

DISK LIBRARY FOR MAINFRAME DLm DLm2500 and DLm8500

ABSTRACT

This white paper discusses the benefits of the 5th Generation Dell EMC Disk Library for mainframe (DLm), release 5.4 which introduces new DLm Gen5 capabilities for the model 8500.

DLm is the ideal mainframe tape replacement solution in terms of scalability and performance. It enables mainframe users to replace their physical tape and virtual tape servers with an integrated virtual tape solution that improves storage utilization and processing performance.

July 2021

Revisions

Date	Description
July 2019	Initial release for DLm 5.1
July 2020	Initial release for DLm 5.2
October 2020	Initial release for DLm 5.3
July 2021	Initial release for DLm 5.4

Acknowledgements

This paper was produced by the following:

Author: Mikhail Bermistov, Justin Bastin, 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 © 2021 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 2021. White paper h12225.9

Table of Contents

EXECUTIVE SUMMARY	4
Introduction	5
Audience.....	5
Tape use cases in the mainframe environment	5
Challenges with physical tape	6
DISK LIBRARY FOR MAINFRAME AND MAINFRAME TAPE USAGE	7
Disk Library for mainframe, DLM2500 and DLM8500.....	7
Disk Library for mainframe product description	8
Virtual tape engines (VTEs)	9
Back-end storage.....	9
Disk Library for mainframe management and support.....	10
Disk Library for mainframe remote replication	10
Data Domain Replicator for both DLM2500 and DLM8500	12
Data recovery.....	12
Flexible Recovery testing	12
DISK LIBRARY FOR MAINFRAME ARCHITECTURE	13
Redundancy	13
RAID 6 data protection	13
Hot spare drive	13
Hot standby controller	13
Deduplication Storage for DLM	14
DLM8500 AND THE USE OF ELASTIC CLOUD STORAGE (ECS) TO REPLACE PHYSICAL TAPE FOR LONG-TERM RETENTION.....	14
GDDR TAPE LEVERAGES GLOBAL VIRTUAL LIBRARY AND DASD AUTOMATED FAILOVER FOR DLM8500	15
CONCLUSION	16

Executive summary

Dell EMC Disk Library for mainframe (DLm) offers IBM Z Systems and Unisys Dorado / Clearpath mainframe customers the ability to replace their physical tape systems, including traditional virtual tape servers such as the IBM TS7700 family and Oracle/STK VSM, with a dynamic virtual tape solution, eliminating the challenges tied to traditional tape-based processing.

Now in its 5th generation, DLm continues to be the industry's fastest and most flexible VTL to enable complete mainframe tape replacement. This flexible system supports a mix of primary and deduplication storage types to support all of the use cases commonly found in mainframe data centers.

DLm release 5.4 introduces the PowerScale and Isilon storage options, along with 3-site GDDR tape for PowerScale and Isilon. Additionally, two usability enhancements have also been added: **Replication Status by VOLSER** which will list all tape volumes that are currently in process of being replicated and provide definitive information regarding which tape volumes are in replication status, and, **Scratch Tape Health Check** which enables setting and triggering of an alert when the number of available scratch tape volumes falls below a user specified number of volumes.

DLm release 5.3 introduced a brand new DLM UI – DLm System Manager. DLm System Manager replaces DLm WebConsole and provides simple system dashboard view, simplified management and enhanced reporting capabilities. It also extended the LTR use case by introducing Amazon Cloud (AWS) support for cost effective LTR data storing.

In combination with PowerMax8000, PowerProtect DD or ECS DLm Release 5.4 continues IBM Transparent Cloud Tiering compatibility released in 5.3 exclusively for DLm8500.

DLm release 5.3 introduced the model DLm2500, a 2U-size standalone or rack-mountable virtual tape system with two 16Gb FICON channels and four 10 Gb Ethernet ports providing an NFS connection to a wide variety of the Dell EMC storage. Internally, the DLm2500 provides up to 256 virtual tape drives. Its graphical user interface simplifies the powerful operational aspects of DLm that has made Dell EMC a leader in the mainframe virtual tape market.

The DLm2500 enables connection to Dell EMC Isilon storage (as well as the newer PowerScale models), PowerProtect DD models and legacy Dell EMC Data Domain models.

DLm release 5.1 introduced PowerMax 8000 fiber channel attach as a backend storage for Disk and Tape Universal Data Consistency and synchronous replication capability using SRDF/S for mainframe tape environments that demand tape and disk consistency for critical workloads and / or synchronous replication of tape data between primary and DR sites. Additionally, R5.1 simplifies data retrieval from Long Term Retention (LTR / archive) when using ECS for LTR. Several other operational enhancements have been made to the LTR capability of prior releases and will be

discussed later. R5.1 also adds an install option into a customer-supplied rack for single frame (cabinet) DLm8500 configurations as well as 3-phase power options.

