

Peer-to-Peer VTS solution

In this appendix, we describe the Peer-to-Peer VTS solution. We give you a short introduction to the following topics:

- Components of a Peer-to-Peer VTS
- Configuration options
- Terminology
- Modes of operation
- Upgrade paths

IBM TotalStorage Peer-to-Peer VTS

The IBM TotalStorage Peer-to-Peer VTS is designed to eliminate all single points of failure and provide higher performance than the stand-alone VTS. IBM TotalStorage Peer-to-Peer VTS couples two separate VTSs into a single image with additional components and features. The two VTSs can be located at the same site or can be geographically separated for disaster tolerance. An IBM Model B18 VTS, Model B10 VTS, or Model B20 VTS can all be used in IBM TotalStorage Peer-to-Peer VTS configurations. Additionally with feature code 4001-4004 (Advanced Policy Management), the PtP VTS configuration can be further tailored to suit individual site requirements. Control is gained over logical volume copy modes (either immediate or deferred). Media errors can also be eliminated as a single point of failure within each VTS with the ability to create dual copies of logical volumes on separate physical media. For details, see the redbook, *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115.

AX0 configuration

The combined libraries require the addition of one to four 3494 CX0 frames which contain either two or four AX0 tape controllers each. In turn, each AX0 tape controller provides a total of four ESCON channel attachments; two attached to the host systems and one attached to each of the Model Bxx VTSs in the Peer-to-Peer configuration. The host systems no longer attach directly to the Model Bxx VTS. A single model AX0 controller presents itself to the host(s) as one 3490E tape control unit with sixteen 3490E tape drives (Figure A-1).

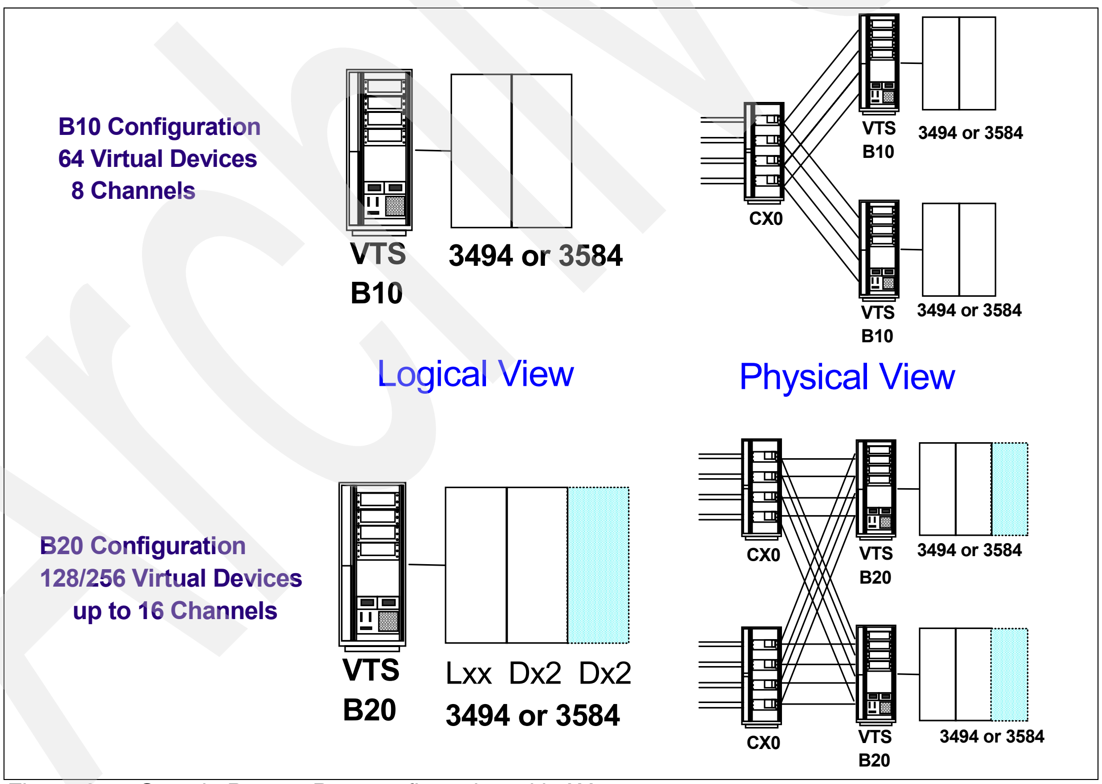

Figure A-1 Sample Peer-to-Peer configuration with AX0

It is important to note that in a mixed Peer-to-Peer (PtP) configuration the Composite Library will assume the attributes of the smaller VTS in relation to virtual drives.

VTC configuration

Following the announcement of the VTC tape controllers, which support FICON, the attachment permutations change slightly. The combined libraries require the addition of one to four 3494 CX1 frames which can be configured with either:

- 4 FICON VTCs
- 2 FICON VTCs and 2 ESCON VTCs
- 2 FICON VTCs or 2 ESCON VTCs with 2 AX0s (requires FC4691)

This flexibility allows for migration from AX0s to the new VTCs. All FICON VTCs in a CX1 frame must have the same feature code. Each VTC controller provides a maximum of four FICON channel attachments; two attached to the host systems and one attached to each of the Model Bxx VTSs in the Peer-to-Peer configuration. The host systems no longer attach directly to the Model Bxx VTS. A single model VTC controller presents itself to the host(s) as one 3490E tape control unit with sixteen or thirty-two 3490E tape drives (Figure A-2).

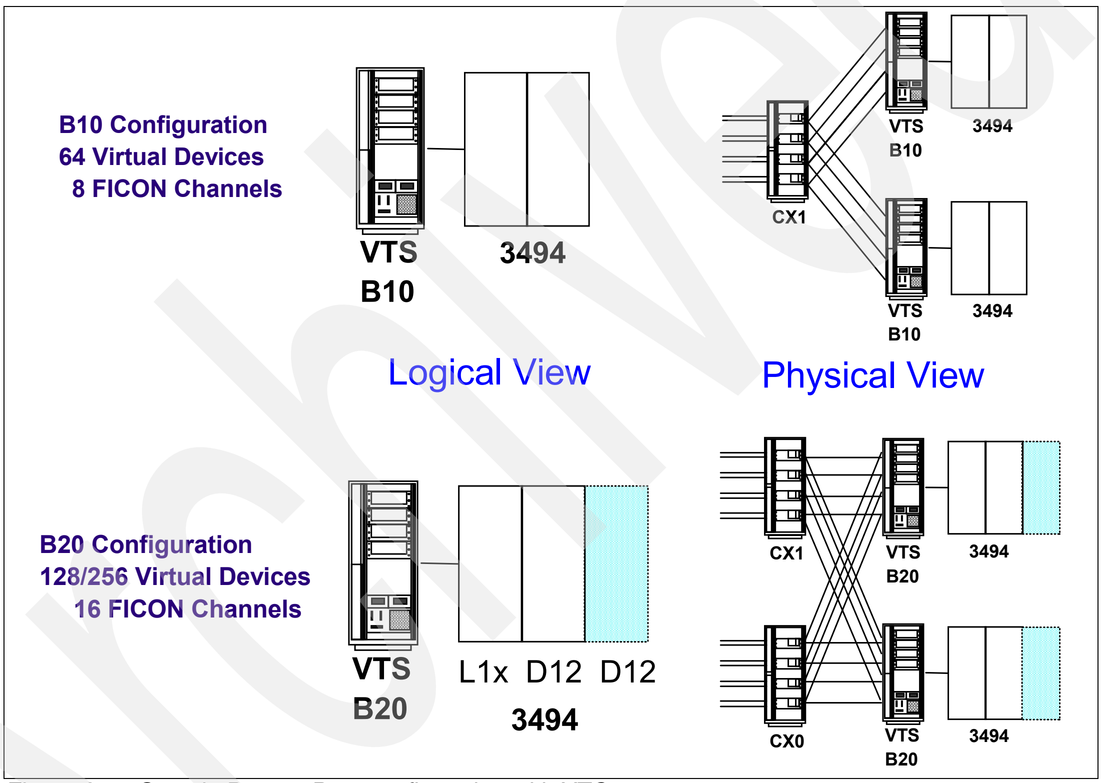

Figure A-2 Sample Peer-to-Peer configuration with VTC

It is important to note that in a mixed Peer-to-Peer (PtP) configuration the Composite Library will assume the attributes of the smaller VTS in relation to virtual drives.

Peer-to-Peer terminology

In this section, we introduce some basic terminology which is unique to the Peer-to-Peer VTS. Terms such as: **Composite Library, Distributed Library, Master VTS, I/O VTS,** and **UI Distributed Library** refer to parts of the logical library as shown in Figure A-3. Whereas, **Immediate Copy, Deferred Copy, and VTS I/O preferencing** refer to functions within the Library Manager (Figure A-3).

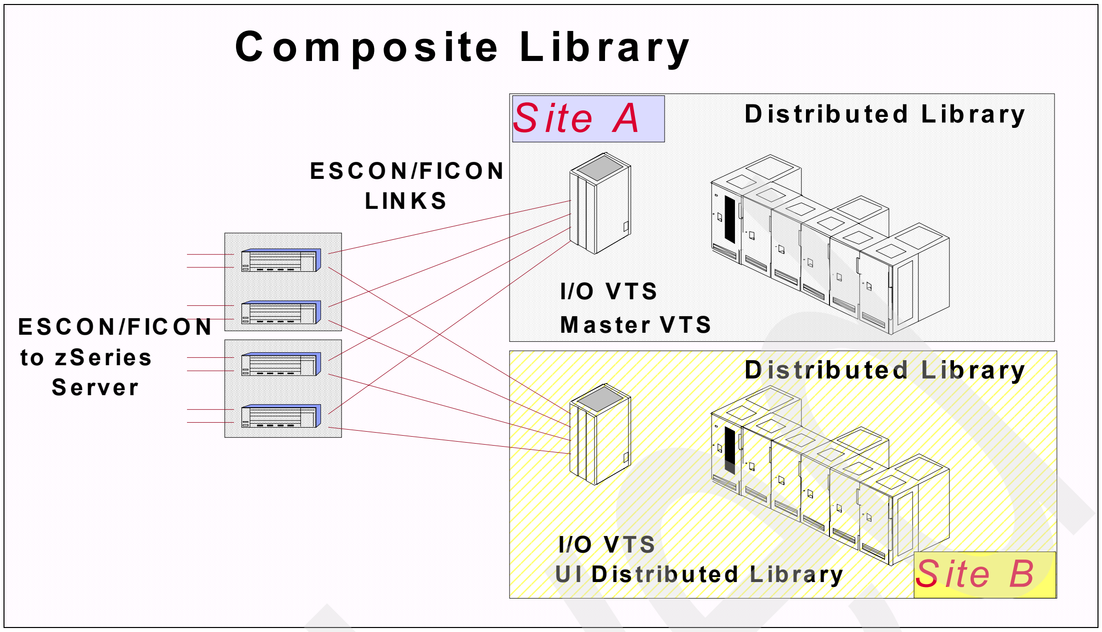

Figure A-3 Peer-to-Peer terminology

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, sixteen or thirty-two 3490E tape control units, giving a maximum of 512 logical 3490E drives (with a model B20). The composite library presents the same image to the z/OS 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 Model AX0 or VTC controllers and they do not need to be defined to the host.

Distributed Libraries: A distributed library (or distributed VTS) is a physical IBM 3494 or IBM 3584/3953 Tape Library in the Peer-to-Peer VTS complex which includes a Model Bxx 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 any single distributed library. On the host, the distributed libraries are only defined to SMS. They are defined using the existing ISMF panels and have no tape devices defined; the tape devices were defined for the Composite library.

Master VTS: The master VTS is one of the Model Bxx 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.

I/O VTS: The I/O VTS is the Model Bxx 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 controller during mount processing (load balancing function). Both VTSs may operate as I/O VTSs for different sets of volumes, this is 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 subsequent mount of the same volume, assuming that the volume is no longer in cache.

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 VTC controllers 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.

Modes of operation

The Peer-to-Peer VTS allows you to select when to create the secondary copy of a virtual volume: immediately during close processing (immediate copy mode), or asynchronously at a later time (deferred copy mode). The immediate or deferred copy mode of operation is configured initially during setup of the VTCs. This can be changed any time afterwards, but only by an IBM System Service Representative (SSR). With the addition of the Advanced Policy Management (APM) feature, the copy mode can be tailored further for a logical volume. You can also create 1-32 pools of stacked media and well as creating duplicate copies of your logical volumes (to the same or different stacked media pools). These functions are discussed in detail within Chapter 4, “Implementation” on page 107.

Immediate Copy: This 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: This 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.

VTS I/O Preferencing: 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 Virtual Tape Controllers (VTC) 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 Model 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.

Dual Copy: With Advanced Policy Management enabled, you have the ability to assign DFSMS management classes to your logical volumes and have a backup copy created. These constructs are then passed to the VTS for actioning as per the constructs defined at either the LM or 3494 Specialist.

Selective PtP Copy: With APM, you can determine to have a logical volume only written to one VTS and no copy created in the other VTS. Through Management Class definition, you can select the copy mode and the VTS to which the single copy should be written.

Volume Affinity / Pooling: With Advanced Policy Management enabled, you have the ability to assign DFSMS storage groups and pool your logical volumes onto separate physical stacked media as per the constructs defined at either the LM or 3494 Specialist.

Upgrading to a Peer-to-Peer VTS

A single VTS, with the exception of a Model B16 VTS, can form a part of a Peer-to-Peer configuration. It may be necessary, however, to upgrade the specific features of any existing VTSs when migrating them into a single Peer-to-Peer complex. Figure A-4 shows potential upgrade paths. See your local IBM representative for details as to specific configurations. See also *IBM Total Storage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115 for details and migration guidelines.

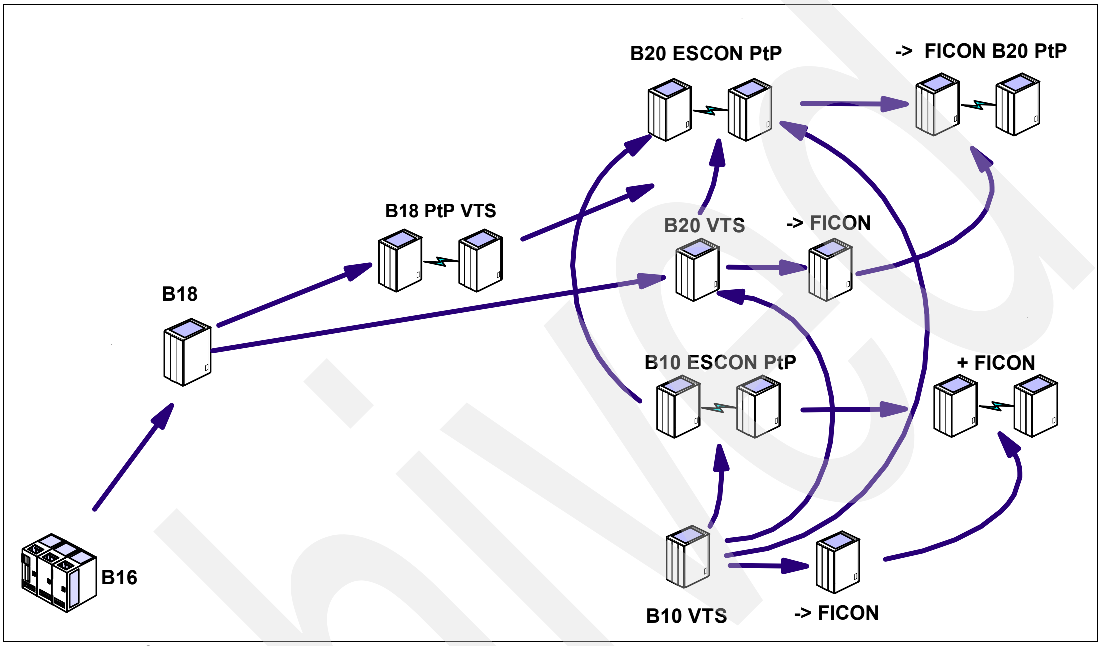

Figure A-4 General upgrade paths

Summary

The IBM TotalStorage Peer-to-Peer Virtual Tape Server (PtP VTS), an extension of the IBM TotalStorage Virtual Tape Server (Model B10, B18 or B20 VTS), builds on a proven base configuration to provide even greater benefits for tape processing and operations. By enhancing data backup and recovery capabilities, the PtP VTS 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 PtP VTS improves system reliability and availability, as well as data access.

Existing Model B18 VTSs cannot be upgraded to a PtP VTS configuration directly. You must upgrade the B18 to a Model B20 VTS first, before you can upgrade it to a PtP VTS configuration by adding the incremental hardware and a second VTS with the Peer-to-Peer Copy features. This protects the investment that many customers have already made in IBM technology.

For more information about IBM TotalStorage Peer-to-Peer VTS, refer to the *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115.

B

VTS Import / Export

Import/Export (I/E) is part of APM which has replaced the I/E functionality of the Advanced Function feature. I/E provides a way to move logical volumes out of a VTS subsystem to physical cartridges and to return them back to the same or another VTS. This function of the VTS removes the “closed store” limitation of the VTS. This function allows customers to move data between Virtual Tape Servers without having to copy the logical volumes back through the host channels or lose the VTS benefits of volume stacking, which exploits the storage capacity of the 3590 cartridge.

In this appendix, we provide information to help you to use the Import/Export operations of the VTS.

> **Important:** In a VTS attached to an IBM 3584 Tape Library, only 3592 drives are supported. If there is a need for the Import/Export function, plan for a separated pool with Economic Enterprise Tape Cartridges (EETC) which provide an uncompressed capacity of 60 GB each. Otherwise, you might waste cartridge space using the JA cartridges with an uncompressed capacity of 300GB, as it is unlikely in most environments that such a large amount of data will be exported in a single export task.

Introduction

If you are uncertain about the ability of your VTS to perform the any Import/Export operations, use the VTS Status panel (Figure 7-24 on page 314) and check the Exp/Imp Capable field. Import/Export is available with both VTS advanced function feature codes:

FC4000	Advanced Function
FC4001-4004	Advanced Policy Management

Export is a host initiated operation which copies one or more logical volumes to physical cartridges and deletes the logical volumes from the VTS. Since exported logical volumes are deleted from the VTS as part of the export process, an exported volume is not a duplicate of the original; it is the only version of an active logical volume. Therefore it is a data move, not a copy operation. Wherever the word “host” is used, it is implied to be an ES/3090, ES/9000, or S/390 processor running z/OS DFSMS System-managed tape with the appropriate support for the Export/Import function. The output volumes are called *exported stacked volumes* and they contain *logical volumes* which have been exported.

Import is a host initiated operation that can be used to copy one or more previously exported logical volumes from exported stacked volumes into a VTS.

You can use Import/Export for the following purposes:

- To move logical volumes to another location for long term archiving or backup purposes
- For workload balancing by moving logical volumes from one VTS subsystem to another
- For data interchange with another VTS subsystem

The support for the Import/Export function is included within DFSMSrmm with PTFs. The functional specifications are made available to other vendors such as:

- Computer Associates (BrightStor CA-1)
- Computer Associates (BrightStor CA-Dynam/TLMS)
- Platinum (AutoMedia, formerly ZARA)
- BMC (Control-T)
- BETA Systems (BETA51)

For BrightStor CA-1 support of the Import/Export function and BrightStor CA-Dynam/TLMS, see “BrightStor CA-1” on page 414 and refer to the respective product documentation.

Configuration and feature requirements

The Import/Export function of the VTS requires the following features and configuration:

- One of these feature codes:
  - The Advanced Functions feature code, FC4000, which has been withdrawn from marketing and is no longer available

    or

  - The Advanced Policy Management feature codes FC4001- FC4004
- Model B18 with feature code 3200 (ESCON High Performance Option) or feature code 3400 (Extended High Performance Option), or a Model B10 or B20 VTS
- A minimum of four 3590 or 3592 tape drives. We recommend that you have a minimum of six drives when using Import/Export. Copy is always done from tape to tape; the Tape Volume Cache is not used or involved. Therefore both operations, export and import, need two physical drives.

- The 3494 Tape Library or 3584 Tape Library must have a convenience I/O station. The high capacity I/O facility cannot be used for Import/Export.

Figure B-1 shows the reason for our recommendation to have a minimum of six tape drives if the Import/Export operations are used:

- One drive may be used for recalls.
- One drive may be used for the copy task.
- Two drives may be used for reclamation.
- Two drives may be used for Import/Export.

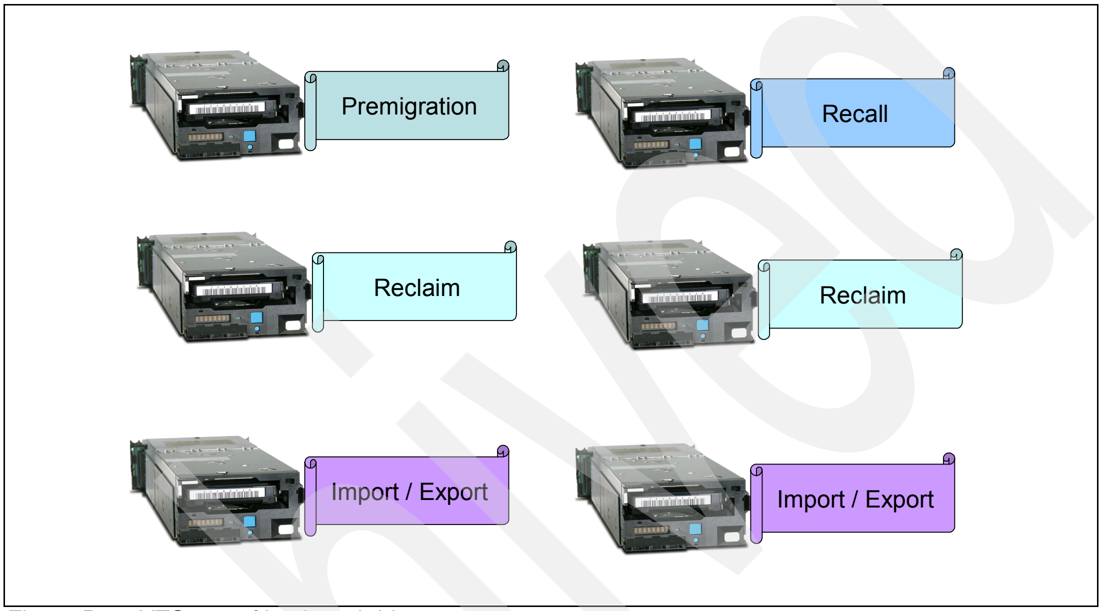

Figure B-1 VTS use of back-end drives

Software requirements

Import/Export is available only in a z/OS DFSMS System-managed tape environment. The software support is provided with z/OS V1R4 and later releases.

The cartridge entry and cartridge eject installation wide exits (CBRUXENT and CBRUXEJC) have been enhanced in support of the Import/Export function. The sample exits provided by DFSMSrmm have been updated accordingly and DFSMSrmm itself has been enhanced in support of Import/Export. If you are using another tape management system product, contact your vendor for details on Import/Export support.

General considerations

Here are some general considerations regarding Import/Export operations:

- A Logical Volume called the Export List Volume, in the VTS performing the Export operation, is used to pass the list of logical volumes to export and to return status to the host about the processing of each logical volume. The list may include logical volumes, that are not in the VTS, which is performing the Export operation. The data records on the volume are written using the EBCDIC character set and compression has to be turned off (TRTCH=NOCOMP).

- An optional destination may be specified for each logical volume to be exported. All logical volumes on an exported stacked volume have the same destination.
- With FC4000, all logical volumes to be exported are copied from the VTS managed storage pool to exported stacked volumes.
- With FC4001, export candidate logical volumes could be assigned to a separate storage pool. The could reduce the amount of stacked volumes mounted for any export process.
- With FC4001, any export function will select as a default a K (EHPCT) type cartridge as an export stacked volume media type. Export will select J (HPCT) type cartridge types:
  - If, in a mixed media library (both Js and Ks), the Export function determines that the exported data will fit within the capacity of a J cartridge.
  - Or, if only J type cartridges are available.
- The exported stacked volumes are written in VTS's export data format.
- All data to be exported must be on stacked volumes. Any data that is resident in the cache is first copied to stacked volumes by the VTS and then removed from the cache (fragmented). This occurs as an integral part of the export operation.
- Exported stacked volumes created on 3590 E1A models cannot be imported into a VTS with B1A models, similarly exported stacked volumes created on a 3590 H1A model cannot be imported into a VTS with E1A or B1A models.
- The Export operation is stopped by the VTS when less than four physical drives are found. The operation in process will be completed for the current exported stacked volume.
- The VTS must have at least one available scratch stacked cartridge for the export operation to be accepted.
- Only one Export/Import process can be active in a VTS subsystem at a time.
- Only one import process can be active in a library at a time.
- Import and export can run concurrently in a library with more than one VTS subsystem.

The exported stacked volume is designed to be imported again by the import function. If there is no VTS with the Import/Export function available and you need access to an exported logical volume, you may copy single logical volumes by using DITTO/ESA plus PTFs to native volumes on native drives. Refer to “DITTO VTS functions” on page 414 for further information.

Specific Library Manager categories associated with Import/Export are:

```
FF12 Export pending category
FF13 Exported category
FF14 Import category
FF15 Import pending category
FF16 Unassigned category
FF17 Export hold category
```

The Import/Export operations, while they can only be initiated from z/OS DFSMS, are not restricted to logical volumes created from the initiating host. In an environment where the VTS is shared by multiple hosts (such as VM/ESA, or SCSI attached hosts), logical volumes belonging to these hosts can be exported from a host running z/OS DFSMS. However, there is no way to track and control the exported logical volumes from a host other than z/OS DFSMS. Import/Export statistics are written in SMF94 records.

