Introduction to Virtual Storage Manager

Version 6.1.0

CRC Update Only

Revision B - October 2006

This edition applies to Version 6.1.0 of the Nearline Control System and Virtual Tape Control System software. Information in this publication is subject to change. Send comments about this publication to:

GLSFS@Stortek.com
or
Global Learning Solutions
Storage Technology Corporation
One StorageTek Drive
Louisville, CO 80028-3256
USA

Please include the publication name, part number, and edition number in your correspondence if they are available.

About this Book

This book introduces the VSM solution. It describes the basic concepts of VSM, including the following:

- "What is VSM?"
- "Key VSM Benefits and Features"
- "How Does VSM Work?"
- "VSM Performance"
- "VSM Compared to HSM/TMM"
- "What Features Did Previous Releases of VTCS and NCS Offer?"
- "What Enhancements Does VTCS/NCS 6.1.0 Offer?"
- "Frequently Asked Questions"
- "VSM Requirements"
- "VSM Configurations"
- "VSM Terminology"

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.

Terminology

The abbreviations used in this book are defined in the last chapter, "VSM Terminology".

Contents

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 6.1.0 (VTCS 6.1.0), which is the MVS host software, the portions of NCS 6.1.0 that support VTCS, and the VTSS.

HSC provides mount and dismount services for physical multi-volume cartridges (MVCs), and NCS, working with VTCS, provides the ability for greater than 16 MVS hosts running MVS/CSC to route data to VSM. **Also note** that NCS/VTCS allows non-MVS/CSC clients to request VTV mounts using HSC and LibraryStation.

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.
- 8 or 16 ESCON 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).

VSM5: Increased Capacity and Throughput

The VSM5 provides greater capacity and throughput than the VSM4, while retaining its advantages over the VSM3 and VSM2. Table 1 summarizes the VSM5 features.

Table 1. VSM5 Features

Feature	Description
Host/Nearlink Interfaces	Up to 32 (FICON only)
RTDs supported	Up to 16 via FICON directors (in 3490-emulation mode only), can be a mixture of the following: 9840B, 9840C, 9940B, T10000.
LSMs supported	9740, 9360, 4410, 9310, SL8500
Host Software	NCS/VTCS 6.0 and above
Maximum VTDs per VTSS	256
Maximum VTVs per VTSS	297,000

Note: PTF L1H12ZT (SWS6100) is required for VTCS 6.1 support of VSM5. This PTF also changes QUERY/DISPLAY VTSS to report the VTSS capacity in gigabytes (Gb).

VSM Software

VTCS, which works as an extension to HSC, 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

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.

VSM Solution vs. Nearline Solution

Figure 1 shows the Nearline architecture with an HSC-controlled library. The VSM solution integrates easily into this open, flexible architecture.

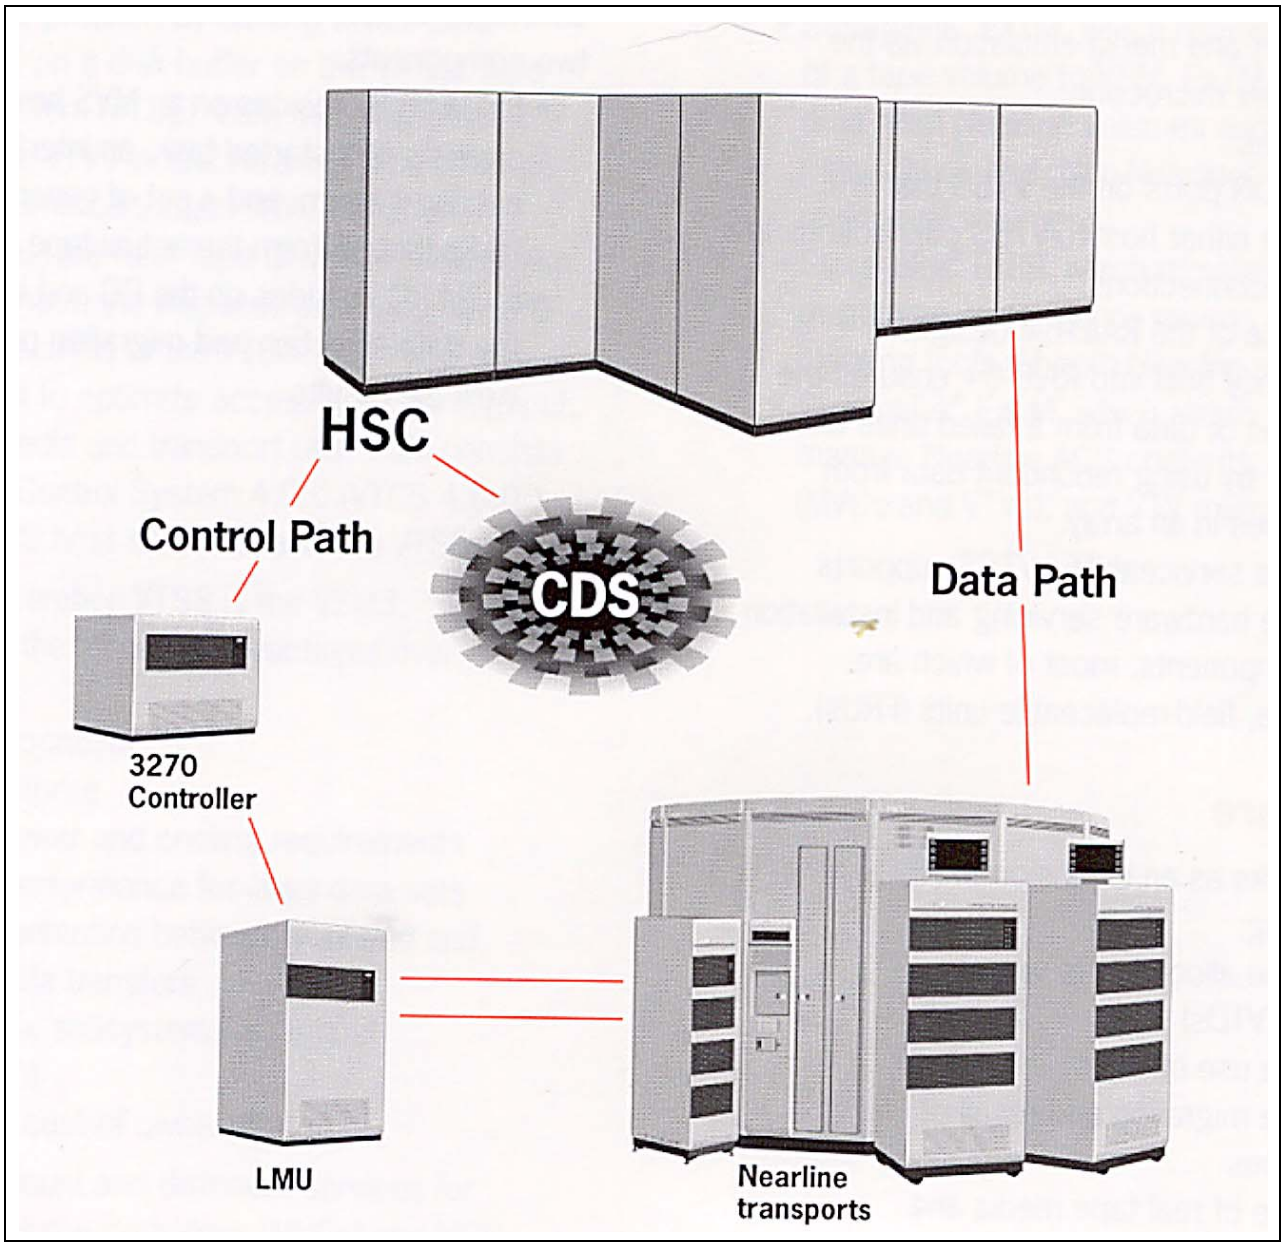

Figure 1. ***Nearline Solution***

Figure 2 shows VSM components integrated into the HSC system:

- The VTSS and its ESCON connection to a real tape drive (RTD). The connection allows migration and recall of VTVs from and to the VTSS.
- Combined ESCON control and data paths for VSM.
- VTCS
- ExPR
- ExLM

