

DM - What's new from Dell Technologies: PowerMax and DLm updates

Robin Fromm
Dell EMC

November 7, 2019
Session **DM**

Place your
custom
session
QR code here.
Please remove
the border and
text
beforehand.



DELL EMC PowerMax & VMAX All Flash Update for IBM Z

Robin Fromm – Global CTO Mainframe Solutions



DELL EMC

PowerMax/VMAX Unique Features

- **Large Global Cache**
 - Advanced Algorithms
 - Partitioning
 - FlashBoost Improving Response Time for Read Miss IO
- **SRDF Is Single Program Product With 4 Operational Modes**
 - SRDF Is A Mirror Enabling IO TO Served Through Replication Link Avoiding DASD Swap Or Site Failover Due To Local Device Or Raid Failure
 - Sync, Async, Adaptive & Active Active (FBA)
 - Dynamically Change Operation Modes
 - Async Automatically Pages To SRP To Ensure Stability
 - Dynamic Volume Expansion With Active Replication
- **Virtual Storage Provisioning And Extreme Space Efficiency**
 - Single Track Allocation 56K CKD, 128K FBA
 - 2:1 Or Better Data Compression For FBA Data
- **TimeFinder SnapVX**
 - Target Less Snapshots Requiring No Addresses
 - Flashcopy Compatibility
 - 256 Consistent PIT Copy Of A Volume
 - One Full Copy Of Storage Can Yield 256 Active Usable Copies Of Data
- **Extreme Space Efficiency Enables New Cost Effective Possibilities**
 - Parallel Application Development And Testing (DevOPS)
 - Multiple Data Recovery Points From Data Corruption Or Destruction Events

VMAX 950F All Flash

- ⇒ 1-8 Engines
 - ⇒ 56Gb/s Infiniband Engine Interconnect
 - ⇒ 72 Intel Broadwell 2.8GHz Cores/Engine
 - ⇒ Up to 256 16Gb/s FICON ports
 - ⇒ Up to 16TB cache (1TB, 2TB DDR4)
-
- ⇒ 120-drive DAE (1920 max)
 - ⇒ Up to 1.7PBu




PowerMax 8000

- ⇒ 1-8 Engines
 - ⇒ 56Gb/s Infiniband Engine Interconnect
 - ⇒ 72 Intel Broadwell 2.8GHz Cores/Engine
 - ⇒ Up to 256 16Gb/s FICON ports
 - ⇒ Up to 16TB cache (1TB, 2TB DDR4)
-
- ⇒ NVMe Drives:
Storage Class Memory & NAND flash
 - ⇒ 24-drive DAE (288 max)
 - ⇒ Up to 1.7PBu



