

Mainframe Data Library with HDS Content Platform (HCP)

This paper provides an overview of support for Hitachi Data Systems' Content Platform offered with the Bus-Tech Mainframe Data Library

Background

The amount of new information and content that large-scale data centers are required to create, distribute, and store digitally continues to grow dramatically. And, in addition to protecting and preserving data, today's corporations are also looking for cost-effective alternatives for keeping archive data readily accessible. The time when it would be considered acceptable to take hours to retrieve archived data is over. Business managers are demanding that data be readily available when needed—even when that data may be months, or even years, old.

In addition, corporations are now required to comply with regulations at both federal and state levels. Compliance impacts how they manage their data with regards to retention, accessibility, integrity, security, and privacy. Examples of regulations include:

- 17 CFR 240.17a-4
- Electronic signatures in Global & National Commerce Act
- Gramm Leach-Bliley Act
- Health Insurance Portability and Accountability Act (HIPAA)
- 21 Code of Federal Regulations Part 11
- California Database Protection Act
- Sarbanes Oxley Act of 2002

The impact of these new requirements is forcing many IBM mainframe data centers to look at both the methods and technology they use to store long-term archive data.

Most mainframe-based archive applications work on the principle that data needs to be readily available for retrieval and access for a relatively short period of time: typically 30, 60, or even 90 days. These applications are designed so that data is kept on a Direct Access Storage Device (DASD) during that time, making it instantly available for queries and background processing requests. When the

data is deemed to be no longer in high demand, it is migrated to less expensive storage (usually tape or optical) where it can be permanently kept at a reasonable cost.

But once moved off of online DASD, retrieval of archived data is, at best, inconsistent. Data retrieval of an individual record or report from tape or optical storage can take anywhere from 10 seconds to several hours, depending on the type of storage being used and the availability of resources at the time of the request.

For example, a typical automatic tape library (ATL) may contain 4,000 to 6,000 tape cartridge volumes yet have only 16 or 32 tape drives on which to mount cartridges. To retrieve a record from a tape volume in an automated tape library (ATL) requires that: a tape drive in the library is available for allocation; the robotic arm in the library is not active doing something else; and the tape needed is not being used by another task in the system. When all those conditions are satisfied, the tape can be mounted and the record retrieved, typically in 10 to 45 seconds. But if any of conditions are not met then the request for data must wait.

And, while traditionally less expensive than mainframe DASD, storing massive amounts of archive data in large-scale automated tape libraries (ATLs) is itself an expensive proposition. ATLs require considerable amounts of floor space, electricity, preventative maintenance, and an ongoing investment in tape cartridges.

Mainframe Data Library

Bus-Tech's Mainframe Data Library (MDL) provides a cost-effective alternative for storing long-term mainframe archive data. The MDL is a family of tape-on-disk controllers for IBM and compatible mainframes. An MDL connects to the mainframe using 4 Gbit FICON or ESCON channel interfaces and emulates standard IBM tape drives, including 3490, 3490, and/or 3590.

The mainframe and its applications see the MDL simply as a collection of tape drives that can be used just like any other tape drive to archive and back up data. But instead of using expensive robotic tape libraries and cartridges, the MDL writes tape volume images to IP-attached open-systems disks. Later, when the mainframe application requests that a tape be mounted for processing, the MDL simply opens the file containing the VOLSER being requested and immediately responds to the mainframe's mount request. The time required to initially access

any record on the tape volume is usually under 1 second. And the MDL never has to wait for a robotic arm to become available.


HCP

Hitachi Data Systems' Content Platform (HCP) delivers a best-of-breed active archive solution designed to meet large-scale storage growth while addressing the need for corporations to preserve, protect, and access data well into the future.

HCP implements a unique storage solution that combines records management and archive sciences practices with industry-leading storage technology. The HCP solution allows organizations to satisfy regulatory and corporate governance mandates and provides rapid access to archived data when needed. The HCP also delivers automated policy management for data in order to ensure content authenticity, retention, and integrity.

MDL and HCP

Bus-Tech and Hitachi Data Systems have partnered to offer enterprise-level IBM mainframe data centers a virtual tape-on-disk solution built on Hitachi's compliant storage. Using this technology, mainframe tape applications can write tape volumes to an MDL where the tape volume images will be stored on HCP.



Furthermore, support for HCP has been incorporated across the entire family of MDL controllers. Mainframe-attached MDL/HCP solutions can scale from a single FICON channel emulating from 16 to 256 tape drives to 12 FICON or ESCON channels providing up to 1,536 tape drives.

