Mainframe Data Library with Hitachi Network Attached Storage (HNAS)

An Alternative for Traditional Mainframe Backup

Executive Summary

Today's mainframe data centers are, for the most part, still highly dependent on tape systems to provide daily / weekly backups for disaster recovery as well as long term data archive. These tape systems are large and expensive; occupying huge amounts of data center floor space, often using tens of thousands of cartridges, requiring on-going preventative maintenance, and intense operations support. And for all that, retrieval and/or recovery of data stored on tape is generally time consuming and unreliable.

As disk technology has advanced, use of low-cost ATA disk storage for data backup and archive has become a popular alternative to tape in the open-systems server market. Virtual Tape Libraries (VTLs) are available from a number of vendors, appearing on the SAN as a robotic tape library while actually storing data on RAID protected disks.

Bus-Tech's Mainframe Data Library (MDL) provides a virtual tape library solution for IBM zSeries or compatible mainframes. Combining Hitachi Network Attached Storage (HNAS) with Bus-Tech's MDL system greatly reduces or eliminates a data center's dependence on tape by transparently replacing high-cost tape sub-systems including automatic tape libraries (ATL) and virtual tape systems (VTS) with new, low-cost storage.

This paper provides an overview of this exciting solution and the benefits it can provide to customers interested in replacing their aging tape and / or optical solutions with a solution designed for the 21st century.

Tape Sub-System Overview

Tape has always been a prominent part of any large mainframe installation. From the early reel tape drives that were manually mounted by human operators to Automatic Tape Libraries (ATLs) to Virtual Tape Servers (VTS) meant to maximize cartridge utilization, mainframe computer operations has always spent large amounts of time and money managing and using large-scale tape systems.

Part # 86-03249-03-001

MDL w/Hitachi HNAS for Tape Replacement

It is not uncommon for even a medium sized mainframe environment to have anywhere from 16 to 64 tape drives and thousands or even tens of thousands of tape cartridges stored both locally with the mainframe or remotely at a disaster recovery site.

ATLs, like the one shown here can have thousands of storage slots each housing a single tape cartridge; providing 100s of Terabytes of capacity. Mechanical, robotic arms move cartridges from their storage slot to an available tape drive in response to a load request sent to the ATL from the mainframe.

And while ATLs are much more efficient than having a human operator physically mount a tape cartridge in response to a message on the operator console, much of an ATLs potential capacity can be wasted when small tape volumes which are a fraction of the cartridge capacity in the library are written to individual cartridges.

Virtual Tape Servers such as IBM's VTS or Storagetek's VSM attempt to address wasted capacity by temporarily caching small tape volumes to internal disk and then stacking multiple small volumes onto a single physical cartridge. But virtual tape servers create their own management problems as volumes expire and tape to tape copies are needed within the library to recover space.

Moving tapes offsite for disaster recovery is almost always a requirement. For manually mounted tape systems (i.e. drives without robotic arms) sending tapes offsite is simply a matter of writing 2 copies of the tape and then sending one of them to the remote site by courier. But for ATLs and VTS sending tapes to a remote site is much more

complicated. ATLs can write two copies of a tape and then export one through an export slot on the library. But virtual tape servers can have 10s or 100s of logical tape volumes on a single cartridge; making it difficult to impossible to send physical cartridges offsite.

Automatic replication of tapes from a primary to secondary site is possible, but generally very expensive. To consider such a proposition implies that both the primary data center and the disaster recovery site must have the same or similar ATL or VTS. Then high-speed and proprietary communication links must be established between the libraries so that volumes written to the primary library can be duplicated to the remote site.

There are two primary uses of tape within the mainframe data center; backup and archive. And, included in both of these is data duplication for offsite storage for disaster recovery.

The requirements for backup and archive are different. Backup data is usually only required for a relatively short time. If a mainframe backs up its primary customer database each day, for example, than typically the current backup is only current and required until the next day when a new backup becomes current. But, even though it may only be needed for a relatively short period of time, it still must be duplicated at the disaster recovery site so that if the primary site is lost the data is available.

On the other hand, archive data, which often no longer exists on primary DASD, may need to be maintained for many years. And managing when data gets migrated from online DASD to an archive medium such as tape is a balancing act based on whether or not the data is likely to be needed for immediate online processing.

