Virtual Storage Manager

Stephen H. Blendermann
Storage Technology Corporation
Louisville, Colorado 80028-6202
steve_blendermann@stortek.com
Voice: +1 303 673-3184
Fax: +1 303 661-8462

Dennis F. Reed
Storage Technology Corporation
Louisville, Colorado 80028-6202
dennis_reed@stortek.com
Voice: +1 303 673-6491
Fax: +1 303 661-8462

Abstract: Storage Technology Corporation (StorageTek®) has developed Virtual Storage Manager™ (VSM) as a solution to the problems of inefficient use of tape media and tape transports. VSM solves the media use problem by stacking multiple user tape volumes onto a single, physical tape cartridge. The stacking operation occurs transparent to the user and the actual movement of data to accomplish volume stacking occurs outboard of host systems being serviced. VSM solves the tape transport use problem by emulating more tape transports than physically exist. VSM satisfies peak concurrent transport demand without requiring additional transports that would not be needed or used at other times. Physical tape transports are attached to an Automated Cartridge System (ACS). The solution is comprised of hardware, microcode, host software and PC software.

History

In the late 1970s, StorageTek developed an outboard Virtual Storage System (VSS). This system consisted of host-based software, a 370/138 class processor running special software known as the Virtual Control Processor (VCP), a Channel Adapter (CA) with built-in data compression for connecting the VCP to the host, and 3350-class DASD and 3420-class tape devices connected to the VCP. The VSS host software system controlled allocation and provided the capability to create virtual volumes. The data was compressed in the CA, re-blocked in the VCP's memory and then written to the attached DASD. The data was automatically backed up to tape and migrated and recalled as necessary.

The technology available at the time did not allow this product to be brought to the commercial market. However, in 1984, a patent[1] was issued to Storage Technology for an outboard storage system that used virtual volumes, built-in data compression, and automatically managed disk space by moving data between the disks and tape.

With the advent of automated tape libraries and fault tolerant disk subsystems developed by StorageTek, the technology is now available to bring a virtual storage system to the commercial market. At the same time, tape cartridge capacities have increased without a corresponding method to utilize the full capacity of the cartridges. Virtual Storage Manager provides the ability for applications to make use of these cartridges.

In 1995, a small team was chartered with building a prototype of a VSS to demonstrate the feasibility of the concept. That team modified Iceberg microcode and HSC software to produce a working system that presented a 3490E tape transport image to the host. The

host code managed the disk buffer and migrated/recalled virtual volumes to/from tape. In the prototype, manual tape transports were used as StorageTek had been automatically mounting cartridges since 1987 and it was not necessary to prove that could be done.

In conjunction with the prototype, requirements for virtual storage system were gathered from several customers around the world. Those requirements were the basis for the initial functions of the system. These functions were presented at a Customer Advisory Board for validation and further refined to ensure that the product met customer needs.

On October 13, 1997, Storage Technology announced the Virtual Storage Manager.

Solution - Overview

VSM is comprised of hardware, microcode, host software and PC software. VSM hardware and microcode present the image of 64 IBM 3490E tape drives to one or more ESCON-attached MVS-based hosts. This hardware includes RAID-6 disk storage that serves as a buffer for holding user-created (virtual) tape volumes (VTV).

