

NCS/VTCS

Introducing NCS/VTCS MVS Software

Version 6.2
E22962-01



April 2011, Revision 01

Submit comments about this document to STP_FEEDBACK_US@ORACLE.COM

Copyright © 2006, 2010, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Preface

Oracle's StorageTek Nearline Control Solution (NCS) is the software that powers StorageTek's Automated Cartridge System (ACS) hardware. NCS is composed software components that work together to make ACS the industry-leading storage solution that it is.

Oracle's StorageTek Virtual Tape Control System (VTCS) is MVS host software, which together with the portions of NCS that support VTCS and the Virtual Tape Storage Subsystem (VTSS), comprises Virtual Storage Manager (VSM).

Audience

This book presents information for anyone involved in purchasing and planning for the VSM solution. This audience could include MIS managers, system programmers, storage administrators, and performance specialists.

Organization of This Publication

Table 1. Organization of This Publication

Chapter 1, “What is NCS?”	Introducing NCS and its software components.
Chapter 3, “NCS Frequently Asked Questions”	NCS FAQs with cross-references to more detail.
Chapter 4, “Using the SMC Client/Server Feature”	Overview and benefits of the SMC Client/Server Feature.
Chapter 5, “NCS/VTCS Policies”	SMC 6.2 introduces a new method of specifying policies for tape allocation and mount requests. Read all about it here...
Chapter 6, “What is VSM?”	Introducing VSM and its hardware and software components.
Chapter 7, “Key VSM Benefits and Features”	Highlights the benefits of the VSM solution.
Chapter 8, “VSM Frequently Asked Questions”	VSM FAQs with cross-references to more detail.
Chapter 9, “How Does VSM Work?”	Overview of how the VSM components work together to create the solution.
Chapter 10, “VSM Performance”	How VSM improves your system’s performance.
Chapter 11, “VSM Advantages versus HSM/TMM”	VSM architecture bests the competition
Appendix A, “NCS/VTCS Terminology”	Definitions of VSM terms

What’s New in this Book?

Revisions D and 01 contain technical updates and corrections.

For Revision C, “[What Enhancements Does VTCS/NCS 6.2.0 Offer?](#)” on page 5 is now updated with information on new SPEs to NCS/VTCS 6.2.

Contents

Preface iii

Audience iii

Organization of This Publication iv

What's New in this Book? iv

1. What is NCS? 1

StorageTek Management Component (SMC) 2

Host Software Component (HSC) 2

HTTP Server 2

Near Continuous Operations (NCO) 3

MVS/CSC and LibraryStation 3

ExPR and ExLM 3

2. What Enhancements Does VTCS/NCS 6.2.0 Offer? 5

SMC 6.2 POLicy Command and TAPEREQ Statement Enhancements 6

VTV Archiving and Reconciliation 7

Display STORCLAS 8

Conditional VTV Replication 9

Large VTV Pages 9

2 GB/4 GB VTVs 9

Maximum of 65000 VTVs per MVC 10

Maximum of 32 RTDs per VTSS 10

Stacked Migrates 10

Synchronous Replication 10

DISPLAY Output Enhancements 11

INVENTORY Utility 11
MVCATTR Control Statement 11
MIGRSEL SCHLIMIT Parameter 12
CONFIG GLOBAL MAXVTVSZ Parameter 12
Tapeless VSM 12
Native IP Replication 12

3. NCS Frequently Asked Questions 15

What is a TapePlex and How is It Defined? 15

Answer: 15

What is the SMC Client/Server Feature and What are Its Benefits? 15

Answer: 15

What are SMC Policies? 16

Answer: 16

How Do I Implement SMC Policies? 16

Answer: 16

Is There An Easy Way to Convert SMC TAPEREQ Statements to the SMC Policy Method?
16

Answer: 16

Is There A Way to Test SMC POLICY and TAPEREQ Lookup? 17

Answer: 17

What About VTCS Management Classes? 17

Answer: 17

What is the UI? 17

Answer: 17

What type of fault recovery does NCS have? 18

Answer: 18

4. Using the SMC Client/Server Feature 19

5. NCS/VTCS Policies 21

Defining SMC Policies 22

Example: SMC Policies for Real and Virtual Media 22

Example: The StorageTek DFSMS Interface and SMC Policies 23

Example: TAPEREQ Statements and SMC Policies 24

VTCS Management and Storage Classes 25

6. What is VSM? 27

VSM Hardware 28

VTCS Software 28

The VSM Offsite Vault Disaster Recovery Feature 28

The Concurrent DR Test Feature 29

7. Key VSM Benefits and Features 31

8. VSM Frequently Asked Questions 33

What benefits does VSM provide? 33

Answer: 33

Does VSM outperform TMM? 33

Answer: 33

What data sets are good candidates for moving to VSM? 33

Answer: 33

What data sets are inappropriate for VSM? 34

Answer: 34

Is there a tool to help identify candidate data sets for VSM? 34

Answer: 34

How do I direct candidate data sets to VSM? 35

Answer: 35

Can VSM and the non-VSM portion of HSC share library resources? 35

Answer: 35

What virtual transports and media does VSM support? 35

Answer: 35

How many VTVs are supported in VSM? 36

Answer: 36

Can VTDs be shared among multiple MVS hosts? 36

Answer: 36

Can RTDs be shared? 36

Answer: 36

What is the VSM Offsite Vault Disaster Recovery Feature? 37

Answer: 37

What is the Concurrent DR Test Utility? 37

Answer: 37

What happens when high AMT is exceeded? 37

Answer: 37

What type of fault recovery does VSM have? 38

Answer: 38

What are Clustered VTSS Configurations and how do they work? 38

Answer: 38

9. How Does VSM Work? 39

How VSM Automatically Manages VTSS Space and Migrates VTVs 41

How VSM Automatically Recalls Migrated VTVs 43

Reclaiming Space on MVCs 44

How VSM Maintains Data Integrity and Availability When Moving VTVs from One MVC to Another 44

How VTCS Manages Scratched VTVs 45

The Import and Export Functions 45

VTV Archiving and Reconciliation 45

Clustered VTSS Configurations 47

10. VSM Performance 49

VSM reduces mount time and reduces number of tape mounts: 49

VSM reduces problem of allocation recovery: 49

11. VSM Advantages versus HSM/TMM 51

How HSM Works 51

Disadvantage of HSM 51

How TMM Works 52

Disadvantages of TMM 52

How VSM Works 52

How VSM Overcomes the Disadvantages of HSM/TMM 52

A. NCS/VTCS Terminology 53

What is NCS?

Nearline Control Solution (NCS) is the software that powers StorageTek's Automated Cartridge System (ACS) hardware. NCS is composed software components that work together to make ACS the industry-leading storage solution that it is. The following sections describe each of these components.

StorageTek Management Component (SMC)

SMC is the interface between IBM's OS/390 and z/OS operating systems and HSC and MVS/CSC. SMC is a required NCS component, and must reside on every MVS host that accesses Nearline real tape and/or VSM virtual tape. SMC runs on both JES2 and JES3 systems and does the following:

- ? Influences tape device allocation (real and virtual).
- ? Intercepts tape management, and operating system mount, dismount, and swap messages and creates a request for this functions and routes the request to either HSC or MVS/CSC.
- ? Coordinates requests among multiple StorageTek *TapePlexes*. These requests can consist of mounts, dismounts, and swaps, and queries (configuration, volume lookup).

A TapePlex (formerly known as a "library"), is a single StorageTek hardware configuration, usually represented by a single HSC Control Data Set (CDS).

A TapePlex can contain multiple ACSs and Virtual Tape Storage Subsystems (VTSSs).

SMC can communicate with any number of TapePlexes, using cross address space facilities to communicate with HSC or MVS/CSC running on the same host, and TCP/IP to communicate with HSC systems executing on other hosts.

Host Software Component (HSC)

HSC does the following:

- ? Maintains the CDS, which contains information about the physical and virtual drives and corresponding media in a TapePlex.
- ? Receives requests from SMC for mounts, dismounts, and swaps, and queries and sends these requests to the LMU, which carries out these operation.
- ? Manages error conditions, such as lost cartridges.

HTTP Server

The StorageTek HTTP Server provides the middle ware to allow communication between the SMC (client) and a remote HSC subsystem (server). The HTTP server executes as a separate subsystem on the MVS host where the remote HSC subsystem resides. The HTTP Server is **not** required on a host where only the SMC is executing.

Near Continuous Operations (NCO)

NCS/VTCS provides Near Continuous Operations (NCO), which means that you can dynamically add, change, and delete hardware components (such as Nearline transports, RTDs, LSMs, and VTSSs) without having to bring NCS/VTCS down and back up again.

For more information about NCO see *VTCS Command and Utility Reference and HSC System Programmer's Guide*.

MVS/CSC and LibraryStation

In NCS 6.2 you can use MVS/CSC to communicate with LibraryStation. However, the next release of NCS will not permit an MVS/CSC client to connect to a LibraryStation server. You can also use LibraryStation to connect non-MVS clients to an HSC server.

ExPR and ExLM

To help manage your Nearline and VSM systems, you can also separately license the following StorageTek software products:

- ? ExPR, which provides PC and mainframe performance reports and capacity planning tools for both Nearline and VSM systems.
- ? ExLM, which allows you to efficiently manage Nearline ACS contents, VSM resources (MVCs and VTVs), and VTV consolidation, migration, and recall.

What Enhancements Does VTCS/NCS 6.2.0 Offer?