VMAX 950F & PowerMax 8000 All Flash Engines




Broadwell 18 Core CPU 2.8GHz

72 CPU Cores per Engine / 576 per System


Up to 2 TB Cache per Engine / 16 TB per System

Up to 32 FICON Ports per Engine / 256 per System



Virtual Matrix enables global memory

Data cache and metadata striped across all directors




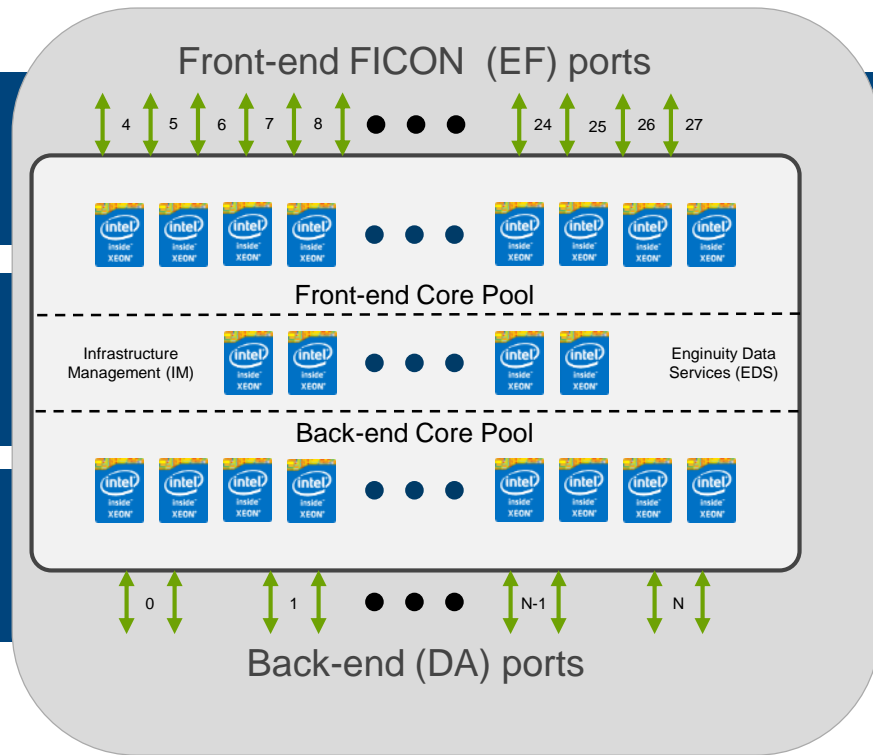
Architected for All Flash

zBricks & zPowerBricks (engine + capacity pack) optimized for multi-core CPUs

Massive increase in per-port performance

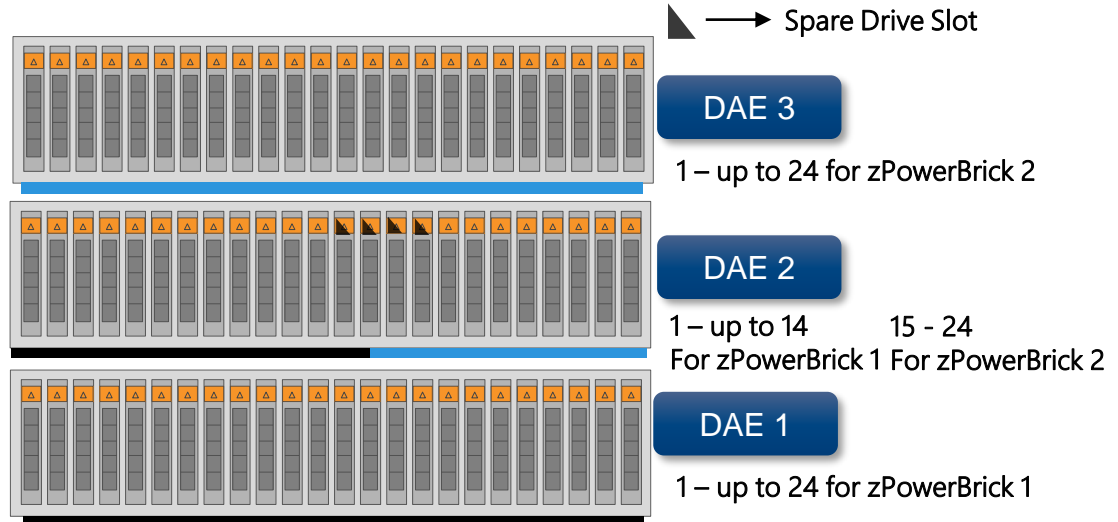
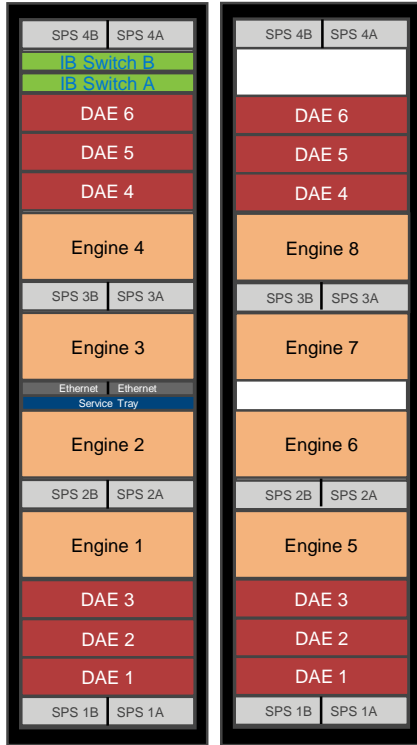
Balanced performance across CPU resources

Optimized for flash performance




PowerMax 8000 System Configuration Details

- Each engine starts with 2 DAEs
- Infiniband switch gets added with 2nd engine
- Each additional engine adds 3rd DAE
 - Each additional even numbered engine gets a 3rd DAE which is shared with the previous odd numbered engine



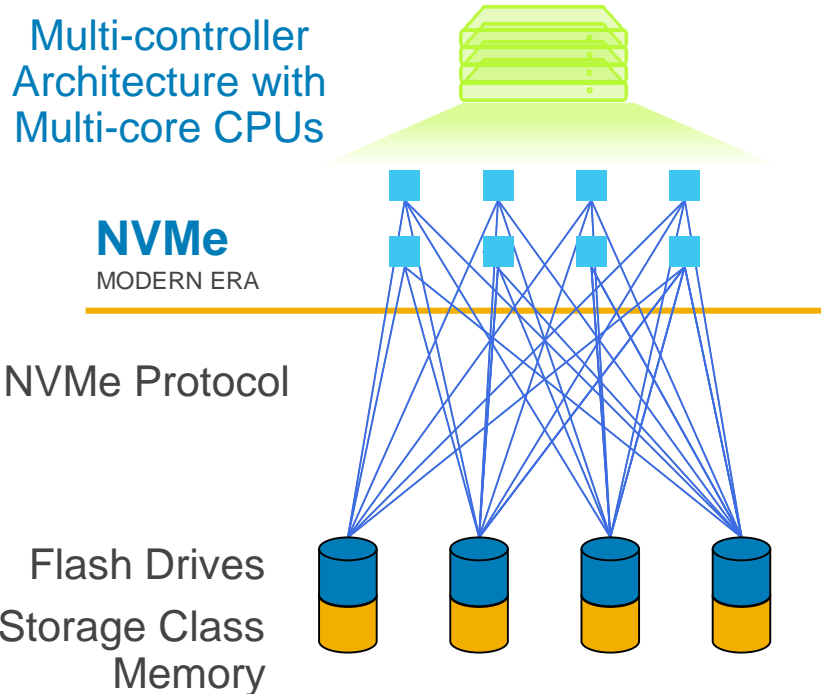
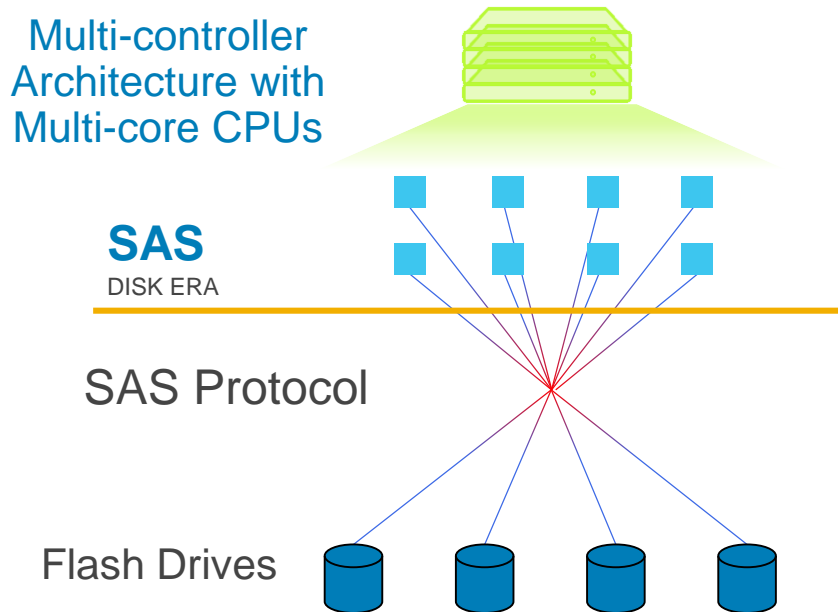
New! FLASHBOOST™ for Mainframe

