BUS-TECH INC.
129 Middlesex Turnpike
Burlington, MA 01803
Phone: 781.272.8200
www.bustech.com

Mainframe Appliance for Storage
White Paper

Bus-Tech's Mainframe Appliance for Storage

Mainframe Appliance for Storage (MAS) is a 'TAPE-ON-DISK' virtual tape controller for IBM zSeries class [or compatible] mainframes running z/OS, OS/390 or VSE. Though it appears to mainframes as an IBM 3494-VTS [i.e. 'Virtual Tape Server'] the MAS is fundamentally different from a VTS. A VTS is still a tape-based storage scheme that uses a disk just to *cache* data that is being written to, or read from, tape. A MAS, on the other hand, is totally disk-based storage system that does not use any tapes – at any point in its operations. The MAS takes the VTS notion to its logical conclusion. Emulating tape operations using low-cost 'open systems' disks accelerates performance, enhances reliability, minimizes data errors and takes up less data center floor space. The MAS thus represents the next generation of VTS technology.

The current version of the MAS supports dual 2Gbps FICON Express mainframe connectivity (as well as ESCON channels); Gigabit Ethernet, fibre channel or Ultrawide SCSI connection to open-system storage; 'tape-mark'-based fast tape access emulation [e.g. fast forward], and support for the highly economical 'Advanced Technology-Attached' (ATA) disk arrays.

This paper introduces the reader to the concepts of 'TAPE-ON-DISK' as well as the solution offered by Bus-Tech and then reviews several of the applications, including state-of-the-art disaster recovery (DR) configurations, where Mainframe Appliance for Storage can provide immediate benefit.

'Tape-on-Disk' Overview

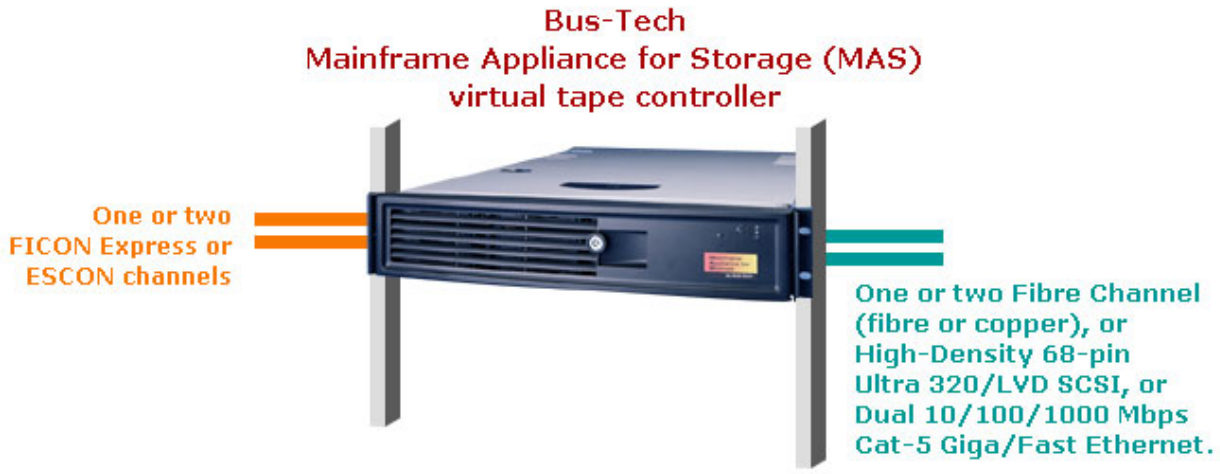

'TAPE-ON-DISK', the most advanced form of VTS technology, emulates physical tape drives and tape volumes using permanently mounted disk storage -- and thus obviates the need for tapes drives and tape cartridges. The disk storage used to store all of the data will be highly reliable, and cost-effective open-systems storage offerings from any of the leading vendors on the market. With VTS, data being written to

a tape volume, rather than being written to a slow tape drive, is actually stored as one of many files on a disk.