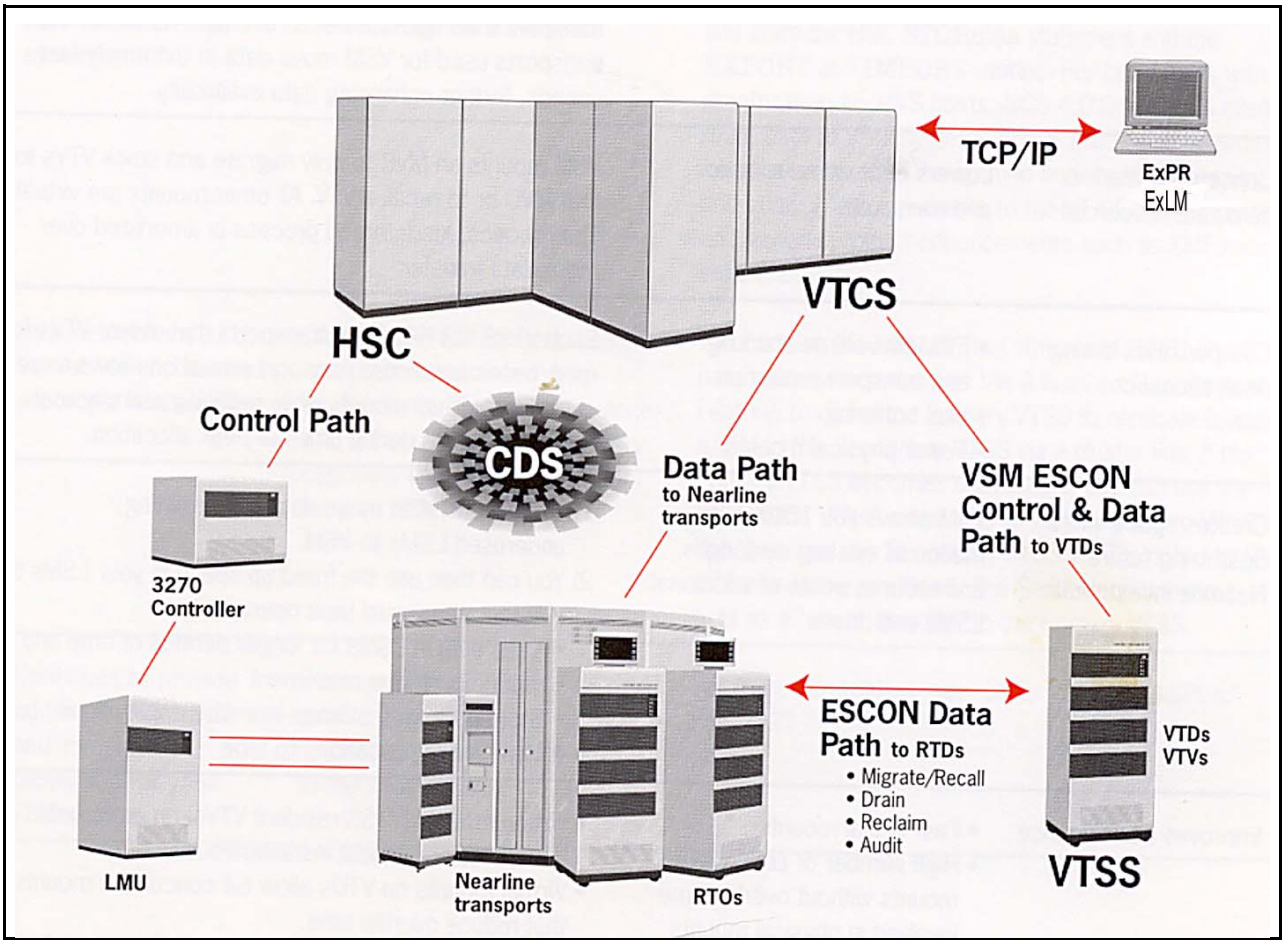

Figure 2. ***VSM Solution***

The VSM Offsite Vault Disaster Recovery Feature

The VSM Offsite Vault Disaster Recovery (DR) Feature is an optional, orderable feature of VTCS 4.0.0 and above. 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

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 Test[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. 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.

The CDRT Utility 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,

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

all data created from the test systems is discarded and the segregated hardware can be re-deployed to the normal production environment.

See Also

- “Key VSM Benefits and Features”
- “How Does VSM Work?”
- “What Features Did Previous Releases of VTCS and NCS Offer?”
- “What Enhancements Does VTCS/NCS 6.1.0 Offer?”
- “Frequently Asked Questions”

Key VSM Benefits and Features

Feature	Benefit	How it Works
VSM5 provides significant enhancements over its predecessors.	The VSM5 provides significant throughput, connectivity, capacity improvements over its predecessors, and it also provides 4x the number of VTDs per VTSS and 3x the maximum VTVs per VTSS compared to VSM2s and VSM3s.	"VSM5: Increased Capacity and Throughput" on page 2
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.
• 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.

Feature	Benefit	How it Works
• Efficient volume stacking and transport emulation plus buffering • Fewer physical mounts	Compensates during peak allocation.	Each VTSS has 64 (for VSM2/VSM3) 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.
VSM allows you 100% automation of existing cartridges and reduces costs of additional LSMs and drives.	Creates space while optimizing future Nearline investments.	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: • 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.
• Fast virtual mounts. • High number of concurrent mounts without overlap time involved in physical mounts.	Improves performance.	• Virtual mounts (VTSS–resident VTVs) on emulated transports are almost instantaneous. • Virtual mounts on VTDs allow 64 (for VSM2/VSM3) or 256 (for VSM4/VSM5) concurrent mounts that reduce overlap time.

Feature	Benefit	How it Works
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 4.0.0, 5.0.0, and 5.1.0 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 33.
VTCS 6.1.0 and NCS 6.1.0 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 6.1.0 Offer?" on page 51.
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 6.

See Also

- "What is VSM?"
- "How Does VSM Work?"
- "VSM Compared to HSM/TMM"
- "Frequently Asked Questions"

VSM Requirements

VTCS 6.1.0 System Software Requirements

Table 2. VTCS 6.1.0 Minimum Software Requirements

Software Description	Minimum Version/Release
Operating System	MVS 5.2.2 and above **Note:** NCS SMC JES3 requires JES3 5.1.1 or higher All versions of OS390
Nearline Control Solution	NCS 6.1 **Note:** • VTCS 6.1.0 requires HSC 6.1.0 and will not run with previous versions of HSC. • You can use the HSC 5.0 and above MERGEcds utility to convert the CDS to VSM Extended Format. The VSM Extended Format CDS **is required for VTCS 6.1.0.** **Note that** after you convert the CDS to VSM Extended Format, you cannot run VTCS 4.0.0 or lower against the converted CDS. • If you are using RMM in an MVS/CSC environment, MVS/CSC can share the tape management catalog with the host(s) running HSC if you have the following installed: • RMM APAR OA03368 • VSM3 microcode N01.00.65 or later or VSM4 microcode D01.00.03 or later. Otherwise, the tape management catalog cannot be shared or VTV scratch mounts will fail.
Expert Performance Reporter (optional software)	ExPR 4.0

Expert Library Manager (optional software)	To use Expert Library Manager (ExLM) with VSM for VTV consolidation using ExLM, ExLM 4.0, HSC 4.0.0, and VTCS 4.0.0. For more information about using ExLM with VSM, see "Using ExLM to Manage VSM Resources" in Chapter 2 of the ExLM *System Administrator's Guide*.

Caution: In a VSM configuration with multiple hosts that share the same HSC CDS, StorageTek strongly recommends that you:

- Install VTCS on all MVS hosts. Installing VTCS on all MVS hosts ensures that these hosts cannot scratch an MVC.

  Note that CONFIG lets you define MVS hosts that are not connected to a VTSS.

- Do *not* use a host that does not have VTCS installed to enter MVCs into an ACS, otherwise these MVCs will be eligible for selection as scratch volumes by any host in the configuration with HSC installed.

VTCS/NCS 6.1.0 SPEs

The VTCS/NCS 6.1.0 SPEs are described in Table 3.

Table 3. VTCS/NCS 6.1.0 SPEs

This SPE...	...is described in...	...and is available for VTCS/NCS 6.1.0 via PTF(s)...
MVC Media Management	"MVC Media Management SPE" on page 53	L1H12DQ and L1H12FS
GLOBAL LOCKSTR	"GLOBAL LOCKSTR SPE" on page 54	L1H12J4 (SWS6100) and L1H12J3 (SOS6100)
DELETSCR utility	"DELETSCR Utility SPE" on page 54	L1H12XT (SWS6100) and L1H12XS (SOS6100)
MGMTclas RESTIME updates	"MGMTclas RESTIME SPE" on page 54	L1H130O

VTCS/NCS 6.0.0 SPEs

The VTCS/NCS 6.0.0 SPEs are described in Table 4.

Table 4. VTCS/NCS 6.0.0 SPEs

This SPE...	**...is described in...**	**...and is available for NCS/VTCS 6.0.0 via PTF(s)...**
MVC Warranty Expired	"MVC Warranty Expired SPE" on page 35	L1H11YP
MVC Media Management	"MVC Media Management SPE" on page 53	L1H12DP (SWS6000) and L1H12Eo (SOS6000)
GLOBAL LOCKSTR	"GLOBAL LOCKSTR SPE" on page 54	L1H12E4 (SWS6000) and L1H12E6 (SOS6000)
DELETSCR utility	"DELETSCR Utility SPE" on page 54	L1H12Q2 (SWS6000) and L1H12Q1 (SOS6000)
MGMTclas RESTIME updates	"MGMTclas RESTIME SPE" on page 54	L1H130N
800 MB VTVs	"800 Mb VTVs" on page 34	L1H11WY **plus** one of the following: • For VSM3s: microcode level N01.00.69.04 **or** microcode level N01.00.70.00 and above • For VSM4s: microcode level D01.00.04.03 **or** microcode level D01.00.06.03 and above

VTCS/NCS 5.1.0 SPEs

The VTCS/NCS 5.1.0 SPEs are described in Table 6.

Table 5. VTCS/NCS 5.1.0 SPEs

This SPE...	...is described in...	...and is available for NCS/VTCS 5.1.0 via PTF(s)...
RECLAIM enhancements	"VTCS RECLAIM Enhancements SPE" on page 37	L1H11EL and L1H11EN
Merge Manifest	"VTCS Merge Manifest SPE" on page 38	L1H11HZ
T9840C support	"T9840C RTD Support SPE" on page 37	• L1H11o5 for VTCS 5.1 • L1H11o6 for HSC 5.1
T9940B support	"T9940B RTD Support SPE" on page 37	• L1H11o5 for VTCS 5.1 • L1H11o6 for HSC 5.1
MVC Warranty Expired		L1H11Y4

VTCS/NCS 5.0.0 SPEs

The VTCS/NCS 5.0.0 SPEs are described in Table 6.

