

VIRTUAL TAPE:

IBM AND STORAGE TEK OFFERINGS - AN UPDATE

BY LLOYD CHRISTENSEN

As this article update demonstrates, both IBM and StorageTek have made numerous enhancements to their respective Virtual Tape Server and Virtual Storage Manager offerings.

In the February 2000 issue of *Technical Support*, I wrote about the concepts involved in virtual tape and presented an overview of the IBM Virtual Tape Server and StorageTek Virtual Storage Manager offerings. This article will describe some of the enhancements that each vendor has provided.

IBM HARDWARE

The first component in the IBM VTS configuration is an IBM 3494 Automated Tape Library (ATL). Each 3494 may have up to two attached Virtual Tape Servers. The original VTS used a model B-16 control unit that was one frame in the 3494 ATL. The B-16 included up to 72GB of cache, two ESCON channels, up to 32 logical tape drives, and a minimum of three and a maximum of six Magstar tape drives. The B-16 also had a limit of 50,000 logical volumes per VTS.

Today, the B-18 controller has replaced the B-16 control unit. The B-18 is not part of the 3494 ATL. You can upgrade an existing B-16-based VTS by adding a B-18 and removing the RISC processors and DASD from the existing B-16 unit. In this upgrade scenario you can use the old B-16 as a cartridge storage frame or a drive frame.

Each IBM 3494 ATL can have up to two attached VTS controllers (a combination of a B-16 and B-18 or two B-18s, but not two B-16s). The standard B-18 has two ESCON channels and two RISC

processors. Tape volume cache comes in the following sizes: 72GB, 144GB, 216GB, or 288GB. The standard B-18 has 32 virtual tape drives. Each of these virtual tape drives can have between a minimum of three and a maximum of six Magstar tape drives. The B-18 also

increased the limit for the number of logical volumes to 250,000 per VTS (500,000 per 3494 ATL).

You can configure the B-18 with the Enhanced High Performance Option (EHPO). The EHPO feature allows the enhanced ESCON channels to perform data compression. Without compression, the ESCON channels now offer tape volume cache to increments of 216GB, 432GB, 648GB, or 864GB per B-18. Expect compression rates of 3:1. The EHPO feature allows the number of ESCON channels to be increased from four and the number of virtual tape drives to be increased to 64. The B-18 with the EHPO feature can have a minimum of three and a maximum of six Magstar tape drives.

IBM has also introduced the Performance Accelerator Feature (PAF). The PAF upgrades the B-18 to four RISC processors. This configuration requires four ESCON channels with the EHPO data compression and comes with tape volume cache of either 432GB or 864GB. This configuration also has 64 virtual tape drives, and requires a minimum of four and a maximum of six Magstar tape drives.

IBM supports either the 3590-B or 3590-E tape drive but not a mix of device types. To convert from the 3590-B to the 3590-E drive requires replacing all of the physical drives at one time. IBM does not yet support the extended length cartridge.

In addition to the Export/Import and Peer-to-Peer features, IBM has announced new Advanced Function Features for tape volume cache management.

IBM SOFTWARE

For an additional cost, IBM provides advanced features including additional disaster recovery and tape volume cache management. The disaster recovery enhancements come in two mutually exclusive forms: Export/Import and Peer-to-Peer VTS.

The Export/Import feature allows the storage manager to have the VTS Library Manager move data from a list of logical (virtual) volumes onto a stacked tape. The 3494 ATL ejects the stacked

cartridge so that it can be taken to another site with a VTS and imported.

The Export process requires that you create a list of the logical volumes to export. The list may contain logical volumes going to multiple destinations. The VTS Library Manager that issued the Export command will then mount an empty stacked tape on one of the VTS' tape drives and mount active stacked tapes as needed for a tape-to-tape copy of all logical volumes going to the same destination. As the output volume is marked full, the logical tapes copied to that stacked tape are deleted from the library manager.

Once all logical volumes are copied to the stacked cartridges and the logical volumes are deleted from the sending VTS, the output cartridge is ejected from the 3494. Your tape management system must be aware of the stacked cartridge so that it can be tracked. The stacked cartridge becomes known as a "Container" volume.