TOTALLY TRANSPARENT TO THE APPLICATIONS

Total software transparency is the defining characteristic of virtual tape systems, and the MAS is no exception. The 'TAPE-ON-DISK' emulation has to be transparent to the host computer, the operating system, and the applications so that they can continue to operate as if they are still dealing with actual tape units. In other words a 'Tape-on-Disk' device must always appear to the host computer as one or more real tape drives; responding to all processing requests exactly as a tape drive would – even though, in reality, there are no tapes being used. The host operating environment and the application software

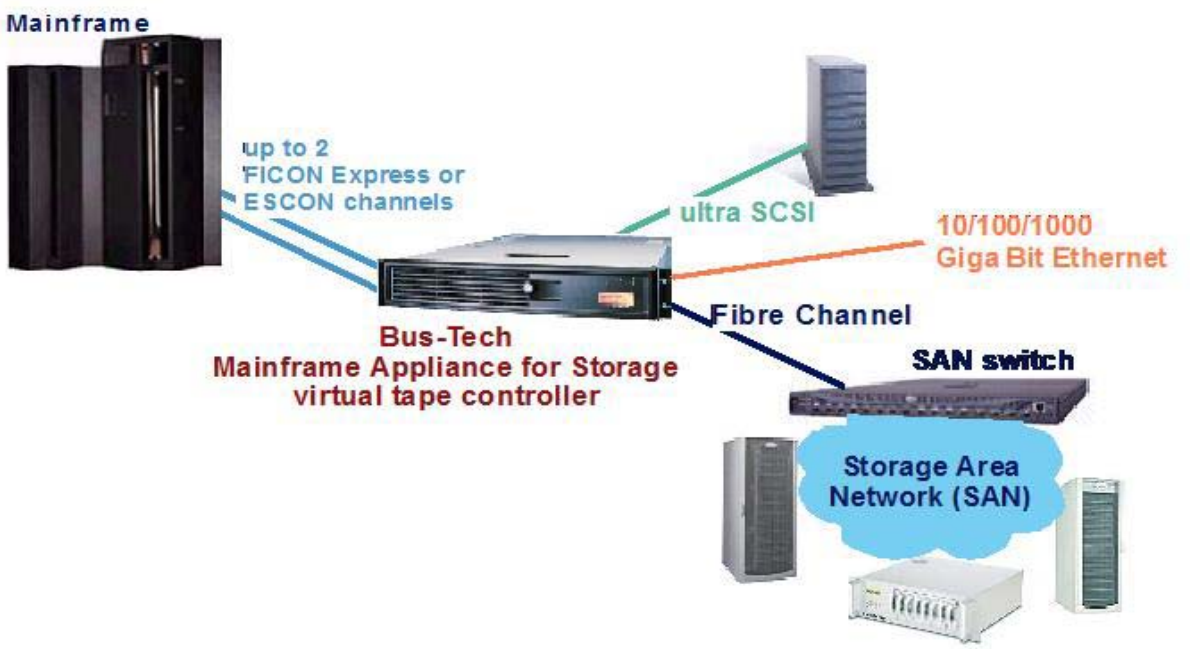

using the MAS 'TAPE-ON-DISK' system cannot be aware that the data is being stored on disk rather than tape. The Bus-Tech MAS, which was introduced in early 2002, and is now being successfully used, 24/7, by very large mainframe customers such as Vanderbilt University and Siemens Dimatic, provides guaranteed 'virtual tape' transparency and resilience.

The Bus-Tech MAS modernizes tape-based mainframe applications. It eliminates the need for tape cartridges and tape drives. With a MAS, IBM 3480/3490-oriented tape operations are emulated using fast, inexpensive, and permanently mounted disk drives. Rather than reading from and writing to slow, error-prone and labor-intensive tape cartridges, mainframe applications can now use fast, reliable, petabyte range disk stores – without any changes to the application software. The MAS, compared to IBM's 3494 systems, can be extremely cost effective with a MAS likely to cost ½ to ¼ of that of an IBM 3494. Detailed

a cost analysis, available from Bus-Tech, will show that for the price of a 20TB, IBM 3494-VTS system, a customer can obtain a 90TB Bus-Tech MAS!

A single, 2U (3.5 x 19 inch) rack-mount MAS can emulate up to 64 IBM 3480/3490 tape dives. Connectivity between a MAS and a mainframe can be realized using one or two 2Gbps FICON Express channels or 17MBps ESCON channels. Disk storage, including the emerging, low-cost ATA drives, can be attached to a MAS using 10Mbytes/sec Ultra 320/LVD SCSI, 2Gbps fiber channel or Gigabit Ethernet connections. The MAS has been tested and certified with all leading, open-system disk platforms.

