

Disk Library for mainframe - DLm6000 Product Overview

Abstract

This white paper introduces the EMC® DLm6000 - the EMC flagship mainframe VTL solution and a member of the EMC Disk Library for mainframe family. The scalability, performance and flexibility of the DLm6000 allows mainframe users to replace their physical tape and virtual tape servers with an integrated tape replacement solution that improves storage utilization and reduces total cost of ownership while improving data availability over physical tape.

October 2011

Copyright © 2011 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate 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 h4193.5

Table of Contents

Executive summary	2
Introduction	2
Audience.....	2
Disk Library for mainframe family overview	3
Dm6000 product description.....	3
Challenges with physical tapes.....	4
Tape usage in the mainframe environment	4
The Dm6000 and mainframe use cases.....	5
EMC Disk Library for mainframe architecture.....	7
Disk Library for mainframe management and supportability.....	9
Disk Library for mainframe remote replication and recovery	10
Flexible Recovery testing.....	11
Conclusion	12

Executive summary

The EMC® Disk Library for mainframe family of solutions offers IBM System z mainframe users 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.

DLM6000 is the EMC enterprise virtual tape solution for replacing physical tape in the mainframe environment. The DLM6000 addresses the challenges of the enterprise data center and delivers industry-leading scalability, performance and flexibility to mainframe tape operations. With up to 2.0 GB/s performance, the DLM6000 provides up to double the performance of competing solutions. DLM6000 offers concurrent mixed storage flexibility by supporting both primary and deduplication storage in a single, manageable solution to address all mainframe tape use cases. The DLM6000 incorporates RAID 6 protected disk storage, hot-standby disks, tape emulation, and hardware compression to meet the data protection requirements of enterprise mainframe data center tape environments.

The DLM6000 is the 3rd generation EMC Disk Library for mainframe system, representing an evolution and maturity of the solution to meet expanding customer requirements.

Introduction

This white paper begins by looking at tape use in the mainframe environment and examines the challenges that physical tape presents. It then describes the DLM6000 tape replacement solution, including benefits, management and support, remote replication and recovery, deduplication, and its architecture.

Audience

This white paper is intended for mainframe storage professionals who are looking to understand how the EMC 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.

Disk Library for mainframe family overview

The EMC Disk Library for mainframe 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.

Disk Library for mainframe systems include one or more virtual tape engines (VTEs) to perform tape emulation operations, mainframe channel connectivity, a variety of internal switches as well as back-end disk storage. Depending on model and configuration, the back-end storage can consist of either primary storage, deduplication storage or both to store the actual tape volumes.

The Disk Library for mainframe family consists of:

- Mainframe Data Library-1000 for Data Domain
- Mainframe Data Library-2000 for Data Domain
- DLm120 supporting up to 95TB of primary storage.
- DLm6000 supporting primary and/or deduplication storage.

This document will focus on the DLm6000.

DLm6000 product description

The base components of the Disk Library for mainframe (VTEs, internal switches, and Access Control Points), all reside within a single cabinet. Additional cabinets may be configured depending on storage capacity requirements. The DLm6000 may be configured with 2 to 6 VTEs depending on the required number of drives and overall system performance requirements.

The DLm6000 is available in 3 different storage configurations:

DLm6000 with one or two EMC VNX7500 storage platform(s)

DLm6000 with one or two EMC Data Domain DD890 deduplication storage system(s)

DLm6000 with one VNX 7500 and one DD890 system

The DLm6000 incorporates the latest virtual tape emulation software referred to throughout this document as EMC Virtuent 7.

Virtuent is a tape-on-disk software package that runs on a base hardware controller which provides two FICON connections to the mainframe. The Virtuent software provides the controller emulation supporting IBM 3480, 3490, or 3590 tape drives. Data that is written to or read from these tape drives by the mainframe is stored and retrieved from either primary or deduplication storage that is configured in the Disk Library for mainframe.

Challenges with physical tapes

Mainframe data centers are highly dependent on tape systems for batch processing, backup/recovery and long term archive. 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 high throughput rates; 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 users 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 will outperform physical tape for “time to first byte” as long as the dataset is resident in cache.

Tape use cases in the mainframe environment

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 centers 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).