Library Manager enhancements

The Library Manager was enhanced with a set of panels and menus to control and manage Import/Export with LM LIC level 526. With LM LIC level 527 or higher (Figure B-3) further enhancements are made and more panels have been added. All panels can be found under *Commands* in the *System management* menu (Figure B-2). The panels are used to:

- Manage unassigned volumes
- Manage import volumes
- Manage insert volumes
- Manage export hold volumes
- Cancel VTS Import/Export
- Display export/import volumes with LM LIC level 527 or higher

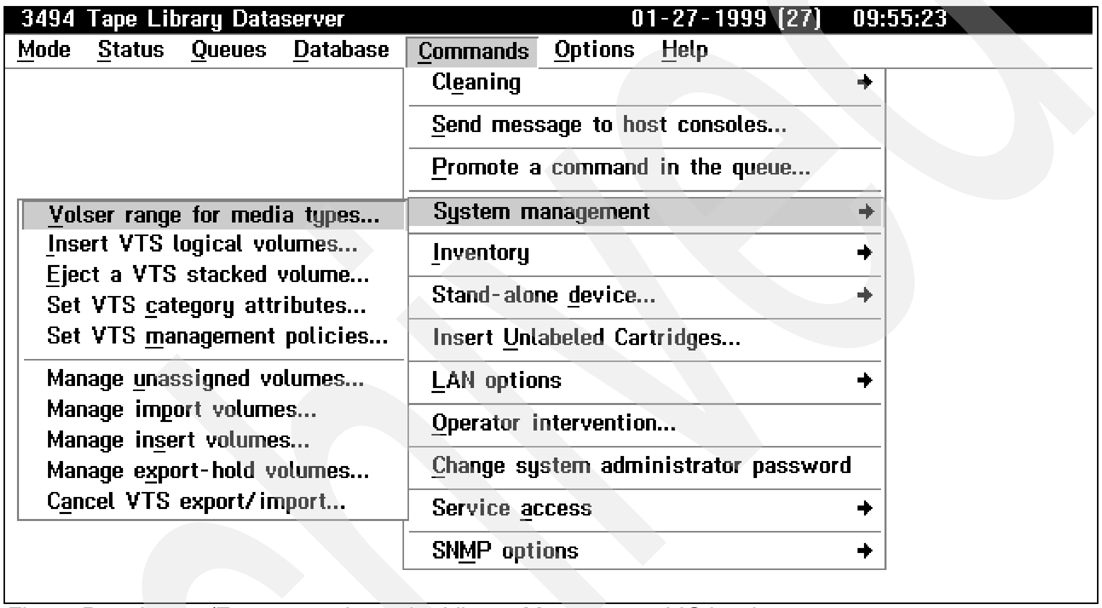

Figure B-2 Import/Export panels on the Library Manager pre LIC level 527

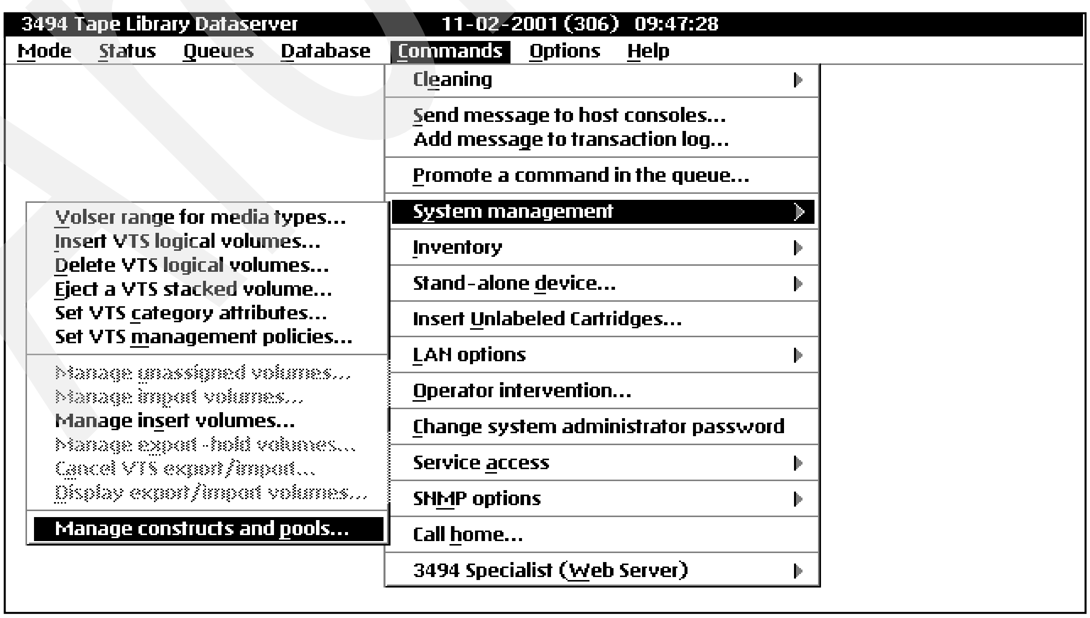

Figure B-3 Import/Export panels on the Library Manager post LIC level 527

To verify pending Import/Export operations in an 3494, the Operational Status panel (Figure B-4) or the host command (see “Run Export function” on page 381) can be used:

```
D SMS,LIBRARY(libname),DETAIL
```

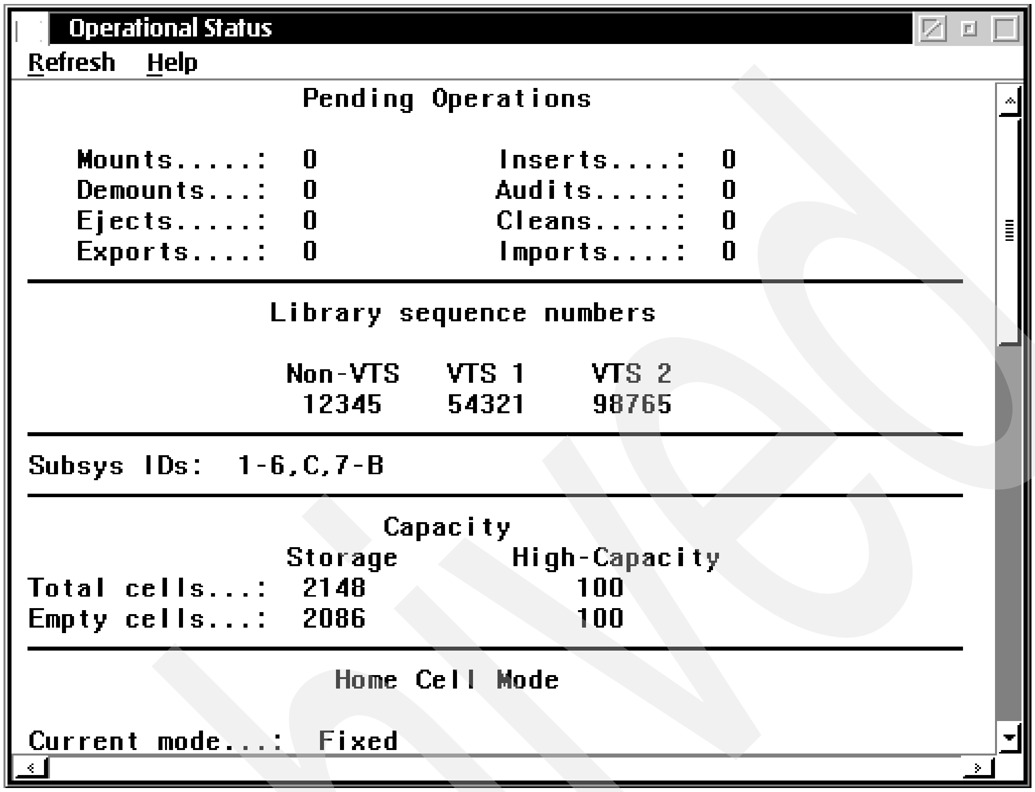

Figure B-4 Operational Status panel

To display volume status for cartridges related to Import/Export operations, the Display VTS Export/Import Volumes panel in Figure B-5 can be used.

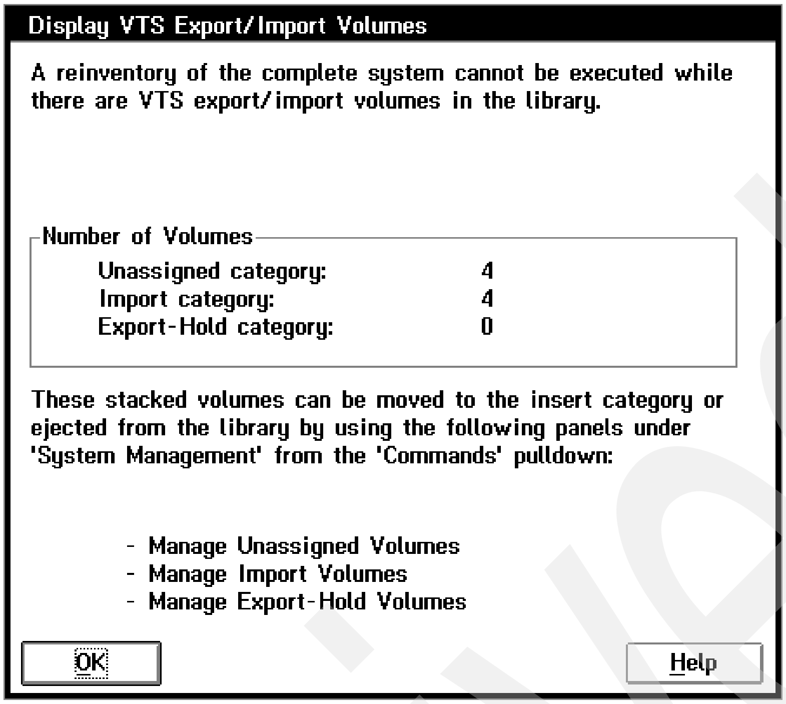

Figure B-5 Display VTS Import/Export volumes

Import/Export terms

Export list volume: This is a logical volume within VTS that contains:

- Export list file
- Export status file

Import list volume: This is a logical volume within VTS that contains:

- Import list file
- Import status file

Exported stacked volume: This is a physical cartridge used as a container:

- It holds the exported logical volumes.
- The volser is reported to tape management system via CBRUXEJC exit.
- It is selected randomly at time of export.

Import and export list volumes

When a host initiates an export or import operation, it has to inform the VTS subsystem which volumes are to be exported or imported. An **export list volume** or an **import list volume** is used for this purpose. The VTS also uses the list volumes to store operation status information. You can access the status file after the operation has completed.

A list volume is a logical volume in the VTS which is executing the Import/Export operation. You need to create the list volume before starting the operation. The volser of the list volume is passed to the VTS when the Import/Export operation is started. The list volumes must have a standard IBM tape label and contain a predefined set of files, as we describe below.

Note: Do not use compaction when writing the import and export list volume files. If compaction is on, data written on a Model B18 VTS EHPO, Model B10 VTS or Model B20 VTS is compressed at the ESCON channel attachment. When the VTS controller internally reads the import or export list file, the data does not pass through the ESCON adapter and therefore does not get decompressed. Either use an appropriate data class, or JCL parameter TRTCH=NOCOMP to prevent the files from being compressed.

In order to ensure that a scratch logical volume is allocated in the target library for the operation, the ACS routines need to have logic to allocate a virtual drive in the target VTS. To accomplish this, you need to have a storage group unique to the VTS library. However, the export and import list volumes do not have to be in SCRATCH status; you can decide to use a predefined volume serial in PRIVATE status to communicate the export and import lists to the VTS.

You can name the import and export list volume files freely; the VTS does not set any requirements on the data set names. Choose the names to fit your naming conventions and, in the case of a scratch allocation, to steer the allocation to the target VTS.

Export operation

The export operation allows you to move logical volumes from the VTS to tape cartridges that are subsequently removed from the VTS. These are the basic steps of an export operation:

- Determine the volumes that are to be removed from the VTS (see “Determine volumes to be exported” on page 376).
- Create an export list volume that contains the list of logical volumes to export (see “Determine volumes to be exported” on page 376).
- Check that you have enough scratch stacked volumes (see “Check for scratch stacked volumes” on page 381).
- Start export. The VTS copies the logical volumes to a set of tape cartridges, and deletes all internal references to the exported logical volumes (see “Run Export function” on page 381 through “OAM messages related to export” on page 386).
- Check export results (see “Analysis of the export status file” on page 387).
- Eject the exported stacked volumes from the library (see “Eject exit” on page 388).

Determine volumes to be exported

Your tape management system keeps track of tape volume vaulting rules and typically provides the information for selecting the volumes that need to be exported to particular destinations.

Run DFSMSrmm storage location management processing (DSTORE) or your tape management system equivalent utility to set the destination for any new volume moves that are required. The required moves would have been determined by normal DFSMSrmm vital record processing or via the DFSMSrmm CHANGEVOLUME subcommand prior to the DFSMSrmm DSTORE processing.

After successful DSTORE processing, you can build a list of volumes to export using DFSMSrmm SEARCHVOLUME command:

```
// EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSTSIN  DD *
 RMM SEARCHVOLUME VOLUME(*) LIMIT(*) LOCATION(vtsname) -
   DESTINATION(dest) CLIST('',',dest') INTRANSIT(N)
```

The resultant output file (EXEC.RMM.CLIST) can be used as input to the export list logical volume file. Use the sample EDGJIMPC to reformat the CLIST output file records to fixed, unblocked format, LRECL 80.

Create the export list volume

Write the required export list volume files on a logical volume in the VTS containing the logical volumes you want to export.

If the volumes you want to export reside in multiple VTS subsystems, you can create one combined export list file and copy it to the export list volumes of each VTS. As export is then run on one VTS subsystem, it skips those volumes that reside in the other VTS subsystems. When finished, you need to combine the export status files from all the export list volumes to check that every volume was successfully exported.

The export list volume

The export list volume is used to pass to the library the list of volumes that you want to export from the VTS and to store status information of the export process (Figure B-6).

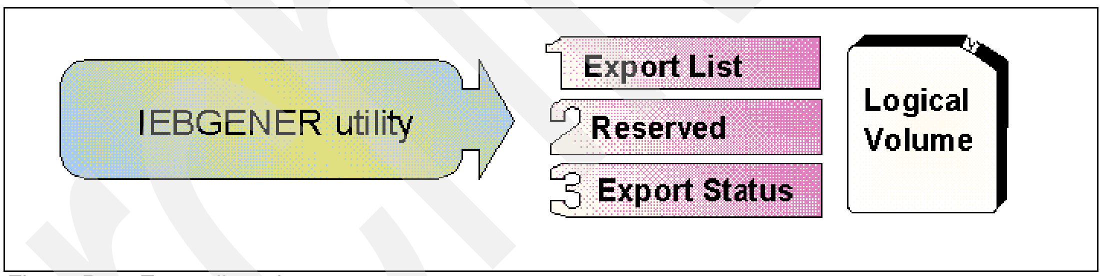

Figure B-6 Export list volume contents

From a VTS perspective, the export list volume is handled as any other logical volume, therefore it might be copied and fragmented, and consequently be recalled upon execution of the export function. It must contain the following three files in this order:

1. Export list file
2. Reserved file
3. Export status file

Export list file:

The export list file contains a list of the volumes to be exported along with their destinations (Figure B-7) and may contain volumes that are not contained in the library in which the export volume is being created.

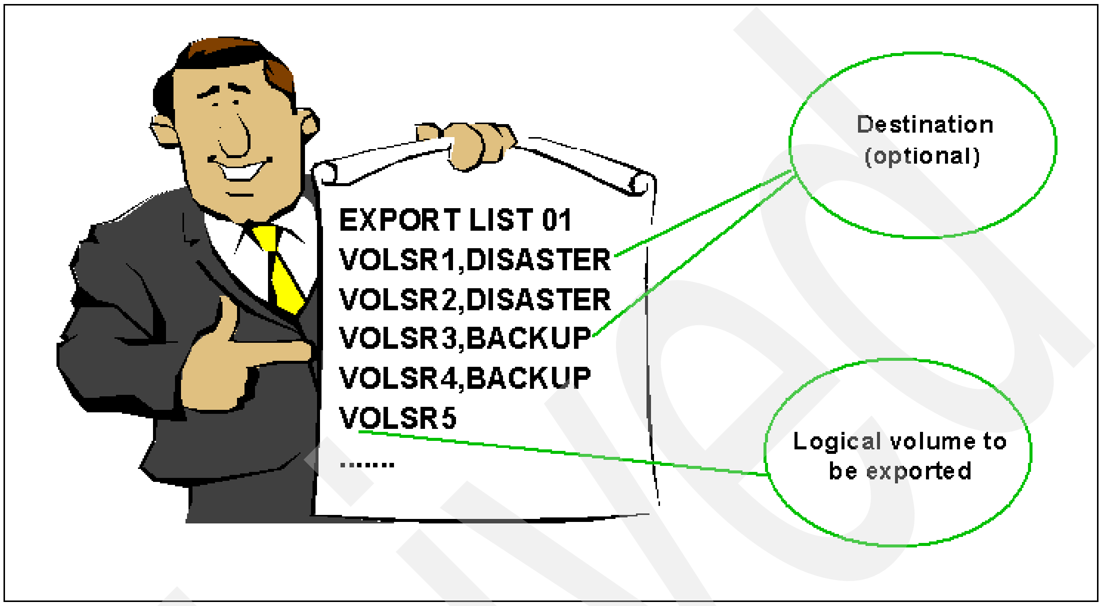

Figure B-7 Export list volume contents

All volumes with the same destination are grouped together on the same set of stacked volumes. A blank destination is valid.

The first record of the export list file is the file identifier and version number and must be specified as illustrated below, starting in the first column:

```
EXPORT LIST 01 <user field>
```

The optional user field is a maximum of 16 characters long, separated with a blank from the version number.

Following the export list file identifier record there is a list of volume records. An export list file record contains two fields, volser and destination, which identify the logical volume to be exported and its destination. The fields must appear in that order separated by a comma. For example:

```
VOLSR2,DISASTER
```

The volser field is 6 characters long.

The destination field for the logical volume is a maximum of 16 characters long. A blank destination is also valid. Use the destination to group the exported logical volumes. Each group of logical volumes to be exported is written on a different set of exported stacked volumes which may then be stored at different locations.

Note: Because the VTS sorts the logical volumes by destination, processing of the exported stacked volumes may be not be completed in the order listed in the export list file.

Reserved file:

The second file on the export list volume is reserved for future use. Even though it is presently not used, you must create this file. It contains one arbitrary record, for example:

```
RESERVED FILE
```

Export status file:

The third file on the export list volume is used for storing status information from the export process. The file identifier and version number is defined in the first record and must be exactly as illustrated below, starting in the first column:

```
EXPORT STATUS 01 <user field>
```

The optional user field is a maximum of 16 characters long, separated with a blank from the version number.

After the export operation is completed, this file contains individual volume export results, including the volser of the exported stacked volume.

Figure B-8 shows sample JCL which writes the required files on the export list volume using a scratch volume:

```
//EXPORT  JOB 1,'EXPORT',MSGLEVEL=(1,1)
//****************************************
//* FILE 1: EXPORT LIST
//****************************************
//STEP1    EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=HILEVELQ.EXPLIST,
//            UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(1,SL),
//            VOL=(,RETAIN),
//            DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1   DD *
EXPORT LIST 01
VOLSR1,DISASTER
VOLSR2,DISASTER
VOLSR3,BACKUP
VOLSR4,BACKUP
VOLSR5
/*
//****************************************
//* FILE 2: RESERVED FILE
//****************************************
//STEP2    EXEC PGM=IEBGENER,COND=(4,LT)
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=HILEVELQ.RESERVED,
//            UNIT=VTS1,DISP=(NEW,KEEP),LABEL=(2,SL),
//            VOL=(,RETAIN,REF=*.STEP1.SYSUT2),
//            DCB=*.STEP1.SYSUT2
//SYSUT1   DD *
RESERVED FILE
/*
//****************************************
//* FILE 3: EXPORT STATUS FILE
//****************************************
//STEP3    EXEC PGM=IEBGENER,COND=(4,LT)
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=HILEVELQ.EXPSTATS,
//            UNIT=VTS1,DISP=(NEW,CATLG),LABEL=(3,SL),
//            VOL=(,,REF=*.STEP1.SYSUT2),
//            DCB=*.STEP1.SYSUT2
//SYSUT1   DD *
EXPORT STATUS 01
/*
```

Figure B-8 Sample JCL to write files on export list volume

You may want to consider designating a dedicated export list volume for use with all exports from a VTS. If no longer needed, it can be deleted. By using a private volume instead of a scratch volume, you can write the export list volume files and start the export process within one batch job.

JCL samples for import and export volume list creation are provided in SYS1.SAMPLIB members:

- CBRSPSIM (import list volume scratch request)
- CBRSPPIM (import list volume private request)
- CBRSPPIP (Sample for Import Policy using a Private Volume)

- CBRSPSXP (export list volume scratch request)
- CBRSPPXP (export list volume private request)

Check for scratch stacked volumes

Before you start export, make sure there are enough scratch stacked volumes in the library to contain the exported logical volumes. To list the numbers of scratch stacked cartridges, find if there is already an Import or Export operation active, and verify the status of the Convenience I/O station, you can issue the following command:

```
D SMS,LIBRARY(libname),DETAIL
```

Output of this command is displayed in Figure B-9.

```
DISPLAY SMS,LIBRARY(LIBVTS1),DETAIL

CBR1110I OAM library status:
TAPE     LIB  DEVICE    TOT ONL AVL TOTAL EMPTY SCRTCH  ON OP
LIBRARY  TYP  TYPE      DRV DRV DRV SLOTS SLOTS   VOLS
LIBVTS1  VL   3494-L10  128 100  50  1443   800    146   Y  Y
----------------------------------------------------------------------
MEDIA       SCRATCH       SCRATCH     SCRATCH**
TYPE          COUNT     THRESHOLD    CATEGORY**
MEDIA1           16            25        0011**
MEDIA2          130            25        0012**
----------------------------------------------------------------------
OPERATIONAL STATE:AUTOMATED
ERROR CATEGORY SCRATCH COUNT:          12
SCRATCH STACKED VOLUME COUNT:         243
PRIVATE STACKED VOLUME COUNT:         400
----------------------------------------------------------------------
Library supports import/export.**
Library supports outboard policy management.**
Host initiated import in process
```

Figure B-9 Display library command listing

> **Note:** In Figure B-9, ** denotes new fields introduced with Advanced Policy Management (FC4001 to FC4004).

If necessary, insert more stacked volumes.

Run Export function

The Export Operation request is initiated by sending a Perform Library Function command that specifies the Library Export order code and the volume serial number of the logical volume that contains the export list file.

To initiate Export, issue the console command:

```
LIBRARY EXPORT,volser
```

Where: *volser* is the volume serial number of the export list volume.

You can initiate the export operation from a batch job using the CBRXLCS program from SAMPLIB which has been updated to support Import/Export functions. Alternatively, you may execute a utility capable of issuing the LIBRARY EXPORT operator command (Figure B-10).

```
//*******************************************
//* WRITE TRANSACTION RECORD IN A TEMP FILE
//*******************************************
//STEP1    EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=&&TEMP,UNIT=SYSDA,DISP=(NEW,PASS),
//            DCB=(RECFM=FB,BLKSIZE=80,LRECL=80),SPACE=(TRK,(1,1))
//SYSUT1   DD *
X  volser
/*
//*************************************************
//* START EXPORT
//*************************************************
//STEP2 EXEC PGM=CBRXLCS
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//INDD     DD &&TEMP,DISP=OLD
//OUTDD    DD DSN=HILEVELQ,TRANSOUT,EXPORT,UNIT=SYSDA,
//            SPACE=(TRK,(1,1)),DISP=(,CATLG)
```

Figure B-10 LIBRARY EXPORT operator command

Column 1 in the input transaction record (INDD statement) contains the transaction code which indicates the requested action. 'X' indicates export. The volser of the export list volume must start in column 4. The CBRXLCS program only initiates the export process without waiting for it to complete. See *z/OS Object Access Method Planning, Installation and Storage Guide for Tape Libraries*, SC26-0427, for additional information on the CBRXLCS program.

Export process flow

The Export Process Flow is described below in Figure B-11.

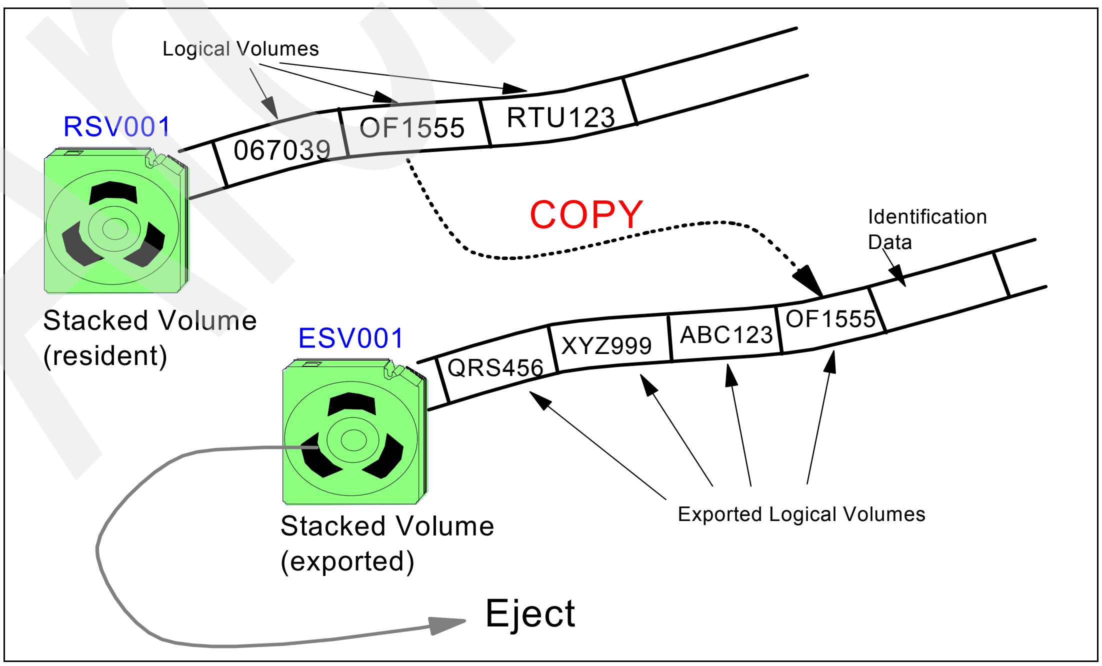

Figure B-11 Export process flow

When export is started, Library Manager first verifies that the export list volume is available, and if so queues the export operation. Export has the same priority as a logical mount. Library Manager also turns on an in-use attribute for the export list volume so that it cannot be used by the hosts while export is running.

When the export operation begins, the VTS first opens and reads the export list and assigns all exportable logical volumes on the list to the export pending category (FF12). A volume is not exportable if it is in use, or not in the VTS executing the export operation. The VTS opens the export status file where it writes records during the export process. A record is written to the status file indicating the reason for a volume not being exportable; the volume is not processed further.

The original category for the logical volume is preserved to allow it to be restored in the event that the volume cannot be successfully exported, or the export operation is terminated or canceled before the volume is exported. Any attempt by a host to mount, audit, or change the category of a volume assigned to the export pending category fails.

The VTS groups the exportable volumes by their destination and processes one destination at a time. Volumes having the same destination are copied to one set of exported stacked volumes. When all volumes for a destination have been processed, the VTS starts exporting volumes for the next destination to a new scratch stacked volume. An exported stacked volume therefore contains only volumes going to the same destination.

The VTS also groups the logical volumes within a destination by the source stacked volume and copies all logical volumes for a destination from a source volume during one mount. An input stacked volume has to be mounted more than once during the export operation, if it contains logical volumes for more than one destination.

The VTS fragments all exportable volumes that are still resident in the TVC: The volumes are copied to tape, if not previously copied, deleted from the TVC and replaced by a stub. The actual export thereafter is a tape-to-tape operation; from the TVC only the fragments are copied on the exported stacked volume. The time required to perform this step depends on the amount of data that has to be copied. Note that exporting volumes which have been fragmented is faster than exporting newly written volumes which have not yet been copied to stacked volumes.

As each exported stacked volume is filled or there are no more logical volume to export for the destination, a 'volume map' file is written on the exported stacked volume that identifies the logical volumes contained on it. This includes each logical volume's volser and its location on the volume.

Note: There is a limitation of 1000 logical volumes per exported stacked volume. Very repetitive data in logical volumes may allow the native cartridge capacity of 3590 cartridges to be exceeded on the exported stacked volume.

After the volume map is written, a file containing the stub information from the Tape Volume Cache for the logical volumes on the exported stacked volume is written on the exported stacked volume. Note that the APM constructs of the logical volumes are copied to the Exported stacked volume together with the logical volume.

When the stub information file has been written, the logical volumes on the exported stacked volume are assigned to the 'Exported' category and the attached hosts are notified. The host passes the information to the tape management system. The Tape Management System is then updated to indicate that the logical volume has been exported and the exported stacked volume it resides on. OAM requests the list of volumes in the exported category from the Library Manager. The list indicates the exported logical volumes and the exported stacked volumes on which they were written. For each volume on the list, OAM takes the following actions:

- OAM calls the CBRUXEJC eject exit and, via the exit, notifies the tape management system of the exported logical volume and the exported stacked volume on which it now resides. The tape management system can then record this information in its database but has no opportunity to fail the subsequent purging of the volume.

Note: VTS Import/Export support today assumes that there is a one-to-one relationship between the tape management system control dataset (CDS) and the tape configuration database (TCDB). So if multiple systems are sharing the TCDB, it is assumed that the tape management system CDS is also being shared. Not all existing customers run with the recommended shared CDS, and as such, the VTS export support will not work with this type of configuration. To support this type configuration, a RC=12 (ignore) option has been added with **APAR OW40969** to the cartridge eject installation exit (CBRUXEJC): If the eject exit is passed an exported logical volume serial number that it does not know about in its CDS, it can now return back with RC=12, leaving the volume in the exported category to be processed by a system that knows about the volume.

DFSMSrmm stores the volser number of the exported stacked volume in the logical volume record's “In container” field (Figure B-12 and Figure B-13).

```
 Panel  Help
-------------------------------------------------------------------------------
EDGPT110             DFSMSrmm Volume Details - A05000
Command ===>

Volume . . . . . . : LOGCL1         Rack number  . . . . . . . :
Media name . . . . : 3490           Status . . . . . . . . . . : MASTER
                                                                 More:     +
Volume type  . . . : LOGICAL        Expiration date  . . . . . : 1999/017
Retention date . . :                Original expiration date . :
Description  . . . :

Data set name  . . :
Media type . . . . : *              Release actions:
Label  . . . . . . : SL               Return to SCRATCH pool . : YES
  Current version  :                  Replace volume . . . . . : NO
  Required version :                  Return to owner  . . . . : NO
Density  . . . . . : *                Initialize volume  . . . : NO
Recording format . : *                Erase volume . . . . . . : NO
Compaction . . . . : *                Notify owner . . . . . . : NO
Attributes . . . . : NONE             Expiry date ignore . . . : NO
Availability . . . :                  Scratch immediate  . . . : NO

Owner  . . . . . . : RMMUSER        Owner access . . . . . . . : ALTER
Assigned date  . . : 1999/012       Assigned time  . . . . . . : 04:44:44
```

Figure B-12   DFSMSrmm Volume Details panel 1

```
Actions pending:
  Return to SCRATCH pool . : NO     Initialize volume  . . . . : NO
  Replace volume . . . . . : NO     Erase volume . . . . . . . : NO
  Return to owner  . . . . : NO     Notify owner . . . . . . . : NO

Location . . . . . : SHELF          Destination  . . . . . . . :
Location type  . . :                In transit . . . . . . . . : NO

In container . . . : STKD01

Storage group  . . :                Home location  . . . . . . : SHELF
                                    Required location  . . . . :
Move mode  . . . . : AUTO           Movement tracking date . . :

Bin number . . . . :                Media name . . . . . . . . :
Old bin number . . :                Media name . . . . . . . . :

Product details:
  Product number . :
  Level  . . . . . :
  Feature code . . :

Enter SCROLL commands for more volume information, or END command to CANCEL.
```

Figure B-13 DFSMSrmm Volume Details panel 2

- OAM requests the Library Manager to purge the logical volume entry from the library inventory.
- Depending on the installation exit, or the EJECT DEFAULT set for the library in SCDS (PURGE or KEEP), OAM either deletes the TCDB entry, or updates the entry to indicate that the volume has been exported.