POSITIVE ROI IN A VERY SHORT TIME

Virtual tape systems typically deliver very fast return on investment by:

- Freeing-up valuable mainframe capacity by significantly expediting the time taken to complete tape dependent batch jobs by eliminating tape mount delays, particularly when it comes to data retrieval [i.e. read] operations – where the MAS always provides near instantaneous, sub 5 second, access compared to the best case, 2 to 3 minutes required with an IBM 3494-VTS.
- Eliminating, the need for expensive tape cartridges – most of which, moreover, are but sparsely used in conventional tape-based systems.
- Greatly minimizes the chances of data loss due to media corruption by storing all mainframe data on error-protected, disk systems as opposed to on error-prone data cartridges as is the case with a traditional VTS system.
- Increases overall RAS [i.e. reliability, availability and serviceability] by not using any mechanical components such as the robotic assemblies used by a VTS.
- Minimizing the need for expensive ‘primary’ DASD storage for fixed-content data by exploiting the very fast, “no-delays-for-tape-mounts”, data retrieval possible with the MAS TAPE-ON-DISK approach.
- Doing away with the high, recurring maintenance costs on tape drives.
- Enhancing operator productivity (and satisfaction) by getting rid of frustrating tape operations related interruptions [e.g. tape intervention requests].
- Reducing the time taken for tape backups by as much as 60%.
- Facilitating disaster recovery scenarios by permitting, high-throughput, unattended remote backups (and retrieval).
- Permitting unattended, no intervention required, ‘TAPE-ON-DISK’ operation that enable the optimization of data center shift staffing, particularly at weekends.

- Significantly reducing tape library costs, in terms of floor-space, racks, equipment and personnel, given that a single open-system storage rack, such as an EMC Celerra NAS Storage Unit, can take the place of over 50,000 IBM 3480/3490 250MB tape cartridges!

- Better utilizing scarce and expensive raised-floor, data center real estate by space-hogging tape drives.
- Uses significantly less power and requires considerably less cooling, thus further reducing the TCO of the MAS compared to that of a VTS.

Since the MAS 'TAPE-ON-DISK' solution, as described above, is totally transparent to the host computer, all of the above benefits can be realized without the need to reengineer and redeploying mission-critical business applications.

Though they share some common technology the Bus-Tech MAS differs from the Bus-Tech zDASD 3990 controller, in that the MAS is a 'TAPE-ON-DISK' solution that enables open-system storage to be used to emulate tape operations, whereas the zDASD 3990 can be thought of as a 'MAINFRAME-DASD-ON-DISK' solution that enables open-system storage to be used in place of expensive mainframe DASD units.

MAS: Principles of Operation

The Bus-Tech Mainframe Appliance for Storage (MAS) is a TAPE-ON-DISK virtual tape controller. It connects IBM System/390 or zSeries class (or compatible) mainframes to high-capacity, open-system disk storage products – making the disk storage appear to mainframe software as if it were tape cartridges. A single MAS can emulate from one to 64, IBM 3480/3490 tapes drives. In a dual FICON or ESCON MAS a feature known as alternate pathing is supported. This feature provides load-balancing and automatic failover between the mainframe connections. A MAS is recognized by OS/390, z/OS or VSE as a logical aggregation of 3480/3490 tape drives. To achieve this, a MAS provides the IBM Hardware Configuration Definition (HCD) utility with its own customized Unit Information Module (UIM) modeled on the 'CBDUS005' definition supplied by IBM.

The MAS UIM ensures that all of the tape drives being emulated by a MAS can be collectively referred to as "VTAPE" in all job control language (JCL) data definition statements. Prior references to tape drives are now replaced with a "UNIT=VTAPE" operand. I/O operations to such "VTAPE" units will be sent to the MAS. Alternatively, the MAS can be managed as a Manual Tape Library (MTL) under the control of SMS management. Please refer to the "Complete Software Configuration Flexibility" section below for all of the various configuration options supported by the Bus-Tech MAS.

When a job with JCL DD statements containing "UNIT=VTAPE" is initiated, the MAS will automatically (and transparently) allocate virtual tape drives [i.e. emulated tapes-on-disk] to that mainframe job. Jobs and tasks running on z/OS, OS/390 and VSE can allocate tape volumes on an emulated virtual tape drives in exactly the same manner they would with a real tape drives. The emulated tape drives are totally

transparent to the mainframe applications. There are no changes required to the application software. The only changes required are restricted to the JCL statements.