Table 6. VTCS/NCS 5.0.0 SPEs

This SPE...	...is described in...	...and is available for NCS/VTCS 5.0 via PTF(s)...
RTV utility adds the VALIDATE parameter and support for ANSI labelled VTVs	"RTV Utility Enhancements SPEs" on page 49	L1H10XX for SWS5000
Merge Manifest	"VTCS Merge Manifest SPE" on page 38	L1H11HY

VTCS/NCS 4.0.0 SPEs

The VTCS/NCS 4.0.0 SPEs are described in Table 7.

Table 7. VTCS/NCS 4.0.0 SPEs

This SPE...	...is described in...	...and is available for NCS/VTCS 4.0 via PTF(s)...	...and for NCS/VTCS 4.1 via PTF(s)...
The Vary VTSS SPE	"VT Vary VTSS Command, Clustered VTSS, and VT QUery Command Enhancements SPEs" on page 45	L1H103H	
Clustered VTSS	"VT Vary VTSS Command, Clustered VTSS, and VT QUery Command Enhancements SPEs" on page 45	L1H10FR for VT QUery CONFIG enhancements. For all others, see "The VTCS/NCS 4.0.0 SPEs are described in Table 7." on page 17.	L1H10FS for VT QUery CONFIG enhancements. For all others, see "The VTCS/NCS 4.0.0 SPEs are described in Table 7." on page 17.
The VTVMAINT utility	"VTVMAINT Utility SPE" on page 49	L1H1086 (SWS4000)	
Remove the restriction of 198 total volume ranges for MVCs and VTVs and the restriction on single volume ranges	"MVC and VTV Volser SPE" on page 49	l1h10aa and l1h10av (SWS4000)	l1h10bk and l1h10bl (SWS4100)
9840/9940 Media Management	"9840/9940 Media Management SPE" on page 47	L1H10JV (SWS4000) and L1H10JW (SOS4000)	L1H10MY (SWS4100) and L1H10N4 (SOS4100)
Support for ANSI labelled VTVs	"ANSI Label Support for VTVs SPE" on page 47	L1H10K6 (SOS4000) **Note:** This SPE is only valid with VSM3s with microcode level 44 or higher.	L1H10K9 (SOS4100) **Note:** This SPE is only valid with VSM3s with microcode level 44 or higher.
MVCMAINT lost MVC	"MVCMAINT Lost MVC SPE" on page 47	L1H10FA (SWS4000)	L1H10FC (SWS4100).
Drain/Reclaim processing	"Drain/Reclaim Processing SPE" on page 48	L1H10KA (SWS4000)	L1H10KB (SWS4100)
Maximum MVCs Concurrently Processed for Reclamation and Drain	"SPE for Maximum MVCs Concurrently Processed for Reclamation and Drain" on page 48	L1H10MB (SWS4000) and L1H10MA (SOS4000)	L1H10o0 (SWS4100) and L1H10NZ (SOS4100)

Table 7. VTCS/NCS 4.0.0 SPEs

This SPE...	...is described in...	...and is available for NCS/VTCS 4.0 via PTF(s)...	...and for NCS/VTCS 4.1 via PTF(s)...
Recall VTVs with Read Data Check	"Recall VTVs Read Data Check SPE" on page 49	L1H10JP (SWS4000) and L1H10JQ (SOS4000)	L1H10JR (SWS4100) and L1H10JS (SOS4100)
RTV utility parameters: •ALLVTVS •LISTONLY •FILEnum •DUMP	"RTV Utility Enhancements SPEs"	L1H08DG for HSC 2.01 systems L1H08EF for HSC 2.1 systems L1H108Y for HSC 4.0 systems	
RTV utility adds the VALIDATE parameter and support for ANSI labelled VTVs		L1H10T3 (SWS4000)	L1H10T4 (SWS4100)
HSC MOUNT command and PGMI enhancements	"NCS 4.0.0 Enhancements for VSM" on page 42	L1H10AX (SOS4000)	L1H10AZ (SOS4100)
Support for the T9940A as an RTD	"T9940A RTD Support SPE" on page 47	L1H1043	
SMF Record 27 enhancements	"NCS 4.0.0 Enhancements for VSM" on page 42	L1H109P	
SMF Record Enhancements (unlink VTV and MVC)	"NCS 4.0.0 Enhancements for VSM" on page 42	L1H10oJ (SWS4000) and L1H10oH (SOS4000)	L1H10NG (SWS4100) and L1H10NH (SOS4100)
Merge Manifest	"VTCS Merge Manifest SPE" on page 38	L1H11HW	L1H11HY

Clustered VTSS Requirements

Table 8. Clustered VTSS Requirements

Component	Requirement
2 VTSSs within a cluster (ESCON Interfaces)	The Primary and Secondary VTSSs can be any combination of VSM3 and VSM4 where the Secondary can be of any capacity. All hosts **must** be at VTCS 5.1.0 or above to enable this feature. For example, all of the following are valid: Primary VSM4, Secondary VSM3 Primary VSM4, Secondary VSM4 Primary VSM3, Secondary VSM3 Primary VSM3, Secondary VSM4 (not recommended) For Bi-Directional Clustering, all hosts must be at VTCS 6.1.0 or above to enable this feature.
2 VTSSs within a cluster (FICON Interfaces)	The Primary and Secondary VTSSs can be any combination of VSM4 and VSM5 where the Secondary can be of any capacity. All hosts **must** be at VTCS 5.1.0 or above to enable this feature. For example, all of the following are valid: Primary VSM5, Secondary VSM4 Primary VSM5, Secondary VSM5 Primary VSM4, Secondary VSM5 Primary VSM4, Secondary VSM5 (not recommended) For Bi-Directional Clustering, all hosts must be at VTCS 6.1.0 or above to enable this feature.
Primary and Secondary VTSS microcode	The Primary VTSS microcode must be at a level that supports sending replicated VTVS. The Secondary VTSS microcode must be at a level that supports receiving replicated VTVS and supports the use of the Secondary as a production VTSS. After the microcode is installed, the Clustering feature must be enabled at **both** the Primary and Secondary VTSS via an options floppy disk. See your StorageTek hardware service representative for details.
VTCS software	• VTCS 5.1.0 (for enhanced clustered support) • For Bi-Directional Clustering, all hosts must be at VTCS 6.1.0 or above to enable this feature. • The Advanced Management Feature (to enable the REPLICAT parameter of the MGMTclas statement)

Nearline Hardware Requirements

Table 9. ***VSM Nearline Hardware Requirements***

Hardware	Requirement
LSMs	Any of the following but 9310, 9740, or SL8500 recommended by StorageTek: 4410, 9310, 9740, 9360, SL8500
Transports and media	VSM RTDs can be a mixture of 9490 (Timberline), 9490EE (Timberline EE), T9840A, T9840B,T9840C, T9940A, T9940B, and T10000 transports. Each VTSS must have a minimum of two library-attached transports for each media type used for MVCs. For example, if your MVCs are STANDARD and ECART, you need a minimum of *either* two 9490s *or* two 9490EEs as RTDs. If your MVCs are STANDARD, ECART, ZCART, and STKIR, you need a minimum of two 9490EEs *and* two 9840s as RTDs. **Note:** Using 9840s as RTDs and using ZCART media with 9490EE RTDs requires HSC 2.1.0. You cannot use 9840s as RTDs, ZCART media, or 9740 LSMs in a configuration where HSC 2.1.0 and HSC 2.0.1 share the CDS. Using T9940s as RTDs requires HSC 4.0.0 and VTCS 4.0.0 with PTF L1H1043 applied. Using T10000 as RTDs requires one of the following: 6.0 with the following PTFs: • L1H12ZI (SOS6000) • L1H12ZJ (SWS6000) 6.1 with the following PTFs: • L1H12ZN (SOS6100) • L1H12ZO (SWS6100)

	Valid media types for the supported RTDs are: 9490: STANDARD, ECART 9490EE: STANDARD, ECART, ZCART T9840A, T9840B, T9840C: T9840A/T9840B and T9840C transports use the **same physical form factor** but **different recording techniques** resulting in the following restrictions: • T9840Cs can read from media written to by T9840As/T9840Bs, but cannot write to T9840A/T9840B media **unless** the entire volume is rewritten from beginning of tape. • T9840As and T9840Bs cannot read from or write to media written to by T9840Cs. T9940A, T9940B: T9940A and T9940B transports use the **same physical form factor** but **different recording techniques** as follows: • T9940Bs can read from media written to by T9940As, but cannot write to T9940A media **unless** the entire volume is rewritten from beginning of tape. • T9940As cannot read from or write to media written to by T9940Bs. T10000: T10000T1, T10000TS For more information, see *VTCS 6.1 Installation and Configuration Guide*.

How Does VSM Work?

Figure 3 shows how VSM migrates data from the VTSS to MVCs, then recalls the data back to the VTSS.