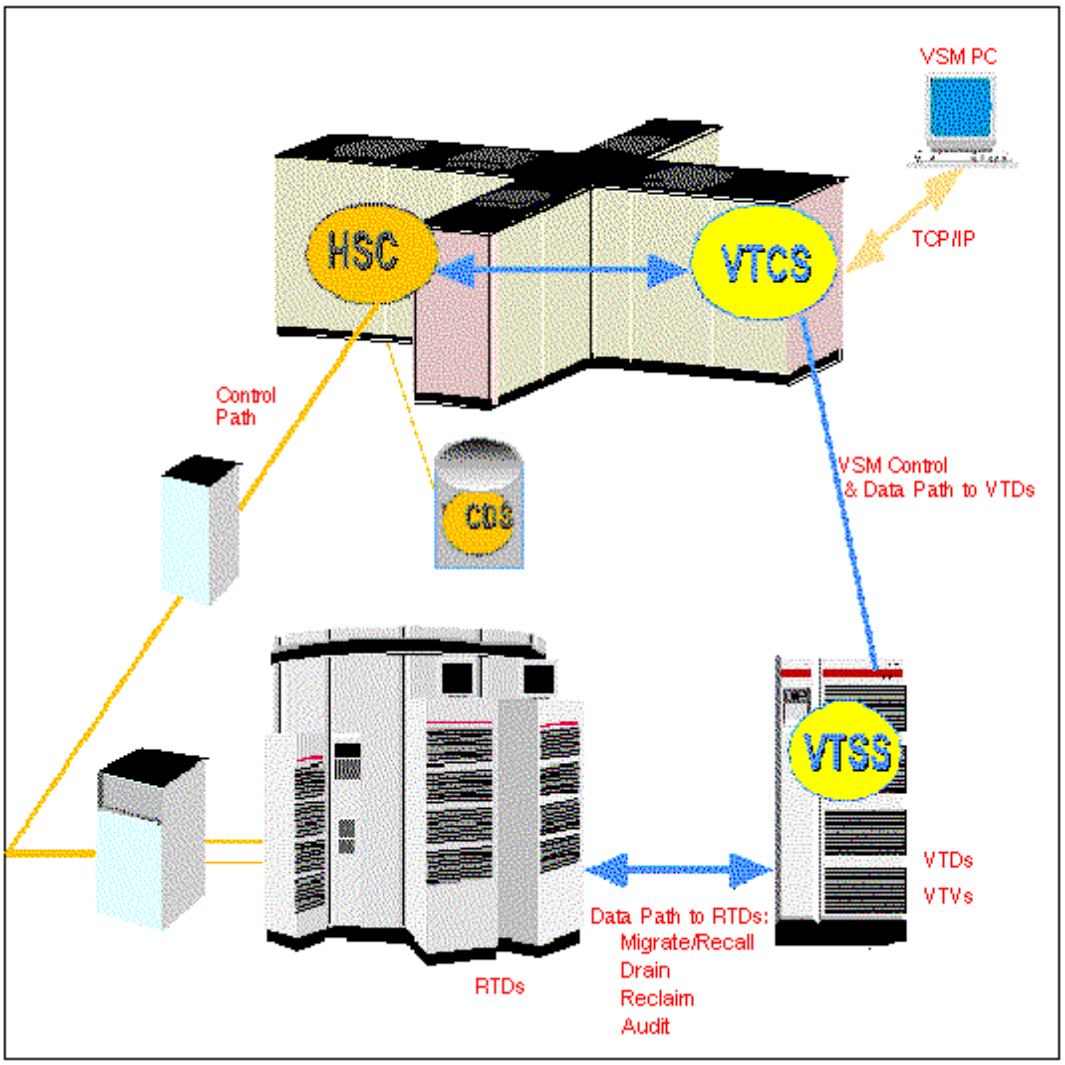

Figure 1. VSM Solution

The VSM host software directs tape jobs to the VSM hardware, influences the allocation of the virtual 3490E tape drives (VTD), and instructs the VSM hardware to perform mounts of virtual tape volumes on the appropriate virtual tape drive (VTD). The VSM host software manages the disk buffer, and instructs the VSM hardware to migrate virtual tape volumes to real tape cartridges, stacking several VTVs onto a single, physical (multi-volume) cartridge (MVC). When needed to satisfy a tape job, VSM host software automatically instructs the VSM hardware to recall a VTV from a MVC back into the VSM hardware disk buffer. VSM host software maintains the locations of all VTVs, whether resident on a VSM hardware disk buffer or a MVC. All potential MVC contention is mitigated by VSM host software. All real tape drives (RTD) used by VSM hardware to satisfy VTV migrate and recall actions, are library-attached to fully automate the solution. VSM PC software provides a Windows 95 or NT platform with a graphical user interface for administration and reporting of VSM functions, including installation, configuration and maintenance

The features of VSM include:

- Efficient utilization of tape media by stacking multiple (user) VTVs onto a physical MVC.
- Emulation of multiple, host-compatible tape control units and drives, backed by a high-performance RAID-6 disk buffer, to satisfy peak tape drive allocation as well as deliver superior VTV read access when resident in a disk buffer.
- Efficient data throughput to minimize the number of physical tape drives needed to satisfy a larger number of emulated tape drives.
- Transparent user access to tape data, regardless of physical data location, either on a disk buffer or a MVC.
- Mapping of user VTV data from emulated, host-compatible devices to other storage device technologies.
- Support for multiple physical cartridge and RTD technologies.
- Outboard (of host) movement of compressed VTV data for migration and recall.
- Increase tape automation capacity. All RTDs are library-attached. The library may be shared with non-VSM activity.

VSM is a natural extension of StorageTek's core technologies: tape automation and device virtualization. The combination of these technologies yields a solution that can dramatically lower the cost of storage and storage management. VSM takes advantage of current, underutilized device and tape cartridge technologies (investment protection), as well as providing a seamless technology migration path for future device and media technologies. VSM further increases the effectiveness of existing tape automation, while saving on precious resources such as floor space, tape racks and device power/cooling.