  If the volume record disposition returned from the exit indicates that the volume's TCDB record should be kept and the shelf location returned from the exit is all blanks, OAM automatically stores the container volser in the shelf location field of the volume record as STACKED=volser. This information can later be used in the volume not in library installation exit (CBRUXVNL) to assist the operator in importing a logical volume. OAM also stores this information if the eject exit is not invoked, the volume record disposition is KEEP and the shelf location in the TCDB record is blank.

  Note that when the eject exit supplied by DFSMSrmm is used, the TCDB entry is always deleted regardless of the PURGE/KEEP setting, as all the required information regarding the exported volume is stored by DFSMSrmm. The DFSMSrmm supplied Volume-Not-in-Library exit (CBRUXVNL) provides this information to the operator so that an import operation can be initiated for the required logical volume.

When all of the volumes in the exported category are purged, the VTS deletes the exported logical volumes from the cache and generates status file records.

Note: Export processing cannot continue until all volumes in the exported category are purged.

The exported stacked volume is then moved to the Export-Hold category in the library and processing continues on to fill another volume for the destination currently being processed. Once all of the logical volumes for a destination have been processed, processing for the next destination begins.

When all logical volumes in the Export List File have been processed, a completion message is broadcast to all hosts attached to the VTS that performed the export process. The completion message results in the generation of a console message that includes the overall completion status of the export operation and metrics about the execution of the operation.

A Library Manager dialog box, **Manage Export-Hold Volumes** (Figure B-17 on page 389), allows an operator to move exported stacked volumes from the Export-Hold category to the Eject category. Exported stacked volumes in the Eject category are moved to the Convenience I/O station by the Library Manager. If an exported stacked volume is needed for an Import Operation, it may be moved to the Import category.

Exported logical volumes do not span physical volumes. If a logical volume does not fit on the current output stacked volume, the VTS changes output volumes and writes the entire logical volume on the next cartridge.

When export has completed, all logical volumes that could not be successfully exported are returned from the export pending category to their original category.

If you are sharing the library between multiple systems, a global resource serialization (GRS) ring allows OAM to serialize the purge processing associated with export. To prevent multiple sharing hosts from attempting to simultaneously process logical volumes in the exported category, OAM sends a SYSTEMS level ENQUEUE around the GRS ring. The resource names are:

```
QNAME - SYSZCBR
RNAME - EXPORTED_CATEGORY_library-name
```

There is no need to include the QNAME and RNAME in the system inclusion RNL, but for documentation purposes you may wish to do so. The `library-name` is the SMS name given to the library when it is defined through ISMF panels. Use of GRS in a shared environment is recommended, but not required.

OAM messages related to export

In addition to the export completion message, status messages are generated during the execution of the export operation. These messages result in the generation of host console messages. See *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280, for further information on host console messages. Messages are generated:

- When export operation has started
- When destination processing begins
- When all volumes on an exported stacked volume have been assigned to the exported category
- As each exported stacked volume is moved into the Export-Hold category
- When the processing for a destination has completed
- When an export operation is suspended
- When an export operation is resumed
- When the time limit of 60 minutes has been exceeded for providing scratch stacked volumes

- When an export operation has been terminated or canceled
- When all export processing has been completed

Some examples for export related messages follow:

Successful completion of the export operation is indicated by the message:

```
CBR3855I Export operation for logical list volume volser in
        library library-name completed successfully.
```

If only some of the logical volumes were successfully exported, the completion is indicated by one of two messages:

```
CBR3856I Export operation for logical list volume volser in
        library library-name completed with exceptions or errors.

CBR3857I Export operation for logical list volume volser in
        library library-name completed with exceptions or errors.
        Statistics for the operation were not available.
```

Messages CBR3855I and CBR3856I include statistics for the operation, for example, the number of volumes and megabytes exported.

When you set up an export operation, you can define multiple export destinations. For each export stacked volume, OAM issues the message:

```
CBR3750 Message from library library-name: message-text
```

Where, *message-text* indicates the destination for the stacked volume. Once export has completed, the tape operator can use these messages to find out where each export stacked volume should be sent.

The following message tells which logical volumes have been exported on a stacked volume:

```
CBR3685I Export processing. Volumes exported from library
        library-name on stacked volume volser
        volser1 volser2 ... volser8
```

`volser1 volser2 ... volser8` are logical volumes that were exported on exported stacked volume `volser`. The message contains a maximum of 48 logical volser numbers. If more volumes were exported on a stacked volume, the message is issued several times.

Analysis of the export status file

Unless the overall completion status of the export operation indicated that no processing exceptions were found, the Export Status File records must be analyzed. The export status file contains an entry for each volume specified on the export list indicating whether the export was successful or not. It is a flat file in EBCDIC format. You can use any standard tool to copy or print it.

The export list may contain volumes that do not exist in this VTS. A message is written in the status file to indicate this, and the process continues with the next volume.

Figure B-16 shows sample JCL to read the export status file.

```
//STEP     EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT1    DD DSN=HILEVELQ.EXPSTATS,
//             UNIT=3490,DISP=OLD,LABEL=(3,SL),
//             VOL=SER=volser
//SYSUT2    DD SYSOUT=*
```

Figure B-14 Sample JCL to read the export status file

The status file might look like as displayed in Figure B-15.

```
EXPORT STATUS 01
VOLSR1,      ,18,*No Data Associated With Logical Volume
VOLSR4,      ,18,*No Data Associated With Logical Volume
VOLSR5,      ,05,*Logical volume not in VTS
VOLSR2,STKD01,00, DISASTER
VOLSR3,STKD02,00, BACKUP
```

Figure B-15 Status file

In the above example, logical volume VOLSR1 was not exported (status code 18) because no data was associated with it, therefore it probably was not used after it had been defined. Logical volume VOLSR5 (status code 05) was not found in this VTS. Logical volume VOLSR2 (status code 00) was successfully exported to export stacked volume STKD01 with destination DISASTER. Refer to *IBM TotalStorage 3494 Tape Library Operator's Guide* for a complete listing of the fields and status codes.

Eject exit

The export function requires that OAM is active, that the library for the export operation is online and operational and that the eject installation exit CBRUXEJC is not disabled. The OAM address space does not have to remain active on the host which initiates the export operation, as long as it remains active in one of the hosts sharing the TCDB. If your tape management system does not supply an eject exit and you use the default provided by DFSMSrmm, the exit will be in bypassed state; this does not prevent the export function from working.

Use the following commands to check the status of the eject exit and if necessary to enable it:

```
DISPLAY SMS,OAM
LIBRARY RESET,CBRUXEJC
```

The output of the command is displayed in Figure B-16.

```
D SMS,OAM
CBR1100I OAM status: 000
TAPE TOT   ONL   TOT   TOT   TOT   TOT    TOT    ONL    AVL    TOTAL
     LIB   LIB   AL    VL    VCL   ML     DRV    DRV    DRV   SCRTCH
       2     2    1     1     0     0     76     43     33     1888
There are also 0 VTS distributed libraries defined.
CBRUXCUA processing ENABLED.
CBRUXEJC processing ENABLED.
CBRUXENT processing ENABLED.
CBRUXVNL processing ENABLED.
```

Figure B-16 Exit Status display

The following command indicates if export operations are in process in the library:

```
DISPLAY SMS,LIBRARY(libraryname)
```

Multiple export operations may run concurrently among the VTS subsystems in a physical library.

Use Library Manager console to eject the exported stacked volumes from the library in order to store them on shelf or to send them off site.

Eject exported stacked volumes

The exported stacked volumes are held in the library in the export hold category until an operator releases them. You can eject one exported stacked volume at a time, wait until export for the destination has completed, or eject all of them when the export operation has completed. Use Library Manager panels to eject the exported stacked volumes (Figure B-17).

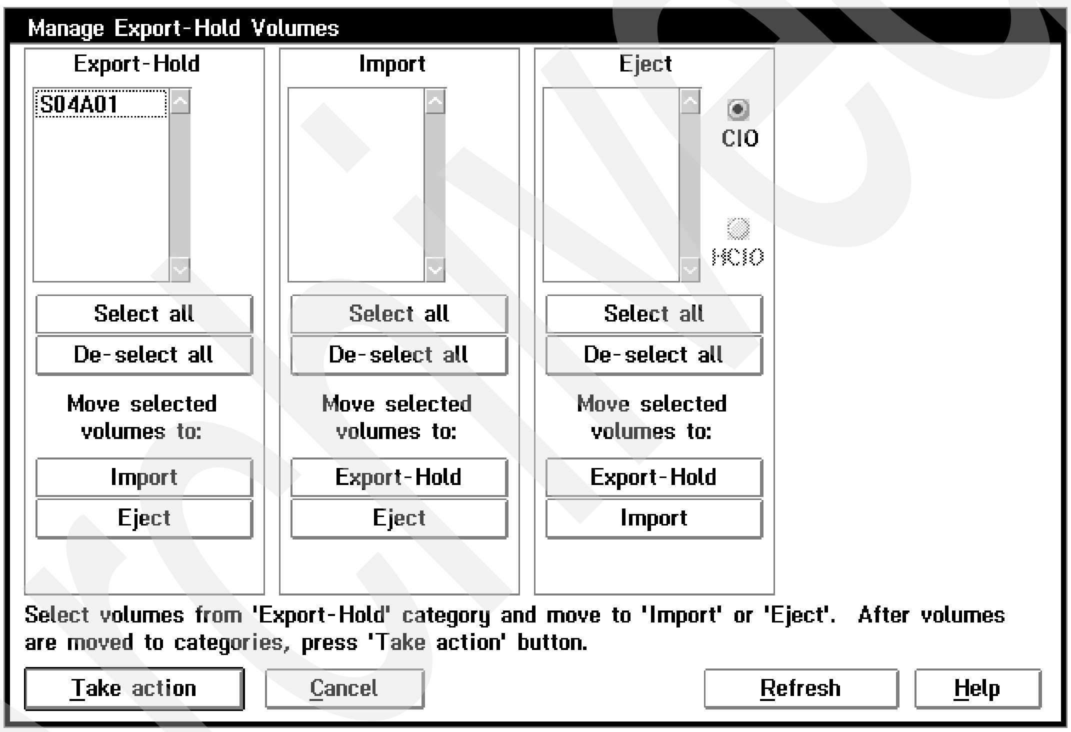

Figure B-17 Manage Export Hold Volumes window

Determine exported volumes

One method to find exported logical volumes and to locate the exported stacked volumes on which they reside is to use the RMM command (RMM SEARCHVOLUME) or the “In container” field on the panel shown in Figure B-18 and Figure B-19.

```
 Panel  Help  Scroll
------------------------------------------------------------------------------
EDGPT010                    DFSMSrmm Volume Search
Command ===>

                                                                  More:     +
Volume  . . . .             May be generic.  Leave blank for all volumes.

Owner . . . . .             Owned by a specific user. Default is your userid

Media name  . .             Limit to a single media name
Pool  . . . . .             or to a particular pool

Status  . . . .
                            Select volume status.  Default is ALL.
Since . . . . .             Volumes assigned since YYYY/DDD
Retention . . .             Volumes retained up to YYYY/DDD
Limit . . . . .             Limit search to first nnnn volumes

Home  . . . . .             Limit to volumes with this home location name
Location  . . .             Limit to volumes in a single location
In container                Physical stacked volser
```

Figure B-18 DFSMSrmm Volume Search panel 1 of 2

```
Volume type . .             ( LOGICAL or PHYSICAL )
Media type  . .             Tape media type ( *, CST, ECCST, HPCT or EHPCT)
Label . . . . .             Tape label standard ( for example SL )
  Current version           Label version number( for example 3 )
  Required version          Label version number( for example 4 )
Density . . . .             Tape recording density
Format  . . . .             Tape recording format ( *, 18TRACK, 36TRACK or
                                                    128TRACK )

Compaction  . .             Limit to volumes containing compacted data
Attributes  . .             Tape special attributes ( NONE or RDCOMPAT )


Destination . .             Limit by destination
Move mode . . .             Limit by move mode ( AUTO or MANUAL )
Intransit . . .             Limit to volumes which are moving ( YES or NO )

The following line commands are available when the list is displayed:
  C - Change volume information            E - Eject volume
  I - Data set information                 L - List volume chain
  O - Owner information                    R - Release volume
  V - Volume information
Enter SCROLL commands for more search values, ENTER to SEARCH or END to CANCEL.
```

Figure B-19 DFSMSrmm Volume Search panel 2 of 2

Canceling export

Depending on the amount of data you are exporting, export can take a long time and you may need to cancel it; for example, to free the VTS resources used by export for other work. The first method of canceling an export operation is to issue this command:

```
LIBRARY EXPORT,volser,CANCEL
```

In this command, *volser* is the volume serial number of the export list volume that was specified when the export operation was started.

You may cancel the export operation at any time. When the export is canceled at the host:

- The VTS completes exporting the logical volume that is currently processing.
- The current exported stacked volume is processed and assigned to the export hold (FF17) category.
- All logical volumes still assigned to the export-pending category are reassigned to the category they were in prior to the start of the export operation.
- Status File Records are written for all logical volumes that had not completed processing.
- A completion message is sent to all hosts indicating that the export operation was canceled via a host request.
- Any volumes that were successfully exported before the cancel command was issued, that is, volumes that were copied to exported stacked volumes and deleted from the VTS, are unaffected.

If you later want to rerun the export operation, you can use the same export list volume without having to modify it. Volumes in the export list file that have already been exported are skipped during the rerun, with a record written in the status file to indicate that it is no longer in the VTS. The VTS overwrites status records from the previous export run by writing new status records to the export status file.

A second method is to cancel an export operation from the Library Manager panel shown in Figure B-20. Canceling the operation through the Library Manager is a hard cancel; it will quickly terminate the operation and will not wait for the host to finish processing the volumes in the exported category. It should be used only when there is no functional host that can issue the cancel.

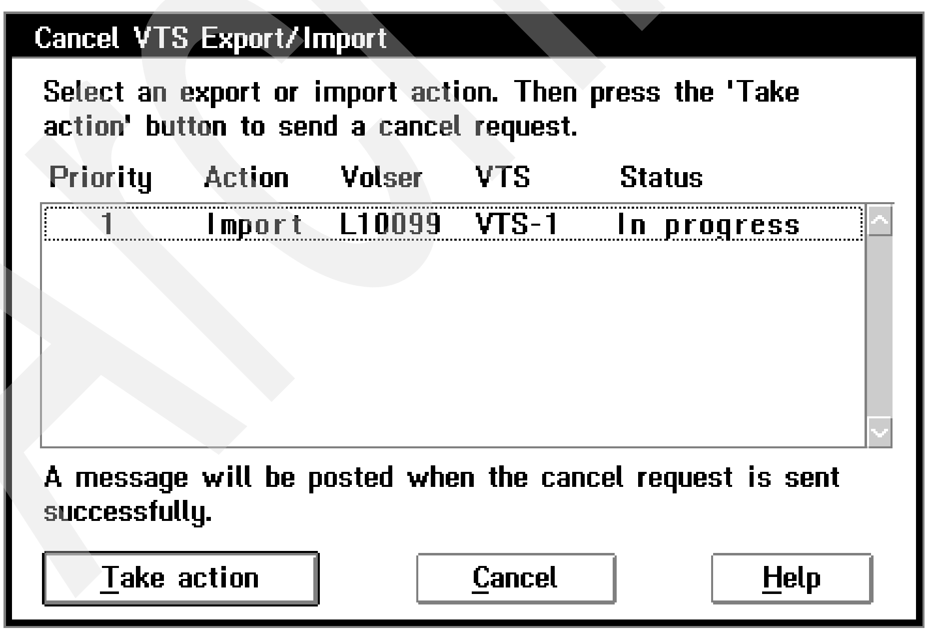

Figure B-20 Cancel Export using Library Manager panels

Depending upon when the cancel is received, the VTS does one of the following:

- It places the logical volumes which had been copied to this incomplete exported stacked volume (and purged by the host processing of the Exported category) into the Insert category for processing when the host becomes available again.
- It returns the logical volumes which would have been exported to their original categories.

Export performance

There are two considerations in defining the performance of an Export Operation — first, the performance of the Export Operation itself; and second, the impact that an Export Operation has on the overall performance of a VTS.

The performance of an Export Operation itself is defined in terms of its total elapsed time. The elements that influence the performance are:

- The number of logical volumes to export:
  - For each logical volume, you should calculate a search time of 20 seconds to locate the logical volume on the stacked volume.
- The amount of data to be copied:
  - It is assumed that data can be moved from source stacked volume to exported stacked volume at a rate of 15 MB/s (3:1 Compression Ratio)
- The number of source stacked volumes to be mounted:
  - For each source stacked volume, there is a nominal processing overhead of 120 seconds to cover robotics mount and demount, load/thread and rewind/unload
- The number of exported stacked volumes created:
  - For each exported stacked volume created, there is an overhead of 120 seconds to write the volume table of contents and the logical stub information files.

Performance Tool

A Spreadsheet to calculate estimated time to accumulate virtual volumes for export named **EXIMCALC.123** is available at:

`http://w3.rmss.tucson.ibm.com/softw/tapetool.htm`

The impact on the overall performance of the VTS can be seen in two primary areas:

- **Data Throughput:** Although export does not move data through the Tape Volume Cache, it uses 3590 SCSI bus bandwidth and VTS processor cycles. The SCSI bandwidth is likely to be insufficient to support maximum data throughput concurrently with an export operation. The throughput bandwidth of the VTS may be reduced slightly.
- **Logical Volume Recall Time:** Since an export operation uses two physical drives of the VTS, fewer drives are available for use to recall logical volumes. This may result in longer logical mount times when the number of pending recalls exceed the number of physical drives less the two used for the export operation.

Import/Export recommendations

We recommend that Import/Export tasks be run during non-peak I/O activity times to minimize contention.

In order to reduce physical mounts required for Export, virtual volumes with similar creation dates should be exported together. This increases the probability that multiple volumes are found on the same physical source stacked volume.

Import operation

The import process allows you to copy logical volumes from exported stacked volumes back to a VTS. The basic steps in an import process are:

- Create an import list volume which contains the list of logical volumes you want to import (see “Create the import list volume files” on page 395).
- Define the logical volumes you are importing to the receiving tape management system if required (see “Define volumes to the tape management system” on page 400).
- Insert the exported stacked volumes (cartridges containing volumes to be imported) into the target VTS library through the convenience I/O station (see “Insert exported stacked volumes in the library” on page 401).
- Assign the exported stacked volumes to the import category (FF14) using Library Manager panels (see “Assign the exported stacked volumes” on page 401).
- Run import. The VTS copies logical volumes from the exported stacked volumes onto stacked volumes already in the library and creates internal references to each imported logical volume (see “Run Import function” on page 402).
- Check import results (see “Analysis of the export status file” on page 387 through “Analysis of the import status file” on page 408).
- Eject exported stacked volumes after import (see “Eject exported stacked volumes after import” on page 408).

Assumptions and definitions

There are several assumptions and definitions for the import operation, as follows.

Individual logical volumes or the entire contents of an exported stacked volume can be imported.

A logical volume, called the Import List Volume, is used to specify the logical volumes to import and to return status information. The Import List File can include either logical volumes, exported stacked volumes pairs or just exported stacked volumes. If only an exported stacked volume is specified, all logical volumes that reside on that volume are to be imported.

The library does not import a logical volume whose volser is a duplicate of one that is already resident in the library.

All data to import is copied to the VTS managed storage pool.

The exported stacked volume, that contains logical volumes to be imported, must be assigned to the import category by the operator prior to the initiation of the import operation

The volumes moved to the import category by the operator remain in the library after an import operation completes. They can be ejected or assigned for use within native 3590 or VTS by an operator through the Library Manager console.

While each exported stacked volume is being processed during an import operation, it is temporarily moved out of the import category to prevent an operator from attempting to eject or reassign the volume to another category.

Import and Export are mutually exclusive in a VTS. Only one Import operation at a time is allowed in a physical library.

Data being imported is not written into the TVC. The stub information file on the exported stacked volume is read and the stub files are written to the TVC.

The import operation is stopped by the VTS when less than four physical drives are found.

If you want to import logical volumes only to make the logical volumes available for reuse by the Tape Management System without importing any data there are two options available. First, specify that the volume contains scratched data and that only the prior volume label data needs to be kept in the VTS. Second, specify that the volume is to be re initialized. None of the data from prior use is kept and the first time the volume is mounted, it contains only the data records of a newly initialized volume.

When you perform an export with APM enabled, all these policies are exported as part the export set. As a result new support has been added to the import function to enable policies defined to be overridden. Additionally volumes exported from a non-APM capable library can have policies assigned at import time when the import target library supports APM.

New command support has also been added. The command known as the Library Manager Policy command enables you to change policy (or construct) names for the objects (in this case tapes).

Here is the format of the command:

```
{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*}]
```

Additionally, the command can be abbreviated to:

```
LI LP,volser,SC=mysc,MC=mymc,DC=*RESET*,SG=mysg
```

{LMPOLICY|LP}

Specifies a request to set one or more of a private volumes storage group to blanks in the TCDB and to the default storage group policy in the library which is also blanks.

volser

Specifies the volume serial number of a private volume which resides in a library with outboard policy management support.

SG={storage-group-name|*RESET*}

Specifies a construct name for the SG keyword. If the request is successful, the construct name becomes the storage group for the volume in the TCDB and the storage group policy in the library. If *RESET* is specified, the request is to set the volumes storage group 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 keyword. If the request is successful, the construct name becomes the storage class policy for the volume in the library. If *RESET* is specified, the request is to set the volume’s management class to the default management class policy in the library which is blanks.

MC={management-class-name|*RESET*}

Specifies a construct name for the MC keyword. If the request is successful, the construct name becomes the management class policy for the volume in the library. If *RESET* is specified, the request is to set the volume's management class to the default management class policy in the library which is blanks.

DC={data-class-name|*RESET*}

Specifies a construct name for the DC keyword. If the request is successful, the construct name becomes the data class policy for the volume in the library. If *RESET* is specified, the request is to set the volumes data class to the default data class policy in the library which is blanks.

With Advanced Policy Management (FC4001-4004) there are further changes to OAM to enable the facilities of APM to be used. The following DFSMS construct names are now passed to enable various outboard functions to be performed by the VTS:

- Storage class
- Storage group
- Management class
- Data class

You can determine the constructs assigned to a particular volume by issuing the command:

```
DISPLAY SMS,VOL(volser)
```

The output of this command is shown in Example B-1.

Example: B-1 New OAM volume status display showing defined construct names

```
D SMS,VOL(SK0008)
 F OAM,D,VOL,SK0008,L=BHAEUSS-Z
 CBR1180I OAM tape volume status: 743
 VOLUME  MEDIA     STORAGE   LIBRARY   USE  W  C  SOFTWARE  LIBRARY
         TYPE      GROUP     NAME      ATR  P  P  ERR STAT  CATEGORY
 SK0008  MEDIA2    SG4PP03   STARLITE  P    N     NOERROR   NOTAVAIL
 -------------------------------------------------------------------
 RECORDING TECH:     36 TRACK        COMPACTION:        NO
 SPECIAL ATTRIBUTE:  NONE            ENTER/EJECT DATE:   2005-04-25
 CREATION DATE:      2005-04-25      EXPIRATION DATE:    2007-04-15
 LAST MOUNTED DATE:  2005-04-25      LAST WRITTEN DATE:  2005-04-25
 SHELF LOCATION:
 OWNER: BHAEUSS
 LM SG: TAPESG2   LM SC: TAPESC2   LM MC: TAPEMC2   LM DC: MEDIA2DC
 -------------------------------------------------------------------
```

Create the import list volume files

Write the required import list file on a logical volume in the VTS where you want to import the logical volumes. Refer to , "Import and export list volumes" on page 375. An import of one or more logical volumes can be initiated independently of DFSMSrmm by creating the import list volume files and requesting the import using the CBRXLCS macro or the LIBRARY command.

You can use the DFSMSrmm SEARCHVOLUME command to build the list of logical volumes with their containing stacked volume and status:

```
RMM SV VOL(*) LIM(*) OWN(*) DEST(vtsname) TYPE(LOGICAL) CLIST
```

With the TYPE(LOGICAL) operand specified, the resulting output file contains:

```
STKD01,LOGCL1,status
```

The resultant output file can be used as input to the import list logical volume file. Use the sample EDGJIMPC to reformat the CLIST output file records to fixed, unblocked format, LRECL 80.

Import list volume

The import list volume is used to pass to the library the list of volumes that you want to import into the VTS, and to store status information of the import process (Figure B-21).

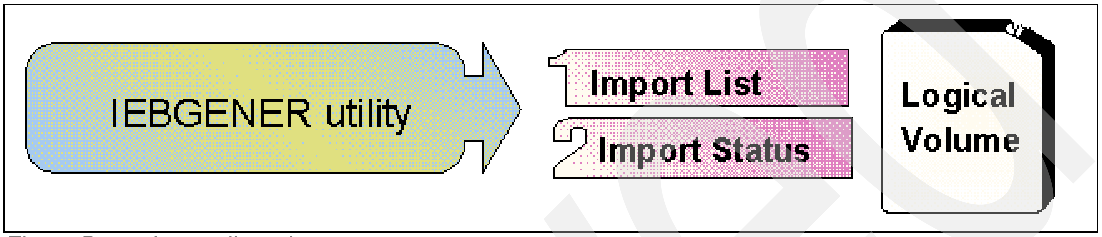

Figure B-21 Import list volume contents 1

From a VTS perspective, the import list volume is handled as any other logical volume, therefore it might be copied and fragmented, and consequently be recalled upon execution of the import function. It must contain the following two files in this order:

Import list file

The import list file contains a list of volumes to be imported. You can import only selected logical volumes or all logical volumes from a stacked volume (see Figure B-22 and Figure B-23).

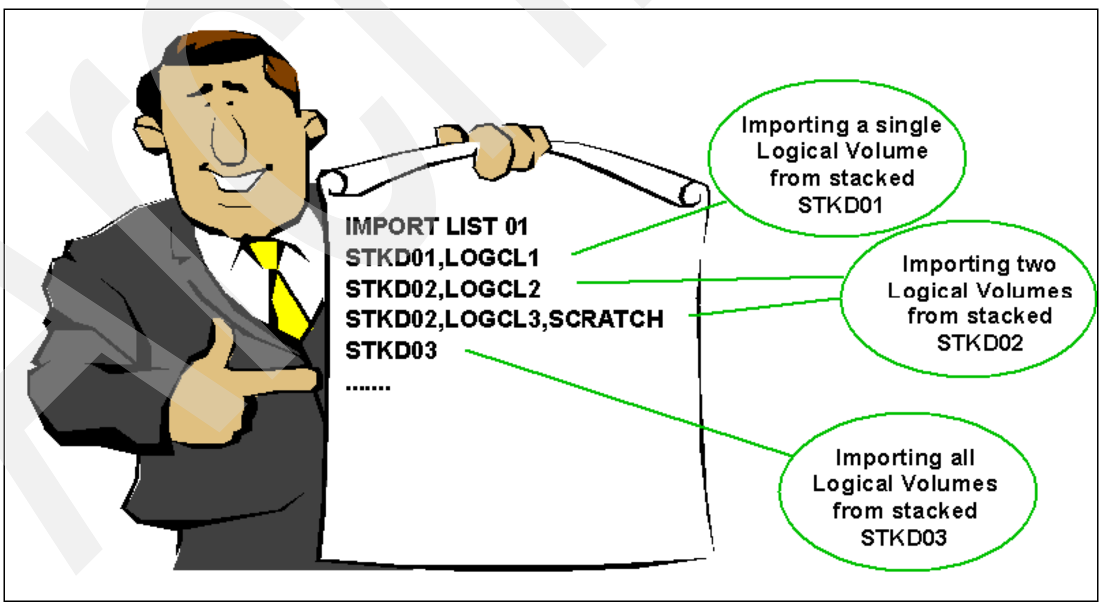

Figure B-22 Import list volume contents 2 without APM

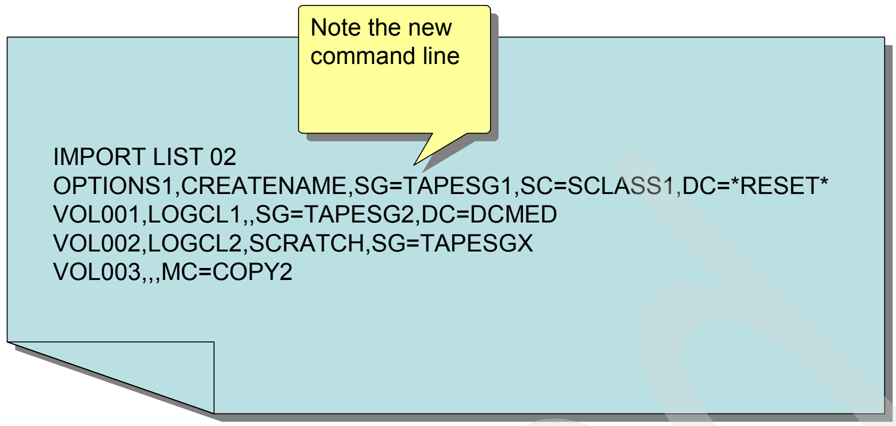

Figure B-23 Import list volume contents with APM

The first record of the import list file is the file identifier and version number and must be specified as illustrated below, starting in the first column.

