



Virtual Tape Systems for IBM Mainframes

A comparative analysis

Virtual Tape concepts for IBM Mainframes

Mainframe Virtual Tape is typically defined as magnetic tape file images stored on disk. In reality it is a little more subtle than that. Most discussions of Virtual Tape focus on the advantages of Virtual Tape versus Real Tape. That begs the question, if it is better to store files on disk, why store them on tape in the first place? Since the earliest versions of OS/360 and DOS/360, programs could easily choose between disk or tape devices for data file storage. If a program was written to use tape, it is generally a trivial matter to change the program to use disk instead. In fact, in most cases the program need not be changed at all, just the JCL (Job Control Language) statements, which describe the input and output files used by the program.

If that is the case, why do we need "Virtual Tape" which uses disk but pretends to be tape? The answer is that in most data centers, there are hundreds or even thousands of jobs that are set up to use magnetic tape. Many of these jobs have been running unchanged for many years. Although changing a single job to use disk is trivial, changing thousands of jobs is quite another matter.

You also might wonder why you would want to change jobs that are working perfectly to use disk instead of tape. There are many reasons to store data files on disk instead of tape. Unattended operation is probably the biggest one. The following paragraphs discuss the advantages of disk and the advantages of tape. For the most part, the advantages of Virtual Tape are the same as the advantages of using disk in the first place, except that existing jobs which have been set up to use tape, can now use Virtual Tape with little or no change. The IBM VTS home page has a very compelling white paper on the virtues of Virtual Tape.

Disk vs. tape for storage

This has been a subject of ongoing debate for as long as there have been computers. For any given job or application, the best choice for storage today might not be the best choice tomorrow. Rapid advances in technology keep changing the grounds rules, i.e... those factors on which such decisions are based. For instance, although the cost of disk is still higher than the cost of tape, as it always has been, the cost of computer operators has continued to increase as the cost of computers and storage has gone down.

Advantages of tape

- **Cost:** The cost per megabyte of tape is still considerably less than disk as long as it is utilized to near its capacity. Most installations have many legacy tape jobs which have been in use for many years since the time when disk was extremely expensive. Even as recently as the early 1990's mainframe disk space was priced as high as \$12,000.00 per gigabyte.

- **Portability:** Tapes are removable and can be taken off site. This makes tape a good choice for backups. Tape has traditionally been used for bulk data interchange applications as well.
- **Scalability:** Tape is infinitely scalable. You can always buy more tapes at a very small incremental cost.

Disadvantages of tape

- **Operator dependence:** Mounting and dismounting of tapes is quite operator intensive, especially where lots of small files are involved. Physical labeling of tapes, storing and retrieving of tapes from racks, also contribute to expensive, inefficient and error prone operation by today's standards. Based on current salary levels, the elimination of one computer operator can pay for multi-terabytes of disk storage each year, year after year.
- **Under-utilization:** In normal practice, tapes often contain only a single relatively small file. This inefficient utilization contributes greatly to the loss of much of the cost per megabyte advantage.
- **Overall throughput:** In general, disk is faster than tape. When tape load, unload and rewind times are considered, along with operator mounting and dismounting times, the total cost advantage is usually lost.
- **Device contention:** Only one program at a time can access a tape drive. Although tape cartridges are very inexpensive, the actual tape drives and control units are not. Additional drives equates to additional floor space and other environmental considerations in addition to the acquisition and maintenance costs for the drives.

Tape management and Automated Tape Libraries

The above disadvantages have been addressed in several ways. Most notably, tape file management and stacking software address the underutilization and tape labeling issues while greatly reducing operator errors. Automatic Tape Libraries are also available to eliminate the dependence on the operator. While these solutions may be cost justified, they are certainly not cheap. Furthermore, while file stacking addresses the underutilization issue, it is at the expense of efficient file access, especially when retrieving a file near the end of the tape.

Advantages of disk storage

- **Share-ability:** Disk devices, also known as DASD (direct access storage devices) can be concurrently accessed by many programs simultaneously with generally little performance impact, effectively eliminating scheduling bottlenecks.

- **Efficient space utilization:** Generally speaking, many, many files can be stored on a single disk storage device with little wasted space.
- **Efficient file access:** Files can be accessed and opened practically instantaneously. It is not necessary to scan past one or even several other files to reach the desired file as it is with tape files.
- **No operator intervention:** Disk storage devices are generally always on line, there is no dependence on the operator. Jobs run basically unattended unless a problem occurs

Disadvantages of disk storage

- **Cost:** It will always be less expensive to buy a few more tape cartridges.
- **Scalability:** Of course you can buy more disk. You just can't buy a 'little bit' more disk.