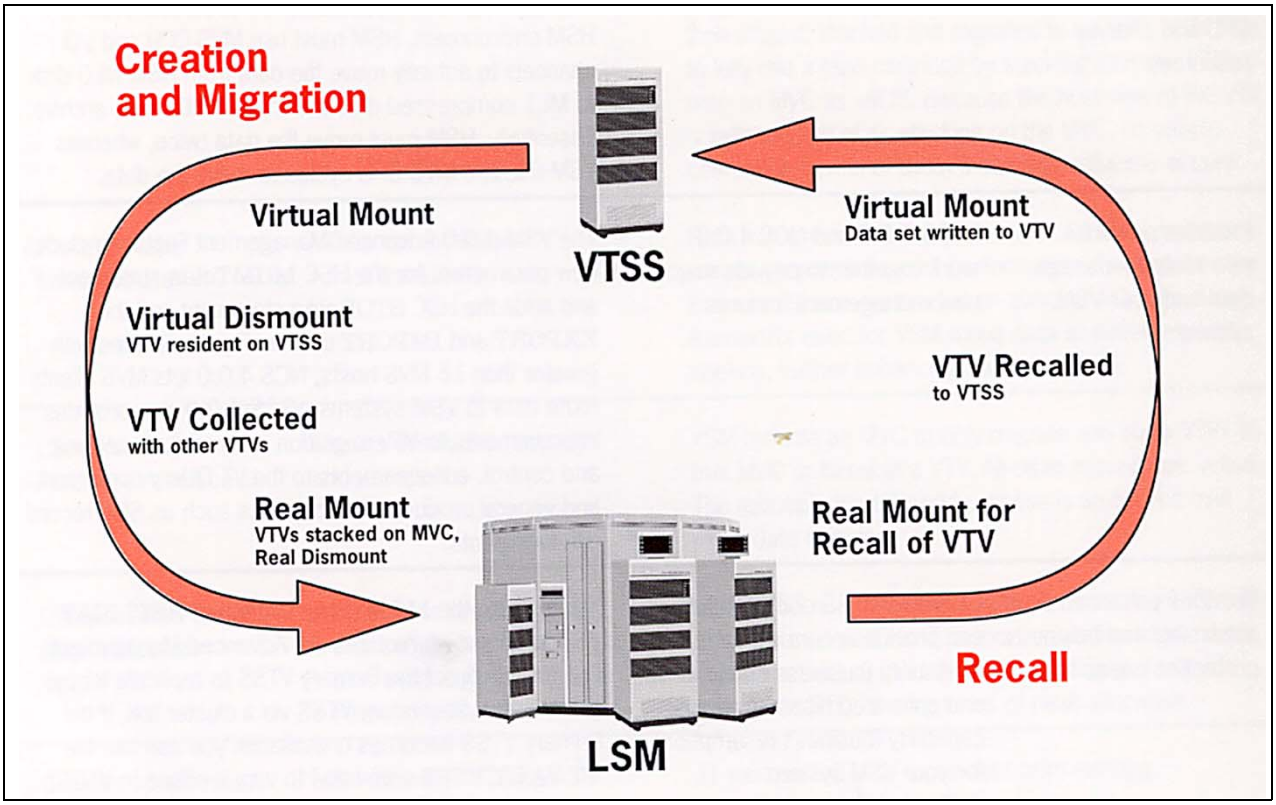

Figure 3. ***VTV Migration/Recall Cycle***

Creation and Migration

The following describes the creation and migration process, using a scratch volume as an example:

Virtual Mount—VTCS mounts a VTV scratch volume on a VTD almost instantaneously. As data is written, it is compressed at an average 4:1 compression ratio.

Virtual Dismount and Residency—The VTV is dismounted from the VTD almost instantaneously. It is stored on the VTSS disk buffer, available for reuse.

VTSS Space Management/VTV Migration—Disk Buffer Utilization (DBU) is the percentage of space used on a VTSS compared to its capacity. VTCS starts automatically migrating 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.

VSM will initially start a number of automigration tasks according to the workload characteristics within the limits of the MINMIG and MAXMIG values, which specify the minimum and maximum concurrent automatic migration, immediate migration, and migrate-to-threshold tasks. 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 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 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.

As an alternative to automatic migration, you can use the MIGRATE utility or VT MIGRATE command 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 above for automigration.

You can also specify that VTCS immediately schedules VTVs for migration on dismount with the IMMEDmig parameter of the MGMTclas statement. Just as with automatic migration and demand migrate-to-threshold, the number of migration tasks is determined as described above for automigration.

Hint: By default, in mixed-media VSM systems, VTV automatic and demand migrations go to MVCs by media type in this order:

1. Standard - 400 Mb
2. ECART - 800 Mb
3. ZCART - 1600 Mb
4. 9840 - 20 Gb
5. 9840C - 40 Gb
6. 9940A - 60 Gb
7. T10000 Sport - 120GB
8. 9940B - 200 Gb
9. T10000 Full - 500 Gb

You can specify the media and ACS preferencing for migration via the Storage Class(es) specified on the MIGpol parameter of the MGMTclas control statement. To optimize recall processing in mixed-media systems, ensure that your MVC pools have at least one media type compatible with each RTD type.

Physical Mount, VTV Stacking to MVC —The VTVs are stacked and written to an MVC through a physical mount, load, and thread. VSM will try to fully use a tape cartridge by moving multiple VTVs onto an MVC until either the tape is full or a user–defined level is reached. Data is written to the tape, and the MVC is unloaded and dismounted.

The user can specify that up to 4 copies of a VTV are written to 4 different MVCs when the VTV is migrated. The MVCs can reside in the same ACS or different ACSs, including an ACS located in an offsite location.

Recall

The following describes the recall process:

Physical Mount for a Recall—When a recall is initiated for a VTV that is not resident in the VTSS, the MVC containing the VTV is physically mounted on the RTD.

VTVs Recalled to VTSS—VTVs are recalled back to the VTSS.

Virtual Mount of Recalled VTV—The recalled VTV is mounted on a VTD in the VTSS.

You can also use the RECALL utility or VT RECALL command to do demand recalls.

Reclaiming Space on MVCs

Reclaiming space on MVCs— You can use the VTCS CONFIG utility to set MVC 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 as an off–hours batch job.

See Also

- "What is VSM?"
- "Key VSM Benefits and Features"
- "VSM Compared to HSM/TMM"
- "What Features Did Previous Releases of VTCS and NCS Offer?"
- "What Enhancements Does VTCS/NCS 6.1.0 Offer?"

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 Compared to 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.

See Also

- “What is VSM?”
- “Key VSM Benefits and Features”
- “How Does VSM Work?”

What Features Did Previous Releases of VTCS and NCS Offer?

Previous releases of VTCS and NCS provided major new features for customers as described in the following sections. Note that many of these features were provided via SPEs to NCS, VTCS, or both. All SPEs are, of course, incorporated into the base product for NCS and VTCS for subsequent releases. For more information, see:

VTCS/NCS 6.0.0 Enhancements

VTCS/NCS 6.0.0 builds on the previous releases to extend the VSM feature/function set. For more information, see:

- "800 Mb VTVs"
- "4 VTV Copies"
- "VTSS Preferencing" on page 35
- "Storage Class Preferencing" on page 35
- "MVC Warranty Expired SPE" on page 35

800 Mb VTVs

Prior to VTCS 6.0, VTVs were restricted to a maximum size of 400 Mb. With VTCS 6.0, you can use the MAXVtvsz parameter of the MGMTclas statement to specify the maximum VTV size in megabytes (400 or 800). This feature allows you to fine tune maximum VTV size for the needs of an application or type of data...especially data that you are migrating to VSM from another vendor. You must use the CDSLEVEL(V6ABOVE) parameter of the CONFIG statement to enable the CDS for this feature.

Note: The MAXVtvsz parameter applies to **only** VSM3s, VSM4s, and VSM5s . **Also note that** 800 Mb support requires VTCS 6.0 **plus** the following:

- L1H11WY **plus** one of the following:
  - For VSM3s: microcode level N01.00.69.04 **or** microcode level N01.00.71.00 and above
  - For VSM4s: microcode level D01.00.04.03 **or** microcode level D01.00.06.03 and above

4 VTV Copies

Prior to VTCS 6.0, you could only duplex VTVs. With VTCS 6.0, you can specify up to 4 copies of a VTV. This feature can be useful in DR solutions or to provide addition protection from data unavailability due to media failure. You must use the CDSLEVEL(V6ABOVE) parameter of the CONFIG statement to enable the CDS for this feature.

VTSS Preferencing

With VTCS 6.0, you can use the VTSSSEL statement to define a VTSS usage rule that applies to the VTSS list and its preferencing specified on a referenced VTSSLST control statement. You can use this feature to preference the VTSS(s) that you want to use for the following functions:

- VTV scratch and specific mounts
- Demand recall
- Reclaim
- Drain
- Audit
- Export
- Consolidation

VTSS preferencing also influence RTD selection for the following functions:

- Demand recall
- Reclaim
- Drain
- Audit
- Export
- Consolidation

Storage Class Preferencing

With VTCS 6.0, you can use the STORSEL statement to define a VTSS usage rule that applies to the Storage Class list and its preferencing specified on a referenced STORLST control statement. Storage Classes control where the data is placed (ACS and media). You can use this feature to influence RTD selection for specific mounts for automatic recalls. You can also use Storage Class preferencing to influence MVC selection for the following functions:

- Demand recall
- Export
- Consolidation

MVC Warranty Expired SPE

VTCS detects media warranty expiration and sets the WARRANTY status to ON. Alternatively, you can use SMF or LOGREG data to detect MVCs approaching end-of-life and use the MVCMAINT to manually set WARRANTY OFF. Knowing that the warranty has expired lets you plan for media replacement before media end-of-life occurs. You can also use the MVCMAINT to set WARRANTY OFF for MVCs erroneously marked as warranty expired.