VSM Hardware

VSM hardware consists of one or more virtual tape subsystems (VTSS™), which are built on the Iceberg® virtual storage platform. Each VTSS:

- Emulates either 32 or 64 virtual 3490E tape drives.
- Emulates standard 3480 cartridges (400 MB capacity after compression). A single VTSS can support up to 100,000 VTVs resident in the disk buffer.

- Supports either 8 or 16 ESCON ports which are configured as host-attached control unit ports or RTD-attached I/O channels.
- Includes a RAID-6 disk buffer with effective user data capacities ranging from 180 GB to over 930 GB.
- Attaches 2 to 8 RTDs, which must be library-attached on a single StorageTek ACS, and can be either TimberLines or RedWoods or a mix of the two.
- Moves data between the VTSS disk buffer and RTDs without using host processing cycles or channel bandwidth.
- Supports all TimberLine® and RedWood® media types, and can be mixed.
- Provides the throughput performance, fault-tolerance and non-disruptive hardware serviceability of the latest generation of Iceberg hardware and microcode.

VSM tape device virtualization, disk buffer and outboard data movement are accomplished on an Iceberg virtual disk array platform with new microcode. This platform provides a robust engine, both in performance and durability, to meet the demands of large tape data processing centers. A VSM solution consists of host software and one or more virtual tape subsystems (Iceberg).

VSM Host Software

Virtual Tape Control Software (VTCS™) is the VSM host software that works as an extension to StorageTek's ACS host software component (HSC). VSM virtual tape processing may be fully integrated with other automated and manual tape processes, including sharing access to ACS robotics with non-VSM activity. VTCS functionality:

- Extends the current HSC database to track all key information concerning VTVs, RTDs, MVCs and VTSSs. This includes maintaining all configuration and operating parameters for the VSM solution.
- Manages the location of all VTVs in the VSM solution.
- Influences MVS allocation of VTDs to satisfy tape job requirements.
- Manages migration and recall of VTVs to/from RTDs
  - ⇒ Insures adequate VTSS disk buffer space available to accommodate new data
  - ⇒ Issues robotic commands via HSC
  - ⇒ Issues migrate/recall orders to appropriate VTSS
  - ⇒ Selects candidate VTVs for automatic migration based on LRU and size.
  - ⇒ Optionally migrates VTVs to 2 MVCs to protect migrated VTVs from media failure
- Manages MVC space reclamation as migrated VTVs expire or are updated.
- Manages recovery of lost HSC control dataset (CDS) information if needed. The CDS itself has multiple levels of data protection as currently provided by HSC.
- Manages recovery from RTD or MVC media errors reported by a VTSS. This includes logging of device and media failures.
- Provides a standalone recovery utility to allow retrieving VTVs from MVCs without benefit of a VTSS.

VTCS requires OS/390 1.2 with Open Edition.

VSM PC Software

VSM PC software provides a graphical user interface on a Windows 95 or NT platform to administer the VSM solution. VSM PC software communicates with VTCS over TCP/IP to the MVS host where VTCS is resident. VSM PC software on a single PC can connect to any of the systems VTCS is installed on. Multiple copies of VSM PC software, on one or more PCs, can be active at the same time. VSM PC software functionality includes:

- Setup of VSM configuration parameters.
  - ⇒ Automatic migration threshold as percent of disk buffer fullness
  - ⇒ Generation of new HSC control statements and parameters
  - ⇒ Definition of MVC cartridge pools and VTV volser ranges
  - ⇒ Definition of maximum depth for stacking VTVs to a MVC
- VSM usage reporting
- Demand migration requests.

VSM Configurations

The VSM solution consists of host (MVS) software, PC software, one or more VTSSs and library-attached tape drives.

A VTSS can be connected to a maximum of 15 MVS hosts systems. Hosts may be attached via ESCON Directors. Each VTSS can be attached to a minimum of two tape drives (RTD) up to a maximum of eight tape drives. If an ESCON Director is used between a VTSS and a RTD, it must be a static connection. A dynamic connection between a VTSS and RTD is not supported in VSM release 1.0. RTDs can be statically shared between MVS hosts and VSM, i.e. a RTD varied online to VSM must be varied offline to MVS, and vice versa.

The VSM solution can support multiple VTSSs shared by multiple MVS hosts.