Below are the typical use cases of tape in a mainframe environment:

- **Backups 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 systems, depending on disaster recovery requirements.
- **Data migration** – Most mainframe users use DFHSM which is a software component within the mainframe environment 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.
- **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 its 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.
- **Work tapes** – 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.

Tape has traditionally provided inexpensive storage for backups, restores, data migration, archival and work tapes as noted above, however tape presents a number of challenges.

The DLm6000 and mainframe tape use

The DLm6000 offers mixed storage flexibility in that it supports both primary and deduplication storage concurrently in a single manageable solution. This allows for tape data to be directed to the appropriate storage based on its intended use, resulting in faster and significantly more efficient storage utilization.

The DLm6000 can achieve over 2GB/s throughput resulting in reduced batch run times, faster backup and restores and near-instantaneous access to DFHSM and archived data.

The following is a description of how the DLm6000 benefits the use cases discussed above:

- **Backups and restores** – backup data is very repetitive which makes it an ideal candidate for a deduplication storage system. Backup workflows can be directed to deduplication storage where data footprint can be reduced by a factor of 10 to 30 to one. This minimizes the amount of storage capacity required to store the data which can significantly reduce overall storage costs and require less bandwidth for replication to a DR site.
- **Data migration** – With the DLM6000, 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 DLM6000, users can go directly from L0 to ML2 and let the DLM6000 compress the data and write it on spinning disk. Recalls of ML2 data from the DLM6000 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 DLM6000 results in faster tape recycles because of no tape stacking and all of the data residing on spinning disk.
- **Data archive** – with the DLM6000, the length of time data must remain on tier-1 storage is greatly reduced. The large number of emulated drives available to the DLM6000 eliminates drive contention and mount requests are typically satisfied in less than one second. Because the data resides on spinning disk and the ability of the DLM6000 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 DLM6000 sooner without compromising recall times and reducing tier-1 storage requirements
- **Work tapes** – The DLM6000 provides over 2GB/sec. throughput and up to 1,536 virtual drives. Drive contention is virtually eliminated and the high-performance of the DLM6000 results in jobs completing sooner. 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.

In many large enterprises the different use cases of tape have required mainframe users to implement multiple storage systems to support all of their tape workloads.

The DLM6000 is the only mainframe VTL solution that can address the full range of mainframe tape workloads with a single, consolidated all-disk system.

EMC 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 DLM6000 with 6 VTEs provides 12 FICON channels to the mainframe host.

Each VTE can support up to 256 total virtual drives. Configured with a maximum of 6 VTEs, a DLM6000 can emulate up to 1,536 virtual tape drives. These tape drives can be shared across a total of 64 active LPARs.

While each VTE operates independently from the others, all the VTEs in a DLM6000 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.

Back-end storage

The DLM6000 is a unique virtual tape library solution in that it can concurrently support both primary and deduplication storage and dynamically direct tapes to the most appropriate storage on a file by file basis. Primary storage is ideally suited for unique data types, such as DFHSM migration, and is available for near instantaneous recalls. The use of deduplication storage is ideal for repetitive backup data, for example 3990 volume dumps from FDR, DFSS and/or CA-DISK. Deduplication of repetitive backups can substantially improve the overall data reduction achieved within the DLM6000 resulting in significant reduction in storage and transmission costs.

VTEs process the arriving mainframe tape volume and write it as a single file on the Disk Library for mainframe back-end storage. Each mainframe tape is stored as a single file whose filename matches the tape VOLSER. Unlike physical tape, the file only consumes as much space that is required to store it, resulting in no wasted storage space. This allows the virtual tape to be easily located and mounted in response to read or write requests, typically within one second.

All disk drives within the DLM6000 are protected with a RAID 6 configuration and hot spare drives for each RAID group.

When writing to deduplication storage, compression is turned off in the VTE. This enables a higher level of data reduction for those applications that can benefit from deduplication. The deduplication storage capability can provide up to 5.7 TB of logical storage capacity based on a conservative 10 to 1 deduplication ratio.

Redundancy

The DLM6000 is designed with significant redundancy to provide for continuous data availability. It includes redundant components such as VTEs, internal switches, Access Control Points, and more. The storage in the DLM6000 utilizes RAID 6 protection and the storage controllers include hot standby drives.