Though MAS functioning is based on IBM 3480/3490 emulation, the operational characteristics of the MAS are in no way restricted by the speeds or cartridge size limitations of real IBM 3480/3490s. The MAS supports virtual cartridge sizes that are greater than 2GB – though this is the standard, default. A user parameter that is included in the JCL statements is provided to permit allocation of virtual cartridge sizes that exceed 2GB. Each MAS can emulate tens of thousands of tape volumes. These tape volumes, per standard IBM JCL conventions, are assigned tape volume serial numbers ranging from “AA0001” to “ZZ9999”. Each volume corresponds to a disk file on the SAN, SCSI or Giga-bit Ethernet attached open-system disk storage product.

INSTANT TAPE MOUNTS AND FAST TAPE ACCESS

If a mainframe application requests a ‘scratch’ tape to create a new volume, the MAS immediately assigns it the next available “XXnnnn” volume serial number and creates a corresponding disk file with that name. It then returns a positive response to the mainframe to indicate the successful completion of that tape mount operation. The whole process takes milliseconds. If an application requests the mount of a specific volume, the MAS checks that the corresponding data set exists on disk and then honors the mount. Again the entire mount operation is completed within milliseconds.

The MAS also supports fast tape access features, in particular the Fast Forward space file and relative block offset commands. The Fast Forward space file command when used with a MAS permits near instantaneous access to the end of the current file being read. This ‘FF’ feature is widely used to add new files to an existing tape volume. If a real tape drive were being used, the tape would have to be mounted and then sequentially read until the EOF mark was reached. The new file would be written following that EOF mark. With the MAS there is no need to sequentially read the file to reach the EOF mark. Instead, the MAS, through the use of pointers, can get to the end of any file within 50 milliseconds. Suffice to say that mainframe applications can get their work done considerably faster.

The relative block offset command is used by mainframe applications when they wish to read a specific block of data from a tape file. It is widely used by applications dealing with fixed content data. If a real tape cartridge were being used, it would have to be first mounted and then sequentially read until the desired data block was encountered. In some cases the tape drive has to first get to the end of the file so as to locate the pointers to the requisite block. In marked contrast, the MAS can always locate any specific block within a data file within 100 milliseconds. The bottom line is that the near instantaneous tape mounts and these fast tape access features considerably speed up and streamline data center operations.

FAITHFUL TAPE EMULATION USING “AWSTAPE”

IBM 3480/3490 tape volumes contain several different types of records including tape marks for separating files on a volume, header records, trailer records, and data records. Each of these records are of different lengths and formats and a single tape volume may contain multiple data files – each with its

headers, trailers, data content and tape marks. In addition, it is possible for records within a given data file to be of varying lengths. Since the MAS guarantees mainframe applications with total tape operation fidelity it has to ensure that the virtual tape volumes that it maintains are no different to those that would be written to an actual 3480/3490 tape cartridge. Given that no changes get made to the mainframe application, the MAS has to ensure that it can reproduce tape volumes during a read operation from a mainframe application in exactly the same manner that the files would have been written to the actual, physical tape in the first place.

MAS relies on a standard IBM file format, known as “AWSTAPE”, to realize the necessary tape emulation fidelity. AWSTAPE is a tape emulation file format for use with disk files. Each AWSTAPE disk file emulates one physical volume and contains an exact replica of the data that would be written to a tape including all tape marks, headers and trailers. Each record written to a virtual tape volume by a MAS is represented on the disk with a block header followed by the actual data. A tape mark is represented by just a block header. The block header includes the exact length of the data that follows, thus allowing the MAS to recreate tape records in exactly the way they would have been stored on a real tape cartridge.

COMPLETE SOFTWARE CONFIGURATION FLEXIBILITY

The Bus-Tech MAS can be used in any z/OS, OS/390 or VSE environment, with any of the standard data storage related access methods including OAM – the Object Access Method used by IBM’s Data Facility Storage Management System (DFSMS). There are no limitations or exceptions. Bus-Tech ensures that any given MAS configuration can be defined in an acceptable manner to the mainframe channel subsystem as well as the operating system.

There are three distinct ways that a MAS can be defined to a mainframe. These being:

1. MAS is defined as a set of real 3480/3490 tape drives
2. MAS is defined as a 3490 Manual Tape Library (MTL) device
3. MAS is defined with a data center specific, unique device type using Bus-Tech’s Unit Information Module (UIM) definitions

As mentioned earlier in this section the MAS provides a true emulation of IBM 3480/3490 tape drives. Therefore a MAS can be defined to a mainframe as a 3480/3490 equivalent. This type of definition, however, is only possible when there are no real 3480 or 3490 units attached to a mainframe. If a mainframe still has one or more 3480s or 3490s attached to it, defining a MAS also as a 3480/3490 could result in potential misallocation errors. If there are no 3480/3490s attached to the mainframe, it is preferable to define the MAS to the mainframe as a 3480/3490 since this will ensure that use of the MAS will be totally transparent to all access methods [e.g. OAM] and applications.

