

DLm™ TRANSFORMS MAINFRAME TAPE! WHY DELL EMC DISK LIBRARY FOR MAINFRAME?

The Business Value of Disk Library for mainframe

OVERVIEW OF THE BENEFITS OF DLm

- DLm is designed to reduce capital and operational expense by replacing physical tape and inefficient virtual tape systems
- DLm enables replaces aging, degrading physical tape with cloud storage for more efficient long-term retention of data
- DLm reduces replication bandwidth expense using the integrated data deduplication capability of PowerProtect DD and legacy Data Domain storage.
- DLm is designed to reduce the operational labor and system overhead of manual failover procedures using field-proven Dell GDDR technology.
- DLm D/R testing read/write snaps won't interrupt production performance.
- DLm is designed to achieve 1 second or less average access times by eliminating tape latency and reducing batch windows consistently, regardless of data age or cartridge size.
- DLm release 7.0 introduces models DLm2700 and DLm8700 which replace existing models DLm2500 and DLm8500. These new models feature new Virtual Tape Engines (VTEs) with 60% more processor cores, 32Gb FICON and 25GbE storage connections.
- DLm release 5.5 introduced attachment of up to four PowerProtect DD or legacy Data Domain Systems or up to two PowerMax models 2500, 8000, 8500. Support for ADABAS backups is also enhanced via AMDD support
- DLm release 5.4 introduced Dell EMC PowerScale and Isilon storage options as well as supporting S3 protocol from DLm8500 to ECS for Long Term Retention
- Reduce CPU MIPS expense via faster HSM recycle times (Eliminate tape drive contention and perform recycles at disk speeds.)

- Reduce data center footprint expenses by replacing large physical tape libraries with DLm (reducing the footprint and associated costs; power, AC, etc.)
- Improve Service Levels (performance, availability), reduced costs, and staff handling tape, which enables better SLAs.

Revisions

Date	Description
July 2019	Initial release for DLm R5.1
July 2020	Initial release for DLm R5.2
October 2020	Initial release for DLm R5.3
June 2021	Initial release for DLm R5.4
July 2022	Initial release for DLm R5.5 and rebranding to Dell Technologies
November 2024	Initial release for DLm R7.0

Acknowledgements

This paper was produced by the following:
 Authors: Chuck Farah, Paul Scheuer
 Support: DLm Product Engineering

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2024 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners, July 2024. White paper, h11879.12

TABLE OF CONTENTS

EXECUTIVE SUMMARY	5
Dell DLm transforms mainframe tape!	5
Introduction	5
Audience	5
DELL DLM - TECHNOLOGY DIFFERENTIATION & LEADERSHIP	6
Support for all mainframe tape workloads	6
WORM compliance (DLm8700 only)	7
Enhanced data protection with RSA encryption (DLM8700 only)	7
Looks like tape drives to the mainframe host	7
No wasted space	7
Sub-second access to tape data	7
Fast locate	8
No single point of failure	8
Tape drives can be shared	8
Use the Cloud to simplify replacing physical tape for long-term retention (DLm8700)	9
GDDR Tape for automated failover pioneered by Dell GDDR (DLm8700)	10
Remote replication	11
Disaster recovery testing	11
HSM L0 to ML2 migration directly	11
Management reporting	12
Tape status management	12
Compression and deduplication	12
PowerMax connectivity (DLm8700)	12
Evolution of mainframe tape technology	13
THE BENEFITS OF DLM	13
Simple operation	13
Quicker access & higher throughput	14
More reliability	14
Improved disaster recovery	14
Significant cost savings	14
Smaller footprint	14

Large scalability 15

Improved security & compliance 15

Enhanced storage efficiency with deduplication..... 15

Reduced risk of litigation or public embarrassment..... 15

Conclusion..... 15

EXECUTIVE SUMMARY

Dell DLm transforms mainframe tape!

Traditional mainframe tape suffers from numerous challenges including:

- Costs – floor space, tape cartridge handling, and CPU cycles used for L0 to ML1 migrations
- Tape problems – physical tape drive and media failures, lost tapes, and numerous points of failure, as well as isolated views of tape volumes
- Slow response time – HSM recycle and time to first byte
- Limited disaster recovery – slow recovery, limited RTO/RPO granularity and cumbersome DR testing

Mainframe tape operations do not have to be limited by the above issues. Dell Disk Library for mainframe (DLm) is a high performance, highly reliable disk-based physical tape replacement that truly transforms mainframe tape environments.