NCS/VTCS 6.2.0 builds on the previous releases to extend the VSM feature/function set. For more information, see:

- ? [“SMC 6.2 POLICY Command and TAPEREQ Statement Enhancements” on page 6](#)
- ? [“VTV Archiving and Reconciliation” on page 7](#)
- ? [“Display STORCLAS” on page 8](#)
- ? [“Conditional VTV Replication” on page 9](#)
- ? [“Large VTV Pages” on page 9](#)
- ? [“2 GB/4 GB VTVs” on page 9](#)
- ? [“Maximum of 65000 VTVs per MVC” on page 10](#)

In addition, NCS/VTCS 6.2 is further enhanced via the following SPEs:

- ? [“Maximum of 32 RTDs per VTSS” on page 10](#)
- ? [“Stacked Migrates” on page 10](#)
- ? [“Synchronous Replication” on page 10](#)
- ? [“DISPLAY Output Enhancements” on page 11](#)
- ? [“INVENTORY Utility” on page 11](#)
- ? [“MVCATTR Control Statement” on page 11](#)
- ? [“MIGRSEL SCHLIMIT Parameter” on page 12](#)
- ? [“Native IP Replication” on page 12](#)

SMC 6.2 POLicy Command and TAPEREQ Statement Enhancements

The SMC POLicy command specifies a tape policy in a single named policy object. You can specify a policy either on a TAPEREQ statement or via the SMC DFSMS ACS routines. This method of creating and specifying tape policies ensures that you get the policy you create with the POLicy command and **only** that policy.

For example, for NCS 6.1 and below, to route volumes created by jobs with jobname PRODSL2 to a VTV mounted on a VTD that VSM selects and specify Management Class MGMTCLS1, you would use the following TAPEREQ:

```
TAPEREQ JOB (PRODSL1) RECT (VIRTUAL) MGMT (MGMTCLS1)
```

With SMC 6.2, to do the same job routing with the same Management Class, you could create the following SMC/HSC statements:

```
POLICY NAME(VSMPOL1) RECT (VIRTUAL) MGMT (MGMTCLS1)  
TAPEREQ JOB (PRODSL1) POLICY (VSMPOL1)
```

To “convert” to the 6.2 policy methodology, you had to create an additional POLicy statement and rewrite my TAPEREQ statement, so what’s the value add? To summarize, the SMC POLicy command specifies a tape policy in a **single** named policy object. If a TAPEREQ statement specifies a policy via the POLicy parameter, SMC then looks up the specified policy. SMC then applies the parameters on that policy, and **only** those parameters, to the allocation and mount process. **No further TAPEREQ lookups are performed.**

VTV Archiving and Reconciliation

You can use the ARCHAge and ARCHPol parameters of the MGMTclas statement to set an *Archive Policy* for VTVs in a Management Class. When the VTV's age exceeds the ARCHAge value, the VTV is eligible for archive per the Storage Class(es) specified on the ARCHPol parameter. The actual archive occurs in one of two ways:

- ? Automatically the next time the VTV is recalled and remigrated.
- ? On demand via the ARCHIVE utility.

Information Lifecycle Management (ILM), a StorageTek storage management strategy, has as its central concept the idea that data should be stored on media that matches its importance to the business enterprise and its pattern of reuse. Active, important data goes on fast-access media, and has multiple copies, while inactive, less important data is archived to high-capacity, cheaper media. Automating this process is the most cost-effective way to manage data storage.

Archiving implements ILM by letting you archive inactive data. Using VTCS archiving, you can move VTVs to different media (for example, from fast-access 9840 media to high-capacity 9940 media) and a different location (for example, from a local ACS to a remote ACS for ejection/vaulting).

You can use the RECONcil utility to reconcile VTV media and location (move VTVs from one Storage Class to another). You typically reconcile VTVs when:

- ? The VTVs are on the wrong media, in the wrong ACS, or both.
- ? An ACS is unavailable for a considerable period of time, then is brought back online. In this case, you would first change the MIGpol parameter on the MGMTclas statement for the affected VTVs to point to a different ACS (and media, if desired). When the original ACS comes back online, you then change the MIGpol parameter on the MGMTclas statement to point to the original ACS, and run RECONcil specifying the updated MGMTclas (or STORclas) statement(s) to move the VTVs to the original ACS.

For more information on VTV archiving and reconciliation, see *VTCS Command and Utility Reference*.

Display STORCLAS

VTCS 6.2 now provides the STORCLAS option on the Display/Query commands and utilities. Display STORCLAS shows:

- ? The characteristics of the Storage Class (ACS, MVC Pool, and Media).
- ? VTVs waiting migration to the Storage Class from any VTSS.
- ? Requirements of the MVCs to be used for migration.
- ? The device type(s) of the RTDs needed to write to the migration MVCs.
- ? Any errors with regard to satisfying the migration requirements.

Also provided with Display STORCLAS are the following:

- ? Greater granularity of message detail for migration failures (1 message replaced by seven messages).
- ? Enhanced validation of MVC Pools and Storage Classes to check for common setup errors.

Taken together, these enhancements give you improved ability to prevent, diagnose, and correct migration failures.

Conditional VTV Replication

VTCS/NCS enhances VTV replication with conditional replication. The new CONFIG GLOBAL REPLICat parameter provides the following value choices:

ALWAYS

The replicate request is added to the VTCS replication queue every time the VTV is dismantled, regardless of whether the VTV was changed while it was mounted (the default).

CHANGED

The replicate request is added to the VTCS replication queue if the VTV:

- ? Was changed while it was mounted **or**
- ? Was only read while mounted but less than the expected number of MVC copies of the VTV exist.

Note – Regardless of the CONFIG GLOBAL REPLICat setting, replication **also** requires that:

- ? The VTV must be dismantled in a VTSS that supports replication **and** there cannot be an identical copy of the VTV in the other VTSS in the Cluster.
- ? In addition to the CONFIG GLOBAL REPLICat value, you **must** specify REPLICAT(YES) on a VTV's Management Class for replication to occur.

Large VTV Pages

NCS/VTCS 6.2 now provides the **VTV Page Size Policy**. You can use the VTVPAGE parameter of the CONFIG GLOBAL statement or the MGMTclas statement to specify the default page size (STANDARD or LARGE) used to store VTV data in the VTSS and on the VTVs. Large page size (which requires that the CDS is at a G level or above) can provide improved performance within the VTSS and for migrates and recalls.

2 GB/4 GB VTVs

For NCS/VTCS 6.2, you can use the MAXVtvsz parameter of the MGMTclas statement to specify the maximum VTV size of a VTV as 2 GB or 4 GB in addition to the existing values of 400 MB and 800 MB. The new larger VTV sizes (which require that the CDS is at a G level or above), coupled with the Large VTV page option, can provide improved performance within the VTSS and for migrates and recalls.

Maximum of 65000 VTVs per MVC

For NCS/VTCS 6.2, you can use the CONFIG GLOBAL MAXVTV parameter to specify a maximum of 65000 VTVs per MVC for a G level CDS. This new maximum value can help you maximize your investment in high-capacity media.

Maximum of 32 RTDs per VTSS

With this SPE, VTCS 6.2 can support a maximum of 32 RTDs per VTSS, which you specify on the CONFIG GLOBAL MAXRTDS parameter. For prerequisites and additional information, see *Installing and Configuring VTCS* and *VTCS Command and Utility Reference*.

Stacked Migrates

With this SPE, you can enable stacked migrates. Enabling the stacked migrates feature can improve migration performance by allowing multiple migrations concurrently to an RTD. For prerequisites and additional information, see *Installing and Configuring VTCS* and *VTCS Command and Utility Reference*.

Synchronous Replication

With this SPE, you can specify whether replications occur synchronously or asynchronously via the CONFIG GLOBAL SYNCHREP parameter. For prerequisites and additional information, see *Installing and Configuring VTCS* and *VTCS Command and Utility Reference*.

DISPLAY Output Enhancements

With these SPEs, DISPLAY output is enhanced as follows:

- ? DISPLAY RTD output for shared RTDs is enhanced as follows:
 - ? The current single line display appears if the global status of the shared RTD applies to all of the connected VTSSs.
 - ? If the global status of the shared RTD does not apply to all of the connected VTSSs, an individual line is displayed showing the status of each VTSS.
- ? The DISPLAY MVCP command output is enhanced for Named MVC Pools to:
 - ? Display the values of the optional parameters that are currently in effect for the Named MVC Pools.
 - ? Produce XML tags showing the values of the optional parameters that are currently in effect for the Named MVC Pools.

For prerequisites and additional information, see *VTCS Command and Utility Reference*.

INVENTORY Utility

The INVENTORY utility lists all VTVs on specified MVCs. Additionally, you can use INVENTORY to:

- ? Request that the contents of the specified MVCs are cross-checked with the information recorded in the CDS.
- ? Terminate the processing based on the end-of-tape position recorded in the CDS for the MVC, rather than at the MVC's physical end-of-tape.
- ? Have the processing terminated for an MVC the first time a VTV is discovered that is incompatible with the information in the CDS.

For prerequisites and additional information, see *VTCS Command and Utility Reference*.

MVCATTR Control Statement

When an error occurs while reading an MVC on an RTD, VTCS may swap the MVC to another RTD to retry the operation. If the swap is for an MVC audit then the selection of the “swap to” device type **excludes** the device type that the error occurred on. Otherwise, the “swap to” device type is restricted to the native device type (device type MVC was created on). However, **note that**:

- ? T9840D devices are the **preferred** device type for reading an MVC created on **any** T9840 device type (T9840A/T9840B/T9840C/T9840D).
- ? T10KB devices are the **preferred** device type for reading an MVC created on **any** T10K device type (T10KA/T10KB).

This enhancement provides the MVCATTR control statement, which lets you specify the “swap to” RTD device type for each MVC media name.

For prerequisites and additional information, see *VTCS Command and Utility Reference*.

MIGRSEL SCHLIMIT Parameter

MIGRSEL SCHLIMIT provides finer control over how Storage Classes are treated for migration, but in the opposite direction. That is, MIGRSEL SCHLIMIT, applies limits on migration resources (RTDs/MVCs) per Storage Class, and you can do so on a VTSS and/or host basis.

For prerequisites and additional information, see *VTCS Command and Utility Reference* and *Installing ELS*.

CONFIG GLOBAL MAXVTVSZ Parameter

You can now specify the maximum VTV size on CONFIG GLOBAL as well as on the MGMTCLAS statement.

For prerequisites and additional information, see *VTCS Command and Utility Reference* and *Installing ELS*.