Back to the Virtual Tape discussion

As has been stated earlier, the case for Virtual Tape vs. Real Tape is basically the same as the case for disk vs. tape but without the need to convert existing tape jobs individually. The point of all this is simply that **all** Virtual Tape systems have the same basic advantages over Real Tape as described above. Interestingly, most virtual tape vendors limit their discussion to those points presented above, comparing their Virtual Tape to Real Tape while ignoring other Virtual Tape Systems. It's as if each VTS vendor thinks they are the only VTS vendor.

The following paragraphs discuss the various types and approaches to virtual tape in contrast specifically to the most popular VTS system which is the Magstar VTS from IBM. The Magstar VTS was introduced in 1997 and has enjoyed great success in large MVS, OS/390 and now Z/OS installations. In general it has been considered much too expensive for use in smaller MVS and VM or VSE shops. Very recently, other VTS systems have begun to emerge, some of which are downward scalable and affordable to the smaller S/390 mainframe community. These systems are based on different technological approaches as described below.

IBM Magstar VTS

This is by far the most sophisticated but also by far the most expensive solution. The Magstar VTS uses dedicated special purpose Raid storage to house the virtual tape images. The IBM white papers extol the virtues of Virtual Tape compared to Real Tape. These advantages are true of all Virtual Tape systems. In particular, the entry level cost for the Magstar VTS makes it feasible only for fairly large mainframe data centers.

Other proprietary Virtual Tape Systems

Storage-Tek and one or two other similar brands of VTS also use proprietary dedicated Raid storage to house the virtual tape images. They are similarly sophisticated and expensive.

Software only, Virtual Tape

These systems intercept I/O requests intended for tape devices and redirect the request to utilize mainframe DASD instead of mainframe tape. These systems are relatively inexpensive and provide unattended operation. Some of these systems use compression to minimize the use of the expensive mainframe DASD. In most cases the mainframe DASD is only used temporarily. The virtual tape files are then stacked and offloaded to real tape thus addressing the under utilization issue. The only drawback to this approach is the trade off of efficient accessibility vs. the cost of expensive mainframe DASD. The more mainframe DASD is made available, the more efficient is the access to the files. However, mainframe DASD is by far the most expensive storage medium.

Open storage Virtual Tape Systems

These systems utilize open network based raid disk storage for storing the virtual tape images. In general, open network storage is much less expensive than mainframe DASD. These systems provide unattended operation, efficient use of storage, and instant accessibility for the raid resident virtual tapes. These systems are generally accompanied by hierarchical storage management systems which eventually migrate some of the virtual tape files, stacking them on large capacity tapes. The concept is the same as above, but the network based storage is generally so inexpensive by comparison, that real tape is only used for long term archival purposes.

Even off site backup and disaster recovery procedures are now possible whereby mainframe backup virtual tapes are compressed and then transmitted (zipped and shipped) off site rather than having physical tapes sent by messenger. Some companies historically have created two sets of backup tapes, one to keep on site and another off site. This virtual tape approach eliminates the need for such duplication.

Comparing the various VTS systems

The following paragraphs contrast the various VTS technologies with respect to their strengths and weaknesses.

Acquisition cost

Least expensive to acquire per emulated tape drive are the “software only” VTS systems. However, the cost per resident on line virtual tape file is quite high due to the use of mainframe DASD. These systems can be quite effective for temporary files and for a small number of fairly small generation data sets. For backup and disaster recovery applications these systems have certain innate limitations. For instance, if mainframe DASD is used for the virtual tape images, then most disaster scenarios would wipe out both the master files and the backups. If the virtual backup tapes must be offloaded to tape anyway, then the original step which created the virtual tape in the first place is somewhat irrelevant. Some software only solutions include transmitting the virtual tape images to a network PC via TCP/IP. This approach puts a major strain on both mainframe CPU cycles and network bandwidth.

At the high end, the Magstar VTS and the STK VSM products are designed and priced for the very large IBM mainframe installation, primarily OS/390 and Z/OS environments. The use of dedicated, proprietary Raid storage contributes to the high acquisition cost. Reliability and serviceability through hardware redundancy is paramount with these systems. Multiple channel paths and sophisticated caching provide excellent performance but make these systems unaffordable to the smaller shop.

Channel attached VTS systems which utilize open networked storage such as the Virtual Tape System from Visara International provide a very affordable entry level price point while at the same time they provide virtually unlimited scalable capacity. The use of open storage and servers allows a wide range of price performance and reliability options to suit the needs of any installation, large or small.