If you still have real 3480/3490s attached to your mainframe, then you will have to define the MAS to it either as a 3490 Manual Tape Library (MTL) device, or an installation specific-device using the Bus-Tech UIM. Which option you choose will depend on the access methods and the applications that will be talking to the MAS. If you plan to use OAM then you will have to define the MAS as a 3490 Manual Tape

Library device. You will also have to use the MTL definition approach if the MAS is going to be used with applications that explicitly check the device type of the unit that they are performing I/O to make sure that they are indeed interacting with a known tape device.

Manual Tape Library was a new tape management scheme introduced by IBM with APAR OW45271. It permits standalone tape drives and the volumes associated with them to be System Managed Storage (SMS) compliant by enabling them to be treated as a logical tape library. SMS is what mainframe data centers use to manage Automated Tape Libraries (ATLs) – the so called tape silos. Thus MTL is a technique that can be used with ATLs. Consequently, I/O operations [e.g. allocations] to MTL devices are managed by SMS as if it were dealing with an ATL, with the only exception being that mount requests for MTL devices are routed to a tape operator console rather than to the robotics unit of an ATL. The MAS can be defined to work as an MTL within the context of SMS. In such situations it would look like an ATL.

If you do not use OAM, or have applications that explicitly check a units device type you may want to define the MAS to your mainframe using the Bus-Tech supplied UIM. This way you can define the MAS as a unique device type. With the UIM scheme, you can have real 3480/3490 tape units working, in parallel, alongside a MAS. The UIM definition scheme will ensure that the operating system will not allocate MAS-based virtual tape volumes to applications requesting the mount of an actual, tape cartridge.

The bottom line is that with these three definition schemes, Bus-Tech can successfully handle any and all mainframe configurations, access method requirements and application specifics. With these schemes at their disposal Bus-Tech will work with you to determine which definition scheme will best fit your needs.

MAS ARCHITECTURE & CONNECTIVITY

The MAS is a 2U (3.5”) high rack mount controller designed to be installed in industry-standard 19” rack units. The MAS comes standard with a single 17Mbytes/sec ESCON mainframe channel, Ultra 320/LVD SCSI port and dual 10/100/1000 Giga bit Ethernet connections.
The MAS has internal, mirrored disks which house the MAS embedded Linux operating system and virtual tape software.

The MAS requires a monitor, keyboard, and mouse, which are used as the MAS operator console. The MAS operator console is used to:

- Configure the MAS Virtual Tape Drives
- Configure the attached direct access storage (disk)
- Monitor the operation of the MAS
- Provide maintenance on the virtual tape library.

The MAS can be enhanced with a second ESCON interface, fibre channel connectors or 2Gbps FICON Express channels. The MAS can support up to two mainframe channels, whether they be ESCON or FICON and multiple open-system disk storage connections. Disk storage, including the emerging, low-cost ATA drives, can be attached to a MAS using Ultra 320/LVD SCSI, 2Gbps fiber channel or Gigabit Ethernet connections.

The Bus-Tech MAS has been tested and certified with all leading open-system disk storage solutions – including the emerging ATA technology disk systems.

Bus-Tech MAS Applications

Now that you have been given an overview of how MAS works, it is time to take a look at how MAS may be applied within the mainframe environment and what potential benefits it might provide.

OPTIMIZING DAILY TAPE BACKUPS

Most of today's mainframe data centers have demanding online processing windows during which mission-critical applications must be available, without any degradation of performance, to networked users. This demand for online availability creates fixed batch processing windows, usually during the night, when all offline activities must be completed. For many data centers, planning and execution of batch processing is dictated by the need to backup large online databases within a fixed processing window.

The instantaneous, totally automatic, tape mount capability of the MAS makes a huge tangible difference when it comes to daily tape backup operations.

Not only does the MAS' 'TAPE-ON-DISK' approach significantly expedite the backup process but it also ensures that operators do not have to scurry around attending to tape drives. On average a MAS will slash the time taken for a daily mainframe tape backup by 60%. This quick completion of data backup, frees up valuable mainframe resource for other, revenue-generating tasks. The automated, unattended backups also enables the streamlining of data center shift staffing. Some of Bus-Tech's MAS customers have achieved a positive ROI within 6 months on just the increased mainframe usage and shift staff optimization – without even having to factor in the cost saving resulting from phasing out tape cartridges and tape libraries.