- Performance feature introduced with VMAX non-mainframe now reintroduced in PowerMax & VMAX All Flash
- Delivers performance acceleration for high demand read intensive workloads
 - VMAX will bypass its internal cache on read miss workloads
 - Cache loading overhead happens asynchronously
 - Performance improvements ~2X on Read Miss
 - Significant savings for All Flash disk response profile

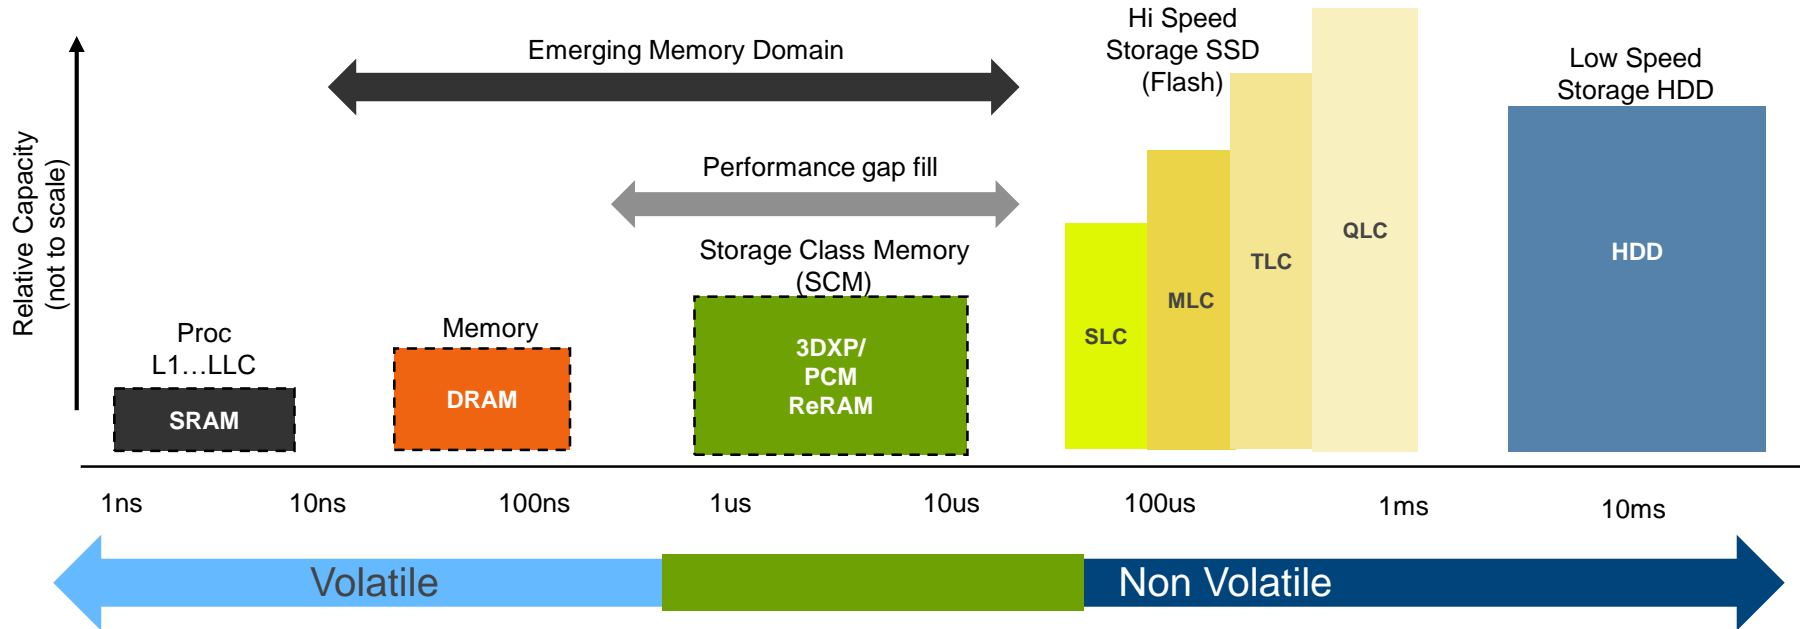


NVMe unlocks the power of next generation media

- Maximizing the performance of multi-controller architectures