A single VTS can only do one Export task at a time. The Export task can only export logical volumes residing within that VTS. If you are considering the Export/Import feature for vaulting data, plan on one minute per logical volume and remember it is single threaded within a VTS. Keep in mind that you will need two VTS tape drives for this process.

The Import process is the reverse of the Export process. You can issue the Library Manager's Import command for the logical volume containing a list of the logical volumes to import. You copy logical volumes on this container to existing stacked cartridges. Then, you eject the container cartridge and return it to the sending site.

The Peer-to-Peer VTS requires the PAF and EHPO channels. The Export/Import feature can not be used in a peer-to-peer VTS. In a peer-to-peer configuration, you have two virtual tape servers residing in different 3494 ATLS within the peer-to-peer distance limitations. The two virtual tape servers are paired, with one being the master VTS and the other being the I/O VTS. You can also call the I/O VTS a User Interface (UI) Library, and each of these libraries is a Distributed Library. Collectively, they are called a Composite Library.

A new hardware frame, the 3494 CX0, which contains two or four 3494 AX0 Virtual Tape Controllers (VTC) is required. The AX0 controller is a RISC-based processor that is used to create logical volume

FIGURE 1: VIRTUAL TAPE STORAGE CAPACITIES

Virtual Tape Storage Capacities				
IBM	Controller	Cache Capacity	Virtual Tape Drives	Virtual Volumes
IBM	B-16	72 GB	32	50,000
IBM	B-18	72/144/216/288 GB	32	250,000
IBM	B-18/EHPO	216/432/648/864 GB	64	250,000
StorageTek	Version	Effective Disk Capacity	Virtual Tape Drives	Virtual Volumes
StorageTek	VSM2-A	180 GB	64	No limit
StorageTek	VSM2-B	324 GB	64	No limit
StorageTek	VSM2-C	468/702/936 GB	64	No limit
StorageTek	VSM3-A	180/320 GB	64	No limit
StorageTek	VSM3-B	460/700/930 GB	64	No limit
StorageTek	VSM3-C	1170/1400/1630/1870 GB	64	No limit

copies and synchronize the coupled VTSes. In the peer-to-peer configuration, each AX0 will have two channels extending from it, one to each attached B-18 control unit. Peer-to-Peer is mutually exclusive of the eight-channel interface feature.

The software that handles most virtual processes, including mount/de-mount of virtual volumes and migration/recall, is Virtual Tape Control System. VTCS is a sub-task of StorageTek's Host Software Component (HSC).

When defining a Peer-to-Peer VTS, you define one VTS as the Master VTS and the other as the I/O VTS. You issue host tape mounts through the Virtual Tape Controller, and the VTC selects the logical volume and logical tape drive. It then selects which distributed library will satisfy the mount. You can assign either of the two distributed libraries to handle any scratch mount. The Distributed Library that satisfies the scratch mount will be responsible for all subsequent references to that logical volume. The User Interface (UI) Distributed Library will be responsible for all insertions for logical volumes.

An implementation option allows the site to distribute the workload between the two distribution libraries. You can specify no preference I/O selection in which the VTC will attempt to balance the workload between them by host read/write/copy

activity and initial mount times (cache and recall times). This option is more common in a local Peer-to-Peer configuration. Alternatively, you can designate one VTS to receive all I/O. This option is most likely to be used in a remote configuration.

There are two copy modes for Peer-to-Peer configurations. One is the Immediate Copy mode. In the Immediate Copy mode, the host application writes a logical volume. At the completion of the rewind/unload CCW, the VTC will make a copy of the logical volume on the alternate virtual tape server. The host I/O is not completed until the copy on the alternate virtual tape server is complete. The Immediate Copy option provides the highest level of data availability but sacrifices throughput.

The alternative to Immediate Copy mode is Deferred Copy. The host I/O is completed upon writing the original logical volume on the primary virtual tape server. The copy operation is queued and completed as a back-end process. This option provides better throughput at the expense of data availability.

In addition to the Export/Import and Peer-to-Peer features, IBM has announced new Advanced Function Features for tape volume cache management. IBM will allow the customer to use the Initial Access Response Time (IART) parameter in the DFSMS Storage Class to indicate a preference for migration of logical volumes. IBM does what they call "Pre-Migration" which is to copy all logical volumes to a stacked tape within minutes of creating/updating a logical volume. Logical volumes are not "migrated" (deleted from the tape volume cache) until necessary.