Introduction

This paper focuses on the technologies within Dell Disk Library for mainframe that differentiates it from other tape solutions in the market and explains why it matters. Its purpose is to explore the technical and financial reasons why Dell DLm can dramatically improve the tape operations for your mainframe environment.

Audience

This paper is intended for Dell customers, Dell sales, Dell Systems Engineers, Dell partners, and anyone else who is interested in learning more about DLm's differentiating technology and the unique advantages that DLm can provide.

Dell DLm - Technology differentiation & leadership

Support for all mainframe tape workloads

Dell Disk Library for mainframe is designed to improve operations for all typical mainframe tape use cases and was first to market in 2008 with a single-vendor 100% disk-based solution.

- **HSM Migration** - identifies data that has not been accessed in some time and moves active data (level zero – L0) to Migration Level One (ML1) someplace else on the tier-1 disk arrays. Mainframe CPU cycles are used to compress the data on writes and to decompress the data during recalls. After data has resided in ML1 for a specified time, it is migrated to Migration Level Two (ML2), which typically resides on tape.
- **Long Term Data retention** – quite often, unchanging fixed content reference data (objects) need to be saved for a decade or more. There are regulatory requirements stipulating data cannot be modified for several years however the data must be preserved. After a pre-determined amount of time, objects are removed from tier-1 storage. Before DLm release 4.5, these volumes were written to physical tape & the process was very labor intensive. Long Term Retention (LTR) provides customers the ability to select data based on their retention requirements and move to the appropriate storage tier optimizing their storage investments.
- **Backup and Recovery** - after a full backup is taken, typically additional backups are incremental and done on a daily basis during the week. Successive incremental backups are full of duplicate data that can benefit greatly from data deduplication.
- **Work Tapes** - Basically, everything else, including batch processing, sort work tapes, etc. Many of these applications do not function properly in a limited cache environment such as VTS or VSM. The unlimited cache of the DLm allows these applications to execute without limitations.

Why DLm for Mainframe Data?

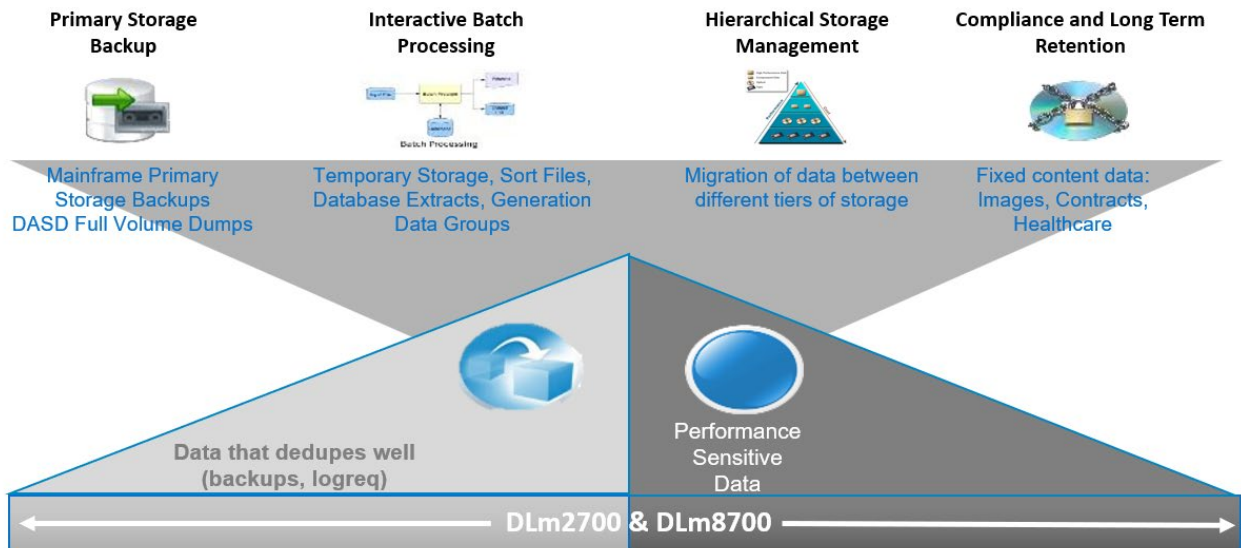


Figure 1 Typical mainframe tape use cases

WORM compliance (DLm8700 only)