For FC4000, use:

```
IMPORT LIST 01 <user field>
```

For FC4001-4004, use:

```
IMPORT LIST 02 <user field>
```

The optional user field is a maximum of 16 characters long, separated with a blank from the version number.

With an Advanced Policy Management (FC4001-4004) enabled VTS, the line following the import list file identifier can be used to identify any Outboard Policy Management (OPM) options. You must start in column one with the OPTIONS1 keyword and the options specified will apply to all volumes in JCL. The options that you have available are:

FAILIMPORT Fails the import of the logical volume if any of its policy names are not predefined in the library.

CREATENAME Automatically creates the policy at the library and assigns the default actions.

USEDEFAULT Assigns default policies for any policy names that are not predefined at the library.

You can override the OPM options specified as part of the export by using the OPTIONS1 keyword. For example:

```
OPTIONS1,CREATENAME,SG=TAPESG1,SC=SCLASS1
```

This statement would set the global policy parameters for all logical volume listed in the JCL to the following:

Storage group	`TAPESG1`
Storage class	`SCLASS1`

For a VTS without APM or where no OPTIONS1 keyword is specified, there is a list of records containing a number of fields which follow the import list file identifier record. The first three fields consist of the exported stacked volume, the logical volser and the import option. The fields are positional and as such must appear in that order separated by a comma. Both the logical volser and the import option are optional. The exported stacked volume field is 1-6 characters long and represents the exported stacked volume containing logical volumes to be imported.

The logical volser field is 1-6 characters long. It identifies the logical volume to be imported. If the field is all blanks, all logical volumes on the specified exported stacked volume are imported. If an import option is specified then the logical volume field must be delimited by a comma prior to the import option keyword.

For example:

```
STKD02,LOGCL3,SCRATCH
STKD02,,SCRATCH
```

The import option allows you to choose how the volume is handled during import. These are the options:

Blank The import option can be omitted (blank) which indicates that the data content of the logical volume is copied into the VTS subsystem and a VTS database entry, a stub entry in the TVC and the Library Manager record for the volume are created. After the import has completed, the logical volume is available for processing in the VTS. If requested for input processing, the newly imported logical volume needs to be recalled from the stacked cartridge.

SCRATCH Indicates that the data content of the logical volume is not copied into the VTS. The SCRATCH option can be specified when the data on the logical volume to import is known to have been expired and is not accessed after the volume is imported into the VTS subsystem. This option can be specified for individual logical volumes or the entire contents of an exported stacked volume. An entry for the logical volume in the VTS database, a stub in the TVC, as well as a Library Manager record are created. After the import has completed, the logical volume can be used as a scratch volume.

INITIALIZE Option can be specified when the logical volume to import is to be initialized and any prior data on the volume discarded. This option can be specified for individual logical volumes or the entire contents of an exported stacked volume. Data will not be copied to a resident stacked volume and stub information with the volume label information will not be written to the cache. Only the Library Manager record for the volume is created. After the import has completed, the volume is like a newly inserted logical volume that has never been used.

All keywords have to be spelled out to be correctly recognized.

If a logical volser is not included with an exported stacked volume, the import option specified applies to all logical volumes on the exported stacked volume.

The import list may include logical volumes or exported stacked volumes that are not found on the exported stacked volumes being used for the import operation.

Use import option 'SCRATCH' or 'INITIALIZE' to prevent unnecessary copying of the data contents of the logical volume. This is useful if the logical volume has expired or returned to scratch after it was exported.

Note: Data contents of logical volumes imported with the SCRATCH or INITIALIZE option cannot be recovered later on, because the exported stacked volume may be reused and overwritten as soon as all logical volumes have been successfully imported.

For an APM enabled VTS, the lines of file records can also contain policy name specifications (Example 8-1). These fields must start after the import option field. These fields conform to the same syntactical rules as the previous fields but are not positional in themselves. Any one of the policy options, or all of them, can be specified in any order. The policy name field is 1-8 characters long:

SG= Storage group name
SC= Storage class name
MC= Management class name
DC= Data class name

Additionally, the policy name of ***RESET*** can be used to assign a policy name of blanks (the default policy name for each category) at the Library Manager.

Important: Any policy name that you define on your file records will override any policies defined by either the **Export** function or the **Options1** keyword.

Example 8-1 Shows some potential file record statements to be used as part of import

```
SVL001
SVL001,LVL001
SVL001,LVL001,SCRATCH
SVL001,,INITIALIZE
SVL001,LVL001,OPTION,SG=NAME,SC=NAME,MC=NAME,DC=*RESET*
SVL001,LOGICAL,,SG=NAME,SC=NAME,MC=NAME,DC=NAME
SVL001,,,SG=NAME,SC=*RESET*,MC=NAME,DC=NAME
SVL001,LVL001,SCRATCH,SG=NAME
SVL001,LVL001,DC=NAME,SC=NAME,MC=NAME
```

Import status file:

The second file on the import list volume is used for storing status information from the import process. The file identifier and version number is defined in the first record and must be exactly as illustrated below, starting in the first column:

```
IMPORT STATUS 01 <user field>
```

The optional user field is a maximum of 16 characters long, separated with a blank from the version number. After the import operation is completed, this file contains individual volume import results. Figure B-24 shows sample JCL which writes the required files on the import list volume using a private volume:

```
//IMPORT  JOB 1,'IMPORT',MSGLEVEL=(1,1)
//***************************************
//* FILE 1: IMPORT LIST FILE
//***************************************
//STEP1    EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=HILEVELQ.IMPLIST,
//            UNIT=VTS2,DISP=(NEW,KEEP),LABEL=(1,SL),
//            VOL=(,RETAIN,,,SER=IMP001),
//            DCB=(RECFM=FB,BLKSIZE=80,LRECL=80,TRTCH=NOCOMP)
//SYSUT1   DD *
IMPORT LIST 01
STKD01,LOGCL1
STKD02,LOGCL2
STKD02,LOGCL3
STKD03
/*
//***************************************
//* FILE 2: IMPORT STATUS FILE
//***************************************
//STEP2    EXEC PGM=IEBGENER,COND=(4,LT)
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=HILEVELQ.IMPSTATS,
//            UNIT=VTS2,DISP=(NEW,CATLG),LABEL=(2,SL),
//            VOL=(,,REF=*.STEP1.SYSUT2),
//            DCB=*.STEP1.SYSUT2
//SYSUT1   DD *
IMPORT STATUS 01
/*
```

Figure B-24 Sample JCL to write required files on import list volume

Note: Import List 01 and Import List 02 are supported by libraries with outboard policy management support; however, only Import List 01 is supported by libraries with no outboard policy management support.

Define volumes to the tape management system

Your tape management system may require that the logical volumes you are importing are defined to the receiving tape management system. If necessary, use the appropriate tape management system commands to define the volumes before you run import. To find out which logical volumes are contained on an exported stacked volume, you need to check the export status file or the CBR3685I messages for the corresponding export operation.

Volumes being imported which are unknown to DFSMSrmm can be automatically added to the control data set in a fashion similar to that supported for native volumes being entered into a system-managed library. DFSMSrmm fails an import of a logical volume, if it has a record for the volume but the volume is not known to be an exported volume. If desired, an exported logical volume can be defined to DFSMSrmm. Use the TYPE and CONTAINER operands of the ADDVOLUME subcommand:

```
AV volser LOCATION(SHELF) CONTAINER(svol) STATUS(MASTER) TYPE(LOGICAL)
```

If you import a logical volume back to the same RMM environment from which it was exported, DFSMSrmm knows the volume as an exported logical volume.

Insert exported stacked volumes in the library

The insert process and the assigned categories are different for 3494s without VTS subsystems and for 3494s with VTS subsystems and the Advanced Function feature installed. Refer to 6.4.4, “Inserting stacked volumes in the VTS” on page 229 for a description of the process.

Use this host command to display the status of the convenience I/O station:

```
D SMS,LIBRARY(libname),DETAIL
```

Verify that the convenience I/O station is in *Import* mode (see , “Run Export function” on page 381).

Assign the exported stacked volumes

When 3590 cartridges are inserted into a 3494 Tape Library that has a VTS with Advanced Function feature installed, all cartridges are assigned to the unassigned category (FF16) in the Library Manager and the operator, using a Library Manager menu, may move the cartridges into the import category, or the insert category, or may eject the cartridge (Figure B-25).

Note: The 3584 Tape Library with 3953 Library Manager supports 3592 drives only.

Cartridges are placed by the Library Manager into insert (FF00) or VTS insert (FF01) categories based upon volume ranges previously defined. Insert the exported stacked volumes in the target library through the convenience I/O station. The Library Manager places the volumes in empty storage cells in the library and assigns them to the unassigned category (FF16). Once you have inserted all of the exported stacked volumes, use the Library Manager menus to move the cartridges from the unassigned category to the import category (FF14) where they must be when the import operation begins.

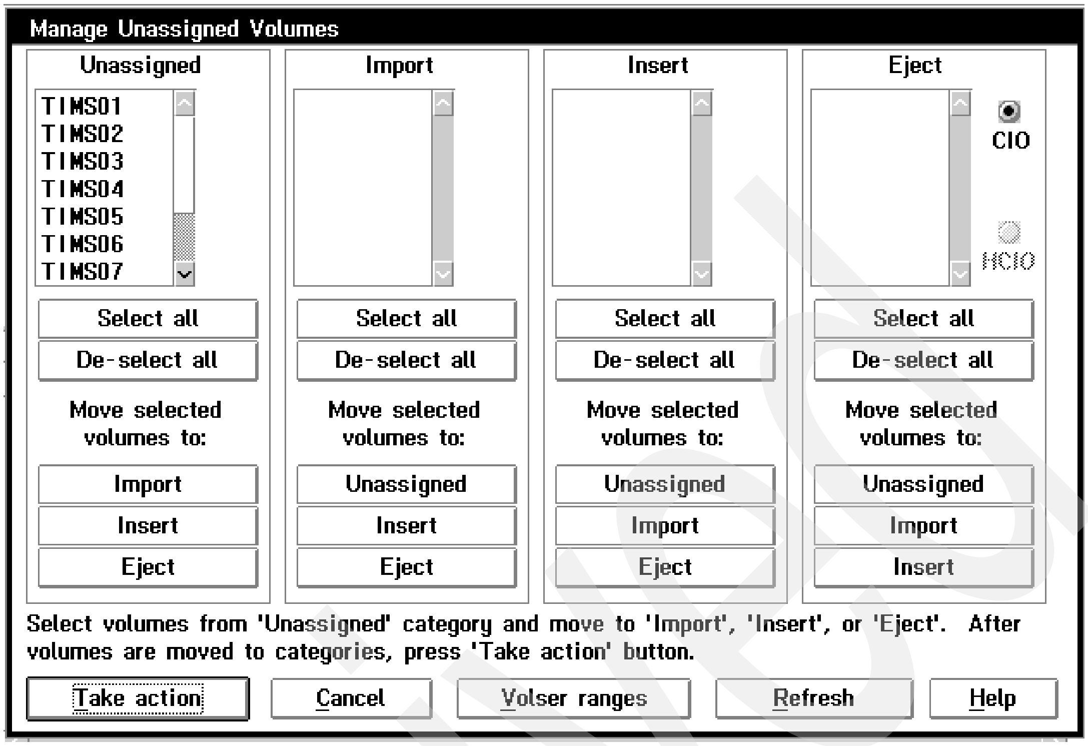

Figure B-25 Manage Unassigned Volumes window

Make sure you do not move an exported stacked volume into the insert category, since you lose all exported logical volumes on it when the VTS uses it as a scratch stacked volume or a native 3590 drive uses it as scratch cartridge.

Note: Cartridges inserted through the high-capacity I/O facility are not assigned to the Unassigned category, but to an insert category, either FF00 or FF01, based upon volume ranges previously defined. Consequently, exported stacked volumes must not be inserted through the high-capacity I/O facility.

Run Import function

Import increases the amount of data in your VTS. Before running import, make sure you have enough empty stacked volumes in the VTS to store the imported logical volumes.

Once you have added all the exported stacked volumes needed for the import into the library and moved them to the import category, you can start the import process. To initiate import, issue the operator command:

```
LIBRARY IMPORT,volser
```

In this command, `volser` is the volume serial number of the import list volume. You can initiate the import operation from a batch job using the CBRXLCS program from SAMPLIB which has been updated to support Import/Export functions. Alternatively, you may execute a utility capable of issuing the LIBRARY IMPORT operator command. See Figure B-26.

```
//*****************************************
//* WRITE TRANSACTION RECORD IN A TEMP FILE
//*****************************************
//STEP1    EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT2   DD DSN=&&TEMP,UNIT=SYSDA,DISP=(NEW,PASS),
//            DCB=(RECFM=FB,BLKSIZE=80,LRECL=80),SPACE=(TRK,(1,1))
//SYSUT1 DD *
I volser
/*
//*****************************************
//* START IMPORT
//*****************************************
//STEP2    EXEC PGM=CBRXLCS
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//INDD     DD DSN=&&TEMP,DISP=OLD
//OUTDD    DD DSN=HILEVELQ.TRANSOUT.IMPORT,UNIT=SYSDA,
//            SPACE=(TRK,(1,1)),DISP=(,CATLG)
```

Figure B-26 Sample JCL to write transaction record in a temp file

Column 1 in the input transaction record (INDD statement) contains the transaction code which indicates the requested action. 'I' indicates import. The volser of the import list volume must start in column 4. CBRXLCS program only starts the import process without waiting for it to complete. See *z/OS Object Access Method Planning, Installation and Storage Guide for Tape Libraries*, SC26-0427, for additional information on the CBRXLCS program.

The import function requires that OAM is active, that the library for the import operation is online and operational, and that the entry installation exit CBRUXENT is not disabled. The OAM address space does not have to remain active on the host which initiates the import operation as long as it is active on one of the hosts sharing the TCDB. Use the following commands to check the status of the exit, and if necessary to enable the exit:

```
DISPLAY SMS,OAM
LIBRARY RESET,CBRUXENT
```

The following command indicates if an import operation is already in process in the library:

```
DISPLAY SMS,LIBRARY(libraryname)
```

Only one import operation at a time can be initiated or be in execution in a library even if there are multiple VTS subsystems in the library.

Import process flow

The steps of the import process flow are described as follows (see Figure B-27).

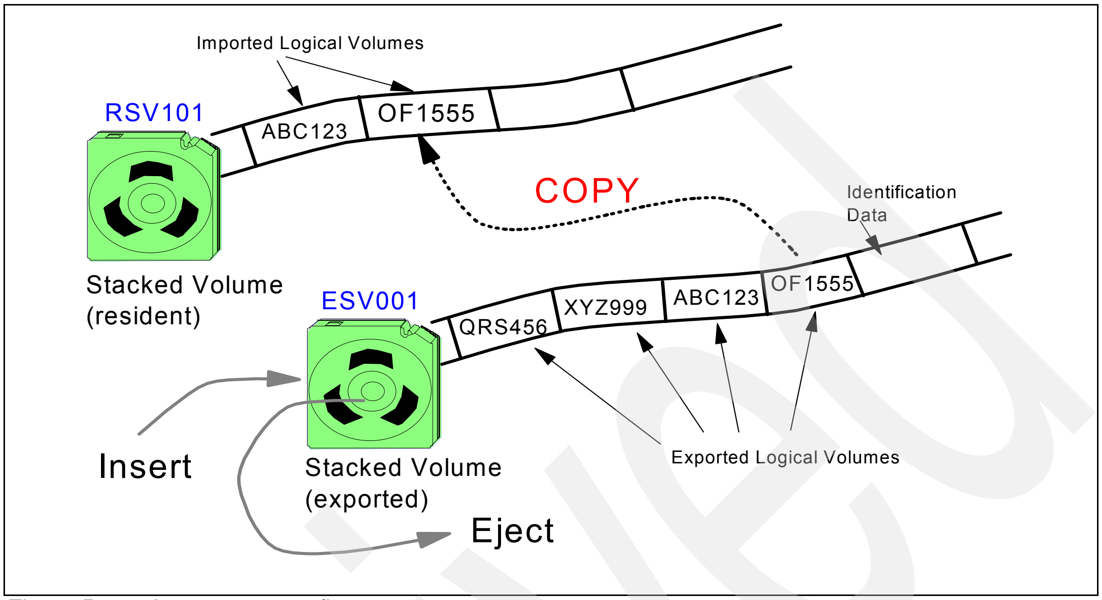

Figure B-27 Import process flow

- First, the VTS checks that the exported stacked volume is in the import category (FF14). If not, a status record is written and the volume is not processed further.
- If the volume is in the import category, the VTS mounts it, reads the volume map file and compares it against the logical volumes specified in the import file list and creates a list of candidate volumes to be imported from the stacked volume. The list contains either all logical volumes on the exported stacked volume, or just the ones specified in the import list file. Import has the same priority as a logical mount. The Library Manager also turns on an in-use attribute for the import list volume so that it cannot be used by the hosts while import is running.
- A Library Manager entry with category set to import pending (FF15) is created for volumes which become candidates for import. If a candidate volume is a duplicate of an existing volume (native or virtual) in the library, it is not imported.
- The data for each volume on the Candidate list is moved from the exported stacked volume to a stacked volume managed by the VTS. If the VTS is APM capable, it also copies the associated APM constructs according to the specifications in the Import List 02 definitions.
- Once all of the data have been moved, a stub file is created in the TVC for each logical volume successfully imported.
- Once all stub files have been written to the TVC, the imported volumes are assigned to the 'INSERT' category and the attached hosts are notified. Volumes that were not successfully imported, due to I/O errors for example, are purged from the Library Manager inventory. OAM then performs entry processing for volumes in the insert category, calling the tape management system from the CBRUXENT exit. Volumes are either accepted and assigned to a private or scratch category, or rejected and purged from the library.
- Tape Management System is informed to update its records for the logical volumes indicating that they are now resident in the library.
- The exported stacked volume is returned to the import category for further disposition by an operator.

- When all logical volumes in the Import List File have been processed, a completion message is broadcast to all hosts attached to the VTS that performed the import process. The completion message results in the generation of a console message that includes the overall completion status of the import operation and metrics about the execution of the operation.
- All volumes which were moved by the operator into the Import category, including exported stacked volumes which contained logical volumes for the import operation, remain in the import category until the operator moves them to the insert category or ejects them.

Data contents of logical volumes with import option “SCRATCH” or “INITIALIZE” is not copied.

After all logical volumes from the exported stacked volume have been copied, the VTS copies the fragments from the file at the end of the volume into the TVC for each successfully imported logical volume. A stub file is not created for logical volumes with the “INITIALIZE” import option.

The process of the Import/Export operation can be monitored at the Library Manager, using the Queues display (Figure B-28).

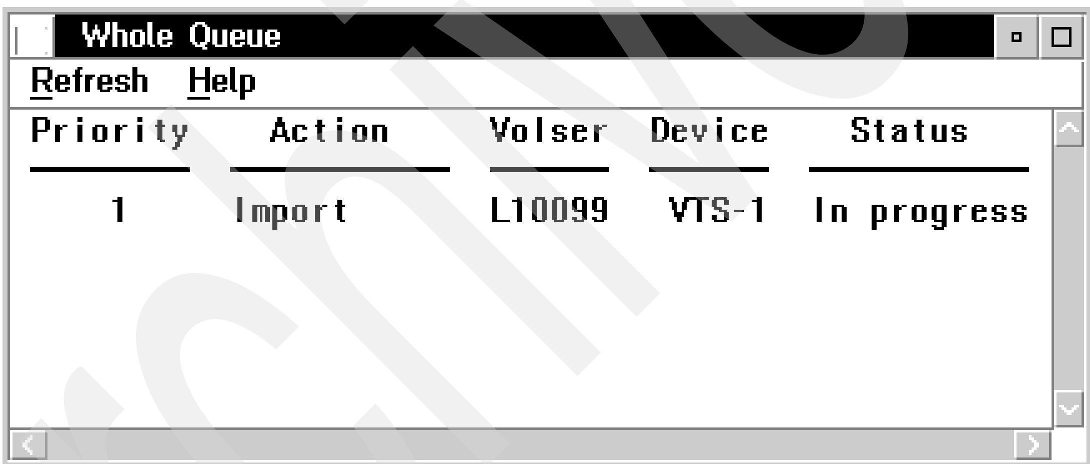

Figure B-28 Queue display with import in progress

> **Note:** When a logical volume is successfully imported, the exported copy of the logical volume is not deleted from the exported stacked volume. You must take care not to import this level of the data again.
>
> DFSMSrmm automatically prevents accidental re-imports of logical volumes. In order to do that, RMM allows import of one of its known volumes only if that volume is identified as being exported on the stacked volume where it is being imported from (identified by an ID in the container field). At import time, RMM removes the container information to prevent re-import of the same volume.
>
> If you want to re-import a logical volume, once it is removed from the VTS, you must remove the volume from the RMM inventory so that it appears to be a foreign volume from another tape management environment, or tell RMM it is back on the stacked volume.

OAM and the Import function

In addition to the import completion message, status messages are generated during the execution of this operation. These messages result in the generation of host console messages. See *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280 for further information. Messages are generated when:

- An import operation has started.
- Import processing has begun for an exported stacked volume
- All logical volumes imported have been assigned to the insert category
- When the processing for an exported stacked volume has completed
- When a time limit has been exceeded for host response
- When an import operation has been terminated or canceled
- All import processing has been completed

Some examples for import related messages follow.

Successful completion of the import operation is indicated by the message:

```
CBR3860I Import operation for logical list volume volser in
         library library-name completed successfully.
```

If only part of the logical volumes were successfully imported, the completion is indicated by one of two messages:

```
CBR3861I Import operation for logical list volume volser in
         library library-name completed with exceptions or errors.

CBR3862I Import operation for logical list volume volser in
         library library-name completed with exceptions or errors.
         Statistics for the operation were not available.
```

Messages CBR3860I and CBR3861I include statistics for the operation, for example the number of volumes and megabytes imported.

With Advanced Policy Management (FC4001-4004), there are further changes to OAM messages. The following messages are new:

```
CBR3625I Entry of volume volser into library library-name failed. Unable to set the
attributes for the volume.
```

Explanation: Volume volser could not be entered into library *library-name*. The attributes for the volume which includes the outboard policy names and the volume's category could not be set. See the secondary error message for a description of the failure.

```
CBR3626I Entry of volume volser into library library-name failed. Unable to obtain the
attributes for the volume.
```

Explanation: Volume *volser* could not be entered into library *library-name*. The attributes for the volume which includes the outboard policy names could not be obtained from the library. See the secondary error message for a description of the failure.

```
CBR3627I Outboard storage group policy storage-group assigned to volume volser in library
library-name failed validation.
```

Explanation: The enter request for volume *volser* into library *library-name* failed because the storage group *storage-group* assigned to the volume at the library is invalid:

1. The storage group is not defined in the active SMS configuration.
2. The storage group is not a tape storage group.
3. The library into which the volume is being entered is not defined to the storage group.

If a storage group does not exist for the volume in the tape configuration data base (TCDB), or is not provided by the cartridge entry installation exit (CBRUXENT), the storage group assigned to the volume at the library is used. This storage group may have been assigned to the volume when it was exported and/or overridden through the import list file. The storage group assigned to the volume at the library may also override any previously existing storage group in the tape configuration data base (TCDB) if it was explicitly specified through the import list file.

CBR3628I Entry of volume *volser* **into library** *library-name* **failed. Outboard policy exceeds the hardware limit.**

Explanation: The enter request for volume *volser* into library *library-name* failed because at least one of the storage management constructs (policy names) is to be created but the maximum number of unique names (255 maximum) for each construct type has already been reached.

CBR3630I Entry of volume *volser* **into library** *library-name* **failed. Library detected duplicate.**

Explanation: The enter request for volume *volser* into library *library-name* failed. The library returned a unit check indicating that there is already a tape volume with this volser in another partition of the non-U/I library of a Peer-to-Peer VTS.

CBR3658I Unable to make the failed eject notification call for volume *volser* **from library** *library-name*.

Explanation: The cartridge eject installation exit (CBRUXEJC) requested that a notification call be made to the exit in case of an eject failure. However, due to a current failure, or a previous failure that resulted in the exit and eject processing being discontinued, the failed eject notification call for volume *volser* from library *library-name* could not be made.

CBR3659I Failed eject notification processing for volume *volser* **from library** *library-name* **failed.**

Explanation: On return from the cartridge eject installation exit (CBRUXEJC) an error was encountered when processing the failed eject notification call for volume *volser* from library *library-name*. This message is issued in conjunction with message CBRxxxxI explaining the cause of the failure.

CBR1086I LIBRARY LMPOLICY results for volume *volser*: *result-text*.

Explanation: The LIBRARY LMPOLICY command was entered for volume *volser*. The result-text *result-text* reports the success or failure of the request as returned from the LCS External Services (CBRXLCS) Change Use Attribute (CUA®) function.

CBR1088I Command rejected. Function not supported in library *library-name*.

Explanation: The operator has entered a command for a library *library-name* that does not support the requested function.

CBR3780I Audit for volume *volser* **in library** *library-name* **cancelled.**

Explanation: The audit for volume *volser* in library *library-name* has been cancelled at the library. An operator at the Library Manager console indicated that the library was to be taken offline. In order for the library to be taken offline, pending operations must either be completed or cancelled. If a state exists at the library that prevents an operation from completing, so that the library can be taken offline, that operation is cancelled.

More details regarding the changes to IMPORT processing and the changes to OAM are covered in “Create the import list volume files” on page 395.

For more information about OAM object tape support, refer to the *z/OS Object Access Method Planning, Installation, and Storage Administration Guide for Object Support*, SC26-0426, and z/OS *Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries*, SC26-0427.

Analysis of the import status file

Unless the overall completion status of the import operation indicated that no processing exceptions were found, the Import Status File records must be analyzed. The VTS generates status records for all logical volumes processed during the import operation. The records indicate the success or failure of the import for each volume. When the import operation has finished, the VTS appends the status records to the import status file on the import list volume.

You can also use the display SMS command to interrogate the information an a successful logical volume import. After a successful import of a logical volume, you will able to check the policy names assigned (Example B-1 on page 395).

Figure B-29 shows sample JCL to read the import status file.

```
//STEP     EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN    DD DUMMY
//SYSUT1   DD DSN=HILEVELQ.IMPSTATS,
//            UNIT=3490,DISP=OLD,LABEL=(2,SL),
//            VOL=SER=volser
//SYSUT2   DD SYSOUT=*
```

Figure B-29 Sample JCL to read the import status file

Eject exported stacked volumes after import

The VTS does not automatically eject the exported stacked volumes at import completion. The exported stacked volumes are held in the library but are not used for storing any logical volumes because they are assigned to the import (FF17) category. You can select options at the Library Manager menus (Figure B-30) to either return the exported stacked volumes to the insert category (to be used for the VTS or native drives) or eject them and then return them to the original site.