Storage Class Memory: Filling the latency gap



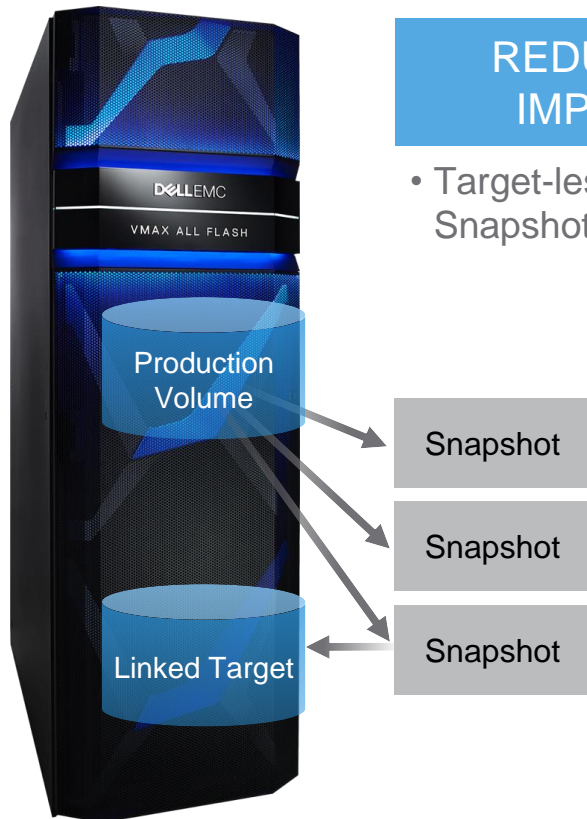
Virtual (thin) Provisioning Concept

- Storage capacity is structured in common data pools
- Thin devices are logical volumes that are provisioned to hosts
- Workload is spread across many disks



Incorporate snaps for data copies

New TimeFinder SnapVX™



REDUCED IMPACT

- Target-less Snapshots

INCREASED AGILITY

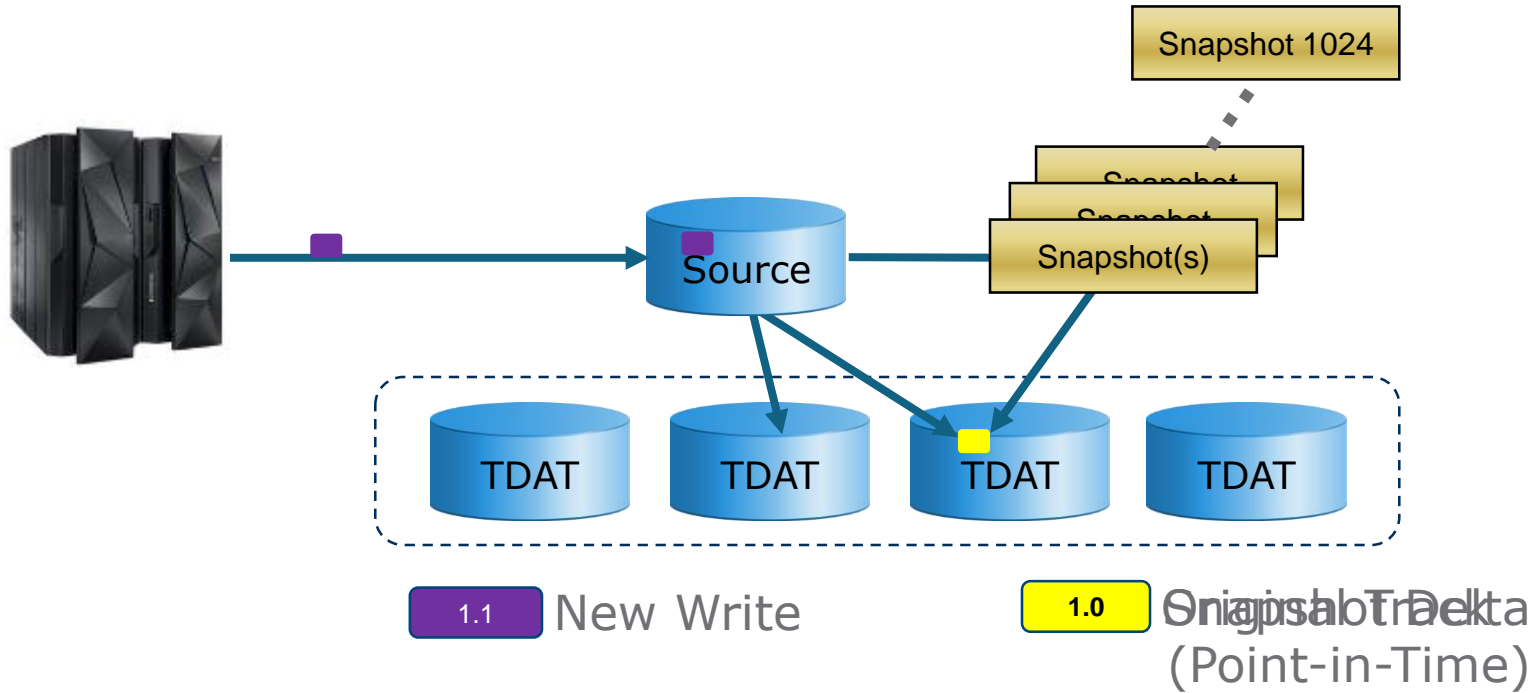
- Up to 256 Snapshots per source
- zDP enables up to 1024 Snapshots per source
- Up to 1024 Linked Targets (snaps and/or clones) per source