This SPE is also available for VTCS 6.0.

VTCS/NCS 5.1.0 Enhancements

VTCS/NCS 5.1.0 builds on the previous releases to extend the VSM feature/function set. For more information, see:

- "Clustered VTSS Enhancements"
- "Alphabetic Volser Range Support"

Note: VSM4 support **is included in** VTCS 5.1.0 base; no additional enabling PTFs are required to connect and configure VSM4s.

Clustered VTSS Enhancements

VTCS 5.1 provides the following enhancements to Clustered VTSS configurations:

- As with the original Clustered VTSS feature, the Primary and Secondary VTSSs must both be VSM3 systems. For VTCS 5.1.0, however, the Secondary can be of any capacity. All hosts **must** be at VTCS 5.1.0 to enable this feature.
- The Secondary can now receive both replicated VTVs from the Primary and non-replicate production workload by any of the standard routing methods (for example, TAPEREQ statements).
- You can now specify the source VTSS for migration of replicated VTVs on the MIGRATE parameter of the HSC STORclas statement. For example, your Management Class specifies replication and duplexing as shown below:

  MGMTCLAS NAME(VSMDR) REPLICAT(YES)
  MIGPOL(LOCAL,REMOTE)

  You can then use the MIGRATE parameter to migrate the source VTV from the Primary VTSS to the local ACS and the replicated VTV from the Secondary VTSS to the remote ACS as follows:

  STORCLAS NAME(LOCAL) ACS(00) MIGRATE(PRIMARY)
  STORCLAS NAME(REMOTE) ACS(01) MIGRATE(SECONDARY)

Alphabetic Volser Range Support

Prior to Version 5.1, HSC and VTCS supported only numeric volser ranges. In 5.1, HSC and VTCS now support alphabetic ranges. Note that the 5.1 support is for alphabetic ranges *only*, not for alphanumeric ranges. An alphabetic volser range consists of a pair of volsers (start volser and end volser) containing an incrementing alphabetic portion of 1 to 6 characters. For example:

00000A-00000Z, ABCAAA-ABCZZZ, 9AA000-9CC000, A00A00-A00M00

VTCS RECLAIM Enhancements SPE

This SPE adds the following parameters to the RECLAIM command/utility:

- ACSid to reclaim the eligible MVCs in a single ACS.
- MVC to reclaim the reclaim the specified MVC(s).

T9840C RTD Support SPE

This SPE lets you use the T9840C transports as RTDs. T9840A/T9840B and T9840C transports use the **same physical form factor** but **different recording techniques** resulting in the following restrictions:

- T9840Cs can read from media written to by T9840As/T9840Bs, but cannot write to T9840A/T9840B media **unless** the entire volume is rewritten from beginning of tape.
- T9840As and T9840Bs cannot read from or write to media written to by T9840Cs.

To ensure media and transport compatibility, you **must use** separate VOLATTR statements to segregate T9840A/T9840B and T9840C media. For more information, see the *VTCS 5.1 Installation and Configuration Guide*.

Note: T9840C support is available **only** for VSM3s, VSM4s, and VSM5s, **not** for VSM2s.

As with other supported RTD device types, you can share the T9840C dynamically between VTSSs, but only manually between VSM and MVS.

T9940B RTD Support SPE

T9940A and T9940B transports use the **same physical form factor** but **different recording techniques** as follows:

- T9940Bs can read from media written to by T9940As, but cannot write to T9940A media **unless** the entire volume is rewritten from beginning of tape.
- T9940As cannot read from or write to media written to by T9940Bs.

To ensure media and transport compatibility, you **must use** separate VOLATTR statements to segregate T9940A and T9940B media. For more information, see the *VTCS 5.1 Installation and Configuration Guide*.

Note: T9940B support is available **only** for VSM3s, VSM4s, and VSM5s, **not** for VSM2s.

As with other supported RTD device types, you can share the T9940B dynamically between VTSSs, but only manually between VSM and MVS.

VTCS Merge Manifest SPE

This SPE lets you merge manifest files created by EXPORT. A merged manifest file can be specified as input to the following:

- IMPORT
- MVCMAINT
- MVCRPT
- VTVRPT

This SPE is also available for VTCS 4.0, 4.1, and 5.0.

NCS 5.0.0 Enhancements for VSM

NCS 5.0.0 provides the following enhancements for VSM:

- For HSC 5.0 and above, the new SCRDEF command lets you dynamically reload VTV scratch subpools.
- For HSC 5.0 and above, the Warn SCRatch, Display SCRatch, and Display THReshld commands are enhanced to let you manage and monitor scratch VTVs.
- For HSC 5.0 and above, you can use the CDS EXPAND command to dynamically expand the CDS.
- For HSC 5.0 and above, you can use the new TRACELKP command to trace HSC definition data sets.
- For HSC 5.0 and above, the Batch API provides additional data in the VTV and MVC records returning from a Batch API Query CDS request.
- For HSC 5.0 and above, you can specify that Nearline transports can only be used as RTDs via the VSMONLY=YES and the ACSDRV parameters of the LIBGEN SLIACS macro or the SET ACS utility.
- NCS/VTCS 5.0 and above does **not** allow allocation of unlabeled tapes to VTVs. Unlabeled VTVs can cause the following for scratch VTV allocation requests:
  - If your JCL specifies a virtual esoteric, the NCS Storage Management Component (SMC) fails the allocation.
  - If you have a default esoteric such as CART and specify allocation to virtual (via TAPEREQ or HSC User Exit), the allocation will go to a non-virtual device.
- HSC 5.0 and above provides two new settings for the LIBGEN DELDISP parameter of the SLILIBRY macro:

  ASCRTCH
  : (All scratch). Both real tape volumes and VTVs are made scratch if they were mounted scratch and the disposition on the dismount message is delete ('D').

  VSCRTCH
  : (Virtual scratch). Only VTVs are made scratch if they were mounted scratch and the delete disposition on the dismount message is delete ('D').

  The current DELDISP settings (SCRTCH and NOSCRTCH) define scratch handling at dismount **only** for real volumes. In an HSC 5.0 system, if DELDISP is set to either of these values, VTVs are **never** scratched at dismount.

  On pre-5.0 HSC systems, acting against a CDS where the DELDISP is set to ASCRTCH, the setting is handled as SCRTCH. Similarly, if the DELDISP has been set to VSCRTCH, the setting is handled as NOSCRTCH.

A LIBGEN and MERGECDS or RECONFIG utility is not required to change the DELDISP setting in an existing system. The DELDISP setting can be changed with the HSC SET utility; for 5.0 it accepts the two new settings. Active systems must be recycled to affect the change for DELDISP. Once set, the DELDISP setting is persistent across HSC initializations.

- MVS/CSC 5.0 and above also provides the new ASCRTCH and VSCRTCH settings for the DELDISP startup parameter which is specified in a sequential file (usually a PDS member) at initialization:

  The current DELDISP settings (SCRTCH and NOSCRTCH) define scratch handling at dismount **only** for real volumes. In an MVS/CSC 5.0 system, if DELDISP is set to either of these values, VTVs are **never** scratched at dismount.

  Each MVS/CSC system can define its own startup parameter file and can have different settings for DELDISP. A recycle of an MVS/CSC system is not necessary to change the DELDISP setting. The MVS/CSC ALTER command can change the setting for DELDISP; for 5.0 it accepts the two new settings. When changing the DELDISP setting via the ALTER command, it goes into affect immediately for that MVS/CSC system. However, if the MVS/CSC is recycled, the DELDISP setting is set to the value defined in the startup parameter file; if omitted it defaults to NOSCRTCH.

VTCS 5.0.0 Enhancements

VTCS 5.0.0 provides the following enhancements:

- The VTCS Programmatic Interface (PGMI).

  VTCS 4.0.0 and below does not allow StorageTek partners and other Independent Software Vendors (ISVs) to interface directly to VTCS. The VTCS 5.0 PGMI provides this interface, which has following benefits:

  - New VTCS functions and updates to functions are immediately made available to the PGMI users.
  - All interfaces to VTCS (commands, utilities and PGMI calls) are now processed through a common interface and have the same syntax.
  - The PGMI provides the following outputs from PGMI calls:
    - Free format output so that a PGMI user can easily reproduce the same output as a VTCS command or utility.
    - Structured output (in XML format) to enable easy programmatic manipulation of the PGMI responses.
  - New feature/function updates to existing VTCS commands or utilities are immediately made available to the PGMI users.
- Enhancements to MVCMAINT, which consist of two new parameters to set MVC error status and “logical eject” status.
- Enhancements to VTVMAINT, which consist of two new parameters to unlink (remove) VTVs from MVCs and set the VTV Management Class.
- The addition of Vary CLInk, which allows you to change Cluster link (CLINK) states. For example, if a CLINK fails or requires service, you can enter a VT Vary CLINK OFFline command to vary the CLINK offline. You enter a a VT Vary CLINK ONline command to vary the CLINK online.
- The RTV utility VALIDATE parameter SPE. The VALIDATE parameter validates that the RTV utility can decompress the specified VTVs without performing the decompression.