The new tape volume cache management feature will allow customers to assign DFSMS storage classes that govern logical volume deletion from the tape volume cache. There are two groupings, Group 0 and Group 1. Data sent to a storage class

with an IART value greater than or equal to 100 seconds fall into Group 0 and data sent to a storage class with an IART value less than 100 seconds fall into Group 1. All volumes within Group 0 will be deleted from the tape volume cache before any volume within Group 1. Within each group, migration is done with a least-recently used algorithm.

IBM has tools available that can help you determine which group to assign and build your DFSMS Storage Class routines. An easy first cut at this would be to review your tape management catalog and compare creating job name to last job name, creation date/time to last used date/time, and creation unit to last unit. If they all match, then assign to Group 0, otherwise assign to Group 1.

STORAGETEK HARDWARE

StorageTek's Virtual Storage Manager version 2, now referred to as "VSM Classic," used the StorageTek 9393 RVA disk subsystem. Each Virtual Tape Subsystem (VTSS) could have up to 16 ESCON channels and could perform up to eight concurrent I/Os. They could connect to any one of a variety of StorageTek robotic tape libraries and could use StorageTek's Timberline, Redwood, and 9840 tape drives.

StorageTek now offers Virtual Storage Manager version 3, which uses the StorageTek 9500 SVA disk subsystem. Like the Classic, each VTSS can have up to sixteen ESCON channels. Unlike the Classic, VSM-3 can perform up to 16 concurrent I/Os. Each VTSS can be connected to the same robotic tape libraries as the Classic. VSM-3 can use Timberline, Redwood, 9840, and the new T9940 tape drives. VSM can use a mixture of device types as real tape drives.

The 9500 SVA disk subsystem uses the Serial Storage Architecture (SSA), whereas the Classic's RVA used SCSI drives. The SSA drives are faster, and offer greater reliability and capacity in a smaller footprint. The VTSS itself has a smaller footprint and reduced environmental requirements. VTSS cache has doubled from 1GB to 2GB and non-volatile storage has been increased from 8MB to 32MB. Improved ICE cards and Channel Interface Processors mean the VSM-3 will see a minimum of 50 percent improvement in performance over VSM-2. A 50 percent read/write workload can expect a bandwidth of 120MB/second.

The VSM comes in three base models (Models A, B and C). Within each model, every VTSS is shipped with the maximum disk buffer capacity; however, the customer determines which feature code is enabled and the Customer Support Engineer authorizes the correct amount of disk buffer capacity. Physical Capacity Control lets you change the total amount of disk buffer capacity available. The CSE can authorize increases in the disk buffer capacity without an IML or outage. The Model A can have either 180GB or 320GB of effective disk buffer capacity, while the Model B can have either 460GB, 700GB or 930GB of effective disk buffer capacity. The Model C can have effective disk buffer capacities of 1.16TB, 1.4TB, 1.63TB, or a maximum of 1.89TB. Physical Capacity Control is just one form of non-disruptive microcode

changes now possible. New microcode not affecting CIP code can be upgraded non-disruptively. Effective cache sizes for IBM and StorageTek offerings are shown in Figure 1.

STORAGETEK SOFTWARE

The software that handles most virtual processes, including mount/de-mount of virtual volumes and migration/recall, is Virtual Tape Control System. VTCS is a sub-task of StorageTek's Host Software Component (HSC). StorageTek has upgraded not only HSC but incorporated it into the Nearline Control Solution (NCS), which also may include MVS/CSC and LibStation. NCS is now at release 4.0. HSC 4.0 and VTCS 4.0 can run on either the VSM Classic or VSM-3.

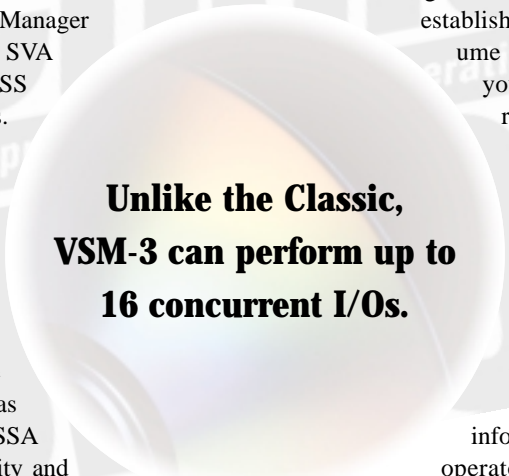
Shortly, after I wrote last year's article, StorageTek announced a new feature: Dual ACS support. With Dual ACS, you can connect a VTSS to two Automated Cartridge Systems in different locations. When you migrate a virtual volume it can be directed to either ACS. Duplexing a virtual volume at migration time provides a migrated copy of the virtual volume in each ACS.