DLm8700, when deployed with PowerProtect DD/Data Domain storage, supports Write Once Read Many (WORM) requirements. This guarantees that the data cannot be modified until a specific expiration date has passed. This feature is critical to help your company meet statutory requirements for SEC or OCC compliance.

Enhanced data protection with RSA encryption (DLm8700 only)

Companies are becoming increasingly aware of the need to protect sensitive customer data regardless of the industry. We have all read or experienced what can happen when customer data is lost by accident or intentionally stolen and manipulated. Enhanced data protection comes from implementing D@RE (Data at Rest Encryption) for PowerMax, PowerScale, legacy VNX and Data Domain storage and KMIP compliant encryption key management.

Looks like tape drives to the mainframe host

The mainframe host views the DLm as physical tape drives. A single Virtual Tape Engine (VTE) emulates up to 256 virtual tape drives (DLm2700) or 512 virtual tape drives (DLm8700). A VTE can access part or all the tape volumes on the back end. The tape volumes (VOLSER) are written to disk as files. Each file will only occupy the amount of space it requires.

No wasted space

The DLm is very efficient with space utilization. The size of virtual cartridges is user definable and can vary depending on intended use. If the application writes out a 50 MB file to a virtual cartridge that is defined at 2 TB, the DLm only consumes enough space on the disk to store the data set (50 MB). This eliminates the need for any tape stacking processing.

Sub-second access to tape data

“Time to first byte”, illustrated in figure 2 below, is the time from the request of the tape volume until the time the information on tape is being read. When using physical tapes, the application must wait for the tape to be mounted. Then, the tape drive needs to perform seek and stage operations before starting to process the information on the tape. Assuming the tape drives are always available, the typical time to first byte with physical tape is around ~13 seconds. Of course, if the drives are busy processing other tapes, there is no way of knowing how long the time to first byte could be.

Because the DLm offers a large number of virtual drives and stores data on spinning disk, mount requests are typically satisfied in less than one second. This eliminates the latency associated with physical tape drives and avoids drive contention.

Faster Performance vs. Tape

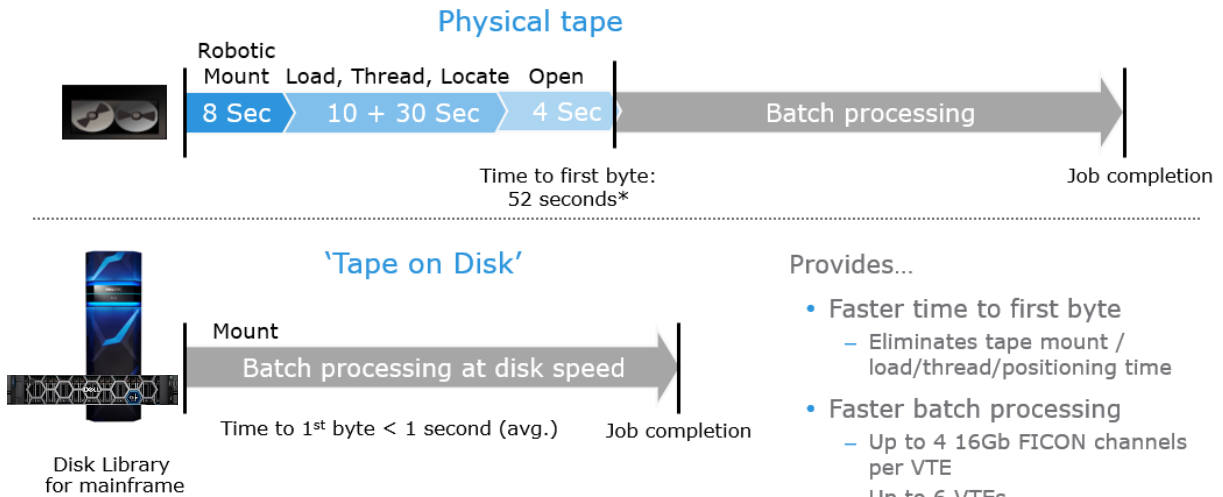


Figure 2 Physical tape vs. DLM access time

Fast locate

Fast locate capabilities allow the DLM to take advantage of the fact that the information resides on disk, which enables random access directly to the required information without the need to read through the entire tape information. When the mainframe host sends a read request, DLM can then access directly to the relative closest block to the required information, taking advantage of the random access method provided by disk. This is done in a transparent fashion to the application and host operating system.