The following are examples of VSM configurations. These examples show simple configurations even though complex configurations are supported. These examples show one or two RTDs, but actual configurations can have up to eight RTDs attached to each VTSS. Each VTSS must have a minimum of two library-attached TimberLine or RedWood tape drives. Each example shows a single LSM, but multiple LSM configurations are supported. All RTD's connected to a specific VTSS must be served by a single ACS. Dual-ported RTDs can be shared by two VTSSs. An ACS can be shared by the VSM solution and non-VSM MVS usage. A LSM can contain MVC cartridges and non-MVC cartridges.

The VSM solution supports the following LSMs via the existing HSC support for these libraries: 4410, 9310 and 9360.

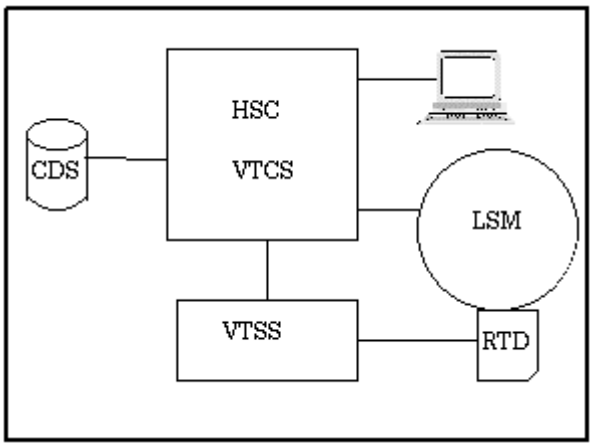

Figure 2. Single Host - Single VTSS

In Figure 2, one VTSS is connected to one host. The VTCS on that host manages all activity of the VTSS.

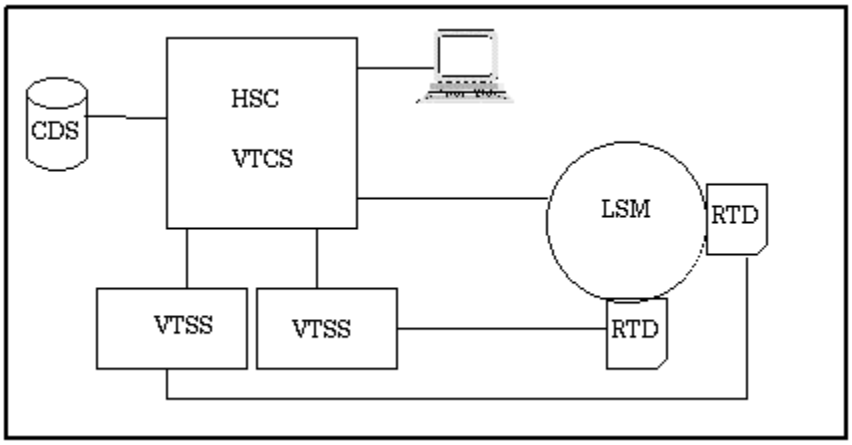

Figure 3. Single Host - Multiple VTSSs

In Figure 3, two VTSSs are connected to one host. The VTCS on that host controls both VTSSs.

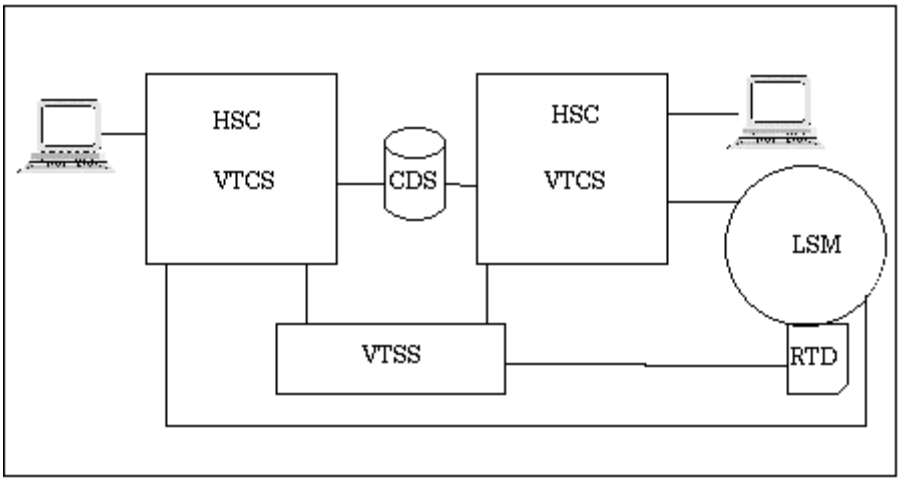

Figure 4. Multiple Hosts - Single VTSS

In figure 4, two hosts share the management of a single VTSS. Each host has a HSC and VTCS, and each host is responsible for managing the VTV requests, including recalls. VTV migration is controlled by the host that first detects the fullness threshold of the VTSS has been exceeded. Up to fifteen hosts can share a VTSS in this configuration.