Tapeless VSM

“Tapeless VSM” basically means that you can have a VTSS without any RTDs directly attached to the VTSS; in the CONFIG deck, there are no RTD statements for the tapeless VTSS. Tapeless VSM applies to VSM4s and VSM5s.

For prerequisites and additional information, see *VTCS Command and Utility Reference* and *Installing ELS*.

Native IP Replication

The VTSS native IP connection feature lets you use TCP/IP protocol to connect two VTSSs for VTV replication. Each VTSS has IFF3 cards with Ethernet ports for connection to the TCP/IP network. Previously, you were limited to ESCON or FICON connections for replication. Using TCP/IP for CLINKs can provide improved replication performance over ESCON or FICON protocols and, if so desired, allows the existing ESCON or FICON ports to be used exclusively for RTD and host connections.

For prerequisites and additional information, see *VTCS Command and Utility Reference* and *Installing ELS*.

NCS Frequently Asked Questions

What is a TapePlex and How is It Defined?

Answer:

A *TapePlex* (formerly known as a “library”, is an actual hardware configuration: the VTSSs, ACSs, real and virtual drives and volumes defined by a single CDS. You define a TapePlex with the SMC TAPEPLEX command. For more information, see [“Using the SMC Client/Server Feature” on page 19](#).

What is the SMC Client/Server Feature and What are Its Benefits?

Answer:

The SMC client/server feature that lets you run SMC only on the client hosts and HSC/VTCS and the HTTP server on one or more server hosts. An SMC client/server configuration:

- ? Reduces the number of hosts on which you run HSC/VTCS,
- ? Eliminates the need for MVS/CSC communicating with LibraryStation, and
- ? Provides failover capabilities when an HSC is recycled for maintenance.

For more information, see [“Using the SMC Client/Server Feature” on page 19](#).

What are SMC Policies?

Answer:

SMC Policies, which are defined via the SMC POLicy command, allow you to specify all policy attributes (including MEDIA, RECTECH or MODEL, ESOTERIC, SUBPOOL, VTCS MGMTCLAS, TAPEPLEX, and IDAX parameters) in a single named object. For more information, see [“NCS/VTCS Policies” on page 21](#)

How Do I Implement SMC Policies?

Answer:

You can specify an SMC Policy on either a TAPEREQ statement or via the SMC DFSMS ACS routines. For more information, see:

- ? [“Defining SMC Policies” on page 22](#)
- ? [“Example: TAPEREQ Statements and SMC Policies” on page 24](#)
- ? [“Example: The StorageTek DFSMS Interface and SMC Policies” on page 23](#)
- ? [“VTCS Management and Storage Classes” on page 25](#)

Is There An Easy Way to Convert SMC TAPEREQ Statements to the SMC Policy Method?

Answer:

Yes! The SMC SAMPLIB member POLCVT02 is a REXX script that helps you to convert existing TAPEREQ statements to use the new SMC POLicy method. The script reads existing TAPEREQ statements and produces the following output:

- ? TAPEREQ statements with existing policy parameters (MEDIA, RECTECH, etc.) replaced by a new POLicy name parameter.
- ? POLicy definitions derived from the original TAPEREQ definitions.

For more information, see *SMC Configuration and Administration Guide*.

Is There A Way to Test SMC POLicy and TAPEREQ Lookup?

Answer:

Yes there is, using the SMCUTRQ utility. After you input POLicy and TAPEREQ lookup criteria, SMCUTRQ performs the lookup as if it were the job or job step with the specified criteria. SMCUTRQ generates a report indicating what TAPEREQ statement by number, within the TREQDEF file, was matched for the specified criteria. If the TREQDEF file is accessible, the actual statement corresponding to the specified sequence TREQDEF TAPEREQ statement is also listed, if matched. For more information, see *SMC Configuration and Administration Guide*.

What About VTCS Management Classes?

Answer:

Because of the advantages of the SMC Policy Method, as described in [“What are SMC Policies?” on page 16](#), StorageTek **strongly recommends** that you convert DFSMS ACS routines that currently return VSM Management Classes to use SMC Policy Names. However, you can still use SMC DFSMS ACS routines to point to VTCS Management Classes as described in [“VTCS Management and Storage Classes” on page 25](#).

What is the UUI?

Answer:

The Universal User Interface (UUI) lets you:

- ? Issue UUI-supported commands from an SMC console via the SMC ROUTE command.
- ? Issue UUI-supported commands from the SLUADMIN or SMCUUI utilities.
- ? Get output for UUI-supported commands and utilities in XML or Comma Separated Values (CSV) format.

What type of fault recovery does NCS have?

Answer:

NCS provides the following facilities for fault recovery:

- ? Multiple copies of the CDS.
- ? Audit of tape libraries.
- ? Option of SMC Client/Server with multiple servers defined.

Using the SMC Client/Server Feature

SMC provides a client/server feature that lets you run SMC only on the client hosts and HSC/VTCS and the HTTP server on one or more server hosts. Using the SMC client/server feature provides the following benefits:

- ? **Reduces the number of hosts on which you run HSC/VTCS.** StorageTek recommends that you execute HSC/VTCS on only two hosts (primary and backup). Running HSC/VTCS on fewer hosts reduces CDS contention and eliminates the need to manage multiple MVS syslog files.
- ? **Communicates with multiple HSC/VTCS TapePlex systems** representing physically different hardware configurations.
- ? **Provides failover capabilities** when an HSC is recycled for maintenance.

For example, in [FIGURE 4-1 on page 20](#), client MVSA is running SMC only, but is connected to two servers running HSC/VTCS and HTTP server.

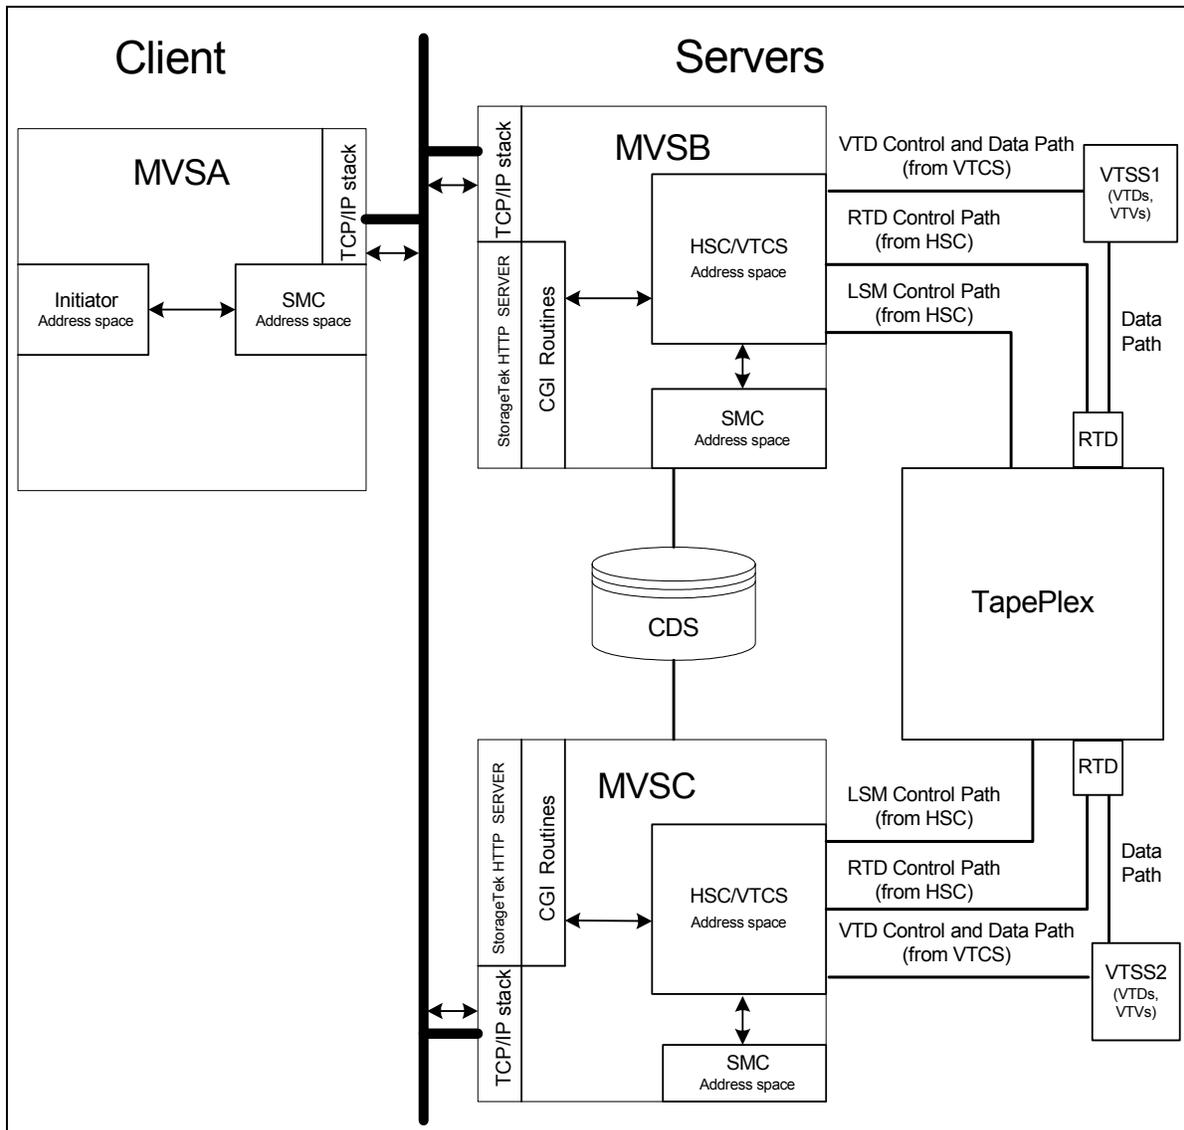


FIGURE 4-1 SMC Configuration: One Client, Two Servers

To make the client/server connections shown in FIGURE 4-1, MVSA a combination of the SMC TAPEPLEX and SERVER commands, for example:

```
TAPEPLEX NAME (SHRLIB)
SERVER NAME (REMPATH1) TAPEPLEX (SHRLIB) HOST (MVSB)
SERVER NAME (REMPATH2) TAPEPLEX (SHRLIB) HOST (MVSC)
```

The TAPEPLEX command lets the client assign a name to the *TapePlex*, or actual hardware configuration: the VTSSs, ACSs, real and virtual drives and volumes defined by the shared remote CDS. The SERVER commands, which point to the *TapePlex* name, complete the connection by specifying the paths to the HTTP server on the server hosts running the HSC/VTCS. For more information on implementing SMC client/server, see the *SMC Configuration and Administration Guide*.