No single point of failure

DLM has no single point of failure. As discussed earlier, each Virtual Tape Engine (VTE) can access all the tape volumes, which ensures that you have access to your tape volumes at all times. The disk drives inside the DLM are protected with RAID 6. Hot spare disk drives are distributed throughout the system to provide additional protection. Additionally, DLM microcode updates and corrective maintenance activities are performed concurrently without the need to interrupt tape operations providing superior availability.

Tape drives can be shared

Many companies need to share their tape drives among several LPARs. When using physical tape libraries, companies usually require external software to enable sharing due to the limited number of tape drives available. With DLM, each of the VTEs can emulate up to 512 tape drives. A single DLM can emulate up to a total of 3,072 tape drives. Coupling two DLM's together, a total of 6,144 drives can share a common view of the tape library. This gives you the ability to define enough tape drives for each of your LPARs and eliminate the need for expensive tape-sharing software. DLM systems can be shared up to 64 LPARs.

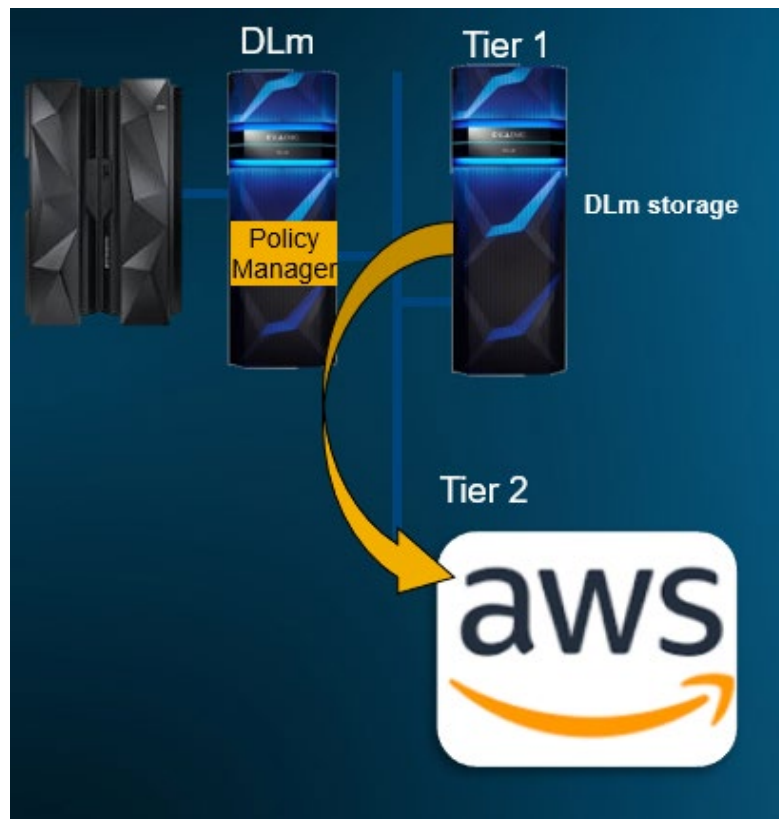
Use the Cloud to simplify replacing physical tape for long-term retention (DLm8700)

Data that need to be stored for extended periods of time, often decades, continue to increase dramatically, and storage administrators are under tremendous cost pressure to store them economically. Until now, physical tape was the only viable option to meet these simultaneous demands. However, the maturity and affordability of the cloud, both public as well as private, offer viable alternatives with considerable benefit compared to physical tape. DLm release 5.3 enabled storage administrators to take advantage of Amazon Cloud (AWS) by using S3 protocol. DLm's built-in policy manager enables storage administrators to plan and automate the movement of volumes between DLm's primary storage and the cloud.