The MAS TAPE-ON-DISK approach also greatly enhances the integrity and validity of the data that is being backed-up. Tape media is renowned for its vulnerabilities that can result in unrecoverable data. With tape-based backups there is also always the danger of accidental tape volume mix-ups that can thwart the recovery of needed data. The inability to totally and accurately recover data that has been backed-up on tapes is a nightmare that continues to confront data center professionals. TAPE-ON-DISK, à la the MAS is a sure-fire solution to this problem, particularly so in that the RAID disk technology that

can be used with the MAS excels in offering cost-effective data protection (and auto-correction) features including affordable and automatic data mirroring.

Once the mainframe backup process is complete online resources can be placed back online. If removable backup is required for offsite storage, the backup of virtual tape volumes can be accomplished using MAS-based backup utilities with no impact to online mainframe applications. A MAS dump utility allows individual virtual tape volumes or the entire tape library to be copied from the disk array to a SAN or SCSI-attached tape subsystem. Such MAS-based backups can be set up to occur automatically at a specific time during the day and make efficient use of tape cartridges by writing multiple virtual tapes to a single cartridge.

Alternatively, if a SAN-based backup server is being used, backup of the virtual tape library can be routinely scheduled as part of a larger SAN-based backup.

Finally, if a real 34xx tape cartridge is required for transfer to another mainframe, simple mainframe-based utilities can be used to read one or more virtual tape volumes and copy them to a real cartridge.

DISASTER RECOVERY

With its Fibre Channel SAN connectivity, Giga-bit Ethernet ports, FICON Express mainframe channels and 'lights-out' unattended mode operation, the MAS is uniquely suited for mainframe disaster recovery scenarios. SAN and FICON each provide significant remote operation capabilities. FICON when used with repeaters can support 100 km [i.e. 60 mile] channel spans at giga-bit data transfer rates. In parallel, IP-based, high-performance wide-area network operations is a hallmark of SAN. Using SAN, Ficon Express or both, a MAS can provide totally automated, high-bandwidth, unattended mainframe data backup configurations between physically distant sites.

Furthermore, as described above, the complete elimination of tape mount and Fast Forward to EOF delays, significantly speeds up the time taken to complete a mainframe backup. Plus the use of disk technology also minimizes the risk of data corruption and media deterioration. With a MAS-centric, open-system disk storage based disaster recovery scenario one is unlikely to be confronted with the unacceptable situation of discovering that the disaster recovery data is not recoverable. All-in-all, more and more mainframe customers are discovering, post 9/11, that the MAS approach is the right and optimum way to realize contemporary disaster recovery scenarios for data centers.

BATCH PROCESSING

The high cost of mainframe DASD in the past dictated that large batch processing mainframe applications were designed to use tapes as their primary data storage and data retrieval mechanism. Batch processing, in general, consisted of reading data from one set of tapes, performing the necessary operations on that data and then writing back the new data onto a new set of tapes. Just the time

demands imposed by tape mounts, tape dismounts and re-mounts meant that these jobs were long drawn and consumed significant operator resources.

A MAS can effortlessly modernize such legacy, batch processing applications without the need for any software changes or application rewrites. The application continues to function as before, but rather than relying on slow, cumbersome and labor-intensive tapes, all I/O can be realized using affordable, high-performance open-system disk storage systems.

By reducing the number of volumes required to hold datasets and eliminating manual operator mount requests, MAS shortens the processing windows of tape-based batch applications. More importantly, these shorter processing windows can be realized without a major re-engineering effort to re-write the applications.

UNIFYING ENTERPRISE DATA

Mainframes are no longer the only mission-critical computing platforms in use at enterprises. With the advent of Web servers, application servers and portals, most mid- to large-size enterprises now rely on a range of platforms for their vital IT needs whether they be Unix/Linux servers, Windows 2000/2003 servers or mini-computers [e.g. IBM iSeries]. Despite this mix of platforms, enterprises favor having a standardized and unified data backup and recovery backup mechanism – typically controlled by a proven data backup software suite such as the IBM Tivoli Storage Manager.

Prior to the MAS it was invariably difficult and expensive to include mainframe data into such a unified, open-systems storage-based backup and recovery solution – given that mainframe backup software is geared towards mainframe-specific tape backups. The MAS is an elegant solution to this problem. With a MAS an enterprise can easily assimilate mainframe data into a unified, enterprise-wide data backup and recovery scheme [replete with disaster recovery where needed] with no changes to the mainframe applications and just minimal changes to the JCL.