Release 5.0 added 16Gb FICON connectivity while doubling the amount of possible FICON connection to 32 (when using 8 Virtual tape engines). Additionally, a new Dell 14G R740xl server and new compression card doubled the bandwidth for even better performance compared to prior DLM models. DLm8500 built on release 4.5 which added cloud-based long-term retention, automated failover and KMIP encryption key manager support.

The DD9800, with up to 1PB (native, no deduplication) of storage capacity, enables growth of the DLm8100's native / logical capacity to 20PB total (assumes 2 DD9800s and 10:1 deduplication of customer data).

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 DLm 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 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.

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, and improve their tape reliability and disaster recovery solutions, while lowering overall TCO.

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

Batch jobs – 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. Batch cycles also occur weekly, monthly, and quarterly. Providing the ability to ensure consistent access times to data is critical to ensuring completion of batch processes within required timeframes.

Backups and restores – Backups in the mainframe environment are executed using IBM utilities such as DFHSM and DFDSS or with software applications from others. Mainframe customers are often required to keep the information for many years to meet compliance and data availability requirements.

Storage management – Most mainframe customers use Data Facility Storage (DFHSM) which is software components within the mainframe environment to automatically manage their tape storage. DFHSM migrate datasets can be moved from primary DASD storage (L0) to secondary DASD in compressed format (ML1) and to tape (ML2).

Disaster recovery – Many customers 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.

Long-term data Retention – All customers have varying time periods for the retention of data stored on virtual or physical tape. It is not unusual for companies with legally-based records retention requirements to ensure data created today can be stored and retrieved for 10, 15, 20 or more years.

Tape has traditionally provided inexpensive storage for batch, backups, disaster recovery, and long-term data retention as noted above, however, tape presents a number of challenges.

Challenges with physical tape

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, batch job cancellations 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 unreliable. Since they are mechanical devices, tape drives naturally wear down over time. They can jam, causing the tape media to be permanently damaged. Head alignment can change, robotic arms can fail or jam and motors can wear out. The net result is that a physical tape drive cannot always access data when required.

Other issues related to physical tape libraries exist, such as poor media utilization. Tapes, however, are fast and can support very fast throughputs; but this speed is achieved only after the tape is mounted and positioned to the first block of data. The problem with tapes is “time to first byte” since tapes must be loaded into 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 their tape drives, which results in long wait times and lengthy time to first byte reads.

Virtual tape servers use a temporary disk cache to store tape data. Tape files are then stacked and written to large format tapes with the primary benefit being better tape media utilization. This means that when information is retrieved from tape, and no longer exist 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 since it does 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. This is especially true for restore operations.

Disk Library for mainframe and mainframe tape usage

As noted in the sections above, traditionally, physical tape has provided inexpensive storage for batch, backups, disaster recovery and long-term archives. However, physical tape presents challenges in today's mainframe environments. DLM addresses the challenges of the enterprise data center by combining RAID 6 protected disk storage, hot-standby disks, tape emulation, hardware compression, and provides both primary and deduplication storage in the same platform. Each of these topics is discussed in more detail in this paper.

Disk Library for mainframe provides both primary and deduplication storage concurrently (see applicable DLM model numbers below), 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. This results in reduced batch runtime, reduced overall batch window times and faster migrations. Data types, such as DFHSM migration, can be directed to primary storage making the data readily available for near-instantaneous recalls, migration times can be reduced significantly as a result. DLM allows for redirecting DFHSM workloads from tier1 storage directly to ML2, avoiding ML1 processing, reducing CPU utilization.

DLM can replicate from one source site to one or two remote sites, the option of replicating all or part of tape data, depending on requirements, can be accomplished and priority can be chosen based on policies for the order of replication. Replication is described in more detail in this paper.

Disk Library for mainframe, DLM2500 and DLM8500

DLM is a tape replacement solution that enables customers to achieve better performance, higher reliability, and significant cost savings by keeping tape information on disk instead of physical tape.

DLM includes one or more virtual tape engines (VTEs) to perform the tape emulation operations and back-end disk storage that stores the tape volumes.