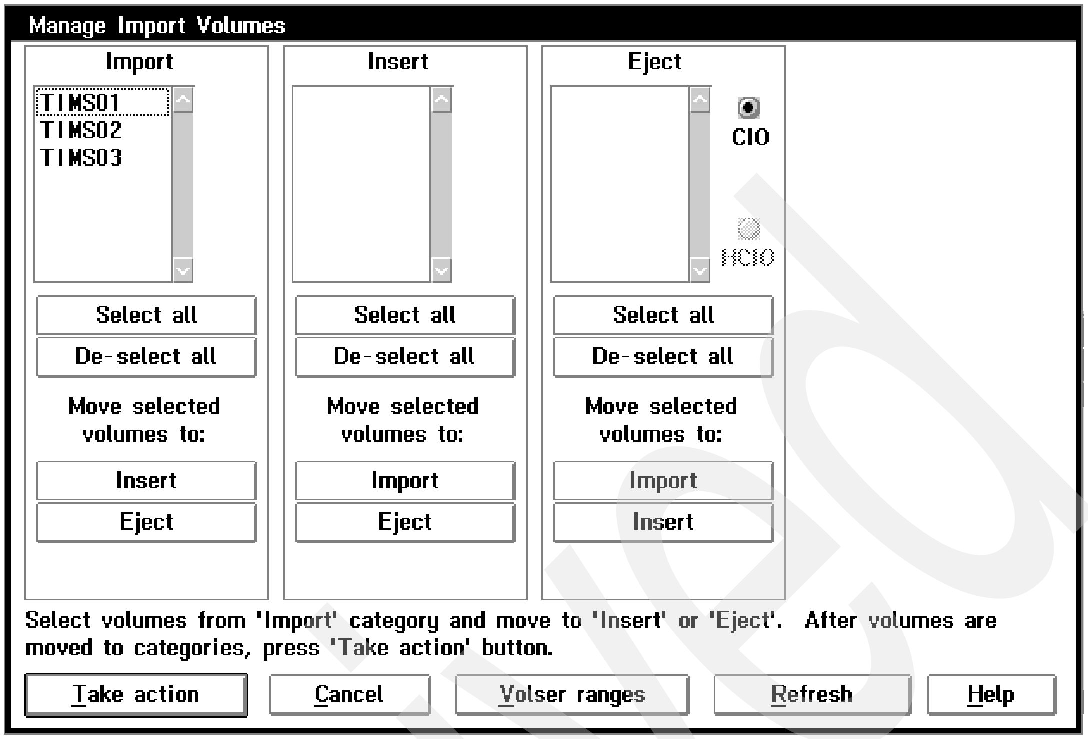

Figure B-30 Manage import volumes window

Canceling import

Depending on the amount of data you are importing, import can take a long time and you may need to cancel it for example to free VTS resources for other tasks. The first method of canceling an import operation is to issue this command from any host attached:

```
LIBRARY IMPORT,volser,CANCEL
```

Where, *`volser`* is the volume serial number of the import list volume containing the list of volumes to be imported.

The VTS accepts a host request to cancel an import operation and processes it as follows:

- Volumes imported successfully prior to the cancel command are indicated by status code of “00” in the Status File List.
- To ensure that the VTS, Library Manager and the host databases remain in sync regarding the logical volumes resident in the library, host processing of those volumes assigned to the insert category already, must complete, i.e. assignment to either eject, private or scratch category, before the cancel becomes effective.
- All logical volumes remaining in the import-pending category are deleted from the library inventory and any stub files written to the cache are deleted.
- Status file records are written for all logical volumes or exported stacked volumes that did not complete the import process prior to the execution of the cancel request.
- A completion message is sent to all attached hosts indicating that the import was canceled.

If you later want to rerun the import operation on the same VTS subsystem, you can use the same import list volume without having to modify it. Volumes on the import list file that were imported during the previous run are found to be duplicates and therefore will be skipped during the rerun, with a record written in the status file to indicate this. The VTS overwrites status records from the previous import run by writing new status records to the import status file.

A second method is to cancel an import operation from the Library Manager panel shown in Figure B-31. Canceling the operation through the Library Manager is similar to the host cancel request. It should be used only when there is no functional host that can issue the cancel.

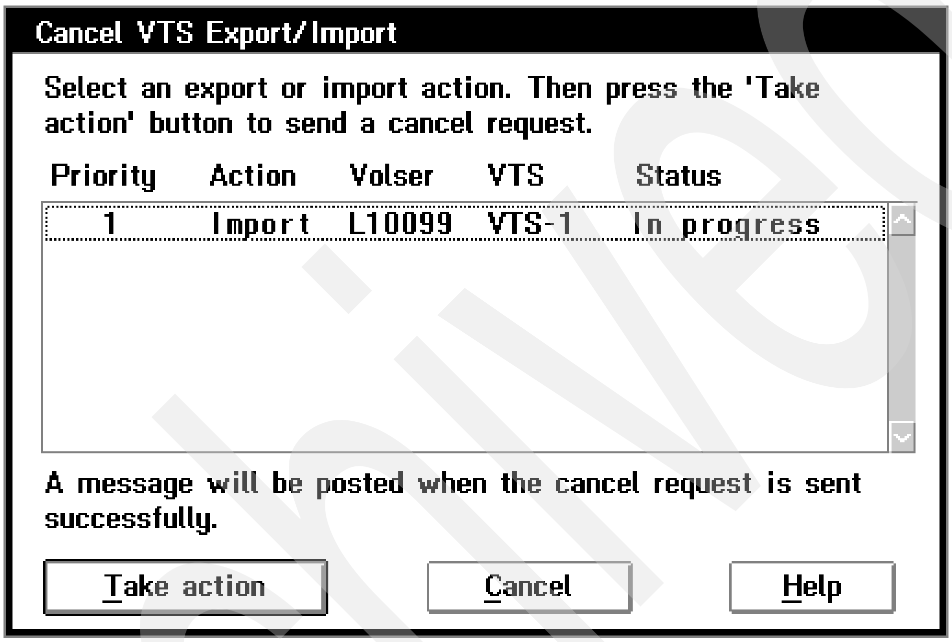

Figure B-31 Cancel import using Library Manager panels

The Library Manager cancel is processed as follows:

- If all logical volumes imported from an exported stacked volume and assigned to the insert category have been processed, cancel processing continues as for a host cancel.
- If this process is not yet finished, volumes that are still assigned to the insert category are left there.
- The status file records for those volumes indicate this condition.
- The rest of the cancel processing continues as for a host cancel.

Handling operator errors

In this section we describe operator errors which may be made during an import operation and the recovery actions.

Error Scenario 1: An exported stacked volume with logical volumes not yet imported has been assigned to the Insert category and the volume serial number falls within the range defined for VTS stacked volumes. **Resultant VTS Action:** The VTS/library adds the volume as a scratch stacked volume. **Recovery Action:** Eject the volume through the console using the **Eject a Stacked Volume** pop-up window.

Error Scenario 2: An exported stacked volume with logical volumes not yet imported has been assigned to the Insert category and the volume serial number **does not** fall within the range defined for VTS stacked volumes. **Resultant VTS Action:** The VTS/library assigns the volume to the insert category for 3590 native use and notify all attached hosts. **Recovery Action:** Eject the volume from the library through host console commands, ISMF or tape management system.

Error Scenario 3: An exported stacked volume has been assigned to the Import category and has not been required for the import operation. **Resultant VTS Action:** The volume remains in the Import category. **Recovery Action:** Use the **Manage Import Volume** pop-up window to eject the volume.

Error Scenario 4: An HPCT volume (not an exported stacked volume) which is needed as a scratch stacked volume or native 3590 volume has been assigned to the Import category. **Resultant VTS Action:** The volume remains in the Import category. **Recovery Action:** Use the **Manage Import Volume** pop-up window to move the volume to the Insert category defined by Volser ranges or eject the volume if it is not desired in the library.

Reuse of exported stacked volumes

Once all of the logical volumes on an exported stacked volume have been imported, the exported stacked volume can be re-used for any 3590 application. No cleanup or special processing is required. However, the exported stacked volume remains in the Import category until operator action is taken. Use the host tape management system to verify that there are no logical volumes contained on a physical volume previously used for export. For example the RMM SEARCHVOLUME command:

```
RMM SV CONTAINER(xxxxxx),
```

In this command, xxxxxx is the volser of the exported stacked volume.

Reuse of the physical volume is accomplished with operator action by using the Library Manager console **Manage Import Volume** pop-up window to move the exported stacked volume out of the Import category or eject it.

Import performance

There are two considerations in defining the performance of an Import Operation — first, the performance of the Import Operation itself; and second, the impact that an Import Operation has on the overall performance of a VTS.

The performance of an Import Operation itself is defined in terms of its total elapsed time. These are the elements that influence the performance:

- **Number of logical volumes to import:** For logical volumes, that do not have the SCRATCH or INITIALIZE import option specified, you should calculate a search time of 20 seconds to locate the logical volume on the stacked volume If all logical volumes on an exported stacked volume have to be imported, then the seek time is eliminated.
- **The amount of data to be copied:** It is assumed, that data can be moved from the exported stacked volume to target stacked volume at a rate of 15 MB/s (3:1 Compression Ratio).
- **The number of exported stacked volumes to be mounted:** For each exported stacked volume, there is a nominal processing overhead of 120 seconds to cover robotics mount and demount, load/thread and rewind/unload. For each exported stacked volume created, there is an overhead of 120 seconds to write the volume table of contents and the logical stub information files.

The impact on the overall performance of the VTS can be seen in two primary areas:

- **Data throughput:** Although import does not move data through the Tape Volume Cache, it uses 3590 SCSI bus bandwidth and VTS processor cycles. The SCSI bandwidth is likely to be insufficient to support maximum data throughput concurrently with an import operation. The throughput bandwidth of the VTS may be reduced slightly.
- **Logical volume recall time:** Since an import operation uses two physical drives of the VTS, fewer drives are available for use to recall logical volumes. This may result in longer logical mount times when the number of pending recalls exceed the number of physical drives less the two used for the import operation. If all the logical volumes to be imported have the SCRATCH or INITIALIZE option specified, only one drive is used.

Exported stacked volumes consolidation

Along time, logical volumes that have been exported will expire and the valid space used on the exported stacked volumes will decrease. It is important to find rules and ways to return exported stacked volumes with a small percentage of active data, as well as to reduce the number of exported stacked volumes. On the other hand, exported stacked volumes contain vital data that are necessary to achieve a correct restart, and therefore cannot be returned to the VTS for consolidation.

The objective of this section is to describe ways of reducing the number of exported stacked volumes outside the VTS.

Links between stacked and logical volumes

Logical volumes that are exported have completely disappeared from the VTS and from the Library Manager. They have also disappeared from the TCDB when the tape management system is RMM. They may or may not be in the TCDB when the tape management system is BrightStor CA-1 or BrightStor CA-Dynam/TLMS. However, the logical volume entries are present and managed by the tape management system, in all cases. Analyzing the tape management system based on container volumes subsequently pointing to the valid (not scratch) logical volumes can give accurate information on the topography of containers.

Such an analysis requires only access to the tape management system data base and can produce a cross reference of logical volume contents and megabytes, crossed by container volumes.

Links between exported logical volumes and in-house volumes

Most of the time, volumes that are exported contain a copy of the data sets that are kept on the creating system. The duplication can be controlled by software such as DFSMShsm that can create two copies of the same data. Alternately, other software or JCL can create a copy of the original data set. The naming rules for the copy are conventions.

When analyzing the topography of exported stacked volumes, we can find the name of the data set in the tape management system, but this name is the name of the copy. We have to find the name of the original data set in order to create a second copy of the data, as well as the adequate data mover.

Returning the container or making a second copy

Once the topography has been done and the original volumes identified, we can derive the amount of data to copy and export. An example of such a list is shown in Figure B-32.

```
STKEX0 0MB
    No valid logical volume

STKEX1 250MB
    DSN1 V00001 V00001 25KB ICEGENER
    DSN2 V00001 V00001 25KB ICEGENER
    DSN3 V00021 V00021 100KB DFDSS
    DSN4 V00023 V00023 100KB IDCAMS

STKEX2 1250MB
    DSN5 V00002 V00052 25KB ICEGENER
    DSN6 V00002 V00052 25KB ICEGENER
    DSN7 V00024 V00043 600KB DFDSS
    DSN8 V00025 V00057 387KB IDCAMS
    ......................

STKEXx yyyyMB

    ......................
```

Figure B-32 Sample coding to derive the amount of data to copy and export

In this list we find the name of the copy volume, the name of the original volume, the size, and the creating program name.

With such a list, you can derive the volser to return to the VTS and the value of initiating a subsequent copy and weight the workload in MB and data set to copy by host software, with the saving in exported cartridges that can be returned to the VTS as scratch.

Export/Import programming support

The Library Control System (LCS) External Services has been enhanced to support Import/Export. It provides the following general-use programming interfaces for Import/Export:

The LCS external service FUNC=EXPORT can be used to identify the export list volume and to initiate the export operation. The interface can also be used to cancel an executing export operation. The syntax of the CBRXLCS Assembler H macro to perform the export function is shown in Figure B-33.

```
CBRXLCS TYPE=TAPE
        ,FUNC=EXPORT
        ,VOLUME=volser
        [,CANCEL=YES | NO]
```

Figure B-33 Sample coding to perform the export function

The LCS external service FUNC=IMPORT can be used to identify the import list volume and to initiate the import operation. The interface can also be used to cancel an executing import operation. The syntax of the CBRXLCS Assembler H macro to perform the import function is shown in Figure B-34.

```
CBRXLCS TYPE=TAPE
        ,FUNC=IMPORT
        ,VOLUME=volser
        [,CANCEL=YES | NO]
```

Figure B-34 Sample coding to perform the import function

The FUNC=QVR external service allows the caller to determine in which tape library a specified volume resides. If VOLINFO=YES is specified, a copy of the available information about the volume as contained in the TCDB and the Library Manager inventory is returned to the caller.

This interface is also extended to provide audit capabilities from a tape management system database with the new optional parameter, LIBNAME. Whereas the ISMF audit capability requires the presence of a TCDB record, this function does not. A tape management system can use this new capability to verify its inventory against the TCDB and LM inventory. It can also be used to track and verify the library residency of volumes that are not managed in the TCDB, such as the stacked volumes in a VTS. See Figure B-35.

BrightStor CA-1 support of the Import/Export function

Import/Exportis supported by BrightStor CA-1 and BrightStor CA-Dynam/TLMS.

Contact Computer Associates for the product information bulletin with the detailed information about VTS Import/Export support.

BrightStor CA-1

We refer you to the “BrightStor CA-1 Administrator and Operator Guide, Chapter 9”. This chapter is titled, “Automated Tape Libraries and Virtual Tape Support”. This is now a 24-page chapter, nearly all dedicated to the 3494 or 3494/VTS environment.

See subchapter “Recommendations and Procedures for IBM 3494 VTS with IMPORT/EXPORT”.

BrightStor CA-Dynam/TLMS

We refer you to the “BrightStor CA-Dynam/TLMS Systems Programmer Guide” chapter 8, which is also titled “Automated Tape Libraries and Virtual Tape Support”. This is almost the same as the BrightStor CA-1 chapter, but has been customized for TLMS users, of course.

DITTO VTS functions

In this section, we describe two of the new features of DITTO/ESA 1.3, Exported Stacked Volume Copy and Exported Stacked Volume List. You can also obtain the updated *Ditto/ESA V1R3 User's Guide and Reference, SH19-8221,* at:

http://www-4.ibm.com/software/ad/ditto/library.html#ug

DITTO overview

IBM Data Interfile Transfer, Testing, and Operations Utility for ESA (DITTO/ESA) is a proven utility for working with data across a wide range of storage media and formats. You can use DITTO/ESA to access, display, edit, migrate and maintain data in files and data sets on any storage device supported in the z/OS, z/VSE, or z/VM Environments.

These new features enable the user to physically copy a logical volume from an exported stacked volume on to another physical cartridge. This can be useful for storing the logical volume in an off site location for disaster recovery purposes or for sending it to an external bureau. A list function is also available which allows you to list the contents of an exported stacked volume which can be used as input into the copy process.

VTS terms

In this section we define several VTS terms.

Exported stacked volume: This is a physical 3590 B1A, E1A or H1A cartridge (containing logical virtual volumes) that was created by using the VTS Export function. See “Export operation” on page 376 for more information on using Export.

Internal stacked volume: This is a physical 3590 B1A, E1A or H1A cartridge used to stack virtual volumes which have been written from the Tape Volume Cache of the VTS. This volume is not in the same format as the exported stacked volume.

Logical volume: This is a tape volume created in the Tape Volume Cache in the VTS. The logical volume may be physically in the TVC, on an exported stacked volume or on an internal stacked volume. The logical volume is equivalent to a physical tape volume in a non-VTS environment.

Requirements

All VTS export volumes are written on devices that support the locate block CCW. DITTO uses the locate block CCW to find the start of a logical volume, therefore, the input device for the exported volume must support the locate block CCW.

New features

In this section we describe some new features.

EVC - Exported Stacked Volume Copy: This is used to copy a logical volume from an exported stacked volume to a physical volume. Up to five logical volumes can be copied in a single step to five physical volumes.

EVL - Exported Stacked Volume List: This is used to list the contents of an exported stacked volume. You can choose a short listing, which prints only the logical volumes on the exported stacked volume, or you can choose a long listing which prints the volume structure of the logical volumes.

How DITTO works

The process uses the standard DITTO interfaces, batch JCL, command mode and full screen mode.

This utility once invoked for COPY requires various parameters which enable it to select the source logical volume from the correct exported stacked volume. These parameters also specify the target output volume required and the device type. See the latest DITTO User's Guide for a full description of required parameters.

The copy utility extracts a logical volume from an exported VTS volume. Therefore, by definition, the exported volume no longer resides in the original VTS where it was created. The VTS export volume can be used as input on a standalone device or in another ATL library. This means that the VTS export volume is physically removed from the VTS, where it was created and loaded into another ATL library, where it is treated/defined the same as any other tape being brought into the library from another system.

The output device can be a 3490 or 3590, but the output volume must be large enough to contain the logical volume you are going to copy. If the output volume is to small then the process fails with an EOV message. You can copy up to five logical volumes during one step and each logical volume must be written to a physical volume, so if you want to copy three logical volumes, then you need three physical output volumes to perform the copy.

The copy process reads the volser (VOL1) of the output volume and during the copy effectively renames the logical volume to the name that is on the output volume. If the output volume has no volser (VOL1) information, then the volser of the logical volume is used.

This copy process is a COPY only process, it does not remove or delete the source volume, hence, the procedure is repeatable if you wanted to make more than one copy of the source.

Figure B-35 shows the DITTO copy of logical volume VOL001 to a unique physical volume PHY001. Only the data is copied from VOL001 and not the volser (VOL1) information. PHY001 now contains the data copied from VOL001. A listing can also be produced to display the contents of the exported stacked volume.

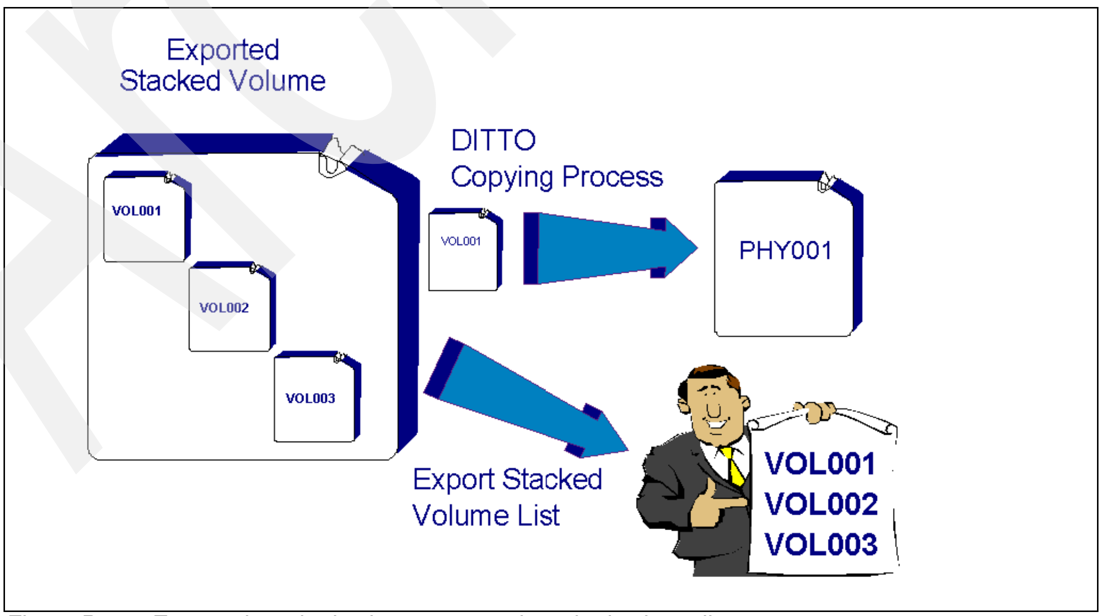

Figure B-35 Exported stacked volume copy and stacked volume list

The exported stacked volume List function scans the VTS export volume and produces one of two reports. These are called the short or long lists. The short list produces a list of volumes contained on the export volume. The long list produces a list of all logical volumes on the export volume as well as the structure of these volumes. It also includes the dataset name on any HDR1 records located. This list can be used in conjunction with the copy command.

As with the copy function, the list function works only with exported stacked volumes and requires various parameters to be set. See the latest DITTO User's Guide for a full description of required parameters.

> **Note:** Only exported stacked volumes created by the VTS export function are supported by the Export Volume Copy function. This function does not support internal stacked volumes that were not created by the VTS export function. Currently, an internal stacked volume is not compatible with an exported stacked volume.

The copy function extracts a block of data from an exported stacked volume and writes the physical block of data on the output volume. A block of data may be user data or HDR, or EOF records. There is no open or close performed for each file on the logical volume. Hence, there is no interface to the tape management systems or to the Offline Control Dataset (OCDS) of DFSMShsm. Also, the z/OS catalog information for individual datasets on the volume are not updated.

A spreadsheet to estimate the time required to perform the DITTO extract operation on exported stacked volumes is named **DITOFEAS.123** and is available at:

http://w3-1.ibm.com/sales/systems/ibmsm.nsf/docnames/tapetools

C

VTS implementation step-by-step

VTS implementation

In this appendix, we document three possible implementation scenarios. These are to be taken as a guide only and should be tailored specifically for your site: These three examples assume that the host system is DFSMS:

1. First-time VTS library installation with LM 527 and VTS 2.26 without Advanced Policy Management (FC4001)
2. Upgrading your current VTS from LIC levels older than LM 527 and VTS 2.26 to the latest VTS and LM LIC levels, but without Advanced Policy Management (FC4001 - 4004)
3. Upgrading your current VTS from LIC levels older than LM 527 and VTS 2.26 to the latest VTS and LM LIC levels with Advanced Policy Management (FC4001- 4004)

Note: Once your VTS has been upgraded to the LM 527 or higher and VTZ 2.26 or higher, additional LIC upgrades are without impact to the LM volume categories.

First time VTS library installation

Here we detail a sample definition procedure for a first time tape library installation without Advanced Policy Management (FC4001):

1. Analyze VTS targeted workloads to determine the appropriate VTS configuration for your environment. The following tools can be helpful in making this determination:
   - BatchMagic

Important: This is the crucial step to configuring your VTS correctly and should be done in conjunction with your IBM Storage Specialist.

2. Check the latest 3494 or 3953 and 3584 Systems Assurance Product Review (SAPR) Guide.

3. Order stacked volume media and/or labels.
4. Set up the HCD (see 4.2.4, “Defining devices through HCD” on page 118).
5. Vary CHPIDs, paths, and VTS drives online.
6. JES3 system considerations (see 4.5.3, “JES3 environments” on page 155).
7. For DFSMSrmm:
   – Create VLPOOLS for logical volumes
   – ADD logical volumes/location to DFRMM
   – Check REJECT PARMs
   – Refresh DFRMM
8. For CA1:
   – If your planned logical volume range has not been defined to your Tape Management System (TMS), you will need to extend the TMC. Also check the amount of free DSNBs available and add if required concurrent with your extend.
   – Check that the **SCRTCH** PARM is set to **YES** in PPOPTION. When a tape becomes scratch or a DSNB record gets deleted, control is given to the security exit TMSUX2S. This is the scratch notification interface between CA1 and the 3494.
   – Depending on the parameters specified in PPOPTION, you may need to initialize the logical volumes after they have been inserted.

> **Note:** Check with the vendor of your tape management product to ensure that the level of software installed on your system will support the 3494/VTS or eventually the 3584/3953/VTS.

9. Modify/check current SYS1.PARMLIB members:
   – **IEFSSNxx**
     - Add or update the OAM1 entry with the name of the initialization module (CBRINIT) executed at IPL.
   – **SCHEDxx**
     - Add the OAM initialization module CBROAM to the system program properties table (PPT).
   – **IGDSMSxx**
     - If you want OAM to start automatically as part of SMS initialization, add the OAMPROC and OAMTASK optional parameters.

> **Note:** If you use an ISV tape management system, it may require that the OAM address space be started after the tape management initialization. In this case, do not start OAM automatically. Check with the vendor of the tape management product.

   – **CONSOLxx**
     - If you want to receive library messages at a specific console, update the CONSOLxx member referenced by IEASYSxx. This console name must also be defined to SMS and is done when you define the library through ISMF.
   – **DEVSUPxx**
     - This member is used to hard partition your VTS by specifying category codes to provide a unique range of tapes for each system. If no category codes are specified, defaults will be assigned.

Attention: We recommend that you do not use the DEVSUPxx default categories. Avoiding the use of the defaults will make future partitioning of the library with other systems easier and more secure. See *Guide to Sharing and Partitioning IBM Tape Libraries*, SG24-4409 for additional information.

- **COMMNDxx**
  - If you want your library automatically brought online after each IPL, add the VARY SMS,LIBRARY command. Another area where you can bring your library automatically online is via ISMF when you define your library. Set the **Initial Online Status** to **YES**.
- **GRSCNFxx**
  - If you are going to share the tape library among two or more systems in an DFSMS complex, a global resource serialization ring can be created to include to include all sharing systems. This allows OAM to serialize the cartridge entry process.
- **LOADxx (optional)**
  - The default data set name for the TCDB is SYS1.VOLCAT.VGENERAL. If you want to use this name, no change is required to the LOADxx member. If you want to use a different HLQ than the defined default, update columns 64 through 71 of the SYSCAT statement with the preferred HLQ.
- **COFVLFxx**
  - Add the volume catalogs to the IGGCAS class definitions where you have other ICF catalogs.
- **ALLOCxx**
  - Add tape automation policies
- **IECIOSxx**
  - Set Missing Interrupt Handler values. This value should be 45 minutes for the VTS logical drives.

Note: For coding examples and more information regarding SYS1.PARMLIB updates, refer to *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632.

10. If your IBM System Service Representative (SSR) has installed your VTS, perform the following Library Manager definitions:
    - Define stacked volume ranges (see “Defining volser ranges for stacked volumes” on page 222).
    - Define fast ready categories (see 4.3.5, “Define Fast Ready categories” on page 133). The Fast Ready Category number you define is the same value as those defined for MEDIA1 or MEDIA2 as specified in the DEVSUPxx of your PARMLIB.
    - Define expired logical volume data policy (see “Expired volume management” on page 245).
    - Define VTS management policies (see 4.3.7, “Define VTS management policies” on page 137):
      - Inhibit reclaim schedule
      - Reclaim threshold percentage
      - Free storage threshold

11. If stacked volume media is available, insert the physical VTS stacked volumes (see "Inserting stacked volumes in the VTS" on page 229).

12. Create TCDB as shown in Example C-1.

Example: C-1 Define TCDB

```
//DEFINECAT JOB
//STEP1    EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
   DEFINE USERCATALOG -
   (NAME(SYS1.VOLCAT.VGENERAL) -
   VOLCATALOG -
   VOLUME(volser) -
   CYLINDERS(30,30))
/*
```

When sizing the TCDB, calculate 275 bytes for each volume entry. For 10,000 logical volumes, 2.75 MB or about 3 cylinders of 3390 disk space is required. Allow sufficient room for growth.

> **Note:** If you are sharing the tape library (non-hard partitioned), use IDCAMS to IMPORT CONNECT the VOLCAT to the other sharing systems.

13. Define security profiles
    - ISMF
    - DFSMS constructs
    - STGADMIN
    - DFSMSrmm
    - z/OS Operator Commands

> **Note:** The following manuals with be helpful in setting up your security environment:
>
> - *z/OS DFSMSdfp Storage Administrator Reference*, SC35-0422
> - *z/OS DFSMSrmm Implementation and Customization Guide*, SC26-7405
> - *z/OS MVS Planning: Operation*, SC22-7601

14. Create the cataloged OAM procedure

15. Start OAM. Upon startup, you may receive the following messages:
    - CBR1115I - No libraries defined to OAM.
    - CBR0094E - OAM has initialized without tape or object support.

These messages are informational and expected. They will be resolved when the library has been defined and the DFSMS SCDS activated.

> **Tip:** Code the RESTART=YES PARM in your OAM procedure. This PARM, upon an activation of the SCDS, automatically restarts OAM. This is important when you have changed tape library related constructs. Without this PARM, you must restart OAM manually after an SCDS activation if you have changed tape library related constructs. This can be accomplished by the following command: **`F OAM,RESTART`**.