NCS/VTCS Policies

SMC 6.2 introduces a new method of specifying policies for tape allocation and mount requests. The SMC POLicy command allows you to specify all policy attributes (including MEDIA, RECTECH or MODEL, ESOTERIC, SUBPOOL, VTCS MGMTCLAS, TAPEPLEX, and IDAX) in a single named object, or SMC **policy**. You can specify an SMC policy on either a TAPEREQ statement or via the SMC DFSMS ACS routines. For more information, see *SMC Configuration and Administration Guide*.

Defining SMC Policies

FIGURE 4-1 on page 20 shows a configuration where client MVSA is running SMC only, but is connected to two servers (MVSB and MVSC) running HSC/VTCS and HTTP server. We defined TapePlex SHRLIB, which consists of MVSB and MVSC.

Example: SMC Policies for Real and Virtual Media

Let's say that MVSA needs four policies for scratch mounts of:

- ? T9840C media from one subpool,
- ? T9940B media from a second subpool,
- ? Production VTVs from a third subpool and all VTVs allocated to VTSS1, and
- ? Test VTVs from a fourth subpool and all VTVs allocated to VTSS2.

The POLICY commands might look like the following:

```
POLICY NAME (P19840C) SUBPOOL (SP1) MEDIA (STK1R) RECTECH (STK1RC)
POLICY NAME (P9840B) SUBPOOL (SP2) MEDIA (STK2P) RECTECH (STK2PB)
POLICY NAME (VTVPOL1) SUBPOOL (SP3) MEDIA (VIRTUAL) IDAXESOTERIC (VTSS1) -
MGMTCLAS (PAYR)
POLICY NAME (VTVPOL2) SUBPOOL (SP4) MEDIA (VIRTUAL) IDAXESOTERIC (VTSS2) -
MGMTCLAS ( )
```

Example: The StorageTek DFSMS Interface and SMC Policies

In the StorageTek DFSMS interface, a Management Class name can now reference an SMC Policy name. StorageTek recommends converting your StorageTek DFSMS ACS routines to return a Management Class to reference an SMC Policy because this method does not require defining esoterics as Storage Groups and allows you to specify additional policy options. You must also specify the SMSDef MGMTPol(ALL) parameter. See *SMC Configuration and Administration Guide* for more information.

For example, let's say that we have the same configuration and policies in effect as described in "Example: TAPEREQ Statements and SMC Policies" on page 24, but we want to direct the data sets using DFSMS ACS routines...as shown in [FIGURE 5-1](#).

```
PROC STORCLAS

  IF &ACSENVIR = 'STKTAP1' THEN
    SET &STORCLAS = 'STKSTORC'
  END
END

=====

PROC MGMTCLAS

FILTLIST PAYR INCLUDE (PAYR*)
FILTLIST BACK INCLUDE (BACK*)
FILTLIST PROD INCLUDE (PROD*)
FILTLIST TEST INCLUDE (TEST*)

  IF &ACSENVIR = 'STKTAP1' THEN
    SELECT
      WHEN (&DSN = &PAYR)
        SET &MGMTCLAS = 'P19840C'
      WHEN (&DSN = &BACK)
        SET &MGMTCLAS = 'P9840B'
      WHEN (&DSN = &PROD)
        SET &MGMTCLAS = 'VTVPOL1'
      WHEN (&DSN = &TEST)
        SET &MGMTCLAS = 'VTVPOL2'
    END
  END
END
```

FIGURE 5-1 Example: The StorageTek DFSMS Interface and SMC Policies

Note – In [FIGURE 5-1](#), the STORCLAS routine assigns a storage class to every call. This ensures that the MGMTCLAS routine is also driven when the ACS routine is processed. If SMSDef MGMTPol is set to ALL, the Management Class name is used to reference an SMC policy at IDAX, allocation, and mount. Otherwise, the Management Class name is interpreted as a VTCS MGMTCLAS.

Example: TAPEREQ Statements and SMC Policies

In “[Example: SMC Policies for Real and Virtual Media](#)” on page 22, we defined separate scratch mount policies for T9840C, T9940B, and virtual media. Next, we want to direct data sets according to their HLQs as follows:

- ? PAYR* to T9840C media,
- ? BACK* to T9940B media, and
- ? PROD* to virtual media (VTVs) and VTSS1.
- ? TEST* to virtual media (VTVs) and VTSS2.

The corresponding TAPEREQ statements are as follows:

```
TAPEREQ DSN(PAYR*) POLICY(P19840C)
TAPEREQ DSN(BACK*) POLICY(P9840B)
TAPEREQ DSN(PROD*) POLICY(VTVPOL1)
TAPEREQ DSN(TEST*) POLICY(VTVPOL2)
```

In addition, you **must** specify the POLICY(ON) parameter on the IDAX command to activate IDAX esoteric substitution.

VTCS Management and Storage Classes

VTCS Management and Storage Classes, which are the bread and butter of a lot of leading edge implementations, do the following:

- ? VTCS Management Classes specify *how* VTCS manages VTVs. The HSC MGMTclas control statement defines a Management Class and its attributes. For example, the DELSCR parameter of the MGMTclas statement specifies whether VTCS deletes scratched VTVs from the VTSS. Management Classes can also point to...
- ? ...VTCS Storage Classes, which specify *where* migrated VTVs reside. The HSC STORclas control statement defines a Storage Class and its attributes.

For example, I already have up and running a VTCS Management Class and accompanying Storage Class for PAYR* data sets that looks like this:

```
MGMT NAME (PAYR) MIGPOL (S9840C)
STORCLAS NAME (S9840C) MVCPOOL (SP1) MEDIA (STK1RC)
```

This combination of Management and Storage Class says “For Management Class PAYR, migrate to 9840C media from MVC Pool SP1 (which equates to HSC subpool SP1).”

Next, I write a POLicy command to point to the Management Class:

```
POLICY NAME (P29840C) MGMTCLAS (PAYR)
```

Finally, I write a TAPEREQ statement for PAYR* data sets to assign Policy P29840C:

```
TAPEREQ DSN (PAYR*) POLICY (P29840C)
```

...which, in turn, assigns Management Class PAYR. For more information on VTCS Management and Storage Classes, see *Beyond the Basics: VTCS Leading Edge Techniques*.

What is VSM?

Virtual Storage Manager (VSM) is StorageTek's virtual storage solution to the problem of inefficient use of tape media and transports. VSM solves this problem by storing virtual tape volumes (VTVs) on a disk buffer on the Virtual Tape Storage Subsystem (VTSS). VSM then migrates (and stacks) the VTVs to real Nearline tape volumes called multi-volume cartridges (MVCs) that are mounted on Nearline Real Tape Drives (RTDs). When needed by the host, if the migrated tape volumes are not VTSS-resident, they are automatically recalled to the VTSS. The VTSS and VTVs allow VSM to optimize access time, throughput, and physical media and transport use. VSM consists of Virtual Tape Control System (VTCS), which is the MVS host software, the portions of NCS that support VTCS, and the VTSS.

HSC provides mount and dismount services for MVCs, and NCS, working with VTCS, provides the ability for greater than 16 MVS hosts running SMC and HTTP server. [FIGURE 4-1 on page 20](#) shows an NCS/VTCS Client/Server configuration. For more information, see the *SMC Configuration and Administration Guide*.

VSM Hardware

The VTSS hardware for VSM is RAID-6+ disk that provides the following:

- ? Tape transport and media emulation via the VTSS controller microcode.
- ? The disk buffer.
- ? ESCON or FICON ports on the VTSS that are used to make host-to-VTSS connections, VTSS-to-RTD connections, or VTSS-to-VTSS connections (for Clustered VTSS configurations).
- ? Fault tolerance of the RAID-6+ design. The redundancy built into RAID-6+ ensures the reconstruction of data from a failed drive onto a “hot spare” by using redundant data from surviving drives in an array.
- ? Nondisruptive serviceability. VTSS supports nondisruptive hardware servicing and installation of major components, most of which are hot-pluggable, field-replaceable units (FRUs).

VTCS Software

VTCS, which works as an extension to HSC/SMC, does the following:

- ? Influences the allocation of virtual tape drives (VTDs)
- ? Manages the use of VTVs
- ? Manages the migration and recall of virtual volumes
- ? Manages use of real tape media and transports used by VSM

The VSM Offsite Vault Disaster Recovery Feature

The VSM Offsite Vault Disaster Recovery (DR) Feature is an optional, orderable feature of VTCS. The VSM Offsite Vault DR Feature consists of the VSM Vault Utilities and the *VSM Offsite Vault Disaster Recovery Guide*. This guide tells how to implement offsite vaulting using the “Pickup Truck Access Method” (PTAM). PTAM is a common DR strategy where data sets requiring safeguarding are migrated to MVCs that are vaulted offsite. VSM adds value to PTAM because of VSM’s unique ability to effectively stack VTVs on MVCs. In addition, the VSM Vault Utilities, plus the power of the IMPORT and EXPORT facilities and the accompanying Manifest File, help you to easily manage your DR vaulting. The Offsite Vault Feature works with the following TMSs:

- ? CA-1
- ? CA-TLMS
- ? Control-M/Tape
- ? DFSMSrmm
- ? Zara (Zara support requires PTF L1H12IE)

The Concurrent DR Test Feature

Customers who use a Disaster Recovery (DR) configuration may want to validate their ability to continue normal production processing before an actual disaster occurs. In fact, certain businesses are required by their auditors to prove the readiness of their business continuance model. Using the Concurrent Disaster Recovery Test¹ (CDRT) software, you can do this validation without purchasing additional Nearline and/or VSM hardware. The DR Test software lets you test your DR environment while concurrently running production work. The software supports a complete parallel test of production systems and applications with simultaneous access to production data by both the production and DR test systems. The DRTEST utility creates a DR test version of the production CDS, which provides the DR test system its own view of the Nearline and/or VSM environment. Additionally, the DRTEST utility and changes to HSC and VTCS programmatically enforce certain functional restrictions during the CDS preparation and actual DR test to attempt to ensure system integrity.