INTER-PLATFORM DATA EXCHANGE

Most mainframe installations today have mixed environments including Unix, Linux, and Windows servers in addition to their zSeries (or equivalent) mainframes. Often there is a need to periodically move bulk data between the mainframe and one of these other platforms. In reality, this is an extension to the multi-platform back-up and recovery scenario described above.

With the proliferation of disparate IT platforms, with mainframes just being one of these, enterprises often have a need to transfer large volumes of data from one platform to another. The MAS with its unsurpassed connectivity can act as a powerful and flexible data switch between mainframes and other systems. Since tape emulation is MAS' forte, mainframes can exchange data with other platforms using standard data transfer utilities which think that they are reading or writing data to standard 3480/3490 tape cartridges.

Data exchange in a network environment, utilizing one of the two available gigabit Ethernet connections, and supporting a flat file format, is very straightforward.

Flat file support will allow a virtual tape drive to output a file as a standard "flat" file rather than in AWSTAPE format. Targeted specifically at data exchange, flat files are not easily re-read by the mainframe but can be processed by other servers using standard utilities like Unix/Linux dd.

The mainframe data is written out as a standard tape operation to the MAS. The MAS, by the use of Samba, is defined as an NFS client and the data is sent to the targeted UNIX platform (NFS mount point) using standard file transfer protocols. In a Windows environment, CIFS is supported in the same manner as NFS.

REMOTE SYNCHRONIZATION

Remote synchronization allows the virtual tape library to be mirrored and periodically synchronized across the Internet for added security, reliability and resilience. Ideal for transferring data between mainframes or for disaster backup/recovery implementations, remote sync will only update data files that have changed since the last sync in order to minimize the amount of data that will have to go over the network at any point in time.

All open-system storage vendors now provide powerful data mirroring utilities, that include remote synchronization capabilities. Any of these data mirroring/remote synchronization schemes can be used with a Bus-Tech MAS-centric solution. The use this type of remote synchronization provides the ultimate in data protection. If a local MAS array should fail, critical system operations requiring virtual tape volumes can be immediately re-directed to the remote network-attached tape library. Re-construction of the local array can be performed concurrently with on-going operation.

How the MAS Differs from an IBM 3494-VTS

The 'TAPE-ON-DISK' MAS represents the latest in VTS technology. It builds upon the 3494-VTS concept and takes the IBM VTS approach to its logical conclusion, exploiting the increasingly affordable, cost economics of modern open systems disk systems. The IBM 3494-VTS is a tape-based system. It uses RAID disks just as temporary cache storage when writing or reading data from tape cartridges [though as with any caching scheme it does also try to serve mainframe data read requests, when possible, directly from data that may still be resident within the cache storage]. IBM 3590/3592 tape cartridges, nonetheless, are the primary, data storage media used by the 3494-VTS – and that is the drawback of this approach.

The Bus-Tech MAS, on the other hand, rather than just using disks for caching, uses open system disk systems as its primary (and only) data storage media. Thus the MAS eliminates the need for tape cartridges and tape drives – though it appears to the mainframe and mainframe software as series of up to 64 IBM 3480/3490 tape drives. Hence why the MAS is known as a 'TAPE-ON-DISK' solution – rather than a VTS, because there is a big difference in these two approaches when it comes to performance,

data integrity, reliability and costs. The MAS by emulating physical tape drives and tape volumes using permanently mounted open systems disk storage is now a far better solution to the technologically-dated VTS approach, *as discussed below*.

Since the IBM VTS only uses disks as temporary cache storage its overall read/write performance is still gated by the performance of tape I/O operations – in particular cartridge retrieval in the case of data access. As with all caching schemes, accessing any data that is not within the cache can be a slow process, and this is definitely the case with the VTS.

In general, the VTS [like the MAS] can quickly respond to 'scratch' mount requests since such new files can be initially created on the disk cache – always assuming, however, that the cache storage is not 'full' due to slow tape operations. But this is certainly not the case when it comes to accessing data from a VTS.

When the mainframe tries to access a file that has previously been written to tape, the VTS has to first retrieve the volume containing the requisite file from the tape library. This first step in this process is to locate the relevant cartridge and physically retrieve it using the 3494 robotic arm [i.e. the so called 'cartridge accessor']. Then a free [i.e. currently unused] tape drive has to be found on which to mount the cartridge.