See Also

- “What is VSM?”
- “Key VSM Benefits and Features”
- “The Clustered VTSS SPE”
- “Key VSM Benefits and Features”
- “Clustered VTSS Requirements”

NCS 4.0.0 Enhancements for VSM

The HSC 4.0.0 FEATures control statement specifies which features you have enabled. The Basic Management Feature (which is included with HSC 4.0.0 base) enables the HSC MGMTclas control statement NAME, ACSlist, IMMEDmig, DUPlex, and DELSCR parameters.The Advanced Management Feature, which is available as a chargeable feature, enables the Basic Management Feature plus:

- The HSC STORclas control statement.
- The MGMTclas control statement MIGpol, RESTIME, CONSRC, CONTGT, and REPLICAT parameters.
- The VTCS 4.0.0 EXPORT and IMPORT utilities.

The Advanced Management Feature, therefore, extends VSM Management Classes and introduces VSM Storage Classes. The EXPORT and IMPORT utilities let you create portable MVCs to move data from one VSM system to another.

Using the Advanced Management Feature, you can:

- Control media selection for migration, reclaim, and consolidation.
- Select the location for consolidation input.
- Specify preferred VTSS residency time for VTVs.
- Group or segregate workloads on MVCs.
- Quickly move data on VTVs and MVCs between VSM systems.
- Create and manage Clustered VTSS configurations.

For complexes with greater than 16 MVS hosts, NCS 4.0.0 lets MVS clients route data to VSM systems. Customers can use the new LibraryStation VIRTACS statement to define a virtual ACS that maps to a Virtual Tape Storage Subsystem (VTSS). The 4.0.0 MVS/CSC TAPEREQ statement adds the keyword of Virtual to the MEDia, RECtech, and MODel parameters to route data to VSM. Customers can also use the new ESOTeric or MGMTclas parameters to route data to VSM with even more control than the Virtual keyword provides.

NCS 4.0.0 and VTCS 4.0.0 allow non-MVS/CSC 4.0 clients to request VTV mounts when you configure the system as follows:

1. Define a virtual ACS using the LibraryStation VIRTACS statement.
2. Define an HSC subpool that contains VTVs.
3. Define a LibraryStation subpool that corresponds to the HSC subpool in Step 2.

The Named MVC Pools feature lets you, via the MVCPOOL NAME parameter, associate a name with specified MVCs that form a subset of the system MVC Pool. You can also use new MVCPOOL parameters to specify MVC reclamation values for a Named MVC Pool. You specify a Named MVC Pool on a Storage Class, and you specify this Storage Class on a Management Class that you specify on your job routing parameters.

The following are also HSC 4.0.0 enhancements:

- The Batch API supports bulk reading of CDS VTV and MVC records.
- The HSC MOUNT command can now mount a scratch or specific VTV on a VTD and optionally assign a Management Class to the VTV.
- The Programmatic Interface MOUNT request now supports an additional parameter of MGMTCLAS to assign a VSM Management Class to the VTV.
- The HSC MERGEcds utility SPE provides the DELVirt parameter. Specifying the DELVirt parameter lets you merge CDSs with virtual records without copying volsers of uninitialized or empty VTVs and MVCs to the target CDS.
- The following SMF record enhancements SPEs:
  - SLSSMF14 - HSC SMF Subtype 14 Record is enhanced to record a VTV dismount request.
  - SLSSMF 27 - HSC SMF Subtype 27 Record is enhanced to record a VTV scratch event.
  - SLSSMF29 - New HSC SMF Subtype 29 Record records a VTV and MVC unlink event

VTCS 4.0.0 Enhancements

VTCS 4.0.0 provided the following enhancements:

VT Vary VTSS Command, Clustered VTSS, and VT QUery Command Enhancements SPEs

The Clustered VTSS SPE enables *Clustered VTSS configurations*, which consist of a *Primary VTSS* and a *Secondary VTSS* connected by one or more Nearlink connections (cluster links). Figure 4 shows a single ACS clustered VTSS configuration and Figure 5 on page 27 shows a dual ACS clustered VTSS configuration. 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.

The Vary VTSS SPE provides the VT Vary VTSS command to vary a VTSS to online, offline, or quiesced state. This command is used while doing maintenance to a VTSS or in the event of a VTSS failure. This SPE also enhances the VT QUery VTSS command output to show the VTSS state. The Vary VTSS SPE is a prerequisite for the Clustered VTSS SPE.

The VTCS 4.0.0 VT QUery command SPE adds the following capabilities:

- Detailed VTSS output, which includes host-to-VTSS access information and NOMIGRAT and NORECLAM settings
- VTD location, device number, and status
- CONFIG parameter settings
- Migration summary and detailed statuses
- Task status
- Lock status

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.

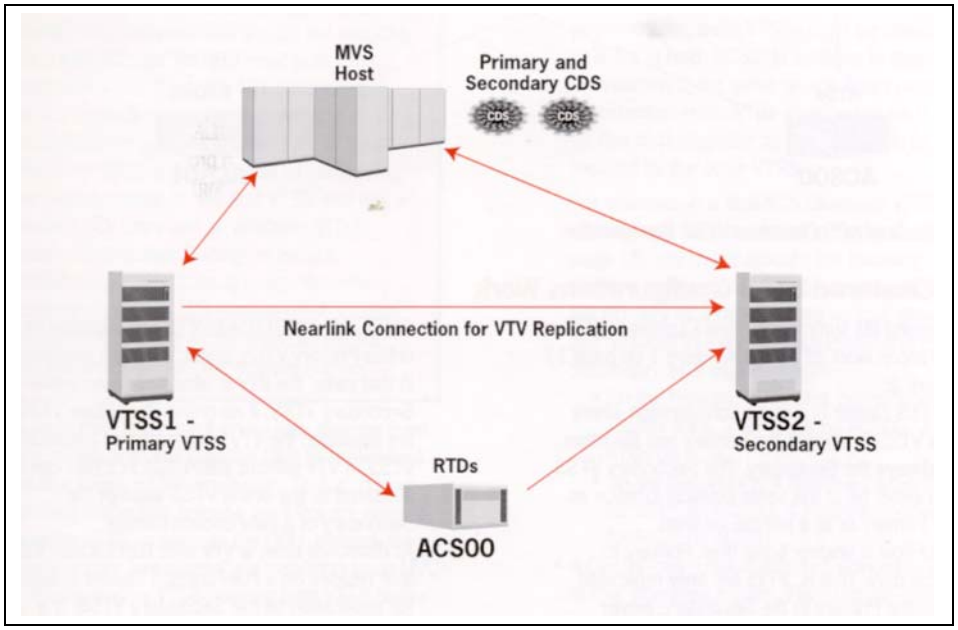

Figure 4. ***Single ACS Clustered VTSS Configuration***

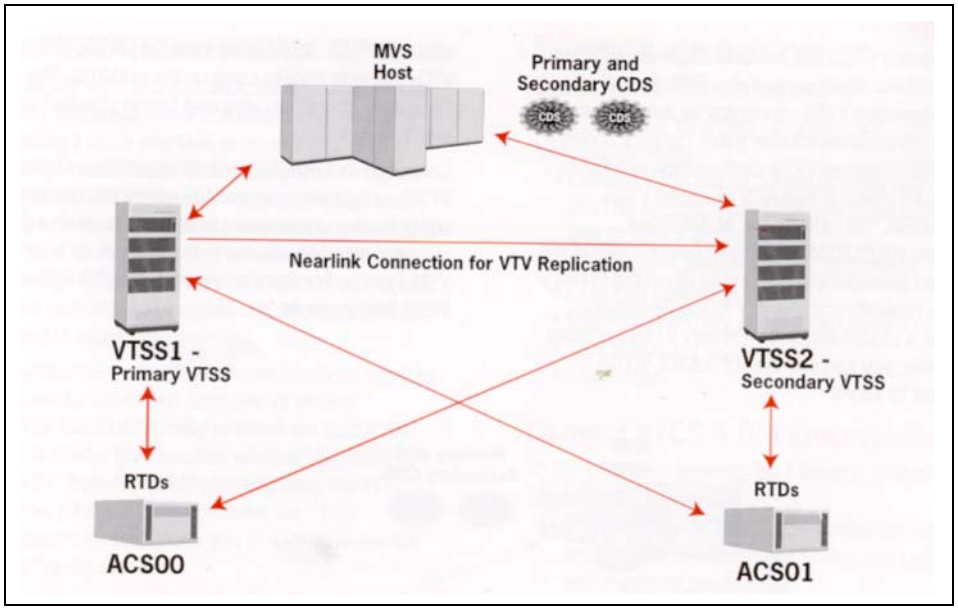

Figure 5. ***Dual ACS Clustered VTSS Configuration***

T9940A RTD Support SPE

This SPE lets you use the T9940A as an RTD. Valid MVC media types are STK2 and STK2P (both currently 60Gb), and these media types are, by default, at the bottom of the migration output MVC hierarchy and at the top of the reclaim MVC output.

As with other supported RTD device types, you can share the T9940A dynamically between VTSSs, but only manually between VSM and MVS.

9840/9940 Media Management SPE

The 9840/T9940 media management SPE includes the following:

- VTCS automatically detects media end-of-life and sets the RETIRED status to ON. Alternatively, you can use SMF or LOGREG data to detect MVCs approaching end-of-life and use the MVCMAINT utility to manually set RETIRED ON.