16. Define the following DFSMS constructs using ISMF:
   - Tape library (Figure C-11 on page 431, Figure C-12 on page 431)
   - Data class (Figure C-1, Figure C-2, Figure C-3 on page 424)
   - Storage class (Figure C-5 on page 426, Figure C-6 on page 427)
   - Management class (Figure C-7 on page 428, Figure C-8 on page 428, Figure C-9 on page 429)
   - Storage group (Figure C-10 on page 430)
17. Create the following ACS routines for DFSMS managed tape:
   - Data class (Example C-2 on page 425)
   - Storage class (Example C-3 on page 427)
   - Management class (Example C-4 on page 429)
   - Storage group (Example C-5 on page 430)

```
                              DATA CLASS DEFINE                 Page 1 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVTS

To DEFINE Data Class, Specify:
  Description ==> Default VTS Data Class
              ==>
Recfm . . . . . . . . . .             (any valid RECFM combination or blank)
Lrecl . . . . . . . . . .             (1 to 32761 or blank)
Space Avgrec  . . . . . .             (U, K, M or blank)
        Avg Value . . . . .             (0 to 65535 or blank)
        Primary . . . . . .             (0 to 999999 or blank)
        Secondary . . . . .             (0 to 999999 or blank)
        Directory . . . . .             (0 to 999999 or blank)
Retpd or Expdt  . . . . .             (0 to 9999, YYYY/MM/DD or blank)
Volume Count  . . . . . . 1           (1 to 59 or blank)
Add'l Volume Amount . . .             (P=Primary, S=Secondary or blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-1 Data Class Define ISMF panel 1 of 4 - first time setup

```
                              DATA CLASS DEFINE                 Page 2 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVTS
To DEFINE Data Class, Specify:

  Data Set Name Type  . . . . . .      (EXT, HFS, LIB, PDS or blank)
    If Ext  . . . . . . . . . . .      (P=Preferred, R=Required or blank)
    Extended Addressability . . . N    (Y or N)
    Record Access Bias  . . . . .      (S=System, U=User or blank)
  Space Constraint Relief . . . . N    (Y or N)
    Reduce Space Up To (%)  . . .      (0 to 99 or blank)
    Dynamic Volume Count  . . . .      (1 to 59 or blank)
  Compaction  . . . . . . . . . . Y     (Y, N, T, G or blank)
  Spanned / Nonspanned  . . . . .      (S=Spanned, N=Nonspanned or blank)



Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-2 Data Class Define ISMF panel 2 of 4 - first time setup

```
                             DATA CLASS DEFINE              Page 3 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVTS

To DEFINE Data Class, Specify:

  Media Interchange
    Media Type . . . . . . .   2         (1, 2, 3, 4, 5, 6, 7, 8 or blank)
    Recording Technology . .   36         (18, 36, 128, 256, 384, E1 or blank)
    Performance Scaling  . .    N         (Y, N or blank)
  Block Size Limit . . . . .             (32760 to 2GB or blank)
  Recorg . . . . . . . . . .             (KS, ES, RR, LS or blank)
  Keylen . . . . . . . . . .             (0 to 255 or blank)
  Keyoff . . . . . . . . . .             (0 to 32760 or blank)
  CIsize Data  . . . . . . .             (1 to 32768 or blank)
  % Freespace CI . . . . . .             (0 to 100 or blank)
              CA . . . . . .             (0 to 100 or blank)

Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-3 Data Class Define ISMF panel 3 of 4 - first time setup

```
                             DATA CLASS DEFINE              Page 3 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVTS

To DEFINE Data Class, Specify:

  Shareoptions Xregion  . . .      (1 to 4 or blank)
               Xsystem  . . .      (3, 4 or blank)
  Reuse . . . . . . . . . . . N    (Y or N)
  Initial Load  . . . . . . . R    (S=Speed, R=Recovery or blank)
  BWO . . . . . . . . . . . .      (TC=TYPECICS, TI=TYPEIMS, NO or blank)
  Log . . . . . . . . . . . .      (N=NONE, U=UNDO, A=ALL or blank)
  Logstream Id  . . . . . . .
  FRlog . . . . . . . . . . .      (A=ALL, N=NONE, R=REDO, U=UNDO or blank)
  RLS CF Cache Value  . . . . A    (A=ALL, N=NONE, U=UPDATESONLY)



Use ENTER to perform Verification; Use UP Command to View previous Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit
```

Figure C-4 Data Class Define ISMF panel 3 of 4 - first time setup

Example: C-2 Data Class ACS routine

```
PROC 0 DATACLAS

 /*------------------------------------------------------------------*/
 /* DEFINE VALID TAPE UNIT ESOTERICS                                 */
 /*------------------------------------------------------------------*/
        FILTLIST TAPE_UNIT INCLUDE(348*,TAPE*,'CART','AFF=SMST','3420',
                                   '3590','3590-1','3490')

        FILTLIST DBASE_BKP   INCLUDE(PROD.DBASE.BKP.**)

        FILTLIST PROD_SMF    INCLUDE(PROD.SMF.**)

          SELECT
             WHEN ((&UNIT = &TAPE_UNIT) OR (&ANYVOL = 'REF=ST'))
                DO
                  SELECT
                     WHEN (&DATACLAS NE '')                /* Allows users to specify  */
                       DO                                  /* data class for tape      */
                         SET &DATACLAS = &DATACLAS
                         EXIT
                       END
                                WHEN (&DSN = &DBASE_BKP)
                       DO
                         SET &DATACLAS = 'DC3590K'
                         EXIT
                       END
                     WHEN (&DSN = &PROD_SMF)
                       DO
                         SET &DATACLAS = 'DCVTS'
                         EXIT
                       END
                  END
```

```
               END
                  OTHERWISE
                   DO
                     SET &DATACLAS = ''
                     EXIT
                   END
          END /* END SELECT */
     END      /* END PROC   */
```

Storage class

```
                              STORAGE CLASS DEFINE             Page 1 of 2
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Class Name  : SCTAPE
To DEFINE Storage Class, Specify:
  Description ==> Default Storage Class
              ==>
  Performance Objectives
   Direct Millisecond Response . . . .          (1 to 999 or blank)
   Direct Bias . . . . . . . . . . . .          (R, W or blank)
   Sequential Millisecond Response . .          (1 to 999 or blank)
   Sequential Bias . . . . . . . . . .          (R, W or blank)
   Initial Access Response Seconds . .          (0 to 9999 or blank)
   Sustained Data Rate (MB/sec)  . . .          (0 to 999 or blank)
  Availability . . . . . . . . . . . . N        (C, P ,S or N)
  Accessibility  . . . . . . . . . . . N        (C, P ,S or N)
   Backup  . . . . . . . . . . . . . .          (Y, N or Blank)
   Versioning  . . . . . . . . . . . .          (Y, N or Blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Page;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-5 Storage Class Define ISMF panel 1 of 2 - first time setup

```
                            STORAGE CLASS DEFINE              Page 2 of 2
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Class Name  : SCTAPE

To DEFINE Storage Class, Specify:

  Guaranteed Space  . . . . . . . . . N           (Y or N)
  Guaranteed Synchronous Write  . . . N           (Y or N)
  Multi-Tiered SG . . . . . . . . . .             (Y, N, or blank)
  Parallel Access Volume Capability   N           (R, P, S, or N)
  CF Cache Set Name . . . . . . . . .             (up to 8 chars or blank)
  CF Direct Weight  . . . . . . . . .             (1 to 11 or blank)
  CF Sequential Weight  . . . . . . .             (1 to 11 or blank)




Use ENTER to Perform Verification; Use UP Command to View previous Page;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-6 Storage Class Define ISMF panel 2 of 2 - first time setup

Example: C-3 Storage class acs routine

```
PROC 0 STORCLAS

 /*-----------------------------------------------------------------*/
 /* DEFINE VALID TAPE UNIT ESOTERICS                                */
 /*-----------------------------------------------------------------*/
    FILTLIST TAPE_UNIT INCLUDE(348*,TAPE*,'CART','AFF=SMST','3420',
                               '3590','3590-1','3490')

    FILTLIST SMS_TAPE  INCLUDE('DCVTS','DC3590K')

     SELECT

       WHEN ((&UNIT = &TAPE_UNIT) OR (&ANYVOL = 'REF=ST'))
         DO
           SELECT
              WHEN (&DATACLAS = &SMS_TAPE)
                DO
                   SET &STORCLAS = 'SCTAPE'
                   EXIT
                END
           END
         END
           OTHERWISE
              DO
                SET &STORCLAS = ''
                EXIT
              END
     END /* END SELECT */
  END  /* END PROC */
```

Management class

```
                          MANAGEMENT CLASS DEFINE               Page 1 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCTAPE

To DEFINE Management Class, Specify:

  Description ==> Default Management Class
              ==>

  Expiration Attributes
    Expire after Days Non-usage  . . NOLIMIT     (1 to 9999 or NOLIMIT)
    Expire after Date/Days . . . . . NOLIMIT     (0 to 9999, yyyy/mm/dd or
                                                 NOLIMIT)

  Retention Limit  . . . . . . . . . NOLIMIT     (0 to 9999 or NOLIMIT)



Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-7 Management Class Define ISMF panel 1 of 3 - first time setup

```
                          MANAGEMENT CLASS DEFINE               Page 2 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCTAPE

To DEFINE Management Class, Specify:

  Partial Release . . . . . . . . . N           (Y, C, YI, CI or N)

  Migration Attributes
    Primary Days Non-usage  . . . .             (0 to 9999 or blank)
    Level 1 Days Non-usage  . . . .             (0 to 9999, NOLIMIT or blank)
    Command or Auto Migrate . . . . NONE        (BOTH, COMMAND or NONE)


  GDG Management Attributes
    # GDG Elements on Primary . . .             (0 to 255 or blank)
    Rolled-off GDS Action . . . . .             (MIGRATE, EXPIRE or blank)

Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-8 Management Class Define ISMF panel 2 of 3 - first time setup

```
                         MANAGEMENT CLASS DEFINE            Page 3 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCTAPE

To DEFINE Management Class, Specify:
  Backup Attributes
    Backup Frequency  . . . . . . . .          (0 to 9999 or blank)
    Number of Backup Vers . . . . . .          (1 to 100 or blank)
      (Data Set Exists)
    Number of Backup Vers . . . . . .          (0 to 100 or blank)
      (Data Set Deleted)
    Retain days only Backup Ver . . .          (1 to 9999, NOLIMIT or blank)
      (Data Set Deleted)
    Retain days extra Backup Vers . .          (1 to 9999, NOLIMIT or blank)
    Admin or User command Backup  . . NONE     (BOTH, ADMIN or NONE)
    Auto Backup . . . . . . . . . . . N        (Y or N)
    Backup Copy Technique . . . . . . S        (P=Conc Preferred, R=Conc
                                               Required or S=Standard)
Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; Cancel to Exit.
```

Figure C-9   Management Class Define ISMF panel 3 of 3 - first time setup

> **Note:** For subsequent management class panels, provide null values.

Example: C-4   Management Class ACS routine

```
PROC 0 MGMTCLAS

/********************************************************************/
/*            DEFINE TAPE DATASET FILTERING CRITERIA               */
/********************************************************************/

 FILTLIST SMS_TAPE  INCLUDE('DCVTS','DC3590K')

 IF (&DATACLAS = &SMS_TAPE) THEN
   DO
     SET &MGMTCLAS = 'MCTAPE'
   EXIT
 END
END /* END PROC *
```

Storage group

```
                         TAPE STORAGE GROUP DEFINE
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Group Name  : VTS1

To DEFINE Storage Group, Specify:

  Description ==> Barry's big storage group
              ==>

  Library Names  (1 to 8 characters each):
    ===> BBBVTS1   ===>           ===>           ===>
    ===>           ===>           ===>           ===>



  DEFINE    SMS Storage Group Status . . Y (Y or N)


Use ENTER to Perform Verification and Selection;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-10 Tape Storage Group Define ISMF panel - first time setup

Notes:

- Select Storage Group Type=TAPE to define a tape storage group.
- Select 'Y' in the SMS Storage Group Status to alter the storage group status of your connected systems or system groups.
- The Library Name you assign in the Define Storage Group panel is matched with the name you use when you define the library (Figure C-11).

Example: C-5 Storage group ACS routine

```
PROC 0 STORGRP

  SELECT
     WHEN (&ANYVOL = 'REF=ST')
        DO
           SET &STORGRP = &STORGRP
           EXIT
        END
     WHEN (&DATACLAS = 'DCVTS')
        DO
           SET &STORGRP = 'VTS1'
           EXIT
        END
     WHEN (&DATACLAS = 'DC3590K')
        DO
           SET &STORGRP = 'ATL1'
           EXIT
        END
 END /* END SELECT */
END  /* END PROC   */
```

18. Define the Tape Library using ISMF, as shown in Figure C-12.

```
                              TAPE LIBRARY DEFINE                  Page 1 of 2
Command ===>

SCDS Name  . : CATIA.SCDS
Library Name : BBBVTS1

To Define Library, Specify:
  Description ==> BBBVTS1 Library Define
              ==>

  Library ID . . . . . . . . . . . 12345     (00001 to FFFFF)
  Console Name . . . . . . . . . .
  Entry Default Data Class . . . . DCVTS
  Entry Default Use Attribute  . . S         (P=PRIVATE or S=SCRATCH)
  Eject Default  . . . . . . . . . P          (P=PURGE or K=KEEP)

  Media Type:          Scratch Threshold     Media1  . . . . 0
    Media2 . . . . 400            Media3  . . . . 0         (0 to 999999)
    Media4 . . . . 0              Media5 . . . . 0         (0 to 999999)
    Media6 . . . . 0              Media7 . . . . 0         (0 to 999999)
    Media8 . . . . 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 C-11 Tape Library Define ISMF panel 1 of 2 - first time setup

```
                              TAPE LIBRARY DEFINE                  Page 2 of 2
Command ===>

SCDS Name  . : CATIA.SCDS
Library Name : BBBVTS1

  Initial Online Status (Yes, No, or Blank):
    SYS1     ===>  YES  SYS2     ===> NO   SYS3     ===>      SYS4     ===>
    SYS5     ===>







Warning:
  When you connect a tape library to a system group rather than a system,
  you lose the ability to vary that library online or offline to the
  individual systems in the system group.  It is strongly recommended that
  the tape library be connected to individual systems only.
Use ENTER to Perform Verification; Use UP Command to View previous Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit
```

Figure C-12 Tape Library Define ISMF panel 2 of 2 - first time setup

19. Translate the DFSMS ACS routines through ISMF.

20. Validate the DFSMS SCDS through ISMF.

21. Test the ACS routines through ISMF.
22. Activate the DFSMS SCDS.

Note: By previously varying the VTS drives online, upon activation of the SCDS, the library will come online automatically. You must have a VTS drive online before the library can come up. To bring a library online manually, issue the following command:

```
VARY SMS,LIB(library-name), ONLINE
```

Note: Once OAM is started, and the library is online, and the host has even a single path initialized for communication with the library, that host is capable of doing cartridge entry processing. The library drives do not have to currently be online for the path to be usable. If this library is being partitioned with other host systems, it is important that each system's tape management CBRUXENT first be coded and linked to prevent the host from entering volumes that do not belong to it.

23. Insert Logical Volumes (see "Inserting logical volumes into the VTS" on page 241)
24. VTS is now ready for use.

Tip: For help in navigating the ISMF panels, as well as detailed information pertaining to the defining of the DFSMS constructs, refer to the following manuals:

- *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632
- *z/OS DFSMSdfp Storage Administrator Reference*, SC35-0422

LM Upgrade to 527 or higher and VTS 2.26 or higher

Here we detail the upgrade process for upgrading from pre-527/2.26 LIVC levels to the latest LM/VTS levels.

Note: This scenario is only valid for 3494-attached VTS systems. IBM 3584/3953 configurations require higher LIC levels initially.

After your SSR has upgraded your 3494/VTS to the new LIC, the following events will occur upon VTS initialization:

1. All VTS stacked volumes are initially set to insert category FF01. This is a one time step at code install.

Note: This may take up to 3.5 hours.

2. VTS performs insert processing on volumes in FF01 and builds pooling, media, and media status (scratch or private) information.
3. Two pools for stacked volumes are now used:
   - Pool 00. Common Scratch Pool
   - Pool 01. Default private pool

Important: All stacked volumes (private and scratch) will have a category of FF04. The previous VTS category for stacked volumes scratch status FF03 is no longer used.

Important: ***Without FC4001 installed, no further action is required. If you have FC4001 installed, perform the following steps to utilize the Advanced Policy Management functions.***

Set up Volume Pooling in your environment

The following example shows how to use Advanced Policy Management in a stand-alone VTS environment. We assume that we will use logical volumes currently assigned to our system. For this example, the criteria will be to:

- Segregate production and testing logical volume data into pools.
- Set up logical volume dual copy.
- Set up early cache removal of logical volumes.

Define Outboard Policy Management Constructs

Through the 3494 Specialist or Library Manager, set up the Outboard Policy Management constructs:

1. Add two storage groups (see 4.4.2, “Defining storage groups” on page 145). For the purpose of this exercise, the storage groups will be named **SGPROD** and **SGTEST**.
   - **SGPROD** will be assigned Primary Pool **02**. This pool will contain logical volume production application data (Figure C-13).
   - **SGTEST** will be assigned Primary Pool **01**. This pool will contain logical volume test application data.

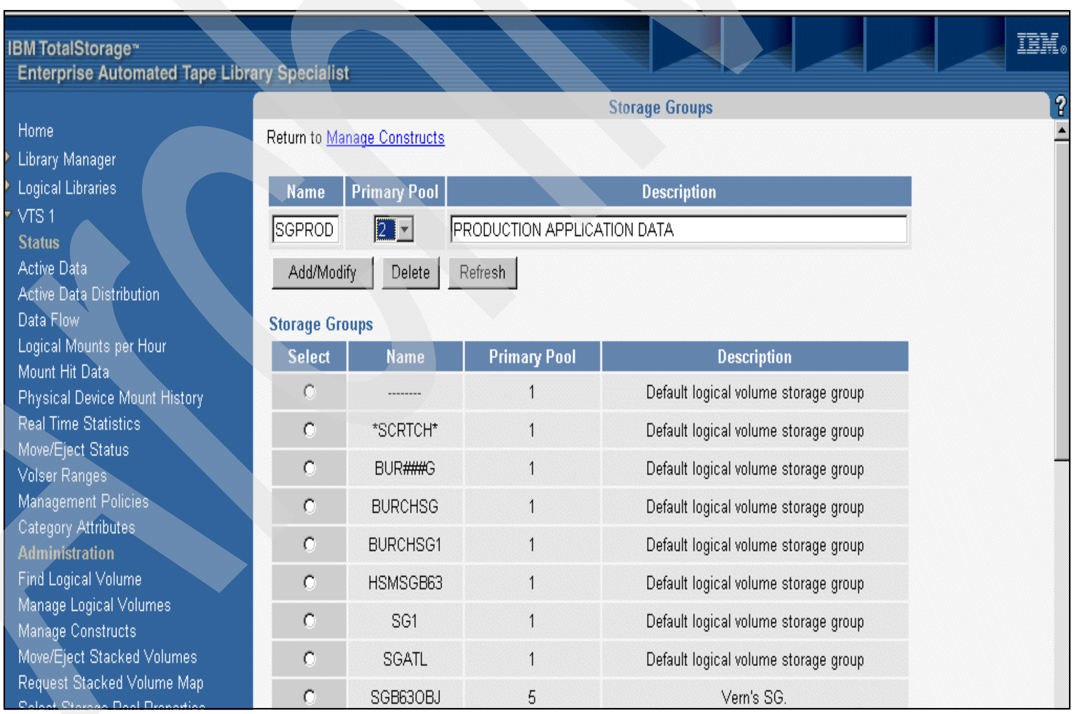

Figure C-13 3494 Specialist Manage Storage Groups - APM setup

2. Repeat the process to create the Library Manager construct for storage group **SGTEST**.

3. Create the DFSMS storage group construct for **SGPROD** through ISMF (as seen in Figure C-14)

```
                         TAPE STORAGE GROUP DEFINE
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Group Name  : SGPROD

To DEFINE Storage Group, Specify:

  Description ==> Production Application Storage Group
              ==>

  Library Names  (1 to 8 characters each):
    ===> BBBVTS1   ===>           ===>           ===>
    ===>           ===>           ===>           ===>



  DEFINE    SMS Storage Group Status . . N  (Y or N)


Use ENTER to Perform Verification and Selection;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-14 Storage Group define ISMF panel - APM setup

4. Repeat the process to create the DFSMS storage group construct **SGTEST**.
5. Create/modify your storage group DFSMS ACS routine (as shown in Example C-6).

Example: C-6 Storage group routine - APM setup

```
PROC 0 STORGRP

  SELECT

     WHEN (&ANYVOL = 'REF=ST')
        DO
           SET &STORGRP = &STORGRP
           EXIT
        END
     WHEN (&STORCLAS = 'SCCACHE')
        DO
           SET &STORGRP = 'SGPROD'
           EXIT
        END
     WHEN (&STORCLAS = 'SCNCACHE')
        DO
           SET &STORGRP = 'SGTEST'
           EXIT
        END
END /* END SELECT */
END /* END PROC */
```

6. Modify the pool properties for pools **01**,**02** and **03** (see “Defining stacked volume pool properties” on page 143). Here, you will need to define the parameters for your pools as described in 4.4.1, “Defining stacked volume pool properties” on page 143.
7. Repeat the same process for Pools 02 and 03.

Note: For the purpose of our exercise, **Pool 03** will be the pool for our logical volume dual copies. Therefore, we do not need to create the DFSMS storage group construct nor make any references to it in the DFSMS storage group ACS routine.

8. Add/modify stacked volser ranges for Pools **01**, **02**, and **03** (see “Defining volser ranges for stacked volumes” on page 222). Assign the following values to the pools:
   - Home pool
   - Media type
9. Insert stacked volumes.
10. Add two storage classes (see 4.4.4, “Creating storage classes” on page 149). For the purpose of this exercise the storage classes will be named **SCCACHE** and **SCNCACHE**. Here you will need to define the TVC preference using the storage class panel as described in 4.4.4, “Creating storage classes” on page 149.
    - **SCCACHE** will be assigned the TVC Preference of ‘**1**’. Logical Volumes assigned to this storage class are removed from the TVC by a Least Recently Used (LRU) algorithm. The logical volumes are copied to tape and deleted when space needs to be made available in the TVC.
    - **SCNCACHE** will be assigned the TVC Preference of ‘**0**’. Logical Volumes assigned to this storage class are removed from the TVC as soon as they are copied to tape. This case uses the largest first algorithm.

Note: Using the above Preference Levels of ‘0’ and ‘1’ overrides any IART values that you may currently have in effect. If you want to utilize current assigned IART values, specify IART as the Preference Level.

11. Repeat the process to create the Library Manager construct for storage class **SCNCACHE**.

12. Create the DFSMS storage class construct through ISMF (as shown in Figure C-15 and Figure C-16).

```
                              STORAGE CLASS DEFINE            Page 1 of 2
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Class Name  : SCCACHE
To DEFINE Storage Class, Specify:
  Description ==> KEEP IN CACHE STORAGE CLASS
              ==>
  Performance Objectives
   Direct Millisecond Response . . . .             (1 to 999 or blank)
   Direct Bias . . . . . . . . . . . .             (R, W or blank)
   Sequential Millisecond Response . .             (1 to 999 or blank)
   Sequential Bias . . . . . . . . . .             (R, W or blank)
   Initial Access Response Seconds . .             (0 to 9999 or blank)
   Sustained Data Rate (MB/sec)  . . .             (0 to 999 or blank)
  Availability . . . . . . . . . . . . N           (C, P ,S or N)
  Accessibility  . . . . . . . . . . . N           (C, P ,S or N)
   Backup  . . . . . . . . . . . . . .             (Y, N or Blank)
   Versioning  . . . . . . . . . . . .             (Y, N or Blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Page;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-15 Storage Class Define ISMF panel 1 of 2 - APM setup

```
                            STORAGE CLASS DEFINE            Page 2 of 2
Command ===>

SCDS Name . . . . . : CATIA.SCDS
Storage Class Name  : SCCACHE

To DEFINE Storage Class, Specify:

  Guaranteed Space . . . . . . . . . N           (Y or N)
  Guaranteed Synchronous Write . . . N           (Y or N)
  CF Cache Set Name  . . . . . . . .             (up to 8 chars or blank)
  CF Direct Weight . . . . . . . . .             (1 to 11 or blank)
  CF Sequential Weight . . . . . . .             (1 to 11 or blank)






Use ENTER to Perform Verification; Use UP Command to View previous Page;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-16 Storage Class Define ISMF panel 2 of 2 - APM setup

13. Repeat the process to create the DFSMS storage class construct **SCNCACHE**.

14. Create/Modify your storage class DFSMS ACS routine as shown in Example C-7.

Example: C-7 Storage class routine - APM setup

```
PROC 0 STORCLAS

 /*----------------------------------------------------------------*/
 /* DEFINE VALID TAPE UNIT ESOTERICS                                */
 /*----------------------------------------------------------------*/
   FILTLIST TAPE_UNIT INCLUDE(348*,TAPE*,'CART','AFF=SMST','3420',
                        '3590','3590-1','3490')

   FILTLIST PROD_DATA    INCLUDE(PROD.**)

   FILTLIST TEST_DATA    INCLUDE(TEST.**)

   SELECT
     WHEN ((&UNIT = &TAPE_UNIT) OR (&ANYVOL = 'REF=ST'))
       DO
         SELECT
             WHEN (&DSN = &PROD_DATA)
                 DO
                   SET &STORCLAS = 'SCCACHE'
                   EXIT
                 END
             WHEN (&DSN = &TEST_DATA)
                 DO
                   SET &STORCLAS = 'SCNCACHE'
                   EXIT
                 END
          END
       END
         OTHERWISE
           DO
             SET &STORCLAS = ''
               EXIT
           END
    END /* END SELECT */
 END  /* END PROC  */
```

15. Add two management classes (see 4.4.3, “Creating management classes” on page 147). For the purpose of this exercise, the management classes will be named **MCTEST** and **MCPROD**. The 3494 Specialist panel for this definition as described in 4.4.3, “Creating management classes” on page 147.

    – **MCTEST** will be assigned a Secondary Pool of ‘**00**’. No dual copy.

    – **MCPROD** will be assigned a Secondary Pool of ‘**03**. Logical volumes assigned to this management class will be dual copied. Stacked volumes defined to Pool 03 will be used.

Note: For a stand-alone VTS, the PTP Copy Control Mode values have no effect.

16. Repeat the process to create the Library Manager construct for management class **MCTEST.** Specify a Secondary Pool value of **0** for no dual copy.

17. Create the DFSMS management class construct through ISMF (see Figure C-17 and Figure C-19)

```
                         MANAGEMENT CLASS DEFINE              Page 1 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCPROD

To DEFINE Management Class, Specify:

  Description ==> PRODUCTION MANAGEMENT CLASS - DUAL COPY
              ==>

  Expiration Attributes
    Expire after Days Non-usage  . . NOLIMIT     (1 to 9999 or NOLIMIT)
    Expire after Date/Days . . . . . NOLIMIT     (0 to 9999, yyyy/mm/dd or
                                                 NOLIMIT)

  Retention Limit  . . . . . . . . . NOLIMIT     (0 to 9999 or NOLIMIT)



Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-17 Management Class Define ISMF panel 1 of 3 - APM setup

```
                         MANAGEMENT CLASS DEFINE              Page 2 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCPROD

To DEFINE Management Class, Specify:

  Partial Release . . . . . . . . . N          (Y, C, YI, CI or N)

  Migration Attributes
    Primary Days Non-usage  . . . .            (0 to 9999 or blank)
    Level 1 Days Non-usage  . . . .            (0 to 9999, NOLIMIT or blank)
    Command or Auto Migrate . . . . NONE       (BOTH, COMMAND or NONE)


  GDG Management Attributes
    # GDG Elements on Primary . . .            (0 to 255 or blank)
    Rolled-off GDS Action . . . . .            (MIGRATE, EXPIRE or blank)

Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-18 Management Class Define ISMF panel 2 of 3 - APM setup

```
                         MANAGEMENT CLASS DEFINE              Page 3 of 5
Command ===>

SCDS Name . . . . . . : CATIA.SCDS
Management Class Name : MCPROD

To DEFINE Management Class, Specify:
  Backup Attributes
    Backup Frequency  . . . . . . . .          (0 to 9999 or blank)
    Number of Backup Vers . . . . . .          (1 to 100 or blank)
      (Data Set Exists)
    Number of Backup Vers . . . . . .          (0 to 100 or blank)
      (Data Set Deleted)
    Retain days only Backup Ver . . .          (1 to 9999, NOLIMIT or blank)
      (Data Set Deleted)
    Retain days extra Backup Vers . .          (1 to 9999, NOLIMIT or blank)
    Admin or User command Backup  . . NONE     (BOTH, ADMIN or NONE)
    Auto Backup . . . . . . . . . . . N        (Y or N)
    Backup Copy Technique . . . . . . S        (P=Conc Preferred, R=Conc
                                               Required or S=Standard)
Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; Cancel to Exit.
```

Figure C-19 Management Class Define ISMF panel 3 of 3 - APM setup

> **Note:** For subsequent management class panels, provide null values.