If all the tape drives are busy, the cartridge will have to be queued until a tape drive becomes available. Once the cartridge has been successfully loaded into a tape drive, the VTS insists on physically reading the entire volume onto its disk cache! The required file is then extracted from the disk cache and forwarded to the mainframe.

IBM's published performance figures for VTS systems [i.e. 'IBM TotalStorage Virtual Tape Storage Performance', by Aare Onton and Jesse L. Thrall, 28 August 2001, available on IBM's Web site] states that the best case response time for this type of read request can be in the order of 2 to 3 minutes. In total contrast, such read operations on a MAS are fulfilled in 5 seconds or less! That is the difference of using disks as the primary data storage media, rather than just as a temporary caching scheme.

Since the MAS is entirely disk-centric, and does not use any tape drives or tape cartridges, it does not suffer, whatsoever, from response time degradations caused by:

- tape cartridge access by the mechanical robotic arm.
- delays waiting for a tape drive to become available for use.
- time taken to physically mount and activate a tape cartridge.
- tape positioning [e.g. fast forward, skip to block offset] delays.

The bottom line here is that the MAS, being unhampered by slow mechanical arms or tape drives, will always deliver significantly faster overall performance than a VTS. The response times for a MAS are, thus, stated and measured in seconds, while those for a VTS tends to be in minutes.

Despite these differences, the 'virtual tape' emulation performed by the 3494-VTS and the MAS are, however, equivalent as far as mainframe operating systems and software is concerned – with both capable of emulating up to 64 virtual tape drives. The fundamental and crucial difference between two, however, being that the MAS never resorts to using actual tape cartridges.

Given the equivalence of the mainframe-end emulation performed by the MAS and the 3494-VTS, the MAS also delivers, without exception, all of the performance and operational advantages associated with the 3494-VTS. The MAS, however, given its use of disks for permanent data storage, goes onto deliver an additional set of advantages that are not possible with a 3494-VTS as described above. In addition to the incontrovertible increases in performance, eliminating tape cartridges, tape drives and mechanical robotic assemblies, without doubt, and at a stroke, enhances RAS, minimizes data loss due to media corruption and slashes maintenance costs. 'TAPE-ON-DISK', in essence, totally modernizes VTS technology, exploiting the latest (and ongoing) advances in open system storage. The table below summarizes the superiority of the MAS approach over that of a 3494-VTS:

MAS EQUALS THE IBM 3494-VTS	**MAS EXCEEDS THE IBM 3494-VTS**
1. significantly reduce the time taken for 'scratch' mounts. 2. eliminates poor utilization of tape cartridge usage. 3. support WORM [i.e. 'write once, read many times'] mode operation.	1. Instant access to previously created files [versus 2(+) minutes on a VTS]. 2. significantly expedites all data transfer and data location [e.g. fast forward, block offset] operations. 3. tangibly reduce floor space requirements. 4. minimize data loss due to media corruption. 5. eliminate the need for tape cartridges and tape drives. 6. eliminate the need for failure-prone mechanical robotic assemblies and tape drives

Summary

The Bus-Tech MAS is a proven solution for modernizing tape-based mainframe applications without the need for rewriting any of the software. Though emulating VTS operations, it totally modernizes the VTS concept by totally eliminating the use of slow and failure-prone tape drives and tape cartridges. It is a high-throughput, 'near-zero' latency, TAPE-ON-DISK virtual tape controller that supports 2Gbps FICON Express mainframe channels, 2Gbps Fibre Channel SAN connections as well as Ultra SCSI attachments. It faithfully emulates IBM 3480/3490 tape unit operation using IBM's prescribed "AWSTAPE" tape emulation file format. The MAS eliminates the need for tape cartridges and obsoletes expensive and space-hogging tape libraries. It can be an extremely cost efficient solution, with a 90TB MAS solution costing the same as a 20TB IBM 3494-VTS!

The MAS delivers near instantaneous tape mounts, Fast Forward to EOF in less than 50 milliseconds and locate any specified block of data within 100 milliseconds. Since all tape volumes are emulated on disks, there is never a need for operator intervention. It thus greatly expedites tape-centric mainframe operations – in particular nightly tape backups. Backups, including disaster recovery, can be realized in unattended mode, thus permitting considerable streamlining of data center shift staffing. Moreover, the use of affordable disk storage, including the emerging ATA technology disks, minimizes the data recovery vulnerabilities brought about by tape media deterioration and tape volume mix-up.