  You can also use the MVCMAINT utility to set RETIRED OFF for MVCs erroneously marked as retired.
- VTCS automatically detects an invalid Media Information Region (MIR) and sets the INVLDMIR status to ON. You can recover the MIR by using either the utility available through the operator panel for the transport or by using the utility available through MPST. After you recreate the MIR, you can use the MVCMAINT utility to set INVLDMIR OFF for the MVC.
- If an MVC is retired or has an invalid MIR, it is depreferenced for migrate processing. For recall processing, an invalid MIR will cause the alternate MVC (if available) to be selected.
- The MVC report, MVC Pool Report, and QUery MVC command report "retired" and invalid MIR status.

ANSI Label Support for VTVs SPE

ANSI label support allows non-MVS/CSC clients, such as the Unisys CSC, to use ANSI labels for VTVs. With ANSI label support, if you define a subpool as ANSI label to HSC, VTCS will mount scratch VTVs as ANSI labeled tapes from this subpool. This feature is only valid, however, for VSM3s with microcode level 44 or higher installed.

MVCMAINT Lost MVC SPE

This SPE lets you set the "lost" status of an MVC, which is required to implement the VSM Offsite Vault DR Feature.

Migration and Recall Enhancements SPE

The VTCS 4.0.0 migration and recall enhancements include the following:

- Changes in migration processing; for more information, see "Creation and Migration" on page 12.
- Changes to the CONFIG VTSS MAXMIG parameter and the new CONFIG VTSS MINMIG parameter to specify the maximum and minimum number of concurrent automatic migration, immediate migration, and migrate-to-threshold tasks for each VTSS.
- The MIGRATE and RECALL utilities, which complement the VT MIGrate and VT RECAll commands by providing a batch interface to do demand VTV migrations and recalls.
- VT SET MIGOPT command that you can use to dynamically set the MAXMIG and MINMIG values for a specific VTSS or for all VTSSs.

Drain/Reclaim Processing SPE

With this SPE installed, VSM maintains data integrity and availability when moving VTVs from one MVC to another during a drain process (MVC drain or space reclamation) 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 duplexed and one MVC is in error or lost, VTCS recalls the VTV from the 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.

SPE for Maximum MVCs Concurrently Processed for Reclamation and Drain

This policy specifies the maximum number of MVCs concurrently processed for reclamation and drain. Use the CONFIG RECLAIM CONMVC parameter to set this policy. Valid values for the CONMVC parameter are 1 to 99. The default is 1. You can also use the following to override the value specified on the CONFIG RECLAIM CONMVC parameter:

- MVCDRAIN utility
- VT MVCDRAIN command
- RECLAIM utility
- VT RECLAIM command

Recall VTVs Read Data Check SPE

By default, VTCS recalls VTVs with read data checks. With this SPE installed, you can specify whether VTCS recalls VTVs with read data checks on the RECALWER parameter of the CONFIG GLOBAL statement.

You can also override the CONFIG GLOBAL RECALWER setting with the following:

- RECALL utility
- VT RECALL command
- MVCDRAIN utility
- VT MVCDRAIN command
- CONSolid utility
- EXPORT utility

MVC and VTV Volser SPE

This SPE removes the restriction of 198 total volume ranges for MVCs and VTVs and the restriction on single volume ranges

CONFIG Utility Enhancement

VTCS 4.0.0 also provided the following enhancements to the CONFIG utility:

- The CONFIG NORECLAM parameter now applies only to reclaim processing (not reclaim and migration processing).
- You can respecify existing VTV and MVC ranges in a CONFIG update run. You can therefore use the DECOM utility output as input to CONFIG without removing existing ranges.

VTVMAINT Utility SPE

The VTVMAINT utility lets you logically dismount VTVs in an offline VTSS. If a VTV is mounted when a VTSS goes offline and a copy of the VTV exists on an MVC, VTCS will not recall the migrated VTV to an alternate VTSS because the VTV is in mounted status on the offline VTSS. In this situation, you can use the VTVMAINT utility to logically dismount VTVs in the offline VTSS (turn off the "mounted" bit in the CDS), then recall the VTV to an alternate VTSS. VTCS records each successful VTV dismount in the SMF14STA field of the SMF Subtype 14 record.

RTV Utility Enhancements SPEs

The RTV utility enhancements include the following new parameters:

- ALLVTVS - converts all VTVs on a specified MVC.
- LISTONLY - lists (but does not convert) the VTVs on a specified MVC.
- FILEnum - lets you specify the VTV to convert by logical data set number on an MVC.
- DUMP - produces a S0C3 abend dump if RTV cannot decompress a VTV.
- VALIDATE - validates that the RTV utility can decompress the specified VTVs without performing the decompression.

The RTV Utility enhancements also allow non-MVS/CSC clients, such as the Unisys CSC, to use ANSI labels for VTVs.

See Also

- “What is VSM?”
- “How Does VSM Work?”

What Enhancements Does VTCS/NCS 6.1.0 Offer?

VTCS 6.1.0 builds on the previous releases to extend the VSM feature/function set. For more information, see:

Bi-Directional VTSS Clustering

Bi-Directional Clusters consist of two *Peer VTSS*s connected by one or more cluster links (CLINKs). Figure 6 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. Bi-Directional Clustering requires an "F" level CDS, which is enabled by specifying V61ABOVE on the CDSLEVEL parameter of the CONFIG utility.

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.

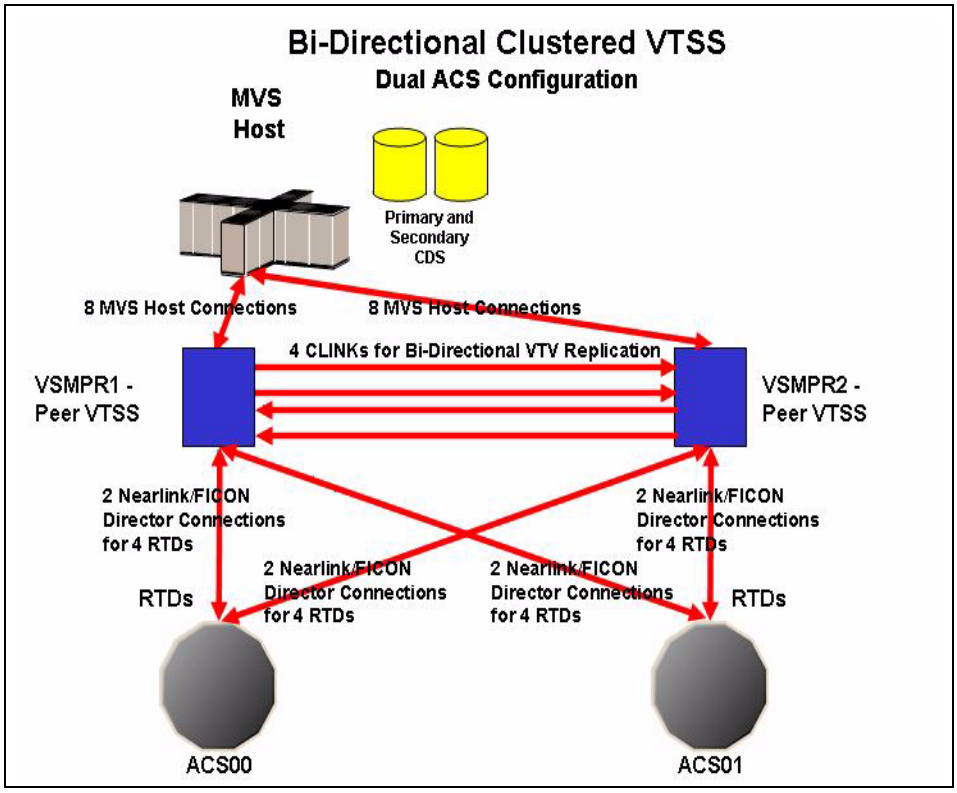

Figure 6. Bi-Directional Cluster attached to a Dual ACS

Near Continuous Operations (NCO)

VTCS 6.1 provides Near Continuous Operations (NCO). Basically, NCO means that you can now dynamically add, change, and delete hardware components, VTVs, and MVCs without having to bring VTCS down and back up again. NCO requires an "F" level CDS, which is enabled by specifying V61ABOVE on the CDSLEVEL parameter of the CONFIG utility. After you have run CONFIG with V61ABOVE, you can add, change, or delete RTDs and VTSSs while the devices are online and the changes take affect immediately without having to recycle VTCS.

Multi-Volume VTVs

Multi-Volume VTV support provides the following enhancements to migrate and recall processing:

- When VTCS auto migration selects candidate VTVs, all VTVs associated with a multiple VTV group that spans a single data set are included within the subsequent migration operation.
- 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.

MVC Media Management SPE

MVC Media Management lets you manage MVCs according to the state of their media as follows:

- You can use the AUDIT utility to audit MVCs with invalid Media Information Records (MIRs).
- You can use the MVCDRAIN utility to drain MVCs with the following conditions:
  - Expired warranty
  - Retired
  - In error
  - Have data checks.

This SPE is also available for VTCS 6.0.

GLOBAL LOCKSTR SPE

This SPE adds the LOCKSTR parameter to the CONFIG utility, where LOCKSTR specifies the name of an MVS Coupling Facility structure to hold VTCS lock data (instead of in the CDS). This configuration can improve performance for large configurations with many hosts. Note, however, that a VTCS locks in a Coupling Facility structure is a solution to the **specific** problem of VTCS causing high I/O demand to the CDS in some configurations, **not** a solution to all CDS performance problems!

This SPE is also available for VTCS 6.0.