CDRT provides a method of segregating a portion of the Nearline/VSM hardware to be used during a test run for displaced production processing at the recovery site. This DR test environment exercises the segregated hardware while normal production processing continues with the non-segregated hardware. The DR test hardware is a minimum of one ACS. Optionally, one or more VTSSs may be employed as DR test hardware. The ACS is shared between the production systems and the DR test systems, and the DR test systems have exclusive use of any DR test VTSS(s) during a DR test. At the end of a DR test, all data created from the test systems is discarded and the segregated hardware can be re-deployed to the normal production environment.

1. Concurrent Disaster Recovery Test (CDRT) was formerly known as “Concurrent VSM Disaster Recovery Test (CVDT)”

Key VSM Benefits and Features

TABLE 7-1 Key VSM Benefits and Features

Feature	Benefit	How it Works
Maximum cartridge use through virtual volume stacking.	Improves media use.	VTV stacking onto MVCs improves media use. Data sets are compressed when written to a VTV, and the VTVs are then staged, stacked, and migrated to an MVC. VSM tries to fully use a tape cartridge by stacking as many VTVs onto an MVC as will fit. Because the host view of the VTV is independent of its stacking on the MVC, no volume contention typical of today's stacking solutions occurs.
<ul style="list-style-type: none"> ? Data compression. ? Fast transports. 	Improves tape transport use.	Data compression in the disk buffer plus fast transports means greater transport performance. Data sent to the transport is compressed at an average 4:1 ratio. The transports used for VSM move data at extremely fast speeds, further enhancing data availability.
Lowers demand on robotics and transports.	Conserves Nearline hardware resources.	VSM only mounts an MVC to migrate and stack VTVs to that MVC or to recall a VTV. All other mounts are virtual. The robotics, load, thread process is amortized over more data transfer.
<ul style="list-style-type: none"> ? Efficient volume stacking and transport emulation plus buffering ? Fewer physical mounts 	Compensates during peak allocation.	Each VTSS has 64 (for VSM2/3) or 256 (for VSM4/VSM5) virtual transports that mount VTVs for read/write operations. Transport emulation allows more concurrent virtual mounts while reducing real physical mounts on RTDs during times of peak allocation.

TABLE 7-1 Key VSM Benefits and Features

Feature	Benefit	How it Works
VSM allows you 100% automation of existing cartridges and reduces costs of additional LSMs and drives.	Creates space while optimizing future Nearline investments.	<ul style="list-style-type: none"> ? 1) You use a data mover utility to move data to VSM from existing LSMs that contain underused media. ? 2) You can then use the freed up space in your LSMs to: <ul style="list-style-type: none"> ? Automate manual tape operations. ? Keep data in LSMs for longer periods of time and reduce cartridge movement. ? Store data from storage-intensive applications. ? Move new applications to tape.
<ul style="list-style-type: none"> ? Fast virtual mounts. ? High number of concurrent mounts without overlap time involved in physical mounts. 	Improves performance.	<ul style="list-style-type: none"> ? Virtual mounts (VTSS-resident VTVs) on emulated transports are almost instantaneous. ? Virtual mounts on VTDs allow 64 (for VSM2/3) or 256 (for VSM4/VSM5) concurrent mounts that reduce overlap time. <p>For more information, see “VSM Performance” on page 49.</p>
<ul style="list-style-type: none"> ? VSM, not the host, performs data movement. 	Reduces management burden and conserves CPU and channel resources.	VSM does data staging and movement, saving valuable MVS CPU cycles and I/O channel resources. In an HSM environment, HSM must use MVS CPU and I/O channels to actually move the data from HSM ML0 disk to ML1 compressed disk and then to ML2 tape archive. Essentially, HSM must move the data twice, whereas VSM uses no MVS CPU cycles to move the data.
VTCS and NCS 6.0 and 6.1 provided major new features for customers.	Provided powerful tools to manage data and your VSM system.	For more information, see “What Features Did Previous Releases of VTCS and NCS Offer?” on page 31 .
VTCS and NCS build on the previous releases to extend the VSM feature/function set.	Continues to provide additional value add and investment protection for your VSM system.	For more information, see “What Enhancements Does VTCS/NCS 6.2.0 Offer?” on page 5 .
The VSM Offsite Vault Disaster Recovery (DR) Solution is an orderable feature of VTCS that helps you safeguard your mission-critical data.	Provides enhanced data availability and business protection capabilities.	For more information, see “The VSM Offsite Vault Disaster Recovery Feature” on page 28 .
Concurrent DR Test (CDRT) is an orderable feature of VTCS that provides a VSM or Nearline DR testing environment.	Allows you to test your DR strategy.	For more information, see “The Concurrent DR Test Feature” on page 29 .

VSM Frequently Asked Questions

What benefits does VSM provide?

Answer:

VSM improves media and transport use and conserves Nearline hardware resources. For information about these benefits and other benefits, see [“Key VSM Benefits and Features” on page 31](#).

Does VSM outperform TMM?

Answer:

Yes. VSM manages all disk-to-tape data migration without using CPU and I/O channel resources. See [“VSM Advantages versus HSM/TMM” on page 51](#) for more information.

What data sets are good candidates for moving to VSM?

Answer:

Types of data sets that are best suited for VSM management include the following:

- ? Sequential data sets

- ? Data sets that would not effectively use the cartridge capacity
- ? Data sets that are *not* created by systems that fill cartridges

What data sets are inappropriate for VSM?

Answer:

Types of data sets that might not be appropriate for VSM management include the following:

- ? HSM archives
- ? Data sets requiring direct access
- ? Data sets that would effectively use the cartridge capacity
- ? Files requiring high performance data transfers

Is there a tool to help identify candidate data sets for VSM?

Answer:

The pre-sales planning tool helps customers identify which data sets would benefit from VSM by analyzing volume and transport activity from real tape activity and from simulated VSM activity. The pre-sales planning tool produces data to help define the optimum VSM solution for your business needs and to identify VSM candidate data sets. Specifically, the tool produces the following:

- ? A VSM candidate data set list
- ? Sizing reports for both real and simulated virtual activity that recommend the following:
 - ? VTSS size
 - ? Number of RTDs
 - ? VTV pool size
 - ? MVC pool size
 - ? Number of channels

How do I direct candidate data sets to VSM?

Answer:

There are two ways in which you can direct data sets to VSM:

- ? SMC TAPEREQ statements, or
- ? The SMC DFSMS interface.

You can also use SMC User Exits or JCL to direct data sets to VSM although StorageTek does not recommend these methods.

Can VSM and the non-VSM portion of HSC share library resources?

Answer:

In general, Nearline jobs cannot directly access MVCs or transports dedicated solely to VSM. However, VSM-owned transports and MVCs are typically only a part of the transports and volumes in an ACS. Physical transports can be shared between HSC and VSM only through static sharing, where the user manually manages the transport sharing. For more information, see [“VSM Configurations” on page 45](#).

What virtual transports and media does VSM support?

Answer:

The VTSS emulates 3490E devices and media.

How many VTVs are supported in VSM?

Answer:

VSM supports a maximum of 100,000 (for VSM2s and VSM3s) or 300,000 (for VSM4s and VSM5s) VTSS-resident VTVs per VTSS, depending on the size of the VTVs, capacity of the VTSS, and size of the CDS. You can have an unlimited number of VTVs that have been migrated to MVCs. The number of VTVs you can migrate to MVCs is limited by the total MVC space you allocate.

Can VTDs be shared among multiple MVS hosts?

Answer:

Yes, VTDs can be shared among multiple hosts using tape sharing products like IBM's GRS or Computer Associates Multi-Image Manager. VTDs are defined to these products just as any other tape drive would be.

Can RTDs be shared?

Answer:

RTDs can be shared in two ways:

- ? Two or more VTSSs can dynamically share one RTD.
- ? An MVS host and a VTSS can share an RTD, but the user must manage the sharing manually.

For more information, see [“VSM Configurations” on page 45](#).

What is the VSM Offsite Vault Disaster Recovery Feature?

Answer:

The VSM Offsite Vault DR optional feature consists of the VSM Vault Utilities and the *VSM Offsite Vault Disaster Recovery Guide*. For more information, see [“The VSM Offsite Vault Disaster Recovery Feature”](#) on page 28.

What is the Concurrent DR Test Utility?

Answer:

Customers who use a Disaster Recovery (DR) configuration may want to validate their ability to continue normal production processing before an actual disaster occurs. In fact, certain businesses are required by their auditors to prove the readiness of their business continuance model. Using the Concurrent Disaster Recovery Test (CDRT)¹ software, you can do this validation without purchasing additional Nearline and/or VSM hardware. The DR Test software lets you test your DR environment while concurrently running production work. For more information, see [“The Concurrent DR Test Feature”](#) on page 29.

What happens when high AMT is exceeded?

Answer:

See [“How VSM Automatically Manages VTSS Space and Migrates VTVs”](#) on page 41.

¹Concurrent Disaster Recovery Test (CDRT) was formerly known as “Concurrent VSM Disaster Recovery Test (CVDT)”.

What type of fault recovery does VSM have?