Consider, for example, a bank's customer data. Bank account data can typically be retrieved in a matter of a second or less if it is stored on DASD. But once it is migrated to tape it may take a minute or more for an ATL or VTS to locate, mount, and retrieve the required information. So bank data processing centers must determine how long to keep a customer's data on DASD based on how often they expect their online banking systems to need the data. The current month's data, or even the last 3 months data, is almost certainly going to be maintained on DASD. Especially considering today's customer online banking applications. But what about data that is 6 months old or 12 months old? When should it be archived to tape? If it is archived to tape too soon then

Internet customers may wait minutes for data retrieval and are almost certain to drop off before the system can respond. But keeping data that is almost never going to be retrieved on DASD increases overall DASD capacity requirements resulting in an increase it the cost per MByte for the data center's storage as a whole.

MDL Overview

As shown in the picture below the Mainframe Data Library (MDL) connects directly to the IBM zSeries (or compatible) mainframe using FICON or ESCON attachments. MDL appears to the mainframe operating system (z/OS, OS/390, VSE, VM, or TPF) as standard IBM tape drives. Existing mainframe-based applications can use MDL tape drives just as they would any real mainframe-attached tape drive. No application changes are required.

MDL – a zSeries Tape-on-Disk Controller

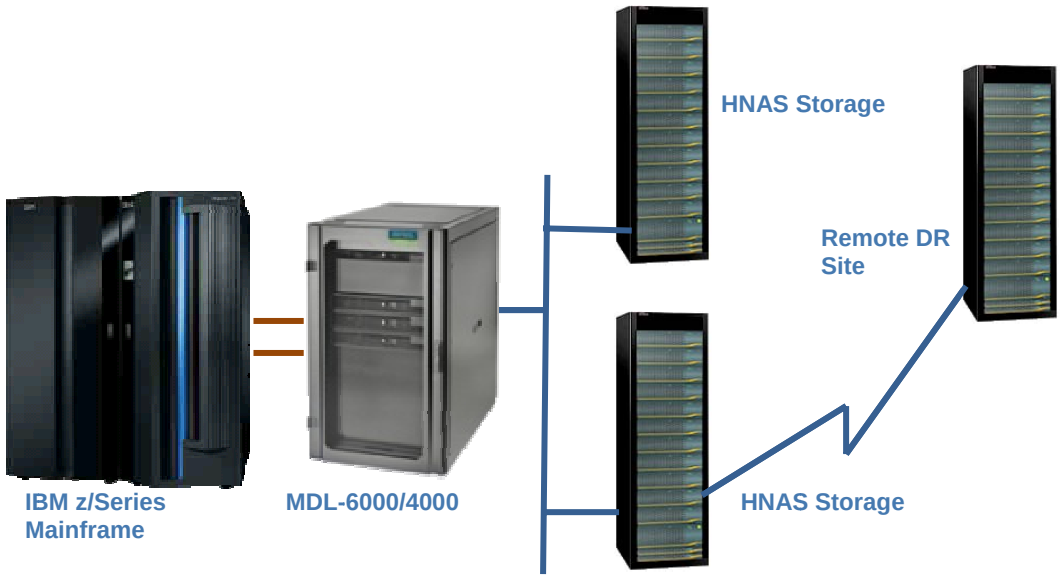

But instead of writing data to physical tape cartridges, MDL uses HNAS to store a tape volume's data as a single file on an industry standard file system. Further, each tape volume stored is named using its mainframe tape Volume Serial Number (VOLSER) allowing it to be easily identified and mounted in response to future read or update requests from the mainframe.

MDL w/Hitachi HNAS for Tape Replacement

The Mainframe Data Library is a scalable tape controller capable of emulating 3480, 3490, or 3590 tape drives. MDL-4000 and MDL-6000 provide multiple independently operating controllers as a single solution providing built in high- availability and redundancy. A fully configured MDL-6000 can provide up to twelve (12) FICON interfaces and emulate a total of 1,536 tape drives.

Configuration of the Hitachi storage is dependent on the storage requirements of the solution. Multiple HNAS devices can provide file systems which can be combined by MDL to form a single virtual library. Additionally, Hitachi storage can provide built-in data replication allowing customers to build mainframe tape-on-disk backup solutions that include remote data replication for disaster recovery. Such an approach eliminates the need for multiple copies of backup tapes and physical transport of tapes to a remote location and can usually be implemented using simple IP connections.

Benefits

The benefits of implementing an MDL virtual tape library (VTL) are considerable. From shortened batch processing windows, to improved data retrieval, to lower operating costs, and electronic disaster recovery, an MDL can quickly pay for itself.