DELETSCR Utility SPE

This SPE provides the DELETSCR VTCS SWSADMIN utility. DELETSCR deletes scratch VTVs from VTSSs and unlinks any migrated VTVs from MVCs. The DELETSCR utility, therefore, provides a batch job alternative to setting DELSCR(YES) on the MGMTclas statement.

With DELETSCR, you can specify scratch VTVs by volser, Management Class, and Scratch Pool. The NOTREF parameter specifies days since last referenced to qualify VTV for deletion and you can use the MAXVTV parameter to specify maximum VTVs deleted. DELETSCR produces a report and a DETAIL option is available.

What are typical uses for DELETSCR? Consider running DELETSCR in the following situations:

- Environments where scratch VTVs are not reused frequently.
- Cleanups after Disaster Recovery.

As with DELSCR(YES), if you run DELETSCR, any data on the VTV is gone…immediately.

This SPE is also available for VTCS 6.0.

MGMTclas RESTIME SPE

Valid values for the MGMTclas RESTIME parameter are 1 to 9999. With this SPE installed, the value 9999 now specifies that the VTVs in this Management Class are resident permanently unless VTSS space management requires VTCS to automigrate the VTV and then delete it from the VTSS.

This SPE is also available for VTCS 6.0.

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 Features and Benefits."

What enhancements does the VSM5 provide?

Answer: The VSM5 provides significant throughput, connectivity, capacity improvements over its predecessors, and it also provides 4x the number of VTDs per VTSS and 3x the maximum VTVs per VTSS compared to VSM2s and VSM3s. For more information, see "VSM5: Increased Capacity and Throughput" on page 2.

Does VSM outperform TMM?

Answer: Yes. VSM manages all disk–to–tape data migration without using CPU and I/O channel resources. See "VSM Compared to HSM/TMM" 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 ESCON channels

How do I direct candidate data sets to VSM?

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

- SMC TAPEREQ statements
- The StorageTek DFSMS interfaces
- A combination of TAPEREQ statements and the StorageTek DFSMS interfaces
- HSC and MVS/CSC User Exits

In addition, you can also change your JCL to direct data sets to VSM although StorageTek does not recommend this method.

What hardware and software does VSM require?

Answer: VSM is supported in MVS 5.2.2 and above (JES2 and JES3 environment) and all versions of OS/390. For more information about requirements, see “VSM Requirements”.

Which LSMs does VSM support?

Answer: Any LSM supported by HSC is supported, although the 9310, 9740, and SL8500 are recommended. Currently the following LSMs are supported:

- 4410
- 9310
- 9360
- 9740
- SL8500

For more information about VSM library support, see “VSM Configurations”.

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”.

What physical transports and media does VSM support?

Answer: VSM RTDs can be a mixture of the following drives and media:

- 9490 (TimberLine)

  Media: STANDARD, ECART
- 9490EE (TimberLine EE)

  Media: STANDARD, ECART, ZCART
- T9840A, T9840B, T9840C

  Media: For more information, see Table 9. on page 20.
- T9940A, T9940B

  Media: For more information, see Table 9. on page 20.
- T10000

  Media: For more information, see Table 9. on page 20.

Each VTSS must have a minimum of two library–attached transports for each media type used for MVCs.

What virtual transports and media does VSM support?

Answer: The VTSS emulates 3490E devices and media. For more information, see "VSM Configurations".

How many RTDs are supported in VSM?

Answer: VSM supports up to 8 RTDs per VTSS (for VSM2/VSM3) or up to 16 RTDs per VTSS (for VSM4/VSM5). For more details about VSM library support, see "VSM Configurations".

How many VTVs are supported in VSM?

Answer: VSM supports a maximum of 100,000 (for VSM2s and VSM3s) or 300,000 (for VSM4/VSM5) 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".

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".

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[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".

What happens when high AMT is exceeded?

Answer: VTCS starts automatically migrating 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. For more information, see "Creation and Migration".

Does VSM support the StorageTek Sysplex solution?

Answer: Yes. For complexes with greater than 16 MVS hosts, NCS lets MVS clients route data to VSM systems. For more information, see "NCS 4.0.0 Enhancements for VSM".

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).

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

- **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 features did previous releases of VTCS and NCS provide?

Answer:

Previous releases of VTCS and NCS provided major features. For the 4.0 releases, for example, the VSM Advanced Management Feature includes new parameters for the HSC MGMTclas statement and adds the HSC STORclas statement and the EXPORT and IMPORT utilities. For complexes with greater than 16 MVS hosts, NCS 4.0.0 lets MVS clients route data to VSM systems. VTCS 4.0.0 also provided improvements to VTV migration and recall processing and control, enhancements to the VT QUery command, and general product enhancements such as SMF record enhancements.

Also note that other SPEs, such as the Clustered VTSS SPE, added value to the VSM solution. The Clustered VTSS SPE provides additional business continuance and business resumption capability for VSM.

VTCS 5.0.0 and NCS 5.0.0 provided enhancements to existing function and also added Vary CLInk, which allows you to change Cluster link (CLINK) states.

VTCS 5.1.0 and NCS 5.1.0 provided VSM4 support, enhancements to Clustered VTSS, and alphabetic volser support.

VTCS 6.0.0 and NCS 6.0.0 provided support for 800 Mb VTVs, 4 VTV copies, and Storage Class and VTSS Preferencing, and MVC Warranty Expired detection.

For more information, see "What Features Did Previous Releases of VTCS and NCS Offer?"

What's new for VTCS 6.1.0?

Answer:

VTCS 6.1.0 builds on the previous releases to extend the VSM feature/function set. For more information, see "What Enhancements Does VTCS/NCS 6.1.0 Offer?" on page 51.

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:

- “VT Vary VTSS Command, Clustered VTSS, and VT QUery Command Enhancements SPEs” on page 45.
- “Clustered VTSS Enhancements” on page 36.
- “Bi-Directional VTSS Clustering” on page 52

See Also

- “What is VSM?”
- “Key VSM Benefits and Features”
- “How Does VSM Work?”
- “VSM Compared to HSM/TMM”
- “What Features Did Previous Releases of VTCS and NCS Offer?”
- “What Enhancements Does VTCS/NCS 6.1.0 Offer?”
- “VSM Requirements”
- “VSM Configurations”

VSM Configurations

VSM consists of host software (VTCS), VTSSs, and use of tape transports (RTDs) attached to an ACS. VSM supports various configurations of VTSSs to hosts and VTSSs to RTDs.

VSM Configurations

VTSS to Host and VTSS to RTD Connections

As described in "VSM5: Increased Capacity and Throughput" on page 2, FICON interfaces have greatly increased the connectivity and throughput of the VSM5.

Multiple LPARs Sharing a Single VTSS

Note that when multiple LPARs share the management of a single VTSS, each LPAR has an HSC and a VTCS, and each LPAR is responsible for managing the VTV requests, including recalls. VTV migration is controlled by the LPAR that first detects that the high AMT is reached or exceeded. The migration process then belongs to that LPAR without the intervention of others. If another migration task is required, it can be run by any LPAR that first detects that the high AMT has been reached or exceeded.

- VTV migration can be run from any LPAR except from those that are specifically excluded from participating in the migration process.
- Each VTSS can be connected to a maximum of 8 RTDs (for VSM2/VSM3) or 16 RTDs (for VSM4/VSM5). The CONFIG utility will not allow fewer than 2 RTDS per ACS. CONFIG cannot check device type, but StorageTek s**trongly recommends** at least two RTDS of each device type in each ACS to which the VTSS is attached....otherwise, you can seriously compromise error recovery and also impact the efficiency of space reclamation.
  - In direct connect mode, each of these RTD connections uses one of the VTSS ESCON ports.
  - In direct connect mode, a dual-ESCON port RTD can attach to two VTSSs.
  - If attaching using an ESCON Director, a single RTD can connect to an ESCON port on multiple VTSSs.
  - If attaching using ESCON Director, an RTD can connect to 64 VTSSs.
- RTDs can be shared in two ways:
  - Two or more VTSSs can share one RTD. VTCS manages the sharing automatically.
  - An MVS host and a VTSS can share one RTD, but the user must manage the sharing manually.

LSM and RTD Configurations

- "Nearline Hardware Requirements" describes the LSM and RTD types that VSM supports.
- The same LSM can have RTDs attached to different VTSSs.
- An LSM can contain both Nearline volumes and VSM MVCs, but these are *not* shared by VSM and conventional Nearline processing.

Virtual Transport, Media, and Volume Support

VSM supports the following device, media, and tape volume emulation:

- VTSS emulation of 64 3490E devices (for VSM2s and VSM3s) or 256 3490E devices (for VSM4s and VSM5s).
- A maximum of 100,000 VTVs per VTSS (for VSM2s and VSM3s) or 300,000 VTVs per VTSS (for VSM4s and VSM5s); varies depending on the size of the VTVs, capacity of the disk buffer, and size of the CDS. An unlimited number of VTVs can reside on MVCs, depending on total MVC space.
- VTSS emulation of standard length cartridges. Physical end of tape occurs when compressed data has been received for a VTV per its size. For example, for an 800 MB VTV, at a 4:1 compression rate, this would be 3200 MB effective capacity.

VSM Terminology

The following table describes the terms and abbreviations associated with VSM.

Term	Description
ACS	Automatic Cartridge System.
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.
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.
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.

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.
MVC Report	This report displays the status of your system's MVCs and has summary and detailed options.
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.
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.
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 or 256 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 64 or 256 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.
VTV Report	This report displays the status of your system’s VTVs.