Answer:

VSM supports the following types of recovery:

- ? **CDS recovery.** The VTCS audit utility can fully recover from the loss of the CDS (which contains a VTV inventory and an MVC list). The MVC audit can also verify the contents of each MVC.
- ? **Multiple VTV Copies.** VSM provides the capability of creating up to 4 copies of a VTV when it is migrated. Each copy is written to a different MVC in separate ACSs if desired. Should an error occur recalling the VTV from one copy, VSM will automatically attempt the recall using another copy.
- ? **VTSS recovery.** Check 0 error recovery for a VTSS is handled so that there is generally no impact to jobs executing with the VTSS (with the possible exception of the job that caused the error condition).
- ? **RTD recovery.** VSM provides error recovery equivalent to that provided by a host system. Specifically, this recovery includes retrying the failed operation on another RTD (for device failures) or using another MVC (for media failures).
- ? **RTV utility.** This utility converts MVC–resident VTVs to data on conventional Nearline volumes. You would typically use this utility if you have no VTSS you can use to recall VTVs.
- ? **Disaster Recovery using the VSM Offsite Vault Disaster Recovery Feature.** For more information, see “The VSM Offsite Vault Disaster Recovery Feature” on page 6.

What are Clustered VTSS Configurations and how do they work?

Answer:

Clustered VTSS configurations provide additional business continuance and business resumption capability for your VSM system. For more information, see *Beyond the Basics: VTCS Leading Edge Techniques*.

How Does VSM Work?

VTCS and HSC work together to provide several methods that you can use to route data to VSM. Each VTSS provides 64 virtual tape drives (VTDs) for VSM2s and VSM3s, and 256 VTDs for VSM4s and VSM5s. VTDs emulate 3490E devices. VSM uses the VTDs to write data sets to virtual tape volumes (VTVs) on the VTSS. The VTSS storage is provided by a RAID-6+ DASD configuration. You specify the VTSS's high and low Automatic Migration Thresholds (AMTs), which control the VTSS space management/VTV migration cycle. Real tape drives (RTDs) write migrated VTVs to physical multi-volume cartridges (MVCs). VTCS controls RTDs (although HSC provides mount and dismount services for MVCs), while HSC controls conventional Nearline tape drives that are not allocated to VSM. By default, VSM migrates a single copy of a VTV. You can, however, use the NCS/VTCS policies to specify that VSM migrates multiple copies of the VTV to separate MVCs.

If the host requests a mount of a VTV that was migrated to an MVC and is not VTSS-resident, VSM automatically recalls the migrated VTV to the VTSS. [FIGURE 9-1 on page 40](#) shows the VTV migration/recall cycle.

Note – VSM supports dynamic sharing of RTDs between VTSSs. **Note, however, that when VTSSs share RTDs, the VTSSs must have access to all the same hosts.**

Also note that VSM does **not** support dynamic sharing of transports between VSM and MVS. That is, a transport cannot simultaneously be online to both MVS and to VSM as an RTD. You can, however, manually vary a transport online to MVS and offline to VSM and vice versa.

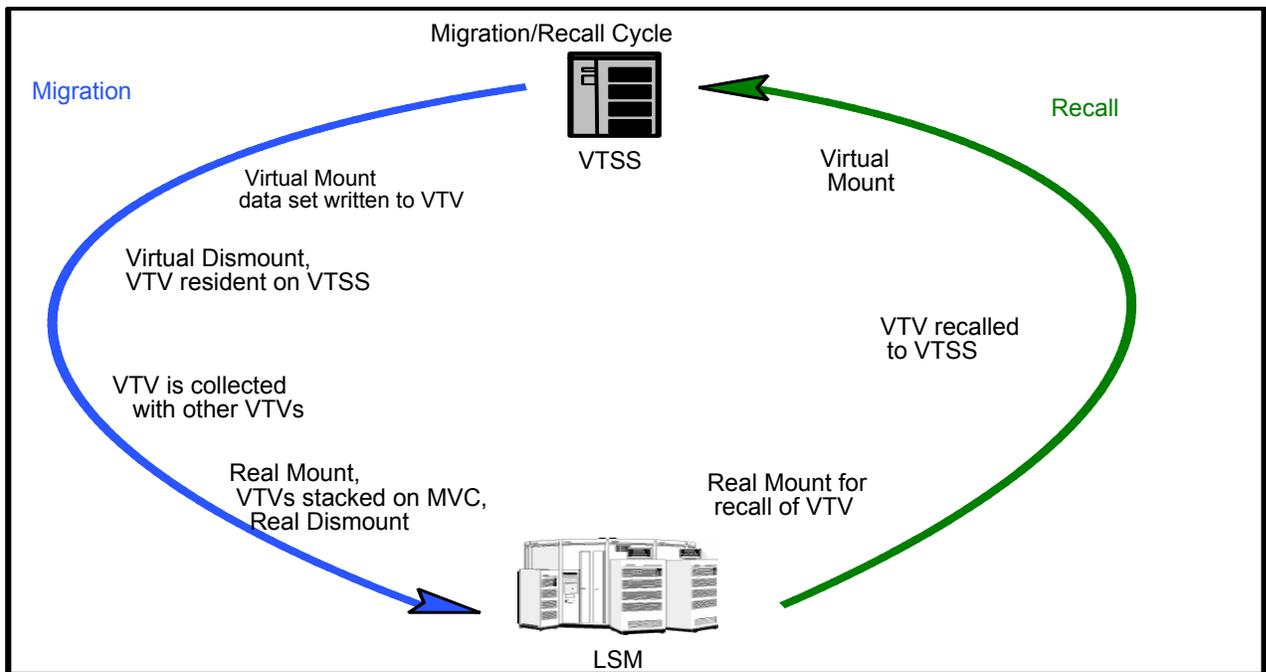


FIGURE 9-1 VTV Migration/Recall Cycle

The following sections provide more information about VTV migration and recall, MVC space reclamation, VTCS and NCS enhancements, VTV consolidation, an overview of how to export and import data, and an overview of Clustered VTSS configurations.

How VSM Automatically Manages VTSS Space and Migrates VTVs

VSM automatically manages VTSS space and migrates VTVs as follows:

1. You route data sets to VSM using SMC Policies that specify virtual media and VTCS Management Class.
2. VTCS mounts a VTV on a VTD, writes the data set to the VTV, then dismounts the VTV, which is now VTSS-resident.
3. Disk Buffer Utilization (DBU) is the percentage of space used on a VTSS compared to its capacity. VTCS starts automigrating VTVs once the DBU reaches or exceeds the High Auto Migration Threshold (HAMT) or the number of VTVs exceeds 97,000 (for VSM2s and VSM3s) or 291,000 (for VSM4s and VSM5s). Automigration continues until the Low Auto Migration Threshold (LAMT) is reached.

Note – VSM will initially start a number of automigration tasks according to the workload characteristics within the limits of the MINMIG and MAXMIG values. During the migration process, VTCS continuously adjusts the number of tasks based on the DBU, the HAMT, and the migration target (the LAMT or the migrate-to-threshold value).

VTCS selects VTVs for automigration in three modes, normal, high and space release:

- ⌚ While the DBU is below or equal to the High Automatic Migration Threshold (HAMT), VTCS operates in *normal mode*. In this mode, VTCS selects VTVs for migration that are least likely to be re-referenced (based on VTV age) but also considers VTV size and any recommendation for the residency time of each individual VTV (RESTIME). VTCS migrates the selected VTVs and deletes the VTSS-resident copies until DBU reaches the Low Automatic Migration Threshold (LAMT).
- ⌚ If the DBU exceeds HAMT but is less than 95%, VTCS switches to *high mode*. As DBU approaches 95%, VTCS progressively more heavily weights VTV size versus age to select VTVs for migration. VTCS continues to consider any RESTIME recommendations. VTCS migrates the selected VTVs and deletes the VTSS-resident copies until DBU reaches the LAMT.
- ⌚ If the DBU reaches 95%, VTCS switches to *space release mode*. In this mode, any VTV that is resident but already migrated is first deleted from the VTSS, after which VTCS only migrates VTVs on a size basis until the LAMT is reached. VTCS does *not* honor any RESTIME recommendations in space release mode.

Also note that when auto migrate selects candidate VTVs, all VTVs associated with a multiple VTV group that spans a single data set are included within the subsequent migration operation.

Demand Migrations

As an alternative to automatic migration, you can also use MIGRATE to do demand migrations.

If you demand migrate specified VTVs, VTCS will start a single migration task. A migrate-to-threshold triggers automigration for the specified VTSS and temporarily sets the LAMT to the specified threshold value. The number of migration tasks is determined as described in [Step 3](#) on page 21.

You can also specify that VTCS immediately schedules VTVs for migration on dismount. Just as with automatic migration and demand migrate-to-threshold, immediate migration spawns one or more migration requests per VTSS within the limits specified by the CONFIG VTSS MINMIG and MAXMIG parameters.

How VSM Automatically Recalls Migrated VTVs

If the host requests a specific mount of a VTV whose most current copy is VTSS resident, VSM mounts the VTSS resident copy of the VTV. Otherwise, VSM automatically recalls the VTV as follows:

1. For VTVs with multiple MVC copies managed by Advanced Management Policies, VTCS selects the MVC to recall from as follows:
 - If one MVC is mounted, VTCS selects that MVC.
 - If neither MVC is mounted, VTCS selects the MVC from the first Storage Class specified by the MIGpol parameter of the MGMTclas statement. If the MVC in the first Storage Class is marked in error, VTCS selects the MVC from the second Storage Class specified by the MIGpol parameter.
2. VTCS selects an RTD connected to the VTSS containing the VTD specified in the mount request.
3. VTCS directs HSC to mount the MVC on the selected RTD.
4. VTCS recalls the entire VTV from the MVC.
5. VTCS directs HSC to dismount the MVC (unless it is needed for another VTV recall request).
6. VTCS mounts the VTV on the VTD.

Note – If a host requests a mount of a data set that spans multiple VTVs, when the mount is issued for any VTV referenced by the data set, VTCS also recalls (if necessary) the next VTV in sequence. VTCS thus attempts to “pre-stage” multiple VTVs referenced by the same data set.

7. The host accesses the VTV. If the host changes the contents of the VTV, VTCS marks the VTV space on the MVC as available for reclamation.
8. At the end of the job, the host unloads the VTD.