Scalability

Most Virtual Tape Systems are not very scalable by their nature. The IBM Magstar and the STK VSM are examples. These systems are certainly not downward scalable due to their high entry level cost. Increasing the amount of dedicated Raid disk storage in small increments is not usually feasible either. Furthermore, the Raid storage is generally so expensive that a relatively small amount is generally purchased simply as a caching front end to an Automatic Tape Library. At any given time, only a small percentage of the virtual tape images are on line in the Raid cache. These on line virtual tapes benefit from the direct accessibility of the Raid storage, while the majority of the virtual tape files have been stacked for efficient utilization on the real, large capacity tapes in the Automatic Tape Library. Remember, the utilization efficiency of the stacking leads to inefficient retrieval of the files further in to the real tape. Also, the

Automatic Tape Library itself is not very scalable. You are normally advised to buy more ATL capacity than you will ever need because when you run out of slots you can't simply just buy "a few more slots".

The software only VTS systems are only as scalable as the extra available mainframe DASD capacity allows.

Open storage based VTS systems like the Visara International Virtual Tape System are extremely scalable by contrast however. They are very affordable at low capacity entry levels, and adding more capacity in small increments is as simple as adding another NAS device, or another server to the network. In fact, reliable open network storage is now so affordable that it is often feasible to have all the virtual tape images or certainly a high percentage of them always on line in Raid storage. If an ATL back end is required, there are many more network compatible brands and models to choose from, for the ATL as well as the Raid Dasd itself. You are not locked in to the few more expensive mainframe compatible brands. The use of open file systems for the virtual tapes also opens the door to the use of any of several brands of HSM (Hierarchical Storage Management) packages as well. The HSM software manages the movement of the virtual tape images between the Raid and an Automatic Tape Library.

Performance

Almost all the Virtual Tape solutions perform well compared with real tape, especially when rewind and operator mount times are considered. Device scheduling, reduced reruns due to operator errors are additional throughput considerations. When comparing the performance of the various VTS systems one to another, an ironic observation becomes apparent. Virtual tape systems perform very well as long as the virtual tape image is in disk cache. When the virtual tape has been migrated off to an Automatic Tape Library, the VTS performance suffers a double whammy. It takes longer to access the virtual tape file, especially if the file has been stacked near the end of a large capacity reel. Secondly, the virtual tape must then be reloaded to the cache in its entirety before it can be opened by the mainframe. The net effect of this is that the ratio of tape files online in cache vs. those which have been offloaded to real tape becomes an important determining factor of overall performance. The irony is that the high end systems are the most adversely affected by this factor because of cost. In fact, reliable, Open Networked Storage has become so inexpensive and so easily extended that you may be able to keep all your tape images on disk. Real tape can then be reserved for long term archival purposes or for files which must be shipped to another data center for instance. Even backup and disaster recovery tapes may not necessarily be required to be offloaded. Current high bandwidth communications, along with compression, snap shot capability, remote mirroring and other state of the art storage technologies enable very affordable off site backup and restore via network transmission, without the need for shipping real tapes by courier.

Summary

Virtual tape has many advantages over real tape, and will most often yield a rapid return on investment. If you have not yet invested in the high end IBM or STK solution, you should investigate an open storage based VTS system. These systems provide very affordable, scalable, reliable and efficient Virtual Tape capability while providing excellent investment protection. A feature comparison of the Bus-tech MAS and the Visara International Vi-4990 VTS follows.

Cost

The Visara International Vi-4990 VTS has a lower entry price point and is generally less expensive for equivalent capacity.

Scalability

The Visara International Vi-4990 VTS offers greater configuration flexibility allowing price, performance, reliability and capacity trade offs tailored to the user's requirements.

Operation

The Bus-tech MAS (and most other VTS systems) are quite dependent on mainframe tape management software. The Visara International Vi-4990 VTS interfaces with any tape management software, but also provides its own built in management functions and GUI operator interface for unrivaled ease of use even without a tape manager.

Compatibility

The Visara International Vi-4990 VTS is built upon the MS Windows platform and file system. This allows incorporation of the widest possible complimentary functionality. The Vi-4990 VTS is therefore compatible with any HSM or other Windows based software applications as well as any tape or disk or any other device which supports MS Windows. The MAS is Linux based, which may limit its compatibility.

Conclusion

Virtual Tape is fast becoming an essential part of the modern data center. Several new products are available which make it more affordable than ever. The Vi-4990 Virtual Tape System from Visara International is one such product which provides affordable and scalable capacity with extraordinary functionality and ease of use.

For more information, contact:

Visit our Internet site at www.visara.com .

For sales and product information, contact us at 888-334-4380 or e-mail us at sales@visara.com .

Or e-mail Bill Phelps, VP of Strategic Planning and Product Marketing, phelps@visara.com .

Or e-Mail Wayne Williamson, Corporate Staff Systems Engineer, waynew@visara.com .

Corporate Offices of Visara International
2700 Gateway Centre Blvd., Suite 600
Morrisville, NC 27560
info@visara.com
919-882-0200 (main)