18. Repeat the process to create the DFSMS management class construct **MCTEST**.
19. Create/modify your management class DFSMS ACS routine like the one listed in Example C-8.

Example: C-8 Management Class routine - APM setup

```
PROC 0 MGMTCLAS

/*******************************************************************/
/*             DEFINE TAPE DATASET FILTERING CRITERIA              */
/*******************************************************************/
    SELECT
       WHEN (&STORCLAS = 'SCCACHE')
           DO
              SET &MGMTCLAS = 'MCPROD'
              EXIT
           END
       WHEN (&STORCLAS = 'SCNCACHE')
           DO
              SET &MGMTCLAS = 'MCTEST'
              EXIT
           END
    END  /* END SELECT */
 END  /* END PROC */
```

20. Define a data class (see 4.4.5, “Creating data classes” on page 150). For the purpose of this exercise the data class will be named **DCVTS1** and can be defined using the panel shown in Figure 4-27 on page 151.

21. Create the DFSMS data class construct through ISMF as shown in Figure C-20, Figure C-21, and Figure C-22.

```
                              DATA CLASS DEFINE                Page 1 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVTS1

To DEFINE Data Class, Specify:
  Description ==>
              ==>
  Recorg  . . . . . . . .              (KS, ES, RR, LS or blank)
  Recfm . . . . . . . . .              (any valid RECFM combination or blank)
  Lrecl . . . . . . . . .              (1 to 32761 or blank)
  Keylen  . . . . . . . .              (0 to 255 or blank)
  Keyoff  . . . . . . . .              (0 to 32760 or blank)
  Space Avgrec  . . . . .              (U, K, M or blank)
        Avg Value . . . .              (0 to 65535 or blank)
        Primary . . . . .              (0 to 999999 or blank)
        Secondary . . . .              (0 to 999999 or blank)
        Directory . . . .              (0 to 999999 or blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

Figure C-20 Data Class Define ISMF panel 1 of 3 - APM setup

```
                              DATA CLASS DEFINE                Page 2 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVTS1
To DEFINE Data Class, Specify:
  Retpd or Expdt  . . . . . .          (0 to 9999, YYYY/MM/DD or blank)
  Volume Count  . . . . . . .          (1 to 59 or blank)
    Add'l Volume Amount . . .          (P=Primary, S=Secondary or blank)
  Imbed . . . . . . . . . . .          (Y, N or blank)
  Replicate . . . . . . . . .          (Y, N or blank)
  CIsize Data . . . . . . . .          (1 to 32768 or blank)
  % Freespace CI  . . . . . .          (0 to 100 or blank)
              CA  . . . . . .          (0 to 100 or blank)
  Shareoptions Xregion  . . .          (1 to 4 or blank)
               Xsystem  . . .          (3, 4 or blank)
  Compaction  . . . . . . . . y        (Y, N, T, G or blank)
  Media Interchange
    Media Type  . . . . . . . 2        (1, 2, 3, 4 or blank)
    Recording Technology  . . 36       (18, 36, 128, 256 or blank)
  Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
  Use HELP Command for Help; Use END Command to Save and Exit; Cancel to Exit.
```

Figure C-21 Data Class Define ISMF panel 2 of 3 - APM setup

```
 DATA CLASS DEFINE              Page 3 of 3
Command ===>

SCDS Name . . . : CATIA.SCDS
Data Class Name : DCVTS1

To DEFINE Data Class, Specify:
  Data Set Name Type  . . . . . .      (EXT, HFS, LIB, PDS or blank)
    If Ext  . . . . . . . . . . .      (P=Preferred, R=Required or blank)
    Extended Addressability . . . N    (Y or N)
    Record Access Bias  . . . . .      (S=System, U=User or blank)
  Reuse . . . . . . . . . . . . . N    (Y or N)
  Initial Load  . . . . . . . . . R    (S=Speed, R=Recovery or blank)
  Spanned / Nonspanned  . . . . .      (S=Spanned, N=Nonspanned or blank)
  BWO . . . . . . . . . . . . . .      (TC=TYPECICS, TI=TYPEIMS, NO or blank)
  Log . . . . . . . . . . . . . .      (N=NONE, U=UNDO, A=ALL or blank)
  Logstream Id  . . . . . . . . .
  Space Constraint Relief . . . . N    (Y or N)
    Reduce Space Up To (%)  . . .      (0 to 99 or blank)
  Block Size Limit  . . . . . . .              (32760 to 2GB or blank)
Use ENTER to perform Verification; Use UP Command to View previous Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit
```

Figure C-22 Data Class Define ISMF panel 3 of 3 - APM setup

22. Create/modify your data class DFSMS ACS routine as shown in Example C-9.

Example: C-9 Data Class routine - APM setup.

```
PROC 0 DATACLAS

 /*-----------------------------------------------------------------*/
 /* DEFINE VALID TAPE UNIT ESOTERICS                                */
 /*-----------------------------------------------------------------*/
     FILTLIST TAPE_UNIT INCLUDE(348*,TAPE*,'CART','AFF=SMST','3420',
                            '3590','3590-1','3490')

     FILTLIST VTS_DATA    INCLUDE(PROD.**,TEST.**)

     SELECT
        WHEN ((&UNIT = &TAPE_UNIT) OR (&ANYVOL = 'REF=ST'))
          DO
            SELECT
               WHEN (&DATACLAS NE '')                /* Allows users to specify  */
                 DO                                  /* data class for tape      */
                   SET &DATACLAS = &DATACLAS
                   EXIT
                 END
               WHEN (&DSN = &VTS_DATA)
                 DO
                   SET &DATACLAS = 'DCVTS1'
                   EXIT
                 END
            END
          END
            OTHERWISE
              DO
                SET &DATACLAS = ''
                EXIT
              END
```

```
     END /* END SELECT */
END  /* END PROC */
```

23. If they have been changed, translate the following DFSMS ACS routines through ISMF:
    - Data class
    - Storage class
    - Management class
    - Storage group
24. Validate the DFSMS SCDS through ISMF.
25. Test the ACS routines through ISMF.
26. Activate the DFSMS SCDS.

D

DEVSERV QLIB command

This appendix lists the syntax and the parameter explanation copied out of the cover letter of the PTFs for APAR OA07505 (see Example D-1).

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. Do not use the command for just testing. It will impact the performance of your VTS since it takes time to take the dump in the hardware.

Example: D-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

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).
```

E

Library Manager volume categories

Table E-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. See DEVSUP discussion in 4.3.5, "Define Fast Ready categories" on page 133.

Table E-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	z/OS DFSMS	Indicates scratch MEDIA1. MEDIA1 is a standard-capacity cartridge system tape.
0002	z/OS DFSMS	Indicates scratch MEDIA2. MEDIA2 is an enhanced-capacity cartridge system tape Type E.
0003	z/OS DFSMS	Indicates scratch MEDIA3. MEDIA3 is the IBM 3590 High Performance Tape Cartridge Type J.
0004	z/OS DFSMS	Indicates scratch MEDIA4. MEDIA4 is the IBM 3590 Extended High Performance Tape Cartridge Type K.

Category (in hex)	Used by	Definition
0005	z/OS DFSMS	Indicates scratch MEDIA5. MEDIA5 is the IBM 3592 Enterprise Tape Cartridge Type JA.
0006	z/OS DFSMS	Indicates scratch MEDIA6. MEDIA6 is the IBM 3592 Enterprise WORM Tape Cartridge Type JW.
0007	z/OS DFSMS	Indicates scratch MEDIA7. MEDIA7 is the IBM 3592 Enterprise Economy WORM Tape Cartridge Type JR.
0008	z/OS DFSMS	Indicates scratch MEDIA8. MEDIA8 is the IBM 3592 Enterprise Economy Cartridge Type JJ.
0009 to 000D	z/OS DFSMS	Reserved
000E	z/OS DFSMS	Indicates an error volume. Volumes in this category are scratch volumes for which the software detected an error during processing.
000F	z/OS DFSMS	Indicates a private volume. Volumes in this category contain user data or are assigned to a user.
0010 to 007F	z/OS DFSMS	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

Category (in hex)	Used by	Definition
008C	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHC
008D	DFSMS/VM including VSE Guest	Indicates that the volume belongs to the VM category SCRATCHD
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

Category (in hex)	Used by	Definition
00B0	Native VSE/ESA	Indicates that the volume belongs to the VSE category SCRATCH16
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.

Category (in hex)	Used by	Definition
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.
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 z/OS DFSMS, 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.

Category (in hex)	Used by	Definition
F00E	BTLS	Indicates a volume in error. Volumes are assigned to the error category during demount if the volume serial specified for demount does not match the external label of the volume being demounted.
F00F to FEFF	-	No assignment to a specific host system. These categories can be dynamically assigned by the Library Manager on request of a host.
FF00	All	Insert category. When a tape volume is added to an automated tape library, the library reads the external label on the volume, creates an inventory entry for the volume and assigns the volume to the insert category. This category 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 as described in 6.7, "Ejecting logical volumes from the VTS" on page 248, or exported as described in "Cache Management exploitation" on page 97.
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 as described in 6.7, "Ejecting logical volumes from the VTS" on page 248, or exported as described in "Cache Management exploitation" on page 97.
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.

F

SMF type 94 record layout

The tape library accumulates statistics over a period of one hour. These statistics represent the activity of the IBM Tape Library resulting from all hosts attached to the library.

At the end of the hour, the system calculates composite statistics about the activity of all devices in the library and writes an SMF type 94 record.

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 in this Appendix 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.

```
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  SMF94ATO      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 AXO
```

```
  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- currency during the last
                                             hour.
  532  (214) UNSIGNED      2  S94OPM_MINVDM  Min Virtual Drives Mounted
                                             con- currency during the last
                                             hour.
  534  (216) UNSIGNED      2  S94OPM_AVGVDM  Avg Virtual Drives Mounted
                                             con- currency 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
```

Crossreference

CROSS REFERENCE

NAME	HEX OFFSET	HEX 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

NAME	HEX OFFSET	HEX 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_AXOVTS0	AB	80	3
S94CLLS0_AXOVTS1	AB	08	3
S94CLLS1_AXOVTS0	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

G

VTS feature codes

In this appendix we describe all feature codes that are related to the installation of a VTS. We group the feature codes according to the frame models to which they belong. Feature codes refer to orderable and chargeable control unit features which are not usually standard on the base machine. If you install additional frames other than a VTS and its attached Model D12, refer to the *IBM 3494 Tape Library Introduction and Planning Guide*, GA32-0279 and to the *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632.

IBM TotalStorage Model B10 VTS and Model B20 VTS feature codes

These feature codes can be installed on the Model B10 VTS and Model B20 VTS frame and determine the configuration of the VTS. The feature codes are used to specify the Remote Support Facility configuration, the size of the TVC, the type and number of channel attachments to be shipped with the Model B10 VTS or Model B20 VTS, and other options. Table G-1 summarizes the feature codes required on the Model B10 VTS and Model B20 VTS.

Table G-1 Model B10 and B20 VTS specifications and specified feature codes

Feature code	Min. / Max. qty	Field / Factory install	Description
0201	0 / 8	Yes / Yes	9 Micron FICON LC/LC 31 meter
0202	0 / 8	Yes / Yes	9 Micron FICON LC/SC 31 meter
0203	0 / 8	Yes / Yes	50 Micron FICON LC/LC 31 meter
0204	0 / 8	Yes / Yes	50 Micron FICON LC/SC 31 meter
0205	0 / 8	Yes / Yes	62.5 Micron FICON LC/LC 31 meter
0206	0 / 8	Yes / Yes	62.5 Micron FICON LC/SC 31 meter
2710	0 / 1	Yes / Yes	Remote Support Facility
2711	0 / 1	Yes / Yes	Remote Support Switch

Feature code	Min. / Max. qty	Field / Factory install	Description
2712	0 / 1	Yes / Yes	Remote Support Attachment
2713	0 / 1	Yes / Yes	Master Console for service
2714	0 / 1	Yes / Yes	Console Expansion
2715	0 / 1	Yes / Yes	Console Attachment
2716	0 / 1	Yes / No	Console Second Modem
3000	0 / 1	Yes / Yes	FICON Enablement
3412	0 / 4	Yes / Yes	Extended Perf. ESCON Channels
3415	0/ 2/4(B10) 0/4/8(B20)	Yes / Yes	Long Wave FICON Channel
3416	0/2/4(B10) 0/4/8(B20)	Yes / Yes	Short Wave FICON Channel
3418	0 / 1 (B10) 0 / 4 (B20)	Yes / Yes	Activate Add'l ESCON Channels
3422	0 / 4	Yes / Yes	SCSI Host Attachment
3704	0 / 1	No / Yes	144/216 GB Disk Storage
3705	0 / 1 (B10) 2 / 4 (B20)[1]	Yes / Yes	288/432 GB Disk Storage
4000	0 / 1	Yes / Yes	Advanced Function
4001	0 / 1	Yes / Yes	Adv. Policy Mgmt up to 250 GB
4002	0 / 1	Yes / Yes	Adv. Policy Mgmt up to 500 GB
4003	0 / 1	Yes / Yes	Adv. Policy Mgmt up to 1000 GB
4004	0 / 1	Yes / Yes	Adv. Policy Mgmt up to 2000 GB
4010	0 / 1	Yes / Yes	Peer-to-Peer Copy Base
4011	0 / 1	Yes / Yes	Peer-to-Peer Copy Increment 1
4012	0 / 1	Yes / Yes	Peer-to-Peer Copy Increment 2
4013	0 / 1	Yes / Yes	Peer-to-Peer Copy Increment 3
4036	0 / 5	Yes / Yes	Increased Logical Volumes
5001	0 / 4	Yes / Yes	4.5 Meter SCSI Cable
5002	0 / 4	Yes / Yes	10 Meter SCSI Cable
5003	0 / 4	Yes / Yes	20 Meter SCSI Cable
5004	0 / 4	Yes / Yes	10 Meter SCSI VHDCI Cable
5235	0 / 2[2]	Yes / Yes	20 Meter SCSI Drive Cables
5238	0 / 1	Yes / Yes[3]	3592 Drives attached to VTS
5264	0 / 2 (B20 only)	Yes / Yes	64 Add'l Virtual Drives
5265	0 / 2 /4 (B20 only)	Yes / Yes	Incremental Virtual Drives

Feature code	Min. / Max. qty	Field / Factory install	Description
9106	0 / 1	Yes / Yes	Attach to RS/6000, pSeries
9109	0 / 1	Yes / Yes	Attach to S/390, zSeries
9201	0 / 99	Yes / Yes	VTS Open System Device Drivers
9210	0 / 1	Yes / Yes	HP-UX Attachment
9211	0 / 1	Yes/Yes	Sun Attachment
9212	0 / 1	Yes / Yes	Attached to Windows
9700	0 / 1	Yes / Yes	No FICON cable from plant
9702	0 / 8	Yes / Yes	Interposer, Double Byte Wide
9798	0 / 8	Yes / Yes	Inline SCSI Terminator
9799	0 / 8	Yes / Yes	VHDCI Cable/Interposer
9986	0 / 1	Yes / Yes	6 ft Chicago Power Cord

Notes:

1. For the Model B10 VTS, one FC3705 may be ordered. For the Model B20 VTS, two or four only of FC3705 must be ordered.
2. For the Model B10 VTS, one FC5235 must be ordered. Factory only for Model B10 VTS
3. Factory only on the 3494 B10 VTS. Factory or Field on the 3494 B20 VTS

For details on the feature codes listed in Table G-1, review the following descriptions.

9 Micron LC/LC 31 Meter FICON Cable (FC0201)

The feature provides a 31 meter long, 9 micron FICON cable with LC/LC connectors.

9 Micron LC/SC 31 Meter FICON Cable (FC0202)

The feature provides a 31 meter long, 9 micron FICON cable with LC/SC connectors.

50 Micron LC/LC 31 Meter FICON Cable (FC0203)

The feature provides a 31 meter long, 50 micron FICON cable with LC/LC connectors.

50 Micron LC/SC 31 Meter FICON Cable (FC0204)

The feature provides a 31 meter long, 50 micron FICON cable with LC/SC connectors.

62.5 Micron LC/LC 31 Meter FICON Cable (FC0205)

The feature provides a 31 meter long, 62.5 micron FICON cable with LC/LC connectors.

62.5 Micron LC/SC 31 Meter FICON Cable (FC0206)

The feature provides a 31 meter long, 62.5 micron FICON cable with LC/SC connectors.

Remote Support Facility (FC2710)

This special feature provides a cable and connectors for connection of Model B10 VTS or Model B20 VTS to an IBM supplied modem to enable remote diagnostic support. This feature should only be ordered for the first remote support product at an installation.

One FC2710, FC2711 or FC2712 must be ordered for each Model B10 VTS or Model B20 VTS.

Remote Support Switch (FC2711)

This special feature provides a switch, cables and connectors for the attachment of the 3494 Model B10 VTS or Model B20 VTS to an IBM supplied modem through the switch for remote diagnostic support. This feature code should be ordered for second unit connection to the modem in an installation. One feature can provide a connection for fourteen support attachment.

One FC2710, FC2711 or FC2712 must be ordered for each Model B10 VTS or Model B20 VTS.

Remote Support Attachment (FC2712)

This special feature provides a cable and connectors for the attachment of the Model B10 VTS or Model B20 VTS to the Remote Support Switch (FC2711) for remote diagnostic support. This feature code should be ordered on the third through the fourteenth unit attached to the Remote Support Switch in an installation site.

One FC2710, FC2711 or FC2712 must be ordered for each Model B10 VTS or Model B20 VTS.

Master Console for Service (FC2713)

This feature provides the Master Console, an Ethernet hub and cable to an attached product for connection to an IBM supplied modem to enable enhanced remote service. One of the feature FC2713, FC2714 or FC2715 must be ordered for each Model B10 VTS or Model B20 VTS. The Ethernet hub provides 14 additional connections for cables supplied with FC2714 or FC2715. One of FC2714 may be connected to the FC2713.

Console Expansion (FC2714)

This feature provides an Ethernet hub and product attachment cable for expanding the number of products to be attached to the Master Console For Service (FC2713). Up to 14 additional connections of FC2714 or FC2715 are provided for product attachment by this feature. One of FC2713, FC2714 or FC2715 must be ordered for each Model B10 VTS or Model B20 VTS. One of FC2714 may be connected to the Ethernet hub provided by this FC2714. A maximum of two of FC2714 may be included in a single Master Console For Service (FC2713) to provide a maximum of 43 product connections.

Console Attachment (FC2715)

This feature provides a cable for connecting a Model B10 VTS and Model B20 VTS to the Ethernet hub provided by FC2713 or FC2714. One of FC2713, FC2714 or FC2715 must be ordered for each 3494 Model B10 VTS or Model B20 VTS. A maximum of 40 of FC2715 may be included in a single Master Console For Service (FC2713).

Console Second Modem (FC2716)

This feature provides a second modem for the IBM TotalStorage MasterConsole. This feature should be only be ordered on 3494 Models AX0, B10, B18, or B20 shipped prior to September 6, 2002, as it is included in any feature FC2713 shipped after that date.

FICON Enablement (FC3000)

This feature enables the attachment of FICON channels to a VTS Model B10 or Model B20. This feature is mutually exclusive with the Peer-to-Peer Copy Base feature (FC4010) or the SCSI Host Attachment feature (FC3422).

Extended Performance ESCON Channels (FC3412)

This special feature provides two Extended Performance ESCON host channel attachments. Up to 64 logical channels can be connected per channel attachment. Two 30m (100 ft.) ESCON cables are provided with this feature. If other cable lengths are required, order group number 3797 or group number 8486 for up to 122m (400 ft.) cables.

For the Model B10 VTS, a maximum of two of this FC3412 may be factory or field installed without FC3418. If FC3418 is installed, only one FC3412 may be installed. If FC3412 is not ordered, FC3422 must be ordered. With FC4010, Peer-to-Peer Copy Base, two of FC3412 must be installed.

For the Model B20 VTS, two or four of this FC3412 may be factory installed and two of FC3412 may be field installed for a maximum of four of FC3412. With FC4010, four of FC3412 must be installed.

Long Wave FICON Channel (FC3415)

Feature #3415 provides one Long Wave FICON host channel attachment. It can be used for attachment of the VTS Model B10 or B20 to host FICON channels. Up to 128 logical channels can be connected per channel attachment. A maximum of four of this feature can be installed on the B10 and eight on the B20. For each FC3415 ordered, one of the cable features FC0201 to FC0206 or FC9700 must be specified.

Short Wave FICON Channel (FC3416)

This Feature provides one Short Wave FICON host channel attachment. It can be used for attachment of the VTS Model B10 or B20 to host FICON channels. Up to 128 logical channels can be connected per channel attachment. A maximum of four of this feature can be installed on the B10 and eight on the B20. For each FC3416 ordered, one of the cable features FC0201 to FC0206 or FC9700 must be specified.

Activate Additional ESCON Channels (FC3418)

This special feature activates two additional Extended Performance ESCON Channels for FC3412 installed on the Model B10 VTS or Model B20 VTS. Two 30 m (100 ft.) ESCON cables are provided with this feature. If other cable lengths are required, order group number 3797 or group number 8486 for up to 122 m (400 ft.) cables.

For the Model B10 VTS, a maximum of one of this FC3418 may be factory or field installed. FC3412 and two of FC3422 are prerequisites for FC3418 on the Model B10 VTS.

For the Model B20 VTS, two or four of this FC3418 may be factory or field installed for a maximum of four of FC3418. FC3412 is a prerequisites for FC3418 and the number of FC3418 must be equal to the number of FC3412.

SCSI Host Attachment (FC3422)

This feature provides two SCSI ports for attachment to the SCSI-2 (Fast, Wide, Differential) 20 MB/s standard interface or the 40 MB/s Differential Wide Ultra SCSI interface on open system hosts. Supported systems with FC3422 are pSeries, Sun, HP, Windows NT and Windows 2000. FC9201 is a prerequisite to FC3422. Any or all of FC9106, FC9211, FC9212, or FC9210 must also be specified.

For the Model B10 VTS, a quantity of two of FC3422 must be installed when one of FC3412 and one FC3418 are installed. Four of FC3422 must be installed when FC3412 is not present.

For the Model B20 VTS, four of FC3422 can be installed only when there are two of FC3412 and two of FC3418.

144/216 GB Disk Storage (FC3704)

This feature provides 216 GB of usable storage capacity (Performance Accelerator function is standard on the Model B10 VTS) for the tape volume cache in the Model B10 VTS. FC3704 is not available with FC3705 installed. FC3704 may be converted to FC3705. A quantity of one of FC3704 is allowed to be ordered for the Model B10 VTS for factory installation only.

288/432 GB Disk Storage (FC3705)

This feature provides 432 GB of usable storage capacity (Performance Accelerator function is standard on the Model B10 VTS and Model B20 VTS) for the tape volume cache in the Model B10 VTS or Model B20 VTS.

For the Model B10 VTS, one FC3705 may be ordered. FC3704 may be converted to FC3705 in the field. FC3705 is not available with FC3704 installed.

For the Model B20 VTS, two or four only of FC3705 must be ordered. A quantity of two of FC3705 may be field installed in a Model B20 VTS.

Advanced Function (FC4000)

This feature activates advanced functions such as Import/Export of logical volumes and Tape Volume Cache management. Tape Volume cache management enables efficient use of cache and potentially improves the likelihood of cache hits.

The Advanced Function feature is supported on z/OS systems with the appropriate software support as described in 3.8, “Software requirements” on page 93.

FC4000 requires minimum of four TotalStorage 3590 tape drives in the associated D12 Frame. In addition, the 3494 must have a convenience I/O station for Import/Export operation.

This feature cannot be installed if FC4010 is already installed and installing FC4010 will disable this feature.

Advanced Policy Management up to 250 GB (FC4001)

This feature provides policy management for tape volume cache residency, volume pooling, dual volume copy, and Peer-to-Peer copy mode control, and the import/export of logical volumes, to further exploit the use of tape and disk subsystems in the VTS. Details of this feature (and including FC4002 to FC4004) are covered within Chapter 4, “Implementation” on page 107.

These are the prerequisites:

- Use of the Import/Export function on a non-PtP VTS requires a Convenience I/O Station (FC5210 or FC5230) on the 3494 Tape Library Base Frame.
- A minimum disk storage capacity of 216 GB is required.
- On the Model B18 a minimum of four 3590 Tape Drives in the associated Model D12 (or FC5502) and the 1 Performance Accelerator feature (FC5236) are required.
- VTS LIC level 2.26 is required.

Advanced Policy Management up to 500 GB (FC4002)

This feature provides policy management for tape volume cache residency, volume pooling, dual volume copy, and copy mode control, and the import/export of 1 logical volumes, to further exploit the use of tape and disk subsystems in the VTS. This feature is required on a 1 VTS with up to 500 GB of disk storage capacity. The Import/Export function may be disabled by an IBM System Service Representative (SSR) using the Library Manager service menus.

Prerequisite: Feature FC4001. A Convenience I/O Station (FC5210 or FC5230) is required on the 3494 Tape Library Base Frame when the Import/Export function is not disabled.

Advanced Policy Management up to 1000 GB (FC4003)

This feature provides policy management for tape volume cache residency, volume pooling, dual volume copy, and copy mode control, and the import/export of logical volumes, to further exploit the use of tape and disk subsystems in the VTS. This feature is required on a 1 VTS with up to 1000 GB of disk storage capacity. The Import/Export function may be disabled by an SSR using the Library Manager service menus.

Prerequisite: Feature FC4001 and FC4002. A Convenience I/O Station (FC5210 or FC5230) is required on the 3494 Tape Library Base Frame when the Import/Export function is not disabled.

Advanced Policy Management up to 2000 GB (FC4004)

This feature provides policy management for tape volume cache residency, volume pooling, dual volume copy, and copy mode control, and the import/export of logical volumes, to further exploit the use of tape and disk subsystems in the VTS. This feature is required on a VTS with up to 2000 GB of disk storage capacity. The Import/Export function may be disabled by an SSR using the Library Manager service menus.

Prerequisite: Feature FC4001, FC4002, and FC4003. A Convenience I/O Station (FC5210 or FC5230) is required on the 3494 Tape Library Base Frame when the Import/Export function is not disabled.

> **Note:** For FC4001 to FC4004, the Import/Export function must be disabled if the Peer-to-Peer Copy Base feature (FC4010) is installed.

Peer-to-Peer Copy Base (FC4010)

This feature enables the IBM TotalStorage Model B10 VTS or Model B20 VTS to function in a Peer-to-Peer VTS configuration. This feature must be installed on each Model B10 VTS or Model B20 VTS that attaches to the four IBM TotalStorage Virtual Tape Controllers in the Peer-to-Peer VTS.

This feature cannot be installed with FC3422. Installing FC4010 will disable already installed FC4000 and FC3418. A VTS with FC4010 can only be used in a Peer-to-Peer configuration.

Peer-to-Peer Copy Increment 1 (FC4011)

This feature extends the Model B10 VTS or Model B20 VTS to function in a Peer-to-Peer VTS configuration when greater than 250 GB usable disk storage capacity is installed. This feature is required on a Model B10 VTS or Model B20 VTS in a Peer-to-Peer VTS when one or more of FC3705 are installed.

Peer-to-Peer Copy Increment 2 (FC4012)

This feature extends the Model B10 VTS or Model B20 VTS to function in a Peer-to-Peer VTS configuration when greater than 500 GB usable disk storage capacity is installed. This feature is required on a Model B10 VTS or Model B20 VTS in a Peer-to-Peer VTS when two or more of FC3705 are installed.

Peer-to-Peer Copy Increment 3 (FC4013)

This feature extends the Model B10 VTS or Model B20 VTS to function in a Peer-to-Peer VTS configuration when greater than 1 TB usable disk storage capacity is installed. This feature is required on a Model B10 VTS or Model B20 VTS in a Peer-to-Peer VTS when four of FC3705 are installed.

Increased Logical Volumes (FC4036)

This feature provides the system reconfiguration necessary to enable addition of more logical volumes to the base number of 250,000. Each order of this feature increases the maximum number of logical volumes supported by 50,000, up to a maximum of 500,000.

Note: Installation of this feature will reduce the amount of available customer storage on the physical volumes. This is because the amount of space reserved for the database used for disaster recovery of a VTS subsystem is increased from 250 MB to 500 MB.

FC5245, Dual Attachment Concentrator, is required when FC4036 is installed.