Disk Library for mainframe DLM2500 options

Component	Details	specifications
Virtual tape engine (VTE)	An individual standalone tape emulator	<ul style="list-style-type: none">• Up to two FICON connections• Up to 256 tape drives• Emulates 3480/3490/3590 tape formats
Storage Array options	Legacy Data Domain: DD9800, DD9500, DD9300 DD6800 or DD6300 for data deduplication	2 TB SAS drives
	PowerProtect DD: DD9900, DD9400, DD6900 for data deduplication	2 TB SAS drives
	All Isilon models	Many combinations of SAS OR All Flash configurations

	All PowerScale models	All Flash storage
--	-----------------------	-------------------

Disk Library for mainframe DLm8500 options

Component	Details	specifications
Virtual tape engine (VTE)	Tape emulation technology consisting of 1-6 "engines"	<ul style="list-style-type: none"> • Up to four FICON connections per VTE • Up to 512 tape drives per VTE • Up to 3,072 tape drives emulated per DLm8500 • Emulates 3480/3490/3590 tape formats
Storage Array options	One or two DD9800s, DD9500s, DD9300s DD6800s or DD6300s for data deduplication	2 TB SAS drives
	One PowerMax8000	24 slot NVMe DAE using 2.5" form factor 1.92 TB, 3.84 TB, or 7.68 TB NVMe Drives
	One or two Isilon models: F800, F810, H400, H500, H600, H5600, A200, A2000	Combinations of SAS OR All Flash configurations
	One or two PowerScale models: F200, F600	All Flash storage
Cloud Options	One PowerMax8000 + one Elastic Cloud storage (ECS)	24 slot NVMe DAE using 2.5" form factor 1.92 TB, 3.84 TB, or 7.68 TB NVMe Drives
	One DD system above + Elastic Cloud storage (ECS)	

Disk Library for mainframe product description

The base components of the DLm are the virtual tape engine (VTE) and its associated storage. The DLm2500 (introduced with release 5.2) is a single VTE with two 16Gb FICON ports capable of attaching to one Dell EMC storage subsystem. The DLm8500 consists of one or two racks housing a series of VTEs and internal switches. all reside within a single VTEC cabinet. DLm may be configured with one to six VTEs depending on the required number of drives and overall system performance requirements. DLm incorporates virtual tape emulation software referred to throughout this document as Dell EMC Virtuent.

Virtuent is the operating system internal to each VTE that runs on a base hardware controller which provides up to four FICON connections to the mainframe on DLm for mainframe. The Virtuent software allows for controller emulation supporting 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 the VMAX, legacy VNX or Data Domain disk subsystems attached to the controller.

Virtual tape engines (VTEs)

A VTE appears to the mainframe operating system as a set of IBM tape drives. Mainframe software applications use the virtual drives of the VTE — specifically IBM 3480, 3490, and 3590 drive types — just as they would any physical tape drive. No application modifications are required to integrate them into an existing mainframe tape environment.

VTEs are connected to the mainframe host via FICON channels. Each VTE of a DLm8500 can be configured with up to four FICON channels. A fully configured DLm8500 (six VTEs) therefore provides up to 24 FICON channels to the mainframe host. A DLm2500 has two FICON channels.

In emulating the IBM 3480 / 3490 / 3590 tape drives, each VTE in a DLm8500 can support up to 512 total virtual drives. Configured with a maximum of six VTEs (eight by RPQ), it can emulate up to 3,072 virtual tape drives (or 4,096 with an approved RPQ). 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 DLm have access to all the tape volumes in the DLm, and any emulated tape drive can access all the tape volumes stored in the DLm. The DLm2500 emulates a total of 256 tape drives.

Back-end storage

VTEs process the arriving mainframe tape volume and write it as a single file on the DLm for mainframe storage. Each mainframe tape is stored as a single file whose filename matches the tape VOLSER. 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 DLm are protected with a RAID 6 configuration and hot spare drives for each RAID group.

When configured with deduplication storage, compression is turned off when writing to disk. This enables a higher level of data reduction for those applications that can benefit from deduplication. The deduplication storage capability can provide up to 20 PB of logical storage based on a mix of typical enterprise data (file systems, databases, email and developer files).

In summary, DLm 8500 supports PowerScale, Isilon, PowerProtect DD, legacy Data Domain deduplication storage, PowerMax and legacy VNX storage. DLm DLm2500 supports PowerProtect DD, legacy Data Domain deduplication storage, Isilon storage and PowerScale storage. The use of deduplication storage (PowerProtect DD and legacy Data Domain) is ideal for repetitive backup data, for example, 3990 volume dumps from FDR, DFDSS and/or CA-DISK. Deduplication of repetitive backups can substantially increase the overall data reduction achieved within DLm configuration resulting in significant reduction in storage and transmission costs. Legacy VNX storage is ideally suited for unique data types, such as DFHSM migration, and available for near instantaneous recalls. PowerMax storage is ideally suited for unique data types that require the robust replication feature set contained within SRDF. Disk Library for mainframe is the only available virtual tape library solution that can concurrently support both deduplication and primary storage and dynamically direct tapes to the most appropriate storage on a tape by tape basis.