EASE OF USE


- User-defined name/version number
- Automatic expiration if desired
- Secure Snapshots

SNAPVX SPACE EFFICIENCY

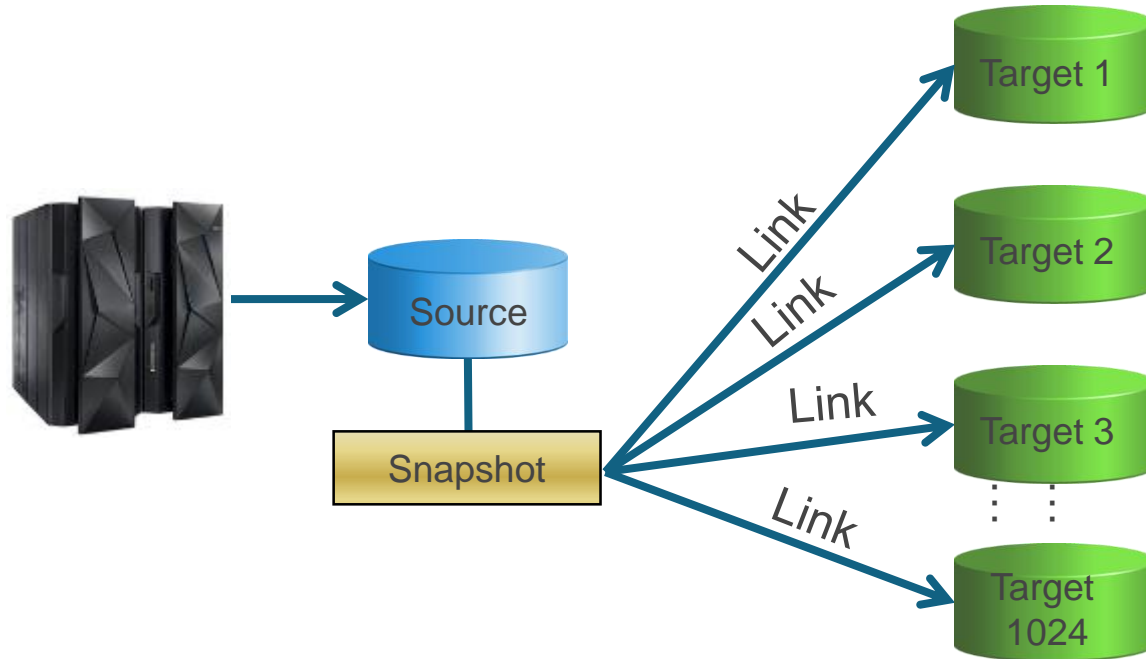
REDIRECT-ON-WRITE AND SNAPSHOT DELTAS



SNAPVX Link Command

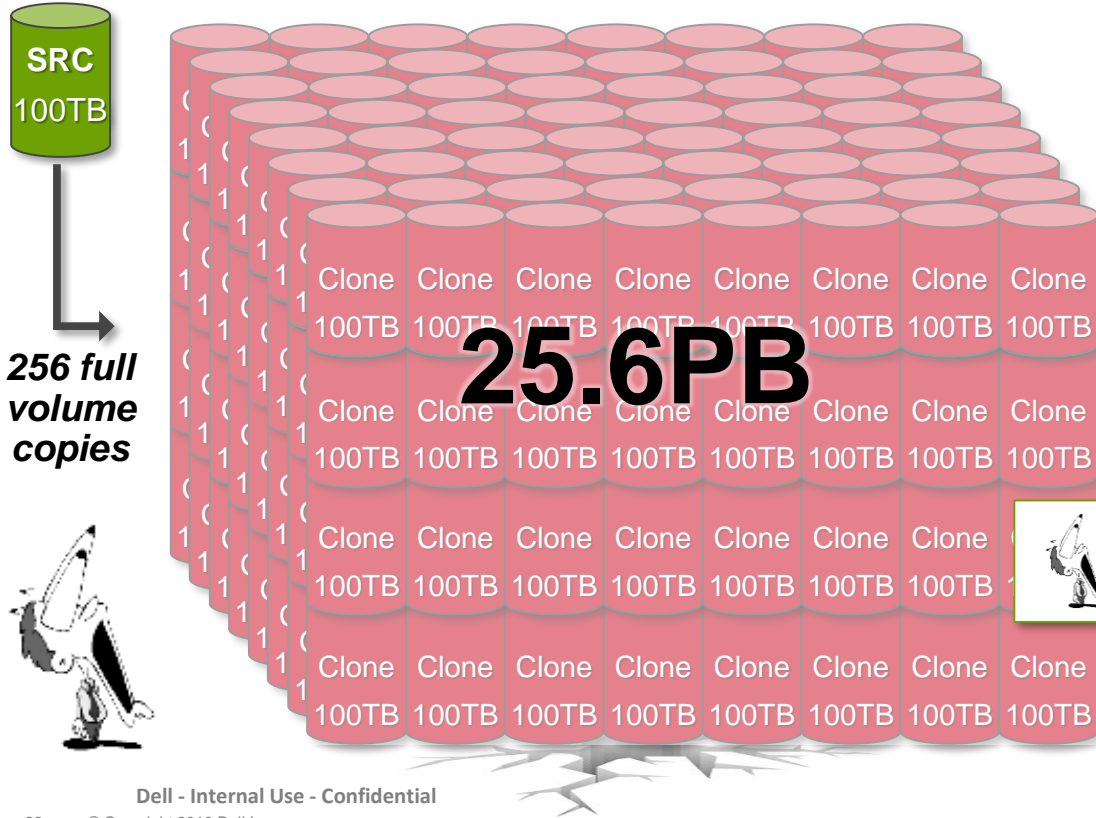


SNAPVX Link Command




Which local replication solution would you choose?