All the features offered for LTR in previous releases are supported by S3 LTR to AWS.

Release 5.1 for the DLm8500 offered several additional long-term retention operational benefits to using LTR engine over the prior release:

- Allow configuring '0' days in policy "move after" field.
- Changed maximum number of "move after" days from 730 (2 years) to 19999 (47 years).
- Migration policy changes can now be reconfigured non-disruptively to DLm
- Added MIGRATE FAILURES LIST and MIGRATE FAILURES CLEAR to examine and clear list of previously failed migrations, so that they can be retried without restarting the VTE .
- Significant performance improvement of the post-migration file-validation process.
- Single, simple, "restore" command to restore tape volumes from ECS Tier 2



GDDR Tape for automated failover pioneered by Dell GDDR (DLm8700)

Today, mainframe tape storage must be as reliable as DASD, and DLm is no exception. Starting with release 4.5, DLm's industry leading High Availability architecture has been enhanced with GDDR (Geographically Dispersed Disaster Restart) technology that has been used to automate failover of Dell DASD for generations. GDDR Tape eliminates the need for a complex, ever-changing compilation of scripts and manual procedures for both DR tests and actual disaster failover. GDDR Tape uses a "heartbeat" to monitor the health of the DLms across sites and alerts the storage administrator if it determines action needs to be taken in the event of an actual outage.



Dell replication software enables network-efficient replication to one or more disaster recovery sites. Data can be encrypted in-flight when being replicated between systems.

Using Dell snapshot technology, storage administrators can perform complete end-to-end DR testing with read/write capabilities on all tape data at the target site. Dell DLm was designed to give storage administrators 100 percent confidence in their disaster recovery (DR) readiness with the least amount of set-up. In addition, replication continues uninterrupted during DR testing. When testing is complete, the snapshot is simply deleted without affecting the existing backup tape volumes.

Remote replication

DLM provides users with the option to perform asynchronous replication between two or three DLM systems. DLM replication is done by utilizing back-end storage IP replication and does not affect the application running on the mainframe host. Only the changed data is replicated to the remote site. You can set RPOs and quality of service levels based on your service level agreements and network bandwidth. Replication can be bi-directional so a source system can also be a target system for remote DLM and vice versa. DLM provides the ability to perform disaster recovery tests while continuing replication and being protected even during the disaster recovery tests.

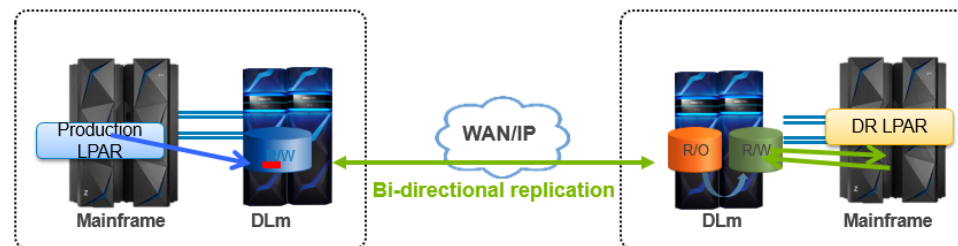
Guaranteed replication ensures a tape volume is replicated from one DLM to another, located locally or remote, before releasing the job for completion. Specific tape volumes can be selected for guaranteed replication, giving the user full control over their system. Guaranteed replication provides a much more “synchronous” approach to replication without adding network overhead when compared to other vendors.

Disaster recovery testing

Mainframe datacenters typically need to perform disaster recovery tests several times a year. These tests can be very complex, requiring DR teams to handle hundreds, if not thousands, of tape volumes. Competitive virtual tape libraries typically require replication to shut down for the duration of the disaster recovery test. This leaves the data center unprotected. With DLM, you do not need to shut down replication and can perform your disaster recovery tests while continuing to be protected. There are two ways you can perform disaster recovery tests with DLM as illustrated in figure 3 below:

- You can mount the DLM volumes at the remote site as read-only and test that tapes can be mounted and all data can be read.
- You can create snapshots of your tape volumes and mount them as read-write at the remote site and perform full disaster recovery tests.

