

Virtual Tape – Getting Rid of a Troublesome Medium

November 2006

Magnetic tape has been the backup medium of choice for decades. Accompanying it, however, is a lot of baggage. In large shops, there can be a floor-full of magnetic tape units and myriad operators. Tapes must be moved to off-site storage and then retrieved in the event of a system failure or data loss. Perhaps magnetic tape's most serious limitation is the hours, days, or weeks that it may take to restore a failed system or to bring up a cold backup.

A recent solution to these problems is virtual tape. A virtual tape system virtualizes magnetic tape cartridges as disk files. Virtual tape cartridges are fast and space-efficient. They may be electronically replicated to off-site storage and to backup sites. Their contents can be written to physical cartridges at the remote storage site if required.

Virtual tape is now an accepted technology with installations at hundreds of sites.

The HP Virtual Tape Server

One virtual tape product is HP's Virtual Tape Server for NonStop systems. It looks exactly like a tape system to the host. No changes are required to the applications or to the operating system to support virtual tape, nor is any change to the host's tape controllers required.

The HP VTS is fault-tolerant and scalable. It can reduce restore times by 50% at a cost commensurate with a magnetic tape subsystem.

Magnetic Tape Operations


To better understand the use of virtual tape, let us look at how magnetic tape is used and at some of its consequences.

Backup

Typically in a data center, the database is protected by periodically writing a full snapshot of the database to magnetic tape. This is a very lengthy procedure, typically requiring that the system be dormant, and is undertaken perhaps weekly.

To minimize the loss of data in the event of a fault, incremental backups are made on a more frequent basis – perhaps daily. In the event of a system failure, the system (or its backup) can be restored by loading the last full backup onto the system from magnetic tape and then by rolling it forward with the incremental backups. All data since the last backup is lost. This defines the

system's RPO, or Recovery Point Objective. The RPO is the organization's goal for the maximum amount of data loss that is tolerable following a failure.



Tape Backup

Periodically, perhaps daily, the tapes are moved to an off-site safe-storage location, from which they can be retrieved if necessary.

In a large data center, there can be a large number of tape drives and a sizeable tape operations staff. The tape operators are responsible for labeling the tape cartridges, creating records of the contents of the tape cartridges, managing retention periods, and so on.

The use of magnetic tape for backup has several problems associated with it, including

- the potential for operator error.
- the cost and maintenance of the many tape drives, including tape silos in many installations.
- the potential for lost tapes in the shipment process to off-site storage.
- the possibility of the destruction of on-site tape cartridges that have not yet been shipped to off-site storage as a result of a fire or other disaster.

Recovery

The time to back up the database is not the problem. The problem is the time to restore the database. Among the restoration steps are:

- accessing the records indicating what data is stored on what cartridges (what if the records have been destroyed in a fire?).
- requesting the retrieval of cartridges from the backup site.
- moving personnel to the backup site if one is to be used.
- loading the last full backup onto the system.
- loading each incremental backup onto the system.

Restoring a system can take days to weeks. During this time, the users of the system are not receiving service. This defines the RTO (the Recovery Time Objective) for the system. The RTO is the organization's goal for the maximum allowable downtime.

There are many things that can go wrong in this process:

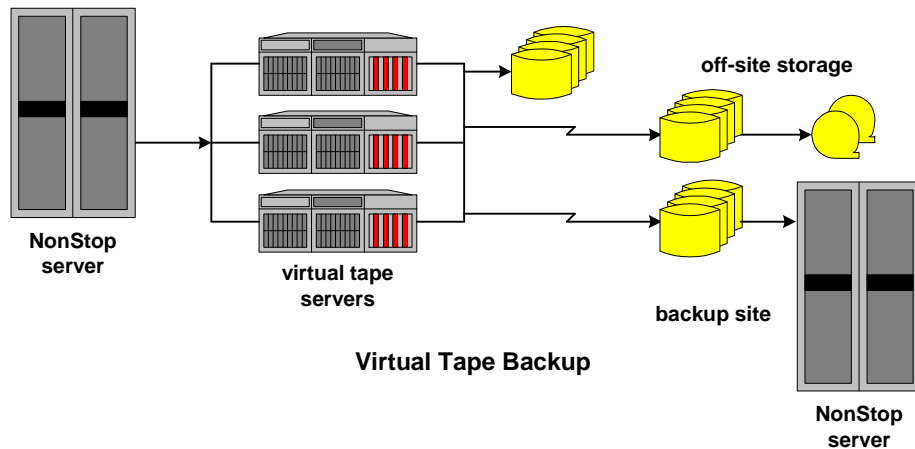
- There may be tapes that are lost.
- There may be tape errors that render a cartridge useless.

- There is further potential for operator error during the restore process.

Enter Virtual Tape

Virtual tape solves or minimizes many of the problems associated with magnetic tape. It basically replaces magnetic tape with disk files that are virtual representations of tape cartridges in a way that is totally transparent to the host system.

Virtual tape servers replace the magnetic tape controllers and present the identical interface to the host system. These servers create virtual tape cartridges, virtual tape drives, and virtual tape libraries on highly reliable disk systems. Disk storage systems are typically either RAID arrays or are enterprise storage area networks.



Backup

With virtual tape, all backup operations are totally automatic and require no operator intervention. Basically, a file which would have been written to magnetic tape is now written to virtual tape instead. The VTS system manages all cartridge labeling, retention management, and all of the other jobs done by operators when magnetic tape was used.

Backups require less time. As a consequence, it may be possible to increase the frequency of full and incremental backups to reduce RPO.