Disk Library for mainframe management and support

DLM works seamlessly with the mainframe host and does not require any mainframe-based code changes to operate. Additionally, clients do not need to change their production operations or production Job Control Language (JCL).

DLM can be managed using DFSMS functionality and supports all tape Channel Commands. 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, often within seconds or minutes, as opposed to hours.

DLM enables customers to manage and query various status and state conditions including the following:

- Customers can perform specific actions on or retrieve information about DLM directly from the mainframe master console. Customers can easily retrieve information such as available space, configuration, scratch count, and more. Customers can use a web-based application, DLM Console, to remotely log in, query and manage DLM 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 Dell EMC Secure Remote Support (ESRS), which enables Dell EMC Customer Support to establish secure IP connectivity to DLM and remotely log in to the system for diagnosing and troubleshooting system issues. In addition, DLM supports Connect Dell EMC, which automatically sends alerts directly to Dell EMC Support.

Disk Library for mainframe remote replication

DLM offers IP-based remote replication for DLM employing both primary and deduplication storage, which uses the customer'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.

DLM DLM8500 replication, when legacy VNX storage is configured, supports multiple target sites (see Figure 1) per source system, which means customers 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. Customers can choose which virtual VOLSERS will be replicated to each remote site.

DLM 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.

Customers can define the Recovery Point Objective (RPO) in minutes or hours and DLM will perform the replication to meet the defined RPO. Customers can also define different RPOs for different VOLSER ranges based on information criticality, which allows them to better tune their system and not overload the network. For example, critical information may have a low RPO (minutes), whereas less critical information can have a higher RPO (hours).

DLM replication also enables the customer to define quality of service (QoS), which optimizes the network traffic to prevent network overload during peak hours.

For replication, the target DLM does not have to have the same configuration as the source DLM. For example, the source DLM can be a high capacity system with eight VTEs, whereas the target DLM can have only two VTEs and lower capacity if it does not require the same throughput and capacity for DR processing as the source DLM.




Figure 1. Disk Library for mainframe remote replication

DLM8500 replication, when PowerMax storage is configured, utilizes Symmetrix® Remote Data Facility (SRDF®) for replication of virtual tape.

SRDF is a business continuance solution that maintains a mirror image of data at the device level in PowerMax primary STORAGE located in physically separate sites. The SRDF product family provides a mirrored data storage solution that allows you to duplicate production site data on one or more local or remote target Symmetrix systems.

SRDF thus provides a recovery solution for component or site failures, reducing backup and recovery costs and significantly reducing recovery time after a disaster.

The SRDF family of software is the gold standard for remote replication in mission-critical environments.

Built for the industry-leading high-end VMAX hardware architecture, the SRDF family of solutions is trusted for disaster recovery and business continuity. The SRDF family offers unmatched deployment flexibility and massive scalability to deliver a wide range of distance replication capabilities. The SRDF family consists of the following options:

- SRDF/S – Synchronous option for zero data exposure loss
- SRDF/A – Asynchronous option for extended distances
- SRDF/Star – Multi-site replication option
- SRDF/CG – Consistency Groups for federated data sets across arrays

SRDF disaster recovery solutions are based on active remote mirroring and dependent-write consistent copies of data maintained at one or more remote locations. A dependent-write is a write operation that cannot be issued by an application until a prior, related write I/O operation completes. Dependent-write consistency is required to ensure transactional consistency when the applications are restarted at the remote

location. SRDF configurations require at least two PowerMax arrays. These arrays are also known as the primary and the secondary system/array. Both sites can be located in the same room, in different buildings within the same campus, or hundreds to thousands of kilometers apart.

The SRDF/Synchronous option maintains a real-time mirror image of data between the arrays. Data must be successfully stored in PowerMax cache at both the primary and the secondary site before an acknowledgment is sent to the production host at the primary site.

The SRDF/Asynchronous option mirrors primary site data by maintaining a dependent-write consistent copy of the data on the secondary site at all times. SRDF/A session data is transferred from the primary to the secondary site in cycles. The point-in-time copy of the data at the secondary site is only slightly behind that on the primary site.