Typical Disaster Recovery Testing with DLM



Read-only mounts

Disk arrays allow instant “read-only” copies
Confirm that tapes can be mounted and all required data can be accessed
No incremental storage capacity required

Snapshots

Disk arrays allow creation of “read-write” snapshot
Confirm operation at the disaster recovery site
Some incremental storage capacity required

Remote replication is uninterrupted during testing

Figure 3 DLM provides two ways to improve DR testing

HSM L0 to ML2 migration directly

HSM migration identifies data that has not been accessed in some time and moves that data from the active data tier (level zero – L0) to a secondary Migration Level One (ML1) disk storage tier and then after a specified period of time, the data is migrated to a tape tier, Migration Level Two (ML2). Mainframe CPU cycles are used to compress the data

on writes and additional cycles for decompression during recalls. As an alternative the HSM application can be set to bypass the L0 to ML1 step and go directly from L0 to ML2 and use the hardware compression feature of the DLm. This results in reducing mainframe CPU cycles and reduces the amount of tier-1 disk space needed.

Management reporting

The DLm Command Processor Facility enables you to submit commands to DLm through your mainframe host in batch. These commands include the ability to retrieve information about DLm and receive the information from the batch job output. You can easily retrieve information such as available space, compression ratios, etc. about your DLm systems through the mainframe host.

Tape status management

DLm manages and maintains tape status log files. These logs are maintained inside the DLm storage array. Using these logs, you can see the status of each tape in the DLm at any given point in time. Because the logs are maintained in the DLm storage, they are also replicated to a remote DLm system. This gives you the ability to know which tapes open and which tapes are closed at the time of replication during a DR event.

Compression and deduplication

All DLm Virtual Tape Engines (VTE) support hardware compression, which typically ranges between 3:1 and 4:1; however, since compression depends on data, higher compression rates are possible. Backups in mainframe environments can be a particular pain point. The percentage of data devoted to backup varies widely by installation but is typically between 20 to 30 percent of the total data. When leveraging deduplication in the storage connected to the DLm, DLm will disable compression processing done at the VTE as to not negatively affect the deduplication process.

Deduplication has a major impact on backup. By reducing the data footprint even by a 10 to 1 factor, backup files use less bandwidth and can get to the disaster recovery site much sooner. DLm leverages Dell PowerProtect DD and legacy Data Domain system's industry leading deduplication technology to maximize storage efficiency and performance.

PowerMax connectivity (DLm8700)

The DLm8700 supports PowerMax storage array models 2500, 8000 and 8500 which utilize SRDF/S (remote replication) and Consistency Groups to ensure Universal Data Consistency between DASD and tape data at identical points in time in production and recovery sites. Tape and disk consistency are essential to applications that rely on multiple datasets (for example, data and log information) being time consistent as well as data consistent. In other words, these datasets are known to be written at the same point in time, with data that correlates precisely with its log or metadata, thus avoiding a potentially lengthy resync between these datasets after a DR event. One example would be DB2 logs and user data.

Synchronous replication between primary and secondary (DR) DLm+PowerMax 8000 sites leverages SRDF/S, which offers improved performance and scalability over IBM's synchronous replication methodology based on "sync points".

The DLm8700 also supports for SRDF/A (asynchronous replication) using Multisession Consistency (MSC) for out-of-region (geographically beyond synchronous distances) locations supporting a three-site STAR configuration. Dell Geographically Dispersed Disaster Restart

(GDDR) product automates disaster restart of applications and systems in mainframe environments in the event of a planned or an unplanned outage.

Evolution of mainframe tape technology

Mainframe tape has evolved since the early days of magnetic reel tape. Dell Disk Library for mainframe has transformed mainframe tape with all its industry leading features and capabilities and represents the evolution of mainframe tape technology as illustrated in figure 4 below.