As the MDL writes new tape volumes (VOLSERs) to the HCP, expiration dates set in standard IBM tape header records are used to define the retention period that HCP should use to protect the volume. Together, the MDL and HCP give data center managers a method to easily move existing tape or optical archive and backup applications to a compliant storage platform without the need for expensive up-front application re-engineering. Simply point your archive applications to MDL-emulated tape devices to immediately get data under regulatory control.

And, as with any MDL-based solution, retrieval of archive data is quick and consistent since HCP is designed to make archive data readily available at any time.

Solution Benefits

The benefits of implementing an MDL backed by Hitachi Data Systems' Content Platform include:

- Reduced tape processing and archive application processing costs
- Improved end-user service levels
- Innovative disaster recovery alternatives
- Mainframe-compliant data storage

Once archive data is placed on an MDL-to-HCP storage solution, cost savings from operating archive applications and their associated tape sub-systems can immediately be realized.

Data that once had to reside on mainframe DASD for months—or even years—in order to be quickly retrieved can be migrated almost immediately to the MDL/HCP environment. This frees valuable, channel-attached DASD for other purposes, lowering overall operating costs for the data center.


Additionally, moving archive data from tape usually maintained in a large ATL to HCP realizes immediate savings in tape processing costs, including reduced floor space requirements; reduced electricity needs; reduced monthly maintenance costs associated with mechanical robotic arms; and reduced consumption of expensive large-volume tape cartridges.

Best of all, while on-going operational costs are reduced, service levels to end users are improved. Retrieval time of archived data is no longer depends on how old the data is. Whether data is a few hours old and still resident on high-speed

mainframe DASD or whether it has already been migrated off to the archive platform, users can retrieve the data quickly. Typically, users are unable to tell if the data is DASD or long-term archive storage when they retrieve it.

Providing offsite data storage for disaster recovery has become a major requirement for most large data centers. And duplicating tape library data in a remote location is no simple task. Alternatives for replicating data often involve some combination of duplicate tape writes, shipping of backup cartridges, and/or expensive tape library mirroring using channel extenders, routers, and duplicate library hardware.

MDL with HCP provides innovative alternatives for satisfying offsite storage requirements for disaster recovery. HCP provides unidirectional cluster-to-cluster data replication specifically intended to deliver a replica of a production HCP cluster to a disaster recovery site. And since the MDL accesses data via IP, access to data at the remote disaster recovery from the primary data center is available.



For organizations that demand backup of their compliant archive disk solution, HCP supports NDMP backup to a physical tape library and supports a choice of leading backup solutions such as Hitachi Data Protection Suite, CommVault Galaxy, NetBackup, and many others. The backup set can be signed, encrypted, and/or compressed.

Whether needing to comply with HIPAA, Saxbanes Oxley, or other federal or corporate compliance mandates, Bus-Tech and HDS Content Archive Platform can help. The MDL with HCP solution provides compliant storage for the IBM mainframe, so data center managers can begin moving their critical archive data onto compliant storage with ease. Data can be protected just by setting an expiration date for tape volumes when they are created and MDL and HCP will take care of the rest.

The MDL reads the expiration date from the first HDR1 record found on new tape volumes it is creating. It will use that date to set the retention period for the volume image stored on the HCP. Once retention is set, it is not possible to delete the tape volume until it has expired. Even if the mainframe tape management system incorrectly scratches a tape before its retention has expired, the volume cannot and will not be scratched. Any attempt to scratch it will be terminated.

Summary

The world of “always online” systems brought about by the public Internet has resulted in the requirement for near-immediate retrieval of data—regardless of its age or frequency of use. For large-scale IBM mainframe data centers, this need for constant and consistent retrieval of data has placed new pressures on the hardware and systems traditionally used to store archive data.

At the same time there is a new awareness and need to provide security and control over data, including the need to satisfy government regulatory requirements that require corporations to be able to provide an audit trail of the data they are storing.

As a result of these developments, corporations need new approaches to data storage that can satisfy compliance issues, provide rapid access, and, at the same time, maintain a reasonable level of expense.

The Bus-Tech Mainframe Data Library allows organizations to replace expensive tape subsystems and inconsistent data retrieval with reliable open-systems disks. Hitachi Data Systems' Content Platform provides an active archive data storage platform that meets compliance requirements and provides rapid access to data—regardless of its age.

By combining mainframe tape-on-disk technology with compliant, rapid-access archive storage, Bus-Tech and Hitachi Data Systems bring to market a truly unique solution. Enterprises that use tape-based archive applications with access methods such as Mobius Viewdirect or IBM OAM can now use compliant, open-systems storage to reduce their ongoing operational costs, improve service level response times, and improve disaster recovery performance. And, since the combined Bus-Tech/Hitachi Data Systems' solution appears to the mainframe as channel-attached tape storage, there is no need to re-engineer your applications in order to get started.