Shortened Batch Windows

When customers implement mainframe VTL solutions an often times unexpected benefit is a reduced batch processing window. For customers replacing non-automated tape drives or even Automated Tape Libraries (ATLs) the reduced time MDL takes to satisfy a mount request can provide significant benefit when a lot of tapes are involved. Additionally, since MDL allows the customer to set their own tape cartridge size, backups that once took multiple volumes and multiple mount requests can now be reduced to a single volume with instantaneous mount resolution. Customers implementing demanding tape applications have realized as much as a 40% reduction in their overall batch processing windows.

Improved Data Retrieval

Retrieval of data stored on traditional tape can be very time consuming. Even when the data is stored in an ATL or VTS, mounting of a tape to retrieve data can require several minutes or more if there are no drives available and the volume is not cached.
And then there is the issue of data reliability. Physical tape can stretch and break especially when the cartridge has been repeatedly used. The inability to read a tape can make data retrieval impossible.

Part # 86-03249-03-001

On the other hand retrieval of data from a VTL solution such as MDL is fast and efficient. Tape mounts are generally instantaneous. And once the tape is mounted retrieval of data from a tape volume can be comparable to DASD retrieval. When applications use tape block locate commands MDL can access any block on the tape in roughly the same amount of time. There is no need to read through the volume sequentially.

Customers who have moved applications from ATL to MDL have found significant improvements in overall retrieval of data using the VTL approach versus waiting for an ATL to physically mount a cartridge.

Improved retrieval of data from a VTL means that archive data can be migrated from DASD to the VTL sooner than it would be migrated to physical tape. Where data may have once been kept on DASD for up to 6 months, it may now only need to be kept on DASD for 15 or 30 days. The result is that DASD resources are freed for other use more quickly lowering the overall cost of storage in the data center.

Lower Operational Expense

Mainframe tape drives, Automated Tape Libraries, and Virtual Tape Servers occupy considerable amount of space within the data center. Additionally physical storage of thousands of cartridges can add to that space requirement. Electrical expense, maintenance contracts, and the ongoing need for new cartridges all work together to make large-scale mainframe tape subsystems an expensive proposition.

An MDL can significantly reduce operational expense. MDL can occupy much less raised floor than a medium or large tape system. Electrical consumption is reduced by the elimination of tape drives and/or robotic arms. The need to continually purchase new tape cartridges can be completely eliminated.
And finally, there is no need to periodically take MDL offline to perform preventative maintenance to mechanical robotic arms and drives.

And finally, there is no need to periodically take MDL offline to perform preventative maintenance to mechanical robotic arms and drives

Electronic Disaster Recovery

In this day of high-speed networks and the global internet, mainframe data centers continue to rely on disaster plans based on physical tape being shipped to remote storage locations. Often, the only alternative is to install very expensive, dedicated proprietary communications links between facilities in order to allow duplication of volumes from one tape library to another.

MDL virtual tape solutions can alleviate or replace both the manual transport of tape and expensive proprietary communications links. Using HNAS replication capabilities, tape volumes stored in a Mainframe Data Library can be automatically duplicated to a remote location over standard IP connections.

In the unlikely event of a disaster, data stored at a remote location can be retrieved just as if it were locally stored in the library. The difficulties of finding a physical tape and transporting it from the disaster site to be processed are replaced with online electronic efficiency.

Summary

Mainframe data centers have always been highly dependent on tape. From the days when mainframes were used for batch processing only to today's mainframe-based online transaction-based systems, tape has always been the lowest cost alternative for backing up and storing data when compared to the high cost of primary DASD.

But today's shared; open-systems disk solutions now allow disk-based backup solutions that offer significant benefit to tape. Over the past several years tape-on-disk and disk-to-disk backup solutions have become a hot technology in the open systems server arena.

But for mainframe data centers tape subsystems continue to prevail.

Bus-Tech and Hitachi Data Systems have collaborated to combine the Bus-Tech Mainframe Data Library (MDL) with Hitachi Network Attached Storage (HNAS) to provide an innovative Virtual Tape Library (VTL) for the mainframe. Appearing to the mainframe as a group of individual tape drives, MDL allows existing tape-based backup applications to store their data directly on Hitachi disk storage. HNAS-based replication capabilities allow expansion of Bus-Tech VTL to provide complete, fully-redundant disaster recovery solutions which can eliminate aging tape systems and advance the data center into the next generation of backup storage systems.

Part # 86-03249-03-001