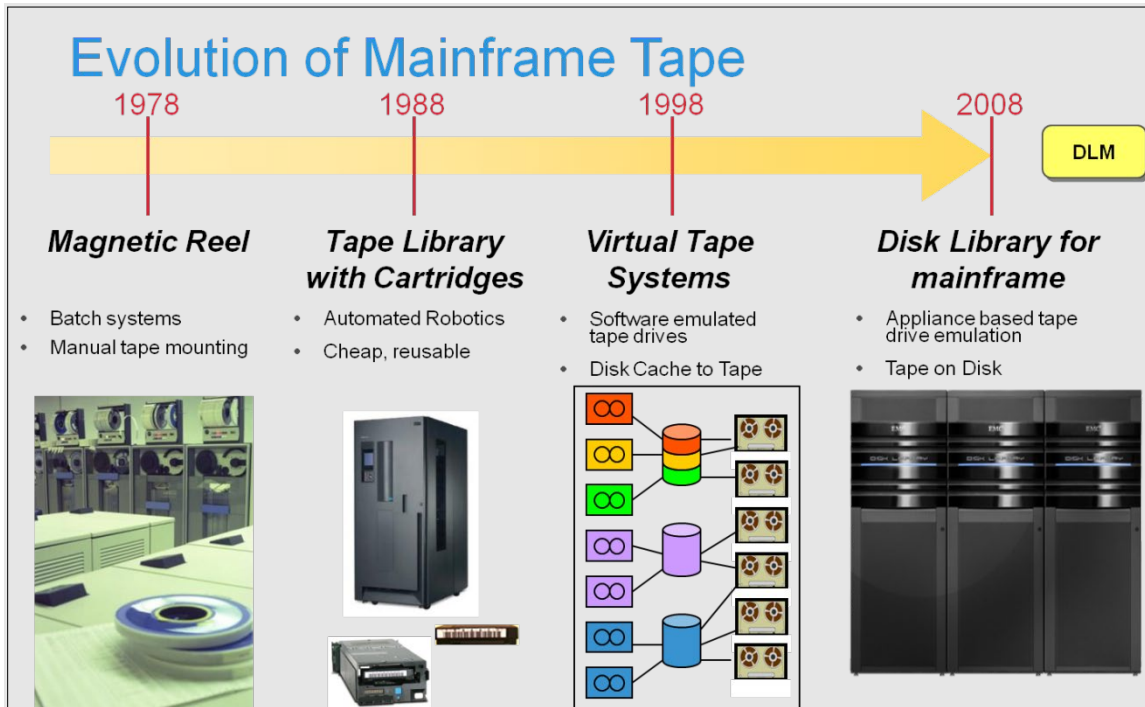


Figure 4 DLm represents the evolution of mainframe tape

The benefits of DLm

Simple operation

Disk Library for mainframe is 100% IBM compatible and emulates true IBM tapes—3480/3490 and 3590—and is transparent to the mainframe host. You do not have to define any started task or any subsystems for DLm to operate. DLm can work seamlessly with any mainframe application such as SMS, HSM, backups, and others and no Job Control Language (JCL) changes are needed. This means DLm is easy to insert into your existing mainframe environment and simple to use because your staff already knows how to use it.

Quicker access & higher throughput

DLM delivers consistently higher performance for both read and write operations at all times regardless of the age of the data because the information is kept on disk. DLM boasts up to 6 GB/s throughput, which is twice as fast as Dell's closest competitor. The typical time to first byte with physical tapes is around ~13 seconds assuming the tape drive is available. With DLM, the time it takes to locate and mount the tape volume averages one second or less. In addition, HSM recycle times are done at disk speeds eliminating contention for tape drives. These high-performance capabilities result in all tape operations being dramatically faster shortening batch and backup windows, providing faster access to data and giving you the ability to provide higher SLAs to your organization.

More reliability

Uptime and system availability remain critically important for mainframe customers. Unlike physical tape subsystems which inherently have multiple points of failure including tape robotics and physical tape drives, DLM was designed to have no single point of failure and allows for non-disruptive code updates for the Virtual Tape Engines and storage. This means that you can depend on tape operations being more stable, more reliable, and help you achieve better SLAs. DLM offers several HA (High Availability) configurations.

Improved disaster recovery

DLM remote replication allows you to set up RPOs and quality of service based on your service level agreements and network bandwidth. With DLM, RTO and RPO can be reduced from days down to hours or minutes. DLM also provides users with the ability to perform 100% of their disaster recovery tests in full read/write mode without interruption of normal tape processing or production replication. It also provides two options to simplify and improve regular disaster recovery testing with read only mounts or using snapshots. No matter which option you choose to perform disaster recovery tests, you do not need to turn off your replication and stay protected at all times. With DLM the files are already at the disaster recovery site where snapshot copies can be made and complete disaster recovery testing can easily be done. In most cases it isn't even necessary to have people onsite to complete the DR testing!