NON-SNAPVX SOLUTION



SNAPVX SOLUTION



- * Based on the following observed avgs
- Snaps every 10 mins
- Observed change rate



99.9%
**LESS COPY
SPACE USED
WITH SNAPVX!!**



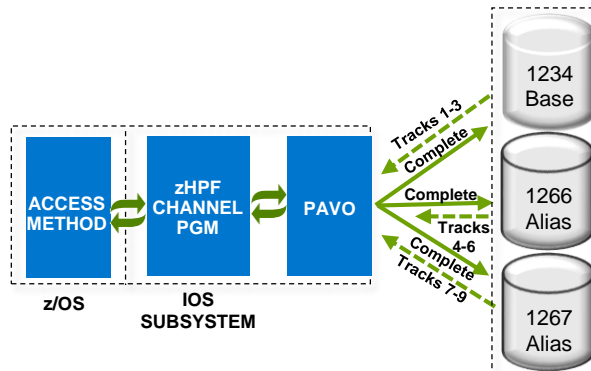
zHPF Optimizers for z/OS

Intercept zHPF I/O and transform it to:


- Create multiple I/O requests
- Execute them parallel
- Reduce latency

PAV OPTIMIZER

UP TO **70%** 
FASTER
RESPONSE TIME



MIRROR OPTIMIZER



Cuts replicated I/O response time in half

For any zHPF performance sensitive data

VMAX and PowerMax IBM Z synergy

- **Multi-Incremental FlashCopy**
- **SuperPAV**
- **PPRC Event Aggregation**
- **Storage Controller Health Messages**
- Enhanced zHPF support
 - List Prefetch
 - Format Writes
 - Bi-Directional transfers
 - BSAM/QSAM support
 - zHPF Extended Distance II
- Query Host Access (ICKDSF VERIFY OFFLINE)
- FICON Enhancements:
 - FICON Dynamic Routing
 - 32K devices per FICON channel
 - Forward Error Correction Codes
 - Read Diagnostics Parameters
- zHyperWrite
- zFBA support
- PPRC SoftFence
- Non-disruptive state save
- 1 TB EAV
- Dynamic Volume Expansion
- D@RE external key manager support
 - IBM SKLM (Secure key Lifecycle Manager) & Gemalto



Disk Library for mainframe DLm8500 release 5.1

Robin Fromm – Global Field CTO Mainframe Solutions




DLm8500 – What Matters - Unique Features

- Continuous Availability In A Single System
- Ability to Read/Write Test %100 of Data Without Disrupting Disaster Recovery
- Superior and Consistent Performance Over The Life Of The System (3X Competitive Offerings) Up to 12 GB/Sec!
- Deduplication Enables Extreme Space Efficiency (8:1 or more reduction)
- Leverage Cloud Object Storage (ECS) For Long Term Retention Of Data

What is Disk Library for mainframe?

- “Virtual” mainframe tape (tape on disk) for all tape use cases
- The first “all flash” virtual tape storage with powermax offering universal data consistency™ and truly synchronous tape (SRDF/S)
- IBM tape emulation, but significantly faster for IBM & unisys mainframes – 3480, 3490, 3590
- Up to 4096 virtual tape drives
- Transparent – looks just like IBM tape
- SAS, Flash & Cloud Object Storage Options
- Dell EMC invented 100% tape on disk
 - QA/tested by dell EMC
 - Developed & manufactured by dell EMC
 - Maintained by dell EMC
 - Professional services by dell EMC



