

EMC Disk Library for Mainframe – DLm2000 Product Overview

Abstract

This white paper introduces the EMC DLm2000 which is the newest member of the DLm family. This system is based upon the same hardware and software technology used on the DLm6000.

The intent of the DLm2000 is to leverage the existing DLm6000 design where possible while also introducing a lower-cost member of the family suited to enterprises with less-demanding requirements.

The DLm2000 system base software is Virtuent 7, the same basic version of emulation software that runs on the DLm6000, but with support for the new hardware configuration.

February, 2012

Copyright © 2012 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation 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 EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

Part Number h10520

Table of Contents

Executive summary	4
Introduction	4
Audience.....	4
Challenge of Physical Tapes in Mainframe Environments	4
DLm2000 Overview	5
Product Description.....	5
Architecture.....	5
Virtual Tape Engines.....	6
DLm2000 Storage.....	6
Continuous Data Availability	7
RAID 6 Data Protection	7
Hot Spare Drive	7
Management and Supportability	8
WORM Tape.....	8
Remote Replication and Recovery.....	9
Flexible Recovery Testing.....	11
DLm2000 Typical Deployment Scenarios	12
Backup and Restores.....	12
Data Migration.....	13
Data Archive.....	13
Work Tapes	13
Conclusion	14

Executive summary

The EMC® Disk Library for mainframe family of products offers IBM System z mainframe customers the ability to replace their physical tape systems, including traditional virtual tape servers such as the IBM VTS and Oracle/STK VSM, with a dynamic virtual tape solution, eliminating the challenges tied to traditional tape-based processing.

DLM2000 is a virtual tape solution for replacing physical tape in the commercial mainframe environment. DLM2000 offers primary storage in a single manageable solution to address mainframe tape use cases. The DLM2000 comes with RAID 6 protected disk storage, hot-standby disks, tape emulation, and hardware compression.

The DLM2000 is a flexible VTL to enable mainframe tape replacement. The DLM2000 is the 3rd generation Disk Library for mainframe systems, representing an evolution and maturity to meet expanding customer requirements.

Introduction

This white paper begins by looking at tape usage in the mainframe environment and examines the challenges that physical tape presents. It then describes the DLM2000 tape replacement solution, including benefits, management and support, remote replication and recovery, and its architecture.

Audience

This white paper is intended for mainframe storage professionals who are looking to understand how the Disk Library for mainframe solution can help improve their mainframe tape operations without having to change their processes or perform a complete update to their tape infrastructure.

Challenge of Physical Tapes in Mainframe Environments

Mainframe data centers are highly dependent on tape systems for production operations as well as protection of their production and development data. This data protection strategy typically has required backup software and physical tapes, accessed either directly in an automated tape library (ATL) or through a caching virtual tape server (VTS). These tape systems often use tens of thousands of cartridges that require a significant amount of administration, physical management, and storage/floor space. Very often physical tapes can be lost or stolen resulting in failed restores and even failed disaster recovery operations.

Tape provides protection for data but the retrieval and/or recovery of the data stored on those tapes can be time-consuming and often unpredictable. As mechanical devices, tape drives naturally wear down over time. Head alignment can change, robotic arms can fail or jam and motors can speed up or slow down. The net result is that a physical tape drive cannot always access a tape when required.

Tapes, however, are fast and can support very fast throughputs; but this speed is achieved only after the tape is mounted and staged. A concern with tape is “time to first byte” since tapes must be mounted onto available tape drives, then the header must be read and the tape must be read sequentially to retrieve the relevant data. In many cases customers need to mount more tapes than the number of tape drives they have, which results in drive contention and lengthy “time to first byte” reads.

Virtual tape servers use temporary disk cache to store tape data. Tape files are then stacked and written to large format tape cartridges with the primary benefit being better tape media utilization. This means that when information is retrieved from tape, and no longer exists in the disk cache, it must be retrieved from physical tape to cache in its entirety before the application can begin to process the data.