As mentioned earlier, all VTEs can access all tape volumes in the system; that is, if one of the VTEs becomes unavailable, a different VTE can access the VOLSER. With the DLM6000, 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 DLM6000 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 and subsequent VTEs 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 the all disk storage on the DLM6000. RAID 6 provides two parity drives. One is horizontal parity of the data for a block location within one block stripe, and the other is diagonal parity, which is unique to RAID 6. Diagonal parity is the parity of the data diagonally across bits in a block stripe. Both horizontal and diagonal parity are completely independent of one another and are contained within a 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 DLM6000 back-end storage configuration.

Hot standby controller

The DLM6000 includes a hot standby storage controller when primary storage is configured. This standby controller can protect up to five active storage controllers. If a storage controller becomes unavailable, the hot standby storage controller will automatically assume the activities of the unavailable storage controller.

Deduplication storage for Disk Library for mainframe

Deduplication storage for the DLM6000 reduces the storage footprint, increases backup application performance, and allows backup data to be retained onsite longer and replicated efficiently for disaster protection. Based on the EMC Data Domain DD890 deduplication storage system, this capability provides:

- Up to 10-30x reduction in the raw storage requirement because of the deduplication process. Note: as with all deduplication systems, this reduction factor can vary greatly depending on the specific data and its retention period.
- Up to 99 percent bandwidth reduction because less data is transported across the IP network due to the deduplication process.
- Continuous recovery verification, fault detection, and healing.
- Flexible scalability options that can scale up to 5.7 PBs of logical storage.

The deduplication storage integrates into the mainframe environment seamlessly by pointing the selected mainframe workloads to a specific storage group, which in turn pass this data to the deduplication storage as opposed to the primary storage.

Disk Library for mainframe 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 Job Control Language (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 client 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 DLm6000 directly from the mainframe master console. Customers can easily retrieve information such as available space, configuration, scratch count, and more.
- Users can use a web-based application, DLm Console, to remotely log in, query and manage the DLm6000 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.

Disk Library for mainframe remote replication and recovery

The Disk Library for mainframe offers IP-based remote replication for the DLM6000 employing both primary and deduplication storage, 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 or performance.

Disk Library for mainframe replication supports multiple target sites (see Figure 1) 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.

Disk Library for mainframe replication is asynchronous with only the changed data being replicated to the remote sites. The Disk Library for mainframe also supports bi-directional replication, which means that the source system can become a target system and a target system can become a source system. For example if you failed over to a DR site for a period of time, data would be replicated back to the production site when connection is re-established.

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, the source Disk Library for mainframe can be a high capacity system with six VTEs, whereas the target Disk Library for mainframe may have only two VTEs and lower capacity if it does not require the same throughput and capacity for DR processing as the production Disk Library for mainframe.




Figure 1. Disk Library for mainframe remote replication

One of the major benefits of deduplication is that replication of deduplicated data typically offers the most economical approach to the automated movement of data copies to a safe site using minimum WAN bandwidth, as considerably less data may have to be moved over the wire because of the reduced footprint of the data.

The [EMC Data Domain Replicator - A Detailed Review white paper](#) provides more detailed information.

As part of a 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. DLM6000 tape volumes can be replicated using IP replication. The fact that the tape volumes are kept on disk and not on physical tape makes the recovery process much faster and can save hours or even days at the remote site. In addition, the Disk Library for mainframe eliminates the risk of losing a tape that might be critical for the recovery process since all tape information is kept on RAID-protected disk.

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 systems 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:

- **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.

Conclusion

DLM6000 enables complete mainframe tape replacement with its ability to support both primary and deduplication storage that can be matched to individual user data and performance requirements. This solution also provides a unique ability to reduce replication bandwidth – a compelling attribute that many mainframe enterprise environments can use to significantly reduce their overall cost of operations. A reduced data center footprint can also result from a DLM6000 deployment when compared to physical tape.

The DLM6000 can provide considerable cost, performance, and availability advantages over existing mainframe physical or virtual tape server solutions.

The DLM6000 is a seamless tape replacement solution that leverages proven EMC hardware and software technology to provide mainframe tape users with the best in performance, scalability and availability 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.