SRDF/A has little or no impact on performance at the primary site as long as the SRDF links contain sufficient bandwidth and the secondary system is capable of accepting the data as quickly as it is being sent across the SRDF links.

The SRDF family of remote replication software offers various levels of Symmetrix-based business continuance and disaster recovery solutions.

Data Domain Replicator for both DLM2500 and DLM8500

Data Domain storage systems incorporate IP-based replication to a remote site. 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. This ensures fast recovery in case of the loss of primary data, the primary site, or secondary store. Data Domain Replicator software provides simple, fast, robust WAN-based disaster recovery for the enterprise.

Data recovery

When implementing a disaster recovery solution in a mainframe environment, customers 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. DLM can become an integral part of a mainframe disaster recovery solution. 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, it eliminates the risk of losing a tape that might be critical for the recovery process since all the tape information is kept on RAID-protected disk and not physical tape.

Flexible Recovery testing

Disaster recovery tests are extremely important in mainframe environments. Many customers 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 customers 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 replication is not interrupted during DR testing and data is always protected. Customers have two options when performing their disaster recovery tests:

- **Read-Only mode:** In this mode, customers mount their tape volumes (VOLSERs) at the remote site as Read Only during the disaster recovery tests. This mode allows customers to read tape data located at the DR site and perform restore operations to check their disaster recovery procedures. Customers may see their tape information change during the disaster recovery tests, as replication continues to update the volumes at the remote site.
- **Read / Write mode:** In this mode, customers perform full disaster recovery tests, including read and write operations, without updating any production tape volumes. Customers 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 system. Full disaster recovery tests can then be performed on these snapshot copies of the tapes. Once the disaster recovery tests are complete, the customer can eliminate the snapshots and free up disk space for future disaster recovery tests.

Disk Library for mainframe architecture

Redundancy

DLM8500 is designed with significant redundancy to provide for continuous data availability. It includes redundant components such as VTEs, internal switches and more. The storage in DLM provides RAID six 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. The customer can load the configuration of an unavailable VTE 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 done.

RAID 6 data protection

Reliability is further enhanced with the implementation of a RAID 6 configuration with all disk storage. 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 DLM back-end storage configuration.

Hot standby controller

The internal storage controller of DLM includes a hot standby storage controller to 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 and issues a call home to the Dell EMC Customer Support Center.

Deduplication Storage for DLm

Deduplication Storage 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 fully integrated deduplication storage system, this capability provides:

- Up to 10x the reduction in the raw storage requirement because of the deduplication process. Note: 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
- RAID 6 protection with hot standby drives
- Flexible scalability options that can scale up to 20 PB of logical storage.

The deduplication storage integrates into the mainframe environment seamlessly by pointing the selected mainframe workloads to a specific range of VOLSERS, which in turn pass this data to the integrated deduplication backend storage system.

DLm8500 and the use of Elastic Cloud Storage (ECS) to replace physical tape for long-term retention

Volumes of 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 5.4 changed the protocol between the DLm and ECS to S3 (it was NFS). 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 offered several additional long-term retention operational benefits to using LTR engine over the prior release:

- a. Allow configuring '0' days in policy "move after" field.
- b. Changed maximum number of "move after" days from 730 (2 years) to 19999 (47 years).
- c. Migration policy changes can now be reconfigured non-disruptively (no Virtuent restart required).
- d. 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 Virtuent (virtual tape engine firmware).
- e. Significant performance improvement of the post-migration file-validation process.

- f. Single, simple, "restore" command to restore tape volumes from Tier 2




Figure 2. Connection of AWS as a 2nd tier of storage for Long Term Retention of tape volumes

GDDR Tape Leverages Global Virtual Library and DASD automated failover for DLM8500

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 EMC 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 disasters. It leverages the Global Virtual Library technology introduced in the previous release of DLM. 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.



Figure 3. Use of GDDR Tape to automate failover of sites

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

Using Dell EMC’s 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 EMC 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.

Conclusion

DLm enables complete mainframe tape replacement with its ability to support both primary and deduplication storage that can be matched to individual customer 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. Reduced footprint is an additional advantage over physical tape.

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

DLm is a seamless tape replacement solution for both physical and competitive virtual tape that leverages many private as well as public cloud providers. GDDR tape uses field-proven GDDR technology for failover and DR test automation to eliminate scripts and cumbersome procedures in run books.

DLm uses proven Dell EMC hardware and software technology to provide mainframe tape customers with the best in performance, scalability and availability over traditional physical tape and tape-based solutions. Customers replacing their tape infrastructure

with disk-based solutions need not change any of their existing mainframe applications or processes.