Virtual tape servers have advantages over physical tape. Virtual tape servers do not deal with the load/unload, device load/eject, seek/rewind, and data transfer operations of tape drives integrated with the robotically controlled cartridge handling system. Virtual tape servers will outperform physical tape for time to first byte as long as the dataset is resident in cache. This is especially true for restore operations.

DLM2000 Overview

Product Description

The EMC DLM2000 is a tape replacement solution that enables users to achieve better performance, higher reliability, and significant cost savings by storing tape information on disk instead of physical tape.

The EMC DLM2000 includes one or two virtual tape engines (VTEs) to perform tape emulation operations, mainframe channel connectivity, a management server, one internal switch, and primary storage to store the actual tape volumes.

Architecture

The base components of the Disk Library for mainframe (VTEs, internal switch, and Access Control Point management server and storage), all reside within a single cabinet. The DLM2000 may be configured with 1 or 2 VTEs depending on the required number of emulated tape drives and overall system performance requirements.

DLm2000 Components




Figure 1. Disk Library for mainframe architecture

Virtual Tape Engines

The VTE appears to the mainframe operating system as a set of standard IBM tape drives. The mainframe manages this set of tape drives as one or more virtual tape libraries. Existing mainframe software applications use the virtual drives of the VTE — specifically IBM 3480, 3490, and 3590 drive types — just as they would any mainframe supported physical tape drives. No application modifications are required to integrate them into the existing mainframe environment.

The VTEs are connected to the mainframe host via FICON channels. Each VTE includes two 4 Gb FICON channels for connection to the mainframe host. A DLm2000 with 2 VTEs provides 4 FICON channels to the mainframe host.

Each VTE can support up to 256 total virtual drives. A DLm2000 configured with a maximum of 2 VTEs can emulate up to 512 virtual tape drives. These tape drives can be shared across a total of 64 active LPARs.

While each VTE operates independently, both VTEs in a DLm2000 have access to all the tape volumes in the Disk Library for mainframe storage system, and any emulated tape drive can access all the tape volumes stored in the Disk Library for mainframe.

DLm2000 Storage

The storage system utilizes a Disk Processor Enclosure (DPE) which includes both Storage Processors (SP) and the first 15 disks. These 15 disks are allocated as follows:

- The first 5 disks make up the ‘vault’ (4 300GB vault drives + 1 Hot Spare)
- The second set of 5 disks are hot spares for the data drives (2TB drives)

- The third set of 5 disks make up the first half of the first RAID Group of data drives. The second half of this first RAID group is located in the first data Disk Array Enclosure (DAE).

The VNX5300 supports from 13 to 143 TB of usable storage via seven expansion DAEs (each with 15 2TB drives). The minimum configuration requires 1 DAE in addition to the drives in the DPE. Each DLM2000 will always have two storage controllers (VNX File Data Movers) configured in the system. Each storage controller has the following I/O:

- 4 x 8Gb FC ports (1 x 4-port card), connection to Storage Processors
- 2 x 10GbE optical ports (1 x 2-port card), internal data path to VTEs
- 4 x 1GbE copper ports (1 x 4-port card), replication ports
- Management ports for internal management network

Continuous Data Availability

The DLM2000 is designed with significant redundancy to provide for continuous data availability. When configured for high-availability it includes redundant components such as VTEs, and more. The storage in the DLM2000 utilizes RAID 6 protection and the storage controllers include hot standby drives.

When configured, both VTEs can access all tape volumes in the system; that is, if one of the VTEs becomes unavailable, the other VTE can access the VOLSER. With the DLM2000, the configuration of an unavailable VTE can be loaded to an alternate VTE to enable access to all tape devices originally accessed by the first VTE. Alternatively, accessing the tape volumes from a different tape drive that is mapped to a 2nd VTE can also be performed.

Software or other equipment upgrades can be performed on the DLM2000 without the need to stop tape processing. For example, in a software upgrade the tape traffic can be diverted from one VTE and then vary the drives offline. The necessary upgrade can be applied and traffic re-introduced to that VTE. The second VTE would be modified in the same way until all updates have been installed.