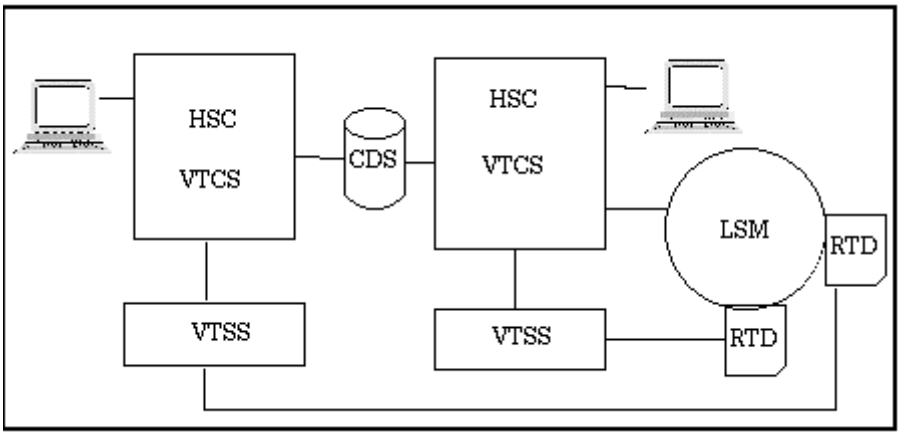

Figure 5. Multiple Hosts - Multiple VTSSs

In figure 5, each host manages its attached VTSS, as in the figure 2 example. If both hosts were attached to both VTSSs, the system would operate like the figure 4 example. Each host would be responsible for its VTV activity, and either host could handle VTV migration on either VTSS.

In this example, the hosts are shown as having separate connections to each VTSS. Each host could be connected to each VTSS and the VTDs shared by those hosts.

Performance

By emulating 64 virtual 3490E tape drives, presenting virtual tape volumes to tape users and caching VTV data on a large disk buffer, VSM optimizes:

- Access time. Virtual mounts occur in less than a second.
- Data throughput. Data is compressed at the Iceberg I/O channel prior to being stored in cache memory, the disk buffer, or migrated to physical tape. Read requests for VTVs resident in the disk buffer are satisfied immediately at disk speeds.
- Physical media utilization. VSM host software stacks multiple VTVs on a single MVC to maximize media utilization, and reduce the number of physical mounts required by the solution.
- Real tape transport requirements and utilization. VSM uses fewer physical drives to satisfy the I/O load handled by the VTDs. VSM requires fewer physical tape mounts to accommodate a larger number of fast, virtual mounts.

Virtual scratch tape mounts are typically satisfied in less than 1 second. The VTSS can support write data rate of 9 MB per second on each of the 8 VTSS data paths at the same time. When a tape mount to read an existing VTV is received, and the data is still resident on a VTSS disk buffer, the mount is satisfied at in less than one second. A VTSS can support read data rate of 10 MB per second on each of the 8 VTSS data paths at the same time.

Conclusion

VSM is another significant advance in StorageTek's successful Nearline® product series. VSM extends the benefits of the unique Iceberg virtualization platform, adds state-of-the-art host software, and dramatically increases the useful storage capacity of existing customer tape libraries. VSM offers equally dramatic improvements in tape job performance while reducing tape management overhead. The benefits that StorageTek Nearline users enjoy today, including leading-edge robotic and tape device technology, are extended for years to come. Seamless integration of future advances in these technology areas is assured. VSM eliminates the constraints of conventional tape subsystems and allows users to fully exploit their existing storage potential.

Futures

VSM hardware and microcode will continue to track and benefit from updates to base Iceberg hardware and microcode. Updates may include:

- Non-disruptive microcode update capability.
- Increased disk buffer capacities.
- Additional channel attachments to non-MVS host platforms.
- Additional RTD device types and associated MVC media types.

VSM Host Software updates may include:

- Full Sysplex support.
- MVC export from VSM and import to VSM.
- Dynamic sharing of RTDs between VSM and MVS.
- User-defined options for controlling how VTVs are grouped on MVCs.
- Support for non-MVS client systems.

References

[1] White, Barry B., United States Patent Number 4,467,421 “VIRTUAL STORAGE SYSTEM AND METHOD”, issued August 21, 1984.

Trademarks