There are no on-site tape drives required (though in some cases, a few may be desirable). Rather, the backup files can be transmitted directly to the off-site storage facility, bypassing the need to transfer any physical magnetic tapes.

In addition, backup files can be transmitted directly to the backup site if there is one and will be immediately available at that site in the event of a failure of the primary site.

It's All About Restore Time

As we said earlier, the amount of time needed to back up a system is not as serious a problem as the time that it takes to restore a system. Backup time may affect the size of the required backup window. Restore time directly affects the operations of the company since no IT services are available until the system is restored. This could take days.

Virtual tape makes many improvements in restore time:

- There is no need to retrieve tapes from an off-site facility. The backup data is already at the site to be restored.
- Restoration from disk can be much faster than restoration from tape.
- Tape reading errors are eliminated. (Disk storage is redundant so that in the unlikely event of a disk read error, the data can be recovered.)
- Operator errors are eliminated since system restoration is automatic.
- If a backup site is used, backup can be started and monitored remotely before staff has arrived at the backup site.

As a result of these time savings, field experience has shown that restore times can be reduced by 50% or more. Days may become hours. Thus, the organization's RTO can be reduced by a similar amount.

Instant DR

HP's Virtual Tape Server has two important modes of operation directed at different problems – Instant DR and AutoCopy.

Instant DR (Instant Disaster Recovery) is aimed at backing up non-transactional databases. Whenever a backup file is closed, the backup data is replicated to the backup sites (local, remote backup, and/or off-site storage).

More importantly, replication is done in a very efficient manner since only the change data is replicated. Typically, the amount of data that has changed in a file between backups is only a small proportion of the file's size. Therefore, replicating just the changes may decrease replication time by orders of magnitude.

Furthermore, incremental backups are not just simply stored to be used in a restore operation later. Rather, they are applied immediately to the latest copy of the backup file so that the restore process only requires that the latest backup copy of the database be loaded. There is no need to apply incremental backups. The change data is retained so that previous backups can be recreated as necessary.

A major advantage of Instant DR is the restoration of lost or corrupted files. With tape, the last backup must be read for its version of the file; and then all of the incremental backups must be searched for updates to that file. This can take hours. With virtual tape, direct access is available to the latest backup of that file. Restoration time is now measured in minutes or seconds rather than hours.

AutoCopy

AutoCopy is used to back up transactional data. This includes full backups and audit trails. Whenever a backup file or an audit trail is closed, AutoCopy will replicate it to the appropriate sites.

VTSPolicy Facility

HP's VTSPolicy facility allows the specification of replication policies. It can be used to direct specific replications to specific sites. The operational staff can command the VTSPolicy facility via EMS messages (NonStop's Event Management System).

Other Characteristics

Other features of the HP VTS include the following:

- *Performance:* Backup and restore performance is greatly enhanced relative to magnetic tape because disk transfers can often be faster, efficient compression is used, and there is no need to deal with off-site transfers, lost tapes, tape errors, or operator errors.
- *Compression:* HP's VTS uses compression to send and to store its data. Compression ratios of 3:1 are common.
- *Scalability:* The HP VTS can scale from two Virtual Tape Servers to over one hundred.
- *Availability:* By using at least two Virtual Tape Servers and highly reliable disk systems, an extremely available backup subsystem is obtained.
- *Encryption:* Data can be encrypted during transit and during disk storage. Gone are the days when unencrypted tapes were lost during transit to or from the off-site storage site. (Citibank once lost four million nonencrypted customer records; Bank of America lost one million such records.) A device from SecureExpress can be used for biometric (fingerprint) authentication.
- *Cost:* The initial cost of a virtual tape subsystem is commensurate with the initial cost of a magnetic tape subsystem. Significant operating savings are made in the cost of magnetic tape unit maintenance, replacement or upgrading of magnetic tape units, and operations personnel.

Some Customer Experiences

Some actual results from customers are very illuminating. They show what benefits virtual tape can bring to a data center.

Gallagher Basset

Gallagher Basset is the largest multiline property and casualty third-party administrator. Over a period of time, its backup operations moved from reel-to-reel to cartridges to virtual tape. Formerly, it had a library of 7,000 to 8,000 tapes and fought constant operator errors. The company was getting significant pressure on its backup window. Restoration took many hours.

When it moved to virtual tape, backup time was reduced by 50%; and the tape library was reduced to 60 tapes.

SunGard Securities Finance

SunGard is a major Application Service Provider (ASP) to the securities lending community. It was performing several nightly backups and used over a hundred cartridges each night. SunGard went to multiple virtual tape servers and RAID disk systems. It used GFS (Global File System) to ensure that all virtual tape servers had access to all data even in the event of a failure.

This move reduced its backup and restore windows by up to 75%. Virtual tape eliminated operator and media errors. SunGard's backup and restore procedures are now completely automated, and allow it to operate its remote site as a "lights out" site.

SunGard reduced its tape usage by 99%. Its daily use of cartridges fell from over one hundred to just one. Its tape library was reduced from 10,000 cartridges to one hundred cartridges.

SunGard achieved significant cost savings in tape drive maintenance, personnel, transportation, and media cost.

The History of HP's VTS

HP's VTS was first developed by Tape Laboratories, Inc., which was recently acquired by Crossroads Systems, Inc.

HP added the Tape Labs Virtual Tape Server to its product line in 2005. The Virtual Tape Server described above is now available as an HP product with full support from HP.

The Bottom Line

Virtual Tape Server success is driven primarily by the higher reliability and shorter times for restore operations. The value of virtual tape is demonstrated by the success of HP's Virtual Tape Server. Today, over 150 units installed around the world at 75 sites provide more than 800 virtual tape drives.