4.5 Meter SCSI Cable (FC5001)

This feature provides one 4.5 m (15 ft.) SCSI cable for host attachment to a IBM TotalStorage VTS with FC3422. It has a HD68 standard 68 pin straight housing connector for host attachment and a VHDCI 0.8 mm connector for VTS attachment.

10 Meter SCSI Cable (FC5002)

This feature provides one 10 m SCSI cable for host attachment to a IBM TotalStorage VTS with FC3422. It has a HD68 standard 68 pin straight housing connector for host attachment and a VHDCI 0.8 mm connector for VTS attachment.

20 Meter SCSI Cable (FC5003)

This feature provides one 20 m SCSI cable for host attachment to a IBM TotalStorage VTS with FC3422. It has a HD68 standard 68 pin straight housing connector for host attachment and a VHDCI 0.8 mm connector for VTS attachment.

10 Meter SCSI VHDCI Cable (FC5004)

This feature provides one cable with VHDCI 0.8 mm connectors on both ends of 10 m length. It attaches a VTS to another VTS.

20 Meter SCSI Drive cables (FC5235)

This feature provides two 20 m (65 ft.) cables with VHDCI 0.8mm connector on one end and HD-68 standard 68-pin straight housing connector on the other end. FC5235 attaches a Model Bxx VTS to up to six 3590 Model B1A, or E1A tape drives in a D12 Frame.

For the Model B10 VTS, on FC5235 must be specified. Factory only for the Model B10 VTS.

For the Model B20 VTS, one or two of FC5235 may be specified. When one of FC5235 is installed or ordered for attachment of 3590 in a D12 Frame.

3592 Drive attached to VTS (FC5238)

This feature provides the Fibre Channel cards in the 3494 B10 or B20 VTS to attach 3592 Tape Drives in an associated 3494 or 3584 Frame to the VTS. It includes the Fibre Channel cables from the VTS to each 2Gb Fibre Channel Switch feature (#3487) in the associated 3494 D22 Frame or 3953 F05 Frame and the appropriate microcode firmware to support that attachment.

64 Additional Virtual Drives (FC5264)

This feature provides an increment of 64 virtual tape drives for a VTS. For the Model B20 VTS, may be ordered only with a quantity of two. FC5264 is not available for the Model B10 VTS.

Note: In a PtP system, adding FC5264 allows each AX0 to support 32 devices.

Incremental Virtual Drives(FC5265)

This feature provides an increment of 64 virtual tape drives for a VTS and should be ordered on a new 3494 B20 VTS instead of FC5264. Each quantity of this feature will allow the VTS to permit an additional 64 virtual devices to be addressed by the host system, up to the maximum number of virtual devices (256) supported by the VTS. When ordered on the 3494 B20 VTS, the quantity must be two or four and indicates the number of virtual tape drives accessible by the host system. When two of this feature is ordered, a PtP VTS configuration of four Virtual Tape Controllers (VTCs) with 32 devices per VTC is supported, for a total of 128 virtual devices. When four of this feature is ordered, a PtP VTS configuration of eight Virtual Tape Controllers (VTCs) with 32 devices per VTC is supported, for a total of 256 virtual devices.

Attach to RS/6000, pSeries (FC9106)

This feature notifies the manufacturing plant that an 3494 frame is to be attached to a pSeries host processor. It is used for documentation. One of FC9106, FC9109, FC9210, FC9211, or FC9212 must be installed.

Attach to S/390, zSeries (FC9109)

This feature notifies the manufacturing plant that an 3494 frame is to be attached to a zSeries host processor. It is used for documentation. One of FC9106, FC9109, FC9210, FC9211, or FC9212 must be installed.

VTS Open System Device Drivers (FC9201)

This feature provides the device drivers for the attachment of the VTS to Open Systems. FC3422 (SCSI Host Attachment) and one of FC9106, FC9109, FC9210, FC9211, or FC9212 must be installed. This feature may be ordered multiple times in order to obtain the latest device drivers for attachment support.

HP-UX Attachment (FC9210)

This feature notifies the manufacturing plant that an 3494 frame is to be attached to a HP host processor. It is used for documentation. One of FC9106, FC9109, FC9210, FC9211, or FC9212 must be installed.

Sun Attachment (FC9211)

This feature notifies the manufacturing plant that an 3494 frame is to be attached to a Sun host processor. It is used for documentation. One of FC9106, FC9109, FC9210, FC9211, or FC9212 must be installed.

Attached to Windows (FC9212)

This feature notifies the manufacturing plant that an 3494 frame is to be attached to a Microsoft Windows 2000 or MS Windows NT host processor. It is used for documentation. One of FC9106, FC9109, FC9210, FC9211, or FC9212 must be installed.

No FICON Cable from Plant (FC9700)

This feature is specified if FC3415 or FC3416 is ordered and a FICON cable is not required.

Interposer, Double Byte Wide (FC9702)

This interposer makes it possible to connect an industry standard high-density 68-pin (HD68) cable connector, typically used on fast/wide devices, to a 68-pin ribbon connector host system adapter such as the pSeries Enhanced SCSI-2 Differential Fast/Wide Adapter/A (FC2412).

Inline SCSI Terminator (FC9798)

This feature provides the required host termination to connect a SCSI cable from HD68 cable connector to an HP Fast/Wide Differential SCSI-2 adapter (A4800A) on an HP V-class system. Cable features FC5001, FC5002, or FC5003 may be used with this terminator when attaching to the V-class server.

VHDCI Cable/Interposer (FC9799)

This feature provides a short cable with a VHDCI 0.8 mm connector for Open System host connection and a HD69 connector which attaches to the cable provided by FC5001, FC5002 or FC5003.

Six-Foot Chicago Power Cord (FC9986)

Feature code 9986 is a specify feature that can be ordered for field and factory installation. If you do not order this feature, which is required only in Chicago, the appropriate power cord is attached at the factory on the basis of the country to which the VTS is designated.

Note that if a Chicago power cord is required, you must specify one for every library frame except the IBM 3494-HA1, which requires two.

D12 feature codes

Table G-2 lists the feature codes on the D12 Frame that are relevant when installing the D12 Frame connected to the VTS. For a complete list of all feature codes available for the D12 Frame, refer to *IBM 3494 Tape Library Introduction and Planning Guide*, GA32-0279.

Table G-2 D12 Frame specifications and specified feature codes

Feature code	Min. / Max. qty	Field / Factory install	Description
4630	0 / 6	Yes / No	Field Install Drive
4663	0 / 6	Yes / No	Replace Model B1A Tape Drive with a Model E1A
5233	1 / 1	Yes / Yes	SCSI Extender
9010	1 / 1	Yes / Yes	Attached to a VTS
9011	1 / 1	Yes / Yes	Additional Drives Support

Feature code	Min. / Max. qty	Field / Factory install	Description
9630	0 / 6	No / Yes	Field Merge Drive (order FC4630 to get drive)
9631	0 / 6	No / Yes	Factory Install B1A
9663	0 / 6	No / Yes	Factory Install E1A

Field Install Drive (FC4630)

This feature provides the necessary hardware to field install an IBM 3590 Model B1A tape drive or IBM 3590 Model E1A tape drive in an already installed D12 Frame. It should be used when the plant has not installed drive-mounting hardware or feature code 9630, or when an odd number of drives is installed in a D12 Frame.

Feature code 4630 is required when the IBM 3590 Model B1A and E1A tape drives are not installed in the factory. One feature code 4630 must be specified for each IBM 3590 Model B1A tape drive or IBM 3590 Model E1A tape drive to be added.

Replace model B1A Tape Drive with a model E1A (FC4663)

This feature provides the necessary hardware to replace an IBM 3590 Model B1A tape drive with an IBM 3590 Model E1A tape drive in an already installed D12 Frame. For every ordered feature code 4663 one feature code 4630 or 9630 or 9631 should be removed from the already installed D12 Frame.

Feature code 4663 is required for every IBM 3590 Model B1A tape drive to be upgraded, so that the plant can send the right installation instructions and required microcode for the VTS and Library Manager.

SCSI Extender (FC5233)

This feature is required to connect the IBM 3590 Model B1A tape drives or IBM 3590 Model E1A tape drives installed in the D12 Frame to the VTS. It provides the hardware to extend the SCSI interfaces between the IBM 3590 Model B1A tape drives or IBM 3590 Model E1A tape drives and the VTS controller. It supports two extended SCSI attachments.

Feature code 5233 is required when a VTS is installed in an 3494. One feature code 5233 must be specified for field or factory installation. Feature code 5233 is not intended for extending SCSI interfaces to other SCSI hosts. Feature code 9010 is a co-requisite that assigns the D12 Frame as a VTS attached D12 Frame.

Attached to a VTS (FC9010)

This feature is required when the D12 Frame attaches to the VTS. This feature code allows the appropriate internal cables to be provided for the installation of the IBM 3590 Model B1A tape drives or IBM 3590 Model E1A tape drives. The individual SCSI cables do not have to be ordered separately for each tape drive.

One feature code 9010 must be specified for field or factory installation.

Additional Drives Support (FC9011)

One FC9011 is required on each D12 Frame that is attached to a Model B10 VTS or Model B20 VTS.

Field Merge Drive (FC9630)

This feature indicates to the manufacturing plant to provide the mounting hardware required to install an IBM 3590 Model B1A tape drive or IBM 3590 Model E1A tape drive in the field. Up to six feature code 9630s can be specified for field installation, depending on the number of tape drives.

Factory Install B1A (FC9631)

This feature indicates to the manufacturing plant to install one IBM 3590 Model B1A tape drive in the D12 Frame. Up to six feature codes 9631 can be specified for factory installation in one D12 Frame.

When the to be ordered D12 Frame is attached to a Model B18 VTS, a minimum of three feature code 9631s must be specified. If specified, this feature code must appear on both the D12 Frame order and on the IBM 3590 Model B1A tape drive.

Factory Install E1A (FC9663)

This feature indicates to the manufacturing plant to install one IBM 3590 Model E1A tape drive in the D12 Frame. Up to six feature codes 9663 can be specified for factory installation in one D12 Frame.

When the to be ordered D12 Frame is attached to a Model B18 VTS, a minimum of three feature code 9663s must be specified. If specified, this feature code must appear on both the D12 Frame order and on the IBM 3590 Model E1A tape drive.

Factory Install H1A (FC9670)

This feature indicates to the manufacturing plant to install one IBM 3590 Model H1A tape drive in the D12 Frame with fibre channel attachment feature (FC9510). Up to six feature codes 9670 can be specified for factory installation in one D12 Frame.

IBM 3590 Model B1A feature codes

The following feature codes can be specified for IBM 3590 Model B1A tape drives that are installed for use with the VTS.

ES/9000, S/3090, S/390 Attachment (FC9000)

One feature code 9000 must be specified per IBM 3590 Model B1A tape drive for field or factory installation.

3590 Model B1A Drive Installation (FC9631)

One feature code 9631 must be specified to indicate to the manufacturing plant that this IBM 3590 Model B1A tape drive is to be factory installed. Feature code 9631 is not required if you order new IBM 3590 Model B1A tape drives to be included in an existing D12 Frame. If specified, this feature code must appear on both the D12 Frame order and on the IBM 3590 Model B1A tape drive.

IBM 3590 Model E1A feature codes

The following feature codes can be specified for IBM 3590 Model E1A tape drives that are installed for use with the VTS.

ES/9000, S/3090, S/390 Attachment (FC9000)

One feature code 9000 must be specified per IBM 3590 Model E1A tape drive for field or factory installation.

Plant Install Drive in 3494 (FC9663)

One feature code 9663 must be specified to indicate to the manufacturing plant that this IBM 3590 Model E1A tape drive is to be factory installed. Feature code 9663 is not required if you order new IBM 3590 Model E1A tape drives to be included in an existing D12 Frame. If specified, this feature code must appear on both the D12 Frame order and on the IBM 3590 Model E1A tape drive.

IBM 3590 Model H1A feature codes

The following feature codes can be specified for IBM 3590 Model H1A tape drives that are installed for use with the VTS.

ES/9000, S/3090, S/390 Attachment (FC9000)

One feature code 9000 must be specified per IBM 3590 Model H1A tape drive for field or factory installation.

Plant Install Drive in 3494 (FC9670)

One feature code 9670 must be specified to indicate to the manufacturing plant that this IBM 3590 Model H1A tape drive is to be factory installed. Feature code 9670 is not required if you order new IBM 3590 Model H1A tape drives to be included in an existing D12 Frame. If specified, this feature code must appear on both the D12 Frame order and on the IBM 3590 Model H1A tape drive.

IBM 3494 Model L1x feature codes

Table G-3 lists the feature codes that are required or can be specified on the L1x Frame frame when an IBM TotalStorage VTS is installed if there are no tape drives in the L1x Frame. For a complete list of all feature codes available for the 3494 L1x Frame, refer to *IBM 3494 Tape Library Introduction and Planning Guide*, GA32-0279.

Table G-3 L1x Frame specifications and specified feature codes

Feature code	Min. / Max. qty	Field / Factory install	Description
2710	0 / 1	Yes / Yes	Remote Support Facility
2711	0 / 1	Yes / Yes	Remote Support Switch
2712	0 / 1	Yes / Yes	Remote Support Attachment
5210	0 / 1	Yes / Yes	Ten Cartridge Convenience I/O Station
5219	0 / 1	Yes / Yes	Token Ring Adapter
5220	0 / 1	Yes / Yes	Ethernet Adapter
5226	0 / 1	Yes / Yes	Remote Library Manager Console
5230	0 / 1	Yes / Yes	30 Cartridge Convenience I/O Station
5232	1 / 1	Yes / Yes	Attachment Concentrator
9002	1 / 1	Yes / Yes	Storage Unit Attach Frames 2 to 8

Feature code	Min. / Max. qty	Field / Factory install	Description
9003	1 / 1	Yes / Yes	Drive Unit Attach Frames 2 to 8
9004	1 / 1	Yes / Yes	Storage Unit Attach Frames 9 to 16
9005	1 / 1	Yes / Yes	Drive Unit Attach Frames 9 to 16
9011	1 / 1	Yes / Yes	Additional Drives Support
9020	1 / 2	Yes / Yes	Model Bxx VTS Attachment
9021	1 / 1	Yes / Yes	Virtual Drive Enhancement
9986	1 / 1	Yes / Yes	6 ft. Chicago Power Cord

Ten Cartridge Convenience I/O Station (FC5210)

This feature allows you to add and remove up to 10 cartridges from the 3494 without interrupting normal operations. If the convenience I/O station is installed, it reduces the cartridge storage capacity in the L1x Frame by 30 cartridges. One feature code 5210 can be ordered for field or factory installation.

Note: The Import/Export function requires a convenience I/O station to be installed. If you plan to use the Import/Export function, you must order either feature code FC5210 or FC5230.

Token-Ring Adapter (FC5219)

Feature code 5219 allows you to connect the Library Manager to a token ring network. Either this feature code or Ethernet Adapter (FC5220) can be ordered for field or factory installation, but not both. If you plan to implement SNMP monitoring, you need to order either feature code 5219 or 5220.

Ethernet Adapter (FC5220)

Feature code 5220 allows you to connect the Library Manager to an Ethernet network. Either this feature code or Token Ring Adapter (FC5219) can be ordered for field or factory installation, but not both. If you plan to implement SNMP monitoring, you need to order either feature code 5219 or 5220.

Remote Library Manager Console (FC5226)

Feature code 5226 allows you to see intervention required messages or perform operator actions without having to be at the site where the 3494 is installed. For additional information about the Remote Library Manager Console, refer to *IBM 3494 Tape Library Introduction and Planning Guide*, GA32-0279.

30 Cartridge Convenience I/O Station (FC5230)

Feature code 5230 allows you to add and remove up to 30 cartridges from the 3494 without interrupting normal operations. If the convenience I/O station is installed, it reduces the cartridge storage capacity in the L1x Frame frame by 80 cartridges. One feature code 5230 can be ordered for field or factory installation.

Note: The Import/Export function requires a convenience I/O station to be installed. If you plan to use the Import/Export function, you must order either feature code FC5210 or FC5230.

Attachment Concentrator (FC5232)

Feature code 5232 provides the hardware for the attachment interface between the Library Manager of the 3494 Tape Library and the VTS that uses the internal LAN connection interface. All cables required for the attachment are included in this feature code. One feature code 5232 must be specified for field or factory installation of the first Model B18 VTS VTS to be installed in a 3494.

Storage Unit Attach Frames number 2 to 8 (FC9002)

Feature code 9002 is used by the manufacturing plant to determine the total length of an IBM 3494 tape library and to provide the correct length accessor and control cables. One feature code 9002 must be ordered for each S10 Frame included in the first eight frames of the 3494.

Drive Unit Attach Frames number 2 to 8 (FC9003)

Feature code 9003 is used by the manufacturing plant to determine the total length of an IBM 3494 and to provide the correct length accessor and control cables. One feature code 9003 must be ordered for each D12 Frame included in the first eight frames of the 3494.

Storage Unit Attach Frames number 9 to 16 (FC9004)

Feature code 9004 is used by the manufacturing plant to determine the total length of an 3494 and to provide the correct length accessor and control cables. One feature code 9004 must be ordered for each S10 Frame included in the second eight frames of the 3494.

Drive Unit Attach Frames number 9 to 16 (FC9005)

Feature code 9005 is used by the manufacturing plant to determine the total length of an IBM 3494 tape library and to provide the correct length accessor and control cables. One feature code 9005 must be ordered for each D12 Frame included in the second eight frames of the 3494.

Additional Drives Support (FC9011)

One FC9011 is required if the L1x Frame has the FC5502 (Drive Unit for Model B18 VTS) attached to a Model B10 VTS or Model B20 VTS.

Model B10 or B20 VTS Attachment (FC9020)

FC9020 notifies the plant of control that a Model B10 or B20 VTS is attached to the IBM L1x Frame to assure the proper cables are supplied for connection to the L1x Frame.

Virtual Drives Enhancement (FC9021)

One FC9021 is required for each Model B10 VTS or Model B20 VTS this feature must also be on the HA1 Frames if that model is included in the library

Media feature codes

Table G-4 lists the feature codes that you can specify to order magnetic tapes along with the 3494 frames.

Table G-4 IBM 3494 magnetic media feature codes

Feature code	Description	Maximum quantity per frame
8002	One 3590 cleaning cartridge	10
8005	One 3490 cleaning cartridge	10
8410	210 3490E data cartridges	Withdrawn
8420	420 3490E data cartridges	Withdrawn
8510	210 3590 data cartridges (HPCT)	1
8520	420 3590 data cartridges	Withdrawn
8610	210 Extended Length data cartridges (EHPCT)	2
9540	No data cartridge	1 *

Note: Cartridges ordered with this feature code will not be labeled or initialized. Initialized and labeled TotalStorage 3590 data cartridges can be ordered through the 3599 machine types. See your IBM representative for further information about ordering TotalStorage 3590 data cartridges independent of a frame or library order.

IBM3584 and 3953 feature codes

All new feature codes for the atachment of the IBM 3584 Tape Library and IBM 3953 Tape Library Manager can be found in the IBM TotalStorage 3584Tape Library for zSeries Hosts: Planning, Implementation, and Monitoring, SG24-6789.

H

Using a VTS on a DFSMS host without a tape management system

Note: While it is possible to connect a VTS to a DFSMS system without a tape management system, it is **not** recommended, due to the number of specific changes required for each and every data set.

A few z/OS customers do not have a tape management system to control tape operations. Rather, the JCL contains the volume serial numbers where each output data set is to be written. In such environments, all tape mounts are specific.

The JCL scheme can be divided into two categories:

- The output data set is cataloged and it is opened for output while the disposition is old.
- The volume list is explicitly coded in the DD statement.

When the output data set grows beyond the capacity of the defined volumes a non-specific mount is requested by Open/Close/EOV and the data set extends on the new volume obtained from the scratch volume list and requested via a non-specific request. If the data set is recataloged, the catalog reflects the supplementary volume, while the JCL volume list will not be automatically extended.

When the data set shrinks, the unused volumes are not automatically removed, neither from the catalog, nor from the JCL DD statement.

The two previous concerns become more important when the tape media capacity evolves greatly, due to technological changes, such as a migration to 3590 or even 3592 native tape drives or a migration to a Virtual Tape Server.

When migrating to a Virtual Tape Server, there is an added performance concern as the output tapes' volumes are not put in the fast-ready category. The objective of this section is to indicate circumventions to avoid unnecessary staging of output tapes.

General concepts

There are a number of items which must be understood before we describe a viable solution. They apply to the entire architecture, software and hardware, of the support of the IBM libraries in an MVS and DFSMS environment:

- **Specific mount request:** The host specifies the volume to be mounted. The specific mount request, generally, is used to access an existing data set residing on the target tape.
- **Non-specific mount request:** The host does not specify the volume to be mounted. The selection of the volume to be mounted is delegated to another component of the architecture. In our case, the Library Manager will select a volume in the scratch category associated to this specific host.

  The interference of the tape management system is limited to accepting or rejecting the proposed volume.
- **Return to scratch:** It is the function of the tape management system to refill the stack of scratch volumes when the data set on volumes have expired. This function can be run daily or only when the stack of scratch volumes goes below a threshold.
- **Scratch and private volume status:** A scratch volume is a volume that does not contain a valid data set. On private volumes, there is at least one data set which has not expired. The status of the volume is retained in the tape management data base.
- **Scratch and private volume category:** The categories belong to the vocabulary of the Library Manager. Every connected host uses a set of categories associated to the various status and nature of the volumes owned and managed by the host.

  Under control of the tape management system, when a volume returns to the scratch status, the category of the volume is updated in the Library Manager data base.
- **Tape Configuration Data Base (TCDB):** The TCDB reflects the contents of all the volumes managed by the host and, logically, in one of the libraries of the system.

  The status of the volume (scratch or private) is recorded in the TCDB, as well as the library name where the volume resides, and the fact of the physical presence of the volume in the library. To preserve the confidentiality of the volume's contents, you cannot issue a specific mount for a volume which is in the scratch status in the TCDB.
- **Fast-ready category:** In a VTS, no volume data is staged in the TVC when the volume is in the fast ready category. Also, an attempt to read any data beyond the VOL1 and HDR1 to HDR8 will end with an I/O error. The fast ready category is associated to the scratch categories of all the connected hosts.

  A volume selected to receive output data set does not need to be staged in the tape volume cache. From the software point of view, this particular use of the volume is formulated in the mount request as a non-specific request commonly designated as scratch mount.
- **Return to Private category:** Once the data have been written to the volume(s), the volume(s) must be removed from the fast ready category, in order to be usable.

Mounting rules

From the list above, “General concepts” on page 516, we can draw the following rules for the VTS operations without tape management system:

- When a specific mount is used for the output data set, the requested volumes must be put in the fast-ready category prior to mounting, to avoid unnecessary staging.

  This can be achieved by executing the CBRSPLCS program which is in SYS1.SAMPLIB.

- When a specific mount is used for the output data set, the requested volumes must be put in the private status in the TCDB.

  This can be achieved by executing an ALTER IDCAMS command against the target volume:

  ```
  ALTER  Vserial VOLUMEENTRY USEATTRIBUT(PRIVATE) +
      STORAGEGROUP('sssssss')
  ```

- Once the data has been written to the volumes, they must be put into the private category.

  This can be achieved by executing CBRSPLCS program which is in SYS1.SAMPLIB. You may execute a preparation procedure, specific to each job and each output tape data set. This preparation procedure will put the volume in the fast ready category in the LM and in private status in the TCDB.

When JES3 is controlling the tape allocations, the disposition NEW or OLD must be considered to determine where, in the job flow, the preparation procedure must be invoked.

The rerun of the job should also be examined for proper positioning of the change category step.

Dynamic allocation of the virtual volume should be reviewed to include the function described above.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 521.

- *IBM TotalStorage Enterprise Tape: A Practical Guide*, SG24-4632
- *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409
- *DFSMShsm Primer*, SG24-5272
- *DFSMSrmm Primer,* SG24-5983
- *DFSMS/MVS V1R5 Technical Guide*, SG24-4892
- *Tivoli Storage Manager Version 3.7: Technical Guide*, SG24-5477
- *IBM Tape Solutions for Storage Area Networks and FICON,* SG24-5474
- *GPFS: A Parallel File System*, SG24-5165
- *GPFS on AIX Clusters; High Performance File System Administration Simplified*, SG24-6035
- *IBM TotalStorage Peer-to-Peer Virtual Tape Server Planning and Implementation Guide*, SG24-6115
- *FICON Native Implementation and Reference Guide*, SG24-6266
- *IBM TotalStorage 3584 Tape Library for zSeries Hosts: Planning and Implementation,* SG24-6789
- *Introduction to SAN Distance Solutions*, SG24-6408

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 TotalStorage 3584 Tape Library Introduction and Planning Guide*, GA32-0469
- *IBM TotalStorage 3584 Tape Library Operator Guide*, GA32-0468
- *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*, GA32-0279
- *IBM TotalStorage 3494 Tape Library Operator's Guide*, GA32-0280
- *IBM TotalStorage Tape System 3590 Introduction and Planning Guide*, GA32-0329
- *IBM TotalStorage Tape System 3590 Operator Guide*, GA32-0330
- *IBM TotalStorage Tape System 3590 Hardware Reference*, GA32-0331
- *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
- *Ditto/ESA V1R3 User's Guide and Reference*, SH19-8221. You can obtain the updated publication at:

  `http://www-4.ibm.com/software/ad/ditto/library.html#ug`
- *Basic Tape Library Support User's Guide and Reference*, SC26-7016
- *DFSMS/VM Function Level 221 Removable Media Services User's Guide and Reference*, SC35-0141
- *IBM TotalStorage Tape Device Drivers Installation and User's Guide*, GC35-0154
- *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:

- DITTO/ESA Home Page:

  `http://www-4.ibm.com/software/ad/ditto/`
- DITTO/ESA Publications:

  `http://www-4.ibm.com/software/ad/ditto/library.html#ug`
- IBM Tape Tools for customers:

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

- IBM Tape Tools for IBMers:

  `ftp://submit.boulder.ibm.com/download/tapetools/`
- IBM Tape Tools for Business Partners:

  `http://testcase-yellow.boulder.ibm.com`
- IBM Tape Storage Publications:

  `http://www.storage.ibm.com/hardsoft/tape/pubs/prodpubs.html`
- Tivoli Storage Manager Home Page:

  `http://www.tivoli.com/products/index/storage_mgr/`
- Tivoli Storage Manager Publications:

  `http://www.tivoli.com/support/public/Prodman/public_manuals/storage_mgr/admanual.html`
- Betasystems Home Page:

  `http://www.betasystems.com`
- Computer Associates Products:

  `http://www.cai.com/products/ca1.htm`
- OpenTech Systems Tape/Copy:

  `http://www.opentechsystems.com/Tape-Copy.htm`
- Innovation Data Processing Home Page:

  `http://www.innovationdp.com`
- Technologic Software Home Page:

  `http://www.technologic.com`

How to get IBM Redbooks

Search for additional Redbooks or Redpieces, view, download, or order hardcopy from the Redbooks Web site:

`ibm.com/redbooks`

Also download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

Redpieces are Redbooks in progress; not all Redbooks become Redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

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

D

S

W

IBM TotalStorage Virtual Tape Server: Planning, Implementing, and Monitoring

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

IBM TotalStorage Virtual Tape Server:

Planning, Implementing, and Monitoring

Improved performance and capacity

Enhanced attachment capabilities

Advanced system management

This IBM Redbook is the seventh update (eighth edition) to the best-selling IBM Magstar VTS Implementation Guide, which was first published in 1997. Since the first VTS installations in May 1997, much has changed in the Virtual Tape Server area. The performance, capacity, and functionality of the VTS has been phenomenally enhanced. The focus of this book is to provide implementation and usage guidance for the latest changes to the IBM TotalStorage Enterprise Automated Tape Library VTS models B10 and B20. Minor additions and corrections have also been applied throughout the book. If you discuss, sell, order, or plan for a Virtual Tape Server, please consult this redbook first!

The IBM TotalStorage Virtual Tape Server (VTS) is integrated with the 3494 Tape Library and with the IBM 3584 Tape Library using the IBM 3953 Library Manager. The host system perceives it as an IBM Tape Library with up to 256 virtual 3490E tape drives and up to 500,000 virtual 3490E cartridges, which can have an uncompressed capacity of 400, 800, 1000, 2000, or 4000 MB. The VTS fully exploits the capability of the IBM TotalStorage 3590 and 3592 tape drives and the capacity of their media.

The VTS introduced the virtual concept into tape subsystems, similar to virtual storage and virtual disk concepts that already have been implemented. The latest changes allow IBM 3592 Tape Drives to attach to the VTS installed with an IBM 3584 Tape Library, thus enhancing the VTS attachment options and performance once again.

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

SG24-2229-07 ISBN 0738492566