RAID 6 Data Protection

Reliability is further enhanced with the implementation of RAID 6 with all disk storage on the DLM2000. RAID 6 provides two parity drives. Diagonal parity is the parity of the data diagonally across bits in a block stripe. RAID 6 also distributes parity among all drives in the RAID group to provide uniform performance.

Hot Spare Drive

A hot spare is a single disk that serves as a temporary replacement for an unavailable disk in a RAID 6 group. Data from the unavailable disk is reconstructed automatically on the hot spare from the parity on the remaining disks in the RAID group, so the data on the device is always accessible. Multiple hot spares are configured depending on the specific DLM2000 back-end storage configuration.

Management and Supportability

Disk Library for mainframe works seamlessly with the mainframe host and does not require any mainframe-based code changes to operate. Additionally users do not need to change their production operations or production JCL.

The Disk Library for mainframe can be managed using DFSMS functionality and supports all Channel Command Words (CCWs) for tape. Therefore DFHSM, backups, and other host applications continue to work without change. Additionally these operations are no longer dependent on a specific tape drive range and tape processing is done at disk speed. This reduces the time it takes for recycle/recall operations to complete.

A Disk Library for mainframe enables users to manage and query various status and state conditions including the following:

- Users can perform specific actions on or retrieve information about the DLm2000 directly from the mainframe master console. Customers can easily retrieve information such as available space, configuration, scratch tape counts, and more.
- Through a set of batch Command Processors, users can perform these same master console functions, and retrieve details of emulated tape libraries and status on data replication sessions.
- Users can use a web-based application, DLm Console, to remotely log in, query and manage the DLm2000 online.
- Disk Library for mainframe supports Simple Network Management Protocol (SNMP), which provides automatic alerts to email accounts or other third-party management tools.

Support is also provided for EMC Secure Remote Support (ESRS), which enables EMC Customer Support to establish secure IP connectivity to the Disk Library for mainframe and remotely log in to the system for diagnosing and troubleshooting system issues. In addition, the Disk Library for mainframe supports ConnectEMC, which automatically sends alerts directly to EMC Support.

WORM Tape

WORM tape is an optional feature of the DLm2000 that protects files from being modified or deleted until a specified retention date has been met. This feature allows you to create a permanent, unalterable set of files and directories to ensure the integrity of the data to meet WORM requirements.

With WORM tape, you are unable to modify or delete files that are in a “locked” state. Additionally, the path to a file in a “locked” state is also protected from modification. This means you cannot delete or rename a directory on a file system if it contains protected files.

The WORM tape feature of the DLm2000 utilizes the underlying File-Level Retention technology of the VNX storage and provides two options for protection: Enterprise and Compliance. Enterprise is for users who are implementing a corporate wide set of regulations. The Compliance version enables a system to satisfy the regulatory requirements of the US Securities and Exchange Commission (SEC) Rule 17a-4(f), German and Austrian Principles of Proper Accounting, and relevant requirements of the Swiss Law of Obligations and Swiss Accounting Regulations.

Remote Replication and Recovery

The Disk Library for mainframe offers IP-based remote replication for the DLm2000 which uses the user's IP network infrastructure and eliminates the need for channel extension equipment. The replication is storage-based and therefore has no impact on mainframe host operations. Disk Library for mainframe replication is asynchronous with only the changed data being replicated to the remote sites.

Disk Library for mainframe replication supports multiple target sites per source system, which means users can replicate their information to different sites. For example, one site can be the disaster recovery site and one site a bunker site for vaulting. Users can choose which virtual VOLSERs will be replicated to each remote site. The DLm2000 supports 1-to-1 or 1-to-2 replication.

The Disk Library for mainframe also supports bi-directional replication, which means that a source system can concurrently serve as a target system. With this feature, the Disk Library for mainframe units at two data centers can provide enhanced data protection for each other. In all implementations where replication will be configured a proper system sizing review should be completed to guarantee that desired performance objectives will be met.

The Recovery Point Objective (RPO) for Disk Library for mainframe can be defined in minutes or hours, depending on user requirements. Different RPOs for different storage groups, for example, could be defined based on information criticality which would allow tuning of systems and ensuring optimized bandwidth utilization. Critical information may have a low RPO (minutes), whereas less critical information can have a higher RPO (hours).