Note – You can also use RECALL to do demand recalls.

Also note that by default, VTCS recalls VTVs with read data checks. For example, if an MVC is defective, only part of a specific VTV may be readable, but VTCS still recalls the readable portion, which creates a “partial” VTV in the VTSS. If a VTV with read data checks contains critical data, you may choose to recall just the readable portion, or you may choose to only recall complete VTVs.

Reclaiming Space on MVCs

Reclaiming space on MVCs— You can use the VTCS CONFIG utility to set MVC space reclamation parameters, and you can use the VTCS RECLAIM command or utility to adjust these parameters. When MVCs reach specified thresholds, VTCS automatically starts MVC space reclamation, reclaiming space one MVC at a time. VSM reclaims space by copying only valid VTVs from the selected MVC to the VTSS, then copying these VTVs back to another MVC with sufficient space. VSM copies only those VTVs placed after the first open space on the MVC, which saves I/O cycles. The space reclamation thus reduces MVC fragmentation and allows the migration to succeed.

If the system's MVCs are highly fragmented, a demand MVC space reclamation can be scheduled (using VTCS RECLAIM) as an off-hours batch job.

Note –

- 2 To optimize recall processing in mixed-media systems, ensure that your MVC pool has at least one media type compatible with each RTD type.
- 2 The MEDIA parameter of the STORclas statement (which defines Storage Classes) specifies a preference list of MVC media types. This list supersedes the default media selection list. For more information, see *ELS Command, Control Statement, and Utility Reference*.

How VSM Maintains Data Integrity and Availability When Moving VTVs from One MVC to Another

During a drain process (MVC drain or space reclamation), VTCS maintains data integrity and availability by not altering the MVC reference in the VTV record until that VTV has been recalled **and** successfully migrated to a new MVC. In addition, VTCS immediately migrates VTVs recalled by a drain process instead of allowing Automatic Migration to migrate the VTVs. If a drain operation does not complete (is cancelled, the system fails, and so forth), you must rerun the drain operation, but all in-transit VTVs will still point to the MVC that was being drained. If a VTV is migrated to multiple MVCs and one MVC is in error, VTCS recalls the VTV from an alternate MVC and immediately migrates the VTV without waiting for subsequent processing of the in-error MVC.

Note that, for both MVC drains and space reclamations, VTCS does a Management Class lookup after the recall phase and honors any Management Class changes.

How VTCS Manages Scratched VTVs

When a VTV is scratched, VTCS sets the status to “scratch” in the VTV record of the CDS. However, VTCS does **not** delete the VTV from the VTSS buffer. Instead, VTCS unlinks all MVC copies except the last one, which is not unlinked until the VTV is reused as a scratch volume as follows:

- ? If the VTV is simplexed, then the single VTV copy is preserved until the VTV is reused as a scratch volume.
- ? If there is more than one copy of the VTV, then copies are deleted as MVCs are reclaimed until there is only one VTV copy left. As described above, this last MVC copy is not unlinked until the VTV is reused as a scratch volume.

The Import and Export Functions

You can use EXPORT and IMPORT to move VTVs on MVCs from one VSM system to another, where each system has its own unique resources, including different CDSs. EXPORT and IMPORT, therefore, give you the ability to create portable MVCs that you can use to move VTVs from one system to another.

VTV Archiving and Reconciliation

You can use the ARCHAge and ARCHPol parameters of the MGMTclas statement to set an *Archive Policy* for VTVs in a Management Class. When the VTV's age exceeds the ARCHAge value, the VTV is eligible for archive per the Storage Class(es) specified on the ARCHPol parameter. The actual archive occurs in one of two ways:

- ? Automatically the next time the VTV is recalled and remigrated.
- ? On demand via the ARCHIVE utility.

Information Lifecycle Management (ILM), a StorageTek storage management strategy, has as its central concept the idea that data should be stored on media that matches its importance to the business enterprise and its pattern of reuse. Active, important data goes on fast-access media, and has multiple copies, while inactive, less important data is archived to high-capacity, cheaper media. Automating this process is the most cost-effective way to manage data storage. Archiving implements ILM by letting you archive inactive data. Using VTCS archiving, you can move VTVs to different media (for example, from fast-access 9840 media to high-capacity 9940 media) and a different location (for example, from a local ACS to a remote ACS for ejection/vaulting).

You can use the RECONcil utility to reconcile VTV media and location (move VTVs from one Storage Class to another). You typically reconcile VTVs when:

- The VTVs are on the wrong media, in the wrong ACS, or both.

An ACS is unavailable for a considerable period of time, then is brought back online. In this case, you would first change the MIGpol parameter on the MGMTclas statement for the affected VTVs to point to a different ACS (and media, if desired). When the original ACS comes back online, you then change the MIGpol parameter on the MGMTclas statement to point to the original ACS, and run RECONcil specifying the updated MGMTclas (or STORclas) statement(s) to move the VTVs to the original ACS.

Clustered VTSS Configurations

Clustered VTSS configurations come in two types:

- 7 *Uni-Directional Clusters*, which consist of a *Primary VTSS* and a *Secondary VTSS* connected by one or more cluster links (CLINKs). FIGURE 9-2 shows a Uni-Directional Cluster attached to a single ACS. The CONFIG CLUSTER statement defines the Cluster, and the CONFIG CLINK statement(s) define this the Cluster as Uni-Directional.

In a Uni-Directional Cluster, you can use the MGMTclas statement REPLICAT parameter (which requires the Advanced Management Feature) to direct the Primary VTSS to *replicate* (copy) a VTV to the Secondary VTSS via a cluster link (CLINK). If the Primary VTSS becomes unavailable, you can use the VARY VTSS to vary it offline to VTCS. You then vary the Secondary VTSS's VTDs online to MVS to continue the workload. The Secondary, therefore, acts as a "warm standby" to the Primary VTSS. **Note that** the Secondary can accept production work as well as replicated VTVs.

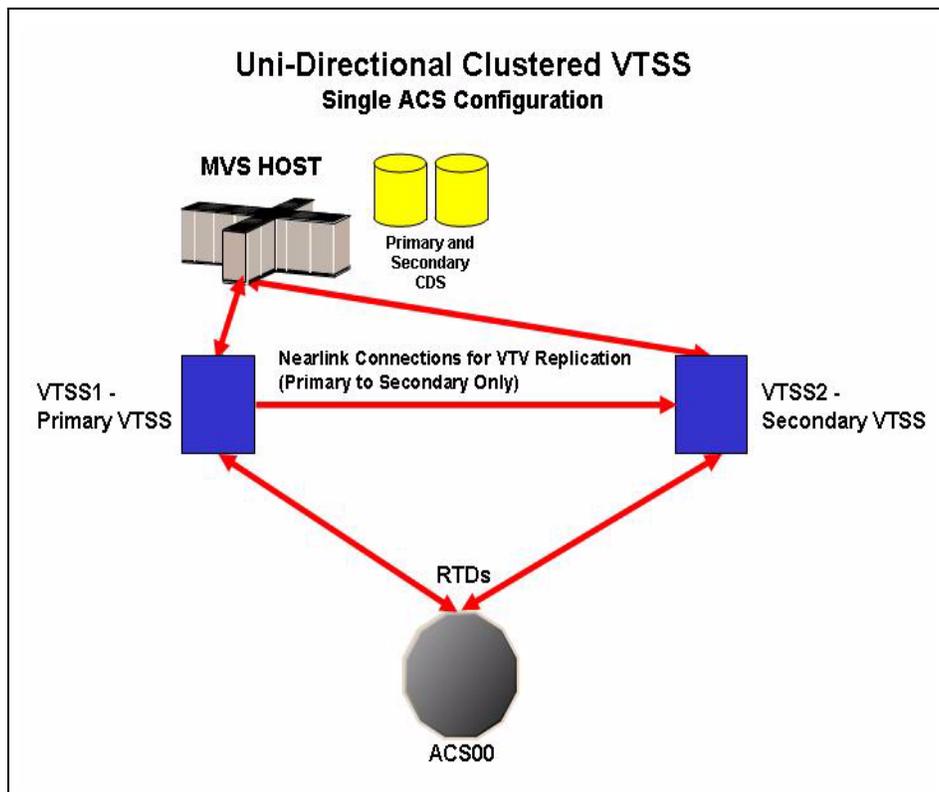


FIGURE 9-2 Uni-Directional Cluster attached to a Single ACS

7 *Bi-Directional Clusters*, which consist of two *Peer VTSSs* connected by one or more cluster links (CLINKs). **FIGURE 9-3** shows a Bi-Directional Cluster attached to a single ACS. The CONFIG CLUSTER statement defines the Cluster, and the CONFIG CLINK statement(s) this Cluster as Bi-Directional.

In a Bi-Directional Cluster, you can use the MGMTclas statement REPLICAT parameter (which requires the Advanced Management Feature) to allow replication *in either direction* from peer to peer across a cluster link (CLINK). If either VTSS becomes unavailable, you can use the VARY VTSS to vary it offline to VTCS and the Peer VTSS takes over the workload. Each Peer VTSS, therefore, acts as a “warm standby” to the other, and each can accept production work as well as send and receive VTVs.

FIGURE 9-3 shows a Peer-to-Peer VTSS Cluster cross-connected to two ACSs for increased redundancy (for example, for enhanced Disaster Recovery capability).