Significant cost savings

DLM provides significant cost savings in many different ways. First, DLM eliminates physical tape associated costs of power, cooling, maintenance, licensing, floor space for physical tape libraries, and tape cartridge recall and storage costs. Second, DLM eliminates the FTE effort associated with loading and unloading tape cartridges, tape storage, and transportation. Because DLM uses hardware compression, you are already saving your data on disk and compacting it based on your migration strategies. Now, by storing your tapes on disk and using compression inside the Disk Library for mainframe, you can eliminate your ML1 requirement altogether. This lets you reclaim the expensive host CPU cycles that were used for on-host compression that ML1 was using and also reclaim the tier-1 storage allocated to ML1. You can also eliminate expensive CPU cycles due to faster HSM recycle times because with DLM there are no tape drive conflicts and HSM recycle is all done at disk speeds.

Smaller footprint

DLM configurations can be as small as a single floor tile depending on the DLM model, the number of Virtual Tape Engines, and the total storage capacity. DLM takes up

significantly less floor space than the physical tape infrastructure it replaces freeing up valuable data center floor space that can be used for other purposes.

Large scalability

DLM can start out very small and inexpensive (DLM2700) and scale up to 10 PBs of logical storage capacity which is almost 4 times as much as our competition, up to 4,096 virtual drives, and virtual tape cartridge sizes up to 16TB. This means a single DLM system is more than capable of meeting the tape operational needs of the most demanding mainframe customers while providing unmatched performance, reliability and improved disaster recovery. And DLM scales simply without the complexity of adding additional subsystems, libraries, network connections, etc. and as it scales in capacity, you are also adding more performance.

Improved security & compliance

Companies are becoming increasingly aware of the need to protect sensitive customer data regardless of the industry. DLM supports most U.S. and International compliance requirements including SEC Rule 17a4(f) and customers can leverage data encryption if desired. DLM gives you the ability to meet your mainframe security and compliance requirements.

Enhanced storage efficiency with deduplication

Deduplication can provide additional storage efficiency and disaster recovery benefits for mainframe backup workloads. DLM with industry leading Dell PowerProtect DD and Data Domain deduplication storage typically achieves 10:1 or better storage reduction efficiencies from deduplication resulting in an even smaller footprint. This also translates to more cost-effective remote replication over IP bandwidth usage. IP replication can go to multiple sites and avoids expensive channel extension equipment.

Reduced risk of litigation or public embarrassment

Eliminating physical tapes with DLM mitigates the risks of public embarrassment or litigation due to lost or stolen tapes simply because there are no physical tapes to get lost or stolen. And DLM in-flight encryption ensures your data is protected as it is replicated from your primary site to your DR site.

Conclusion

After reading this paper you should have a better understanding how Dell Disk Library for mainframe can dramatically improve your mainframe tape operations.

To summarize, Dell DLM will help you:

- Provide long-term retention of tape volumes, replacing physical tape with Dell ECS private cloud solutions for DLM8700.
- Automate failover with field-proven GDDR technology for the DLM8700
- Reduce batch processing windows.
- Improve disaster recovery RTO/RPO.
- Simplify disaster recovery testing.
- Eliminate CPU compression cycles used for ML1 and reclaim that space.
- Save operational expenses for power, cooling.
- Free up valuable data center floor space.

- Improve SLAs to your organization
- Improve reliability of all tape operations with PowerScale, PowerProtect DD and legacy Isilon, Data Domain High Availability storage as well as PowerMax 2500, 85000 and 8000 (for the DLM 8700) .
- Overcome tape drive sharing limitations
- Eliminate physical tape handling
- Eliminate the risk of lost or stolen tape cartridges
- Reduce backup windows.
- Reduce HSM recycle time.

If you would like to know more about Dell Disk Library for mainframe, refer to the [Disk Library for mainframe data sheet](#) and the [Disk Library for Mainframe spec sheet](#)