With VTCS 4.0 comes some new features. Among these is a Management class feature that enables the data center to establish desired residency times for a virtual volume to stay on the disk buffer. For example, you can assign a Management class with a residency time of eight hours. VTCS will delete volumes from the disk buffer that have been on the disk buffer for greater than their respective residency times before deleting those that have been on the disk buffer for less than their respective residency time. The residency time feature on a Management class is part of the Advanced Management Feature.

Additionally, StorageTek has added more information to its operator commands. New operator commands are available for functions previously only done through the VTCS Config utility such as setting the number of maximum and minimum migration tasks or changing the migration thresholds for a specific VTSS. StorageTek added a Migration utility that performs the same function as the Migrate command via batch. The utility version can migrate based on cataloged data set names in addition to volser. You can now also perform recalls in batch via the RECALL utility, thus allowing a data center to pre-stage volumes.

StorageTek has enhanced its management of the Multi-Volume Cartridge (MVC). Customers may now specify multiple ranges of MVCs and have data sets migrated to specific MVC ranges. This feature is similar in concept to a tape management system's use of scratch subpools. For more information on how to do this, refer to the *VTCS 4.0 Installation, Configuration, and Administration Guide under STORCLAS and MGMTCLAS*. Using multiple MVC pools requires the Advanced Management Feature.

StorageTek provides the CONSOLIDATE utility. This utility will allow the customer to copy specified virtual volumes, or a VTCS Management class, to an MVC that can then be ejected from the library. The virtual volume, however, remains on the disk buffer and can later be migrated — unlike the IBM Export utility



**Unlike the Classic,
VSM-3 can perform up to
16 concurrent I/Os.**

discussed earlier. VTCS will only keep track of two copies of a virtual volume. A CONSOLIDATED copy counts as one of the two. The consolidated MVC may be ejected from the ACS.


An extension to CONSOLIDATE is the Export utility. Export creates a Manifest file of the virtual volumes being copied to the Export volume. The Export volume is later ejected from the ACS. The exported copy does not count as one of the two copies of a virtual volume tracked by VTCS so it can be used in conjunction with duplexed virtual volumes. Again unlike the IBM Export utility, the original virtual volume remains within the VTSS. The exported MVC can then be processed at a disaster recovery site by using the associated IMPORT utility and naming the appropriate Manifest file.

I view the StorageTek Consolidate and Export/Import utilities in the following way: You can use the Consolidate utility to create a second copy of virtual volumes on a consolidated MVC that you will store in a rack within the tape library for recovering

from a tape media failure while not requiring additional silo tape slots. You can use the Consolidate utility to recover from a disaster if you also recover the HSC Control Data Set and have an identical configuration. On the other hand, the Export/Import utility would provide a way to copy virtual volumes onto MVCs that you will store at a remote location (offsite storage vendor or disaster recover site). At recovery time, and at a recovery site with VSM, you would issue an Import against the required Manifest files only.

SUMMARY

IBM and StorageTek have each added capacity and performance to their respective virtual tape offerings through improved hardware and microcode. They have enhanced their management of virtual volumes by giving the customer more control over the migration and deletion of virtual volumes from the disk buffer. Both companies have added export capability and

other features to improve disaster recovery planning. Although the two vendors do differ quite a bit in how they have improved performance, view the management of virtual volumes, and in their disaster recovery processes, they are both committed to the virtualization of tapes, data storage costs, and the floor space reduction made possible by virtual tape solutions. 



NaSPA member Lloyd Christensen is a Principle Storage Consultant with StorageTek's Professional Services group. He lives in Granbury, Texas, and has been a NaSPA member since 1992. You can contact Lloyd via email at Lloyd_Christensen@StorageTek.com.