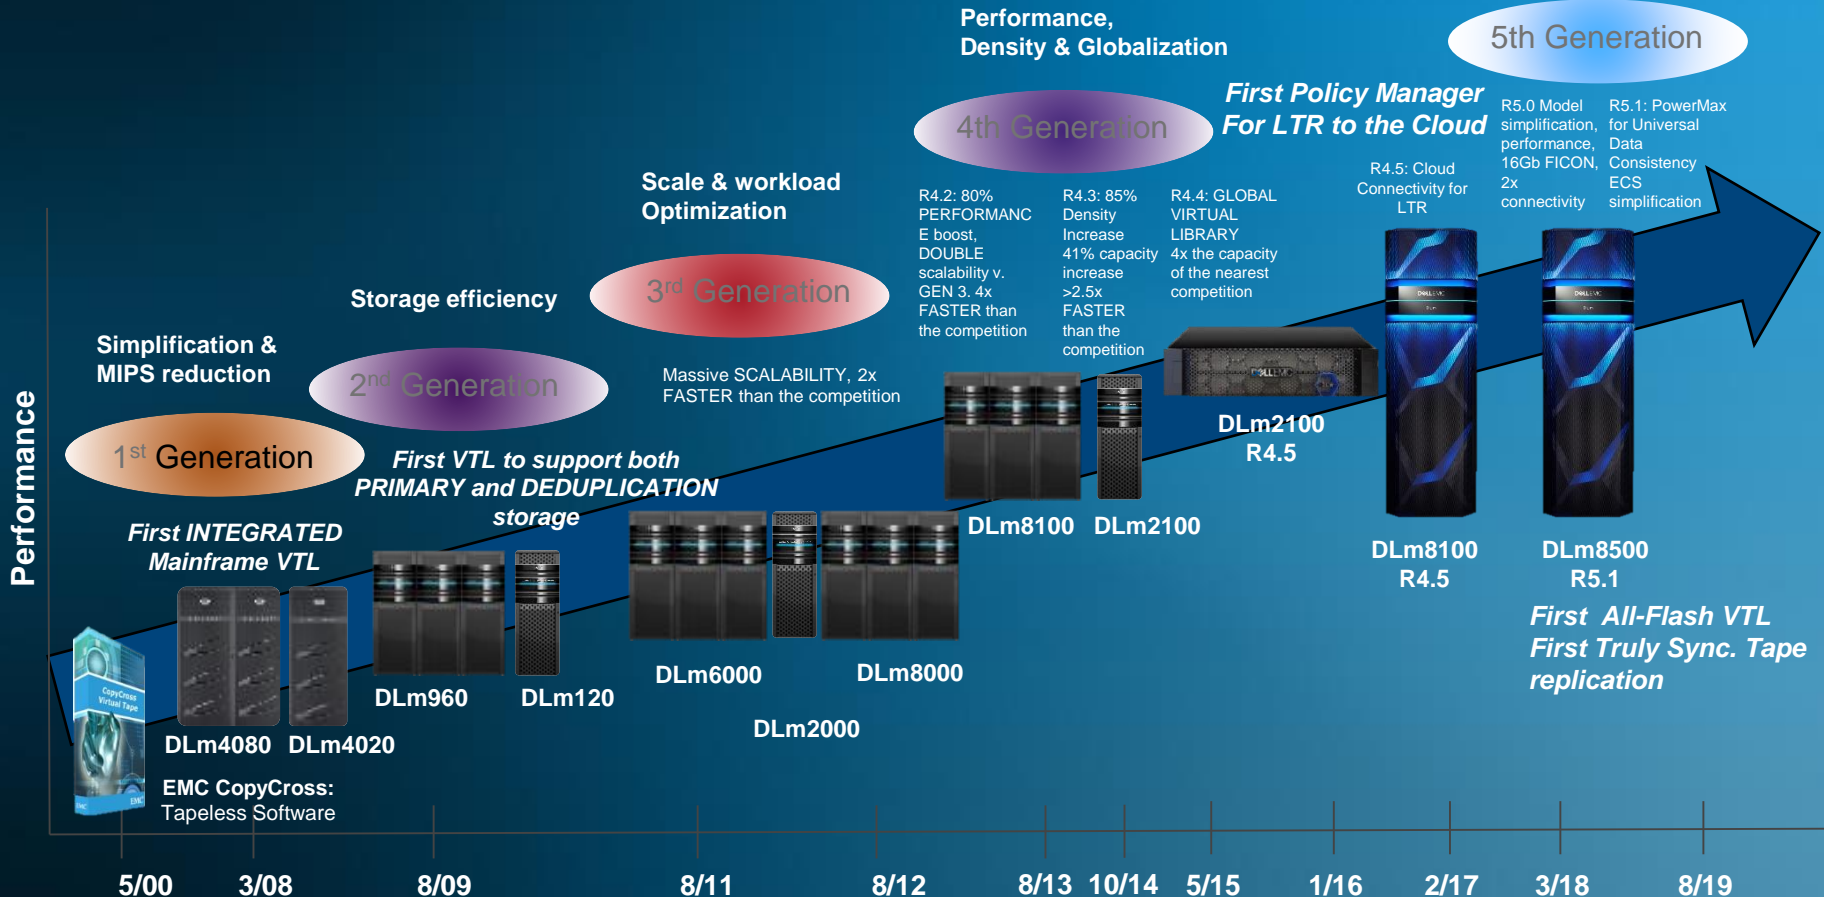
Single-cabinet configuration



Multi-cabinet configuration

Dell EMC's Mainframe Virtual Tape Innovation History


Performance,
PowerMax & Cloud



First All-Flash VTL
First Truly Sync. Tape replication

DLM Starts With Dell EMC's Industry Leading Storage & Cloud Offerings

And allows you to leverage them all



New DLm8500 Release 5.1

Enhances DLm8500 release 5.0

Single-frame Solution

- 1-2 VTEs
- Up to 1024 Virtual Tape Drives
- Up to 3 GB/Sec Bandwidth
- 1-2 1Gb management switch
- 1-2 10Gb Data switch
- Deduplication Storage Options
 - DD6300
 - DD6800 (HA optional)
 - DD9300 (HA optional)
 - DD9800 (HA optional)
- Storage sharing with IBMi (iSeries), open systems
- Long-term tape retention cloud (ECS)



Multi-frame Scale out Solution

- 1-8 VTEs
- Up to 4096 Virtual Tape Drives
- Up to 12 GB/Sec Bandwidth
- 2 1Gb management switch
- 2 10Gb Data switch
- Deduplication Storage Options
 - DD6800-HA
 - DD9300-HA
 - DD9800-HA
- Long-term tape retention cloud (ECS)
- PowerMax8000 high performance DASD




DLm 5.1 “under the covers”




Dm8500 Virtual Tape Engine – 2x All Around

- 16Gb FICON ports into each VTE - Up to 32 per Dm Array – twice that of Gen 4
 - 4 port Virtual Tape Engine using Dell Servers
 - Up to 512 Virtual Tape Drives per VTE – twice that of Gen 4 VTEs
 - 1,500MB/Sec Performance Rating – twice the performance of Gen 4 VTEs
- New and faster Dell server for the Dm VTE Appliance