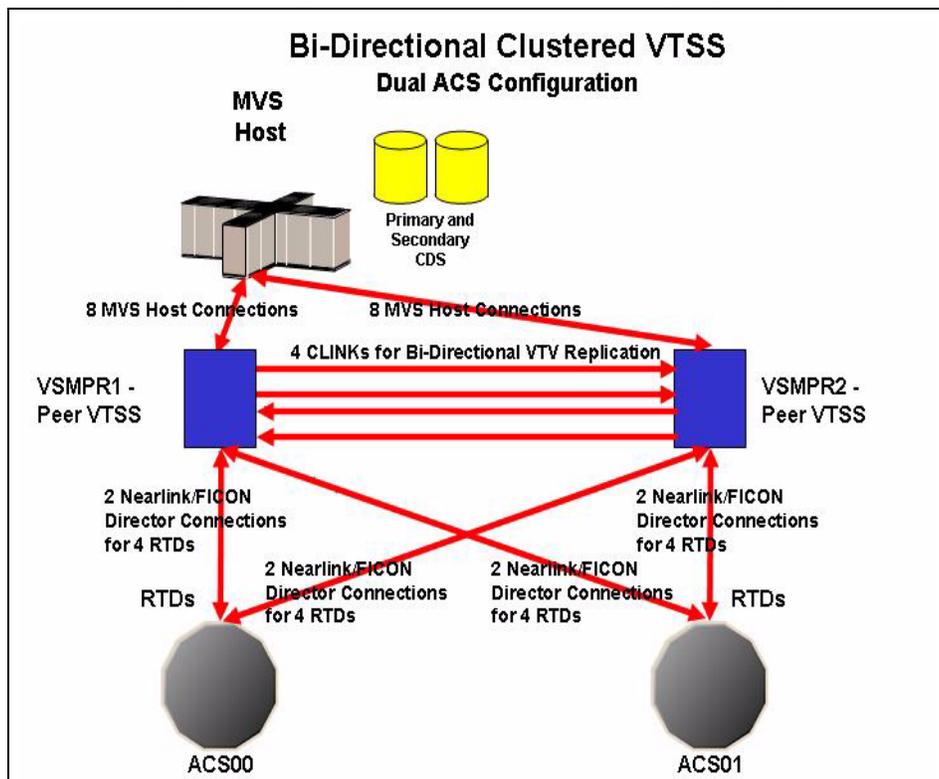


FIGURE 9-3 Bi-Directional Cluster attached to a Dual ACS

Compared to a non-clustered configuration, a Clustered VTSS configuration can provide enhanced data availability (business continuance) and enhanced disaster recovery capability (business resumption) for your VSM system.

VSM Performance

VSM reduces mount time and reduces number of tape mounts:

In VSM, virtual mounts occur almost instantaneously because no physical action is required. The number of real tape mounts is reduced because data is written to the buffer. Data is buffered in the VTSS, and mount requests are serviced at disk speeds. No robotic or mechanical tape I/O is required.

Mount times are reduced in another way. Concurrent jobs in conventional Nearline transports involve overlap time as multiple mounts wait for robot and transport allocation and completion of the mount, load, and thread process. Overlap time results in lost production time. Virtual mounts using VTDs allow concurrent mounts without the overlap time required for conventional Nearline transports.

With 64 VTDs (for VSM2s and VSM3s) or 256 VTDs (for VSM4s and VSM5s) per VTSS, virtual mounts can save thousands of conventional mounts daily.

VSM reduces problem of allocation recovery:

Allocation recovery occurs during times of peak drive use when conventional transports are allocated, and a job is waiting for a drive to be freed by another job. Allocation recovery may continue for several hours with many jobs hung up waiting for transports. In a VSM system, each VTSS has 64 VTDs (for VSM2s and VSM3) or 256 VTDs (for VSM4s and VSM5s), which are serviced by up to 8 RTDs per VTSS (for VSM2s and VSM3s) or up to 16 RTDs (for VSM4s and VSM5s). VTDs, therefore, greatly reduce allocation recovery problems. In addition, when VSM is used to handle small tape files or files with low I/O rates, the real tape system can operate much more efficiently.

VSM Advantages versus HSM/TMM

How HSM Works

HSM automatically manages files on disk to optimize storage. HSM management consists of the following migration levels:

- ? In the first level of migration, ML0, active files reside on disk that are monitored by HSM. After a user-defined period of time during which HSM-managed data has not been accessed (read or written) by any application, HSM moves the data to the second level of migration.
- ? In the second level of migration, ML1, HSM compresses and compacts the data to disk volumes dedicated to HSM. At this level data is unusable to any application except HSM. Data remains in ML1 until either it is referenced for read/write operations or the user-defined time limit for no access expires.
 - ? If the data is referenced for read/write operations, it is restored to the user's disk area in original condition and the ML0 time period begins again.
 - ? If the no access time limit expires, the data is collected along with other unreferenced ML1 data sets and moved to the third level of migration.
- ? In the third level of migration, ML2, the collected data sets are written to HSM tape volumes. The data remains in ML2, managed as a tape volume. If it is referenced by a user or an application, HSM initiates a tape volume recall. The tape volume is mounted and the specified data is restored to the user's disk area in original condition. The HSM migration cycle starts again.

Disadvantage of HSM

By moving data from one area of disk to another and from disk to tape, HSM consumes considerable CPU and channel resources.

How TMM Works

TMM is a methodology operating under SMS. It optimizes use of tape media by reducing tape mounts. HSM collects the data that SMS has redirected to disk and migrates it to tape.

Specifically, TMM works in two ways:

- ? It redirects small tape data sets to reserved disk volumes at the first migration level in HSM.
- ? On an hourly basis TMM uses HSM to move all TMM selected data off the disk buffer to ML2 tape, thus effectively stacking the “tape” data sets on ML2 tape volumes.

Disadvantages of TMM

- ? Because HSM manages the TMM data sets in the SMS disk, the data is written twice; this uses extra CPU and channel resources.
- ? If the TMM disk pool fills up, TMM tape jobs will fail with disk space allocation errors.

How VSM Works

VSM uses a disk buffer to emulate tape transports. It migrates data directly from the VTSS to real tape.

How VSM Overcomes the Disadvantages of HSM/TMM

VSM consumes no CPU cycles or channel resources in the movement of data from disk to tape because VSM, not the host, manages and provides resources for data movement.

NCS/VTCS Terminology

The following table describes the terms and abbreviations associated with NCS/VTCS.

TABLE A-1 NCS/VTCS Terms and Acronyms

Term	Description
ACS	(1) StorageTek's Automatic Cartridge System. (2) A multi-LSM configuration.
ACS routine	An SMS term, referring to automatic class selection routine. Not to be confused with the HSC term, ACS, referring to Automatic Cartridge System.
AMT	Automatic migration threshold. AMT values are user-defined percentage values that determine when virtual tape volume migration begins and ends. VTV migration begins when the VTSS buffer reaches the high AMT and ends when the buffer reaches or falls below the low AMT. These thresholds apply to all VTSSs.
CAP	An assembly that allows several cartridges to be inserted into or ejected from an LSM without human entry into the LSM.
CDS	Control data set. The HSC database. In addition to the current information in the CDS, VSM keeps all its persistent data in the CDS as well.
Clustered VTSS Configuration	A configuration that consists of a Primary VTSS and a Secondary VTSS connected by one or more Nearlink connections (cluster links). You can use the MGMTclas statement REPLICAT parameter (which requires the Advanced Management Feature) to direct the Primary VTSS to replicate (copy) a VTV to the Secondary VTSS via a cluster link. If the Primary VTSS becomes unavailable, you can use the VT VARY VTSS command to vary it offline to VTCS. You then vary the Secondary VTSS's VTDs online to MVS to continue the workload. The Secondary, therefore, acts as a "warm standby" to the Primary VTSS.
DBU	Disk buffer utilization. The ratio of used to total VTSS buffer capacity.

TABLE A-1 NCS/VTCS Terms and Acronyms

ExPR	Expert Performance Reporter. ExPR collects performance data and generates reports about StorageTek Nearline ACSs and VTSS status and performance.
HSC	Host Software Component.
HSM	Hierarchical Storage Manager.
LMU	A Library Management Unit that controls one or more LSMs.
LSM	A Library Storage Module that contains storage cells for cartridges, drive panels, and CAPs.
Migration	The movement of data from the VTSS to the RTD where VTVs are stacked onto MVCs. Migration is initiated by VSM when high AMT levels are reached. VTVs are selected for migration based on use and size: the least recently used and the largest VTVs are selected first. VSM provides the ability to migrate VTVs on demand and to migrate multiple copies of the VTV.
MVC	Multi-volume cartridge. The physical cartridge in the LSM that contains one or more VTVs or no VTVs, but has been identified as a volume that can be selected for VTV stacking. This data is stored in the CDS.
Recall	The movement of VTVs back to the VTSS from the MVC. VSM provides the ability to recall VTVs on demand.
Reclaim	Refers to MVC space reclamation. VTCS uses the amount of fragmented free space on the MVC and the amount of VTV data that would have to be moved to determine if space reclamation is justified. VSM provides the ability to reclaim MVCs on demand.
RTD	Real Tape Drive. The physical transport controlled by VSM/HSC. The transport has a data path to a VTSS and may optionally have a data path to MVS or to another VTSS.
SMC	StorageTek Management Component, which is the interface between IBM's OS/390 and z/OS operating systems and HSC and MVS/CSC.
SMS	System Managed Storage
TMM	Tape Mount Management
VSM	Virtual Storage Manager. A storage solution that virtualizes volumes and transports in a VTSS buffer in order to improve media and transport use. The hardware includes VTSS, which is the disk buffer, and RTDs. The software includes VTCS, an HSC-based host software, and VTSS microcode.

TABLE A-1 NCS/VTCS Terms and Acronyms

VTCS	Virtual Tape Control System. The primary host software that controls activity and information about VTSSs, VTVs, RTDs, and MVCs. This software operates in the same address space from HSC, and communicates closely with HSC.
VTD	Virtual Tape Drive. A transport in the VTSS that emulates a physical 3490E to MVS. The data written to a VTD is really being written to disk. The VTSS has 64 VTDs that do virtual mounts of VTVs.
VTSS	Virtual Tape Storage Subsystem. The disk buffer containing virtual volumes and transports. The VTSS is a StorageTek RAID 6+ hardware device with microcode that enables emulation of 32 or 64 transports. The RAID device can read and write “tape” data from/to disk, and can read and write the data from/to an RTD.
VTV	Virtual Tape Volume. The “cartridge” whose volume number is known to the MVS catalog and the TMS (Tape Management System) as a tape data set.