Disk Library for mainframe replication also enables the user to define quality of service, which optimizes the network traffic to prevent network overload during peak hours.

For replication, the target Disk Library for mainframe does not have to have the same configuration as the source Disk Library for mainframe. For example, a source DLm2000 could be configured with two VTE's and more storage capacity, whereas the target DLm2000 may have only one VTE and lower capacity if it does not require the same throughput and capacity for DR processing as the production DLm2000.

The replication capabilities of the DLm2000 include data replication using the VNX Replicator. VNX Replicator replicates the data periodically and asynchronously. It is configured to periodically create a snapshot of the local DLm storage/filesystems and then asynchronously

transfer the data at the time of the snapshot to the remote DLm. Tape data written subsequently by the host at the source site after that snapshot is replicated in the next snapshot. If replication functions become unavailable before that next snapshot is sent, job processing at the target site will need to recreate those changes if they are needed.

For most situations and applications, this type of periodic and asynchronous replication is adequate. However, for some critical applications where the data is required to be on the remote DLm before specific local job steps can proceed, this type of asynchronous replication may not be adequate.

DLm2000 provides an enhancement to the VNX Replicator capabilities called guaranteed replication. This function ensures VNX Replicator completes the replication of a tape VOLSER to the remote site every time the mainframe issues two consecutive write tape marks on a tape. In this event, guaranteed replication causes the VTE to withhold acknowledgements of tape close from the local mainframe until the VTE confirms that the local Storage Controller has completed replicating the tape volume to the remote VNX.

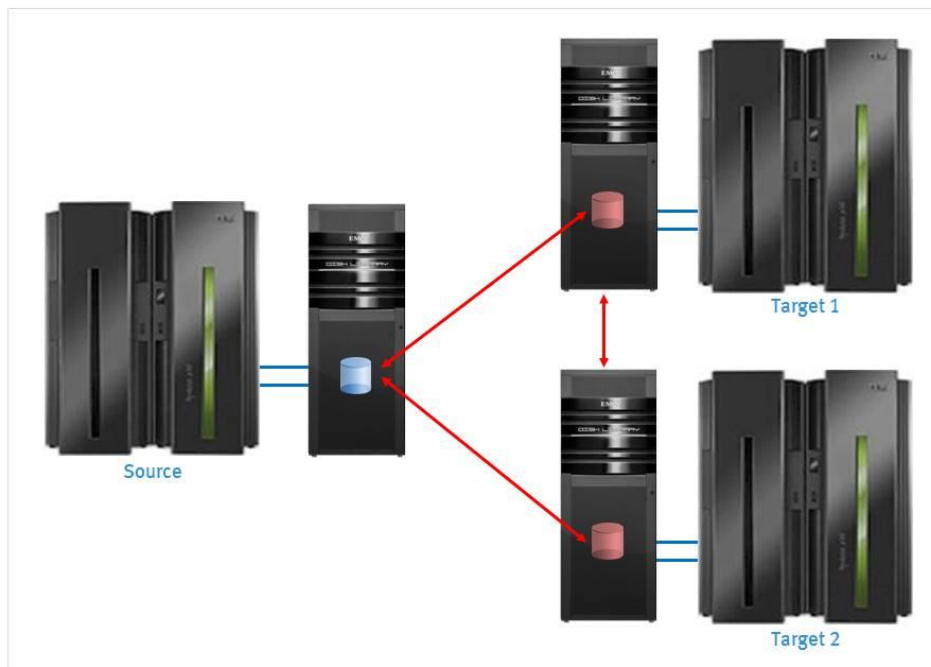


Figure 2. A typical replication scenario

As part of an overall disaster recovery solution in a mainframe environment, users must make sure that their DASD, ICF Catalog, and any tape management datasets, such as tape management catalog (TMC) and the Tape Control Database (TCDB), are replicated to the disaster recovery site.

The fact that the tape volumes are kept on disk and not on physical tape in itself makes the recovery process much faster and can save hours at the remote site. Also, DLm2000 systems automatically replicate tape mount status information which, used in conjunction with the TMC and TCDB, accelerate the resumption of operations at the remote site.

Flexible Recovery Testing

Disaster recovery tests are extremely important in mainframe environments. Many users perform these tests several times a year to ensure that in a time of need their procedures are up to date and that they can successfully recover in minimal time.

The disaster recovery tests often take several days and can require users using remote replication with traditional virtual tape servers to turn off replication, and thus be unprotected and exposed to major data loss during the period of the tests.

With Disk Library for mainframe, replication is not interrupted during DR testing and data is always protected. Using Disk Library for mainframe capabilities, users have two options when performing their disaster recovery tests, while replication continues without interruption:

- **Read only mode:** In this mode, users mount their tape volumes (VOLSERs) at the remote site as read only during the disaster recovery tests. This mode allows users to read tape data located at the DR site and perform restore operations to check their disaster recovery procedures. Users may see their tape information change during the disaster recovery tests, as Disk Library for mainframe replication continues to update the volumes at the remote site(s) during the testing period.
- **Read / Write mode:** In this mode, users can perform full disaster recovery tests, including read and write operations, without updating any production tape volumes. Users can use this mode for disaster recovery tests by taking snapshots of the required tape volumes and mounting them as Read / Write on the target Disk Library for mainframe system. Full disaster recovery tests can then be performed on these snapshot copies of the tapes. Once the disaster recovery tests are complete, the snapshots can be scratched to free up disk space for future needs.

Dm2000 Typical Deployment Scenarios

As the quantity of information continues to grow along with the demand to keep more data available for longer periods of time, most mainframe data center managers are looking for ways to shorten batch and backup durations, improve recall and restore times, improve their tape reliability and disaster recovery solutions, while lowering overall Total Cost of Ownership (TCO).

The performance, scalability, and flexible storage options of the Dm2000 enable it to support a broad range of use cases.

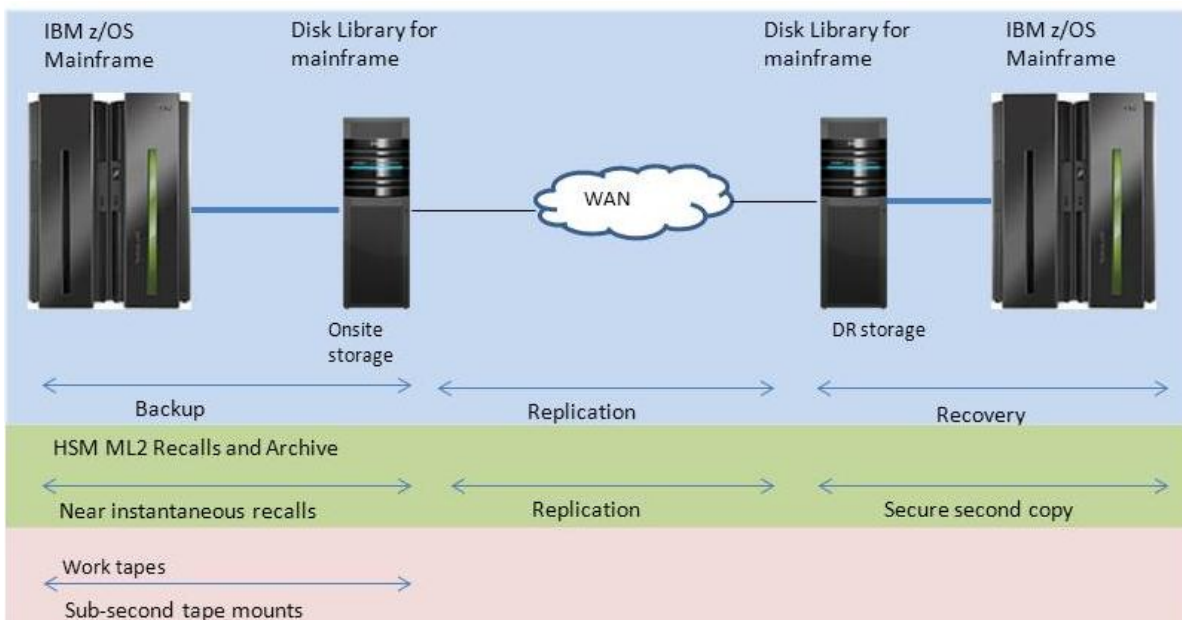


Figure 3. Disk Library for mainframe use cases

Backup and Restores

Backups in the mainframe environment are typically executed using IBM utilities such as Data Facility Hierarchical Storage Management (DFHSM) and Data Facility Data Set Services (DFDSS). Mainframe users are often required to keep the information for many years to meet compliance and data availability requirements. Many users use tape to migrate information between sites for disaster recovery. This usually requires that tapes either be shipped offsite or replicated between traditional virtual tape servers, depending on disaster recovery requirements.

Data Migration

Many mainframe users use DFHSM to automatically manage their storage assets. DFHSM datasets can be migrated from primary tier-1 DASD storage (L0) to another location on tier-1 DASD in compressed format as Migration Level One (ML1) and finally to tape as Migration Level Two (ML2). Recalls of ML2 information from tape require that the tape be mounted and sequentially searched to locate the desired data set. This results in unpredictable response times that can take minutes to satisfy.

With the DLm2000, users can bypass ML1 processing without compromising recall times. When data is migrated from L0 to ML1, the mainframe uses CPU cycles for compression and decompression and stores the information on tier-1 DASD. With the DLm2000, users can go directly from L0 to ML2 and let the DLm2000 compress the data and write it on spinning disk. Recalls of ML2 data from the DLm2000 are fast and consistent and usually are satisfied in less than two seconds. Significant savings in both CPU usage and tier-1 storage costs can be realized. Additionally the DLm2000 results in faster tape recycles because of no tape stacking and all of the data residing on spinning disk.

Data Archive

In a mainframe environment, data archive refers to fixed content data, such as bank statements or insurance policies, which must be maintained in their original state for years. This data is initially kept on tier-1 DASD storage to provide fast and predictable response times for data queries. After a period of time, when access to the information is less frequent, the data is migrated to tape. Accessing this data from tape can take minutes and render the information unavailable for online applications.

With the DLm2000, the length of time data must remain on tier-1 storage is greatly reduced. The large number of emulated drives available to the DLm2000 greatly reduces drive contention and mount requests are typically satisfied in less than one second. Because the data resides on spinning disk and the DLm2000 has the ability to locate specific sections of the tape quickly, data requests are usually satisfied within one second. Because of this data can be migrated to the DLm2000 sooner without compromising recall times and reducing tier-1 storage requirements.

Work Tapes

These can include batch processing, database log files, temporary files for sorting, etc. Production batch jobs are performed daily where it is common practice to find “old master in / new master out” processing and large sequential datasets to be tape-resident. Often the most active DB2 and IMS log(s) are kept on DASD while less active data is kept on tape for recovery purposes. In addition some applications log data directly to tape. Recovery from these log files usually depends on drive availability and can take significant time.

By providing faster throughput and up to 512 virtual drives, the DLm2000 virtually eliminates drive contention. The performance of the DLm2000 can result in faster job completion. Because the data is on spinning disk, mount requests are typically satisfied in less than one second and data is readily available for the step to step nature of batch processing.

Conclusion

The DLm2000 is an excellent tape replacement solution for the commercial mainframe user that leverages proven EMC hardware and software technology to provide mainframe tape users with considerable cost savings along with the many benefits described throughout this paper over traditional physical tape and virtual tape server solutions. Users looking to upgrade or replace their tape infrastructure with disk-based solutions need not change any of their existing mainframe applications or processes.

The DLm2000 is built on the latest Disk Library for mainframe generation architecture that provides both scalable front-end mainframe connectivity and backend storage supporting the EMC VNX storage platform.

The DLm2000 also provides high data availability through the use of RAID 6 data protection and hot spare drives.

The DLm2000 also provides a reduced data center footprint when compared to physical tape.

By leveraging the latest advances in EMC hardware and software technologies, the DLm2000 enables mainframe users to minimize their storage, and improve their disaster recovery capabilities when compared to conventional tape systems.