Data Domain Connectivity for High Performance Extreme Space Efficiency



- Data Domain
- High Availability Option
- Inline Deduplication 8:1 or More data reduction
- @DARE Encryption Option

Data Domain High Availability

ACTIVE Controller



STANDBY Controller



- High availability of backup, archive, and recovery data on Data Domain ensures operational continuity to minimize downtime for users and processes.
- HA configurations are supported on DD6800, DD9300, DD9800 and the legacy DD9500
 - Delivering business continuity for both Large and Midsized Enterprises

Details of DLm 5.1 enhancements

PowerMax 8000 attachment


- Block Storage (supports mainframe and open environments)
- Universal Data Consistency™ (between PowerMax & DLm)
- Synchronous or asynchronous tape protection (SRDF/S, SRDF/A)

New Cloud capability and simplification


- Ability to move data on demand to the cloud tier
- Single, simple, “restore” command

Additional Enhancements

- SNMP V3 network security
- Additional installation options
 - Customer supplied rack
 - 3-Phase power available

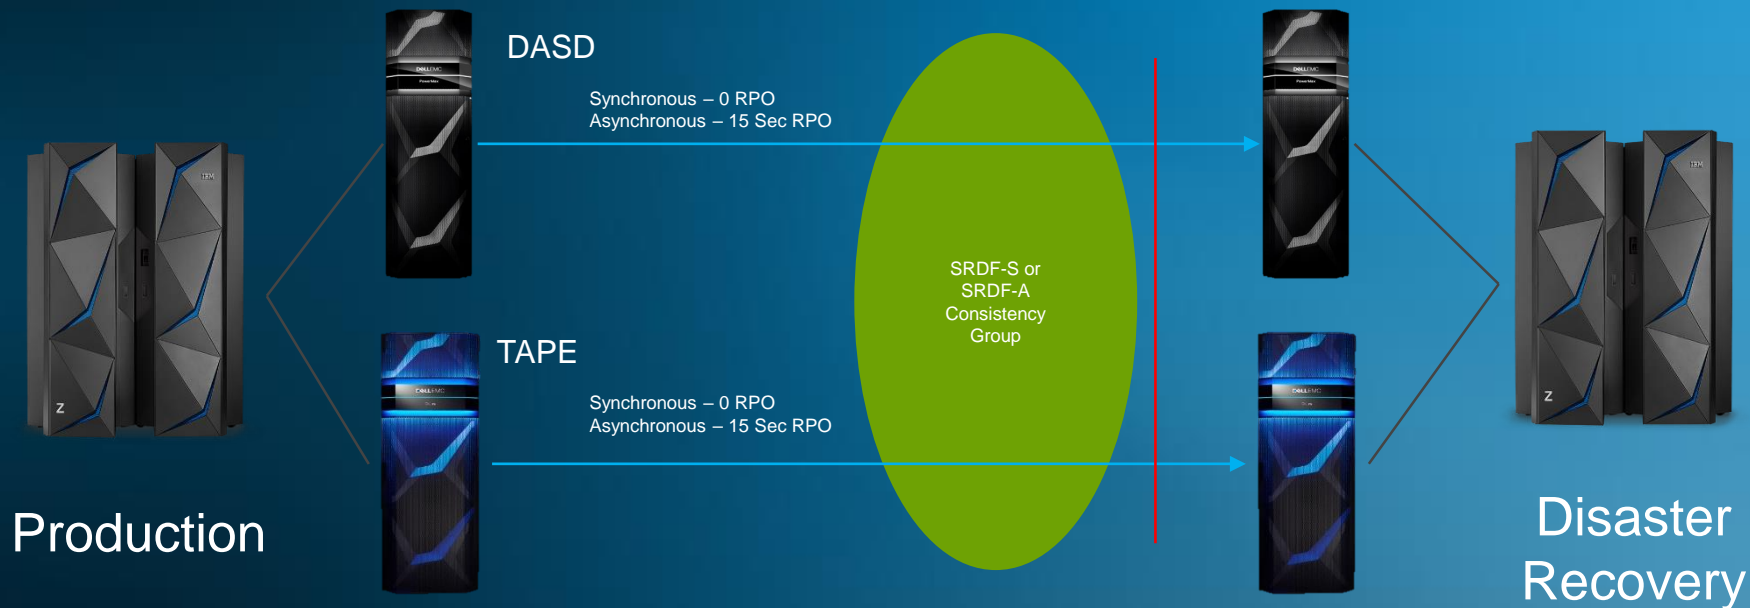


DASD and Tape Data Consistency Issue



- Missing Data
- DASD Files And System Catalogs Are Ahead Of TAPE Files

DASD and Tape Data Consistency Issue - Solved!




- DASD Files, Tape Files And System Catalogs Are Synchronized
- No Missing Tape Data
- No Performance Impact

Why use PowerMax 8000 with DLm?

- When Universal Data Consistency™ is needed to provide data consistency between PowerMax 8000 primary storage (DASD) & DLm (TAPE) on PowerMax for applications that demand it:
- Examples:
 - DFHSM ML2 Migration – DASD File is scratched and catalogs are updated before the TAPE file is replicated resulting is missing ML2 data at the DR location
 - Batch TAPE Files – Jobs or Job Steps terminate and catalogs are updated before the TAPE file is replicated – resulting is missing TAPE data at DR location
- SRDF Synchronous and Asynchronous replication is supported



Dm8500 Single Rack DD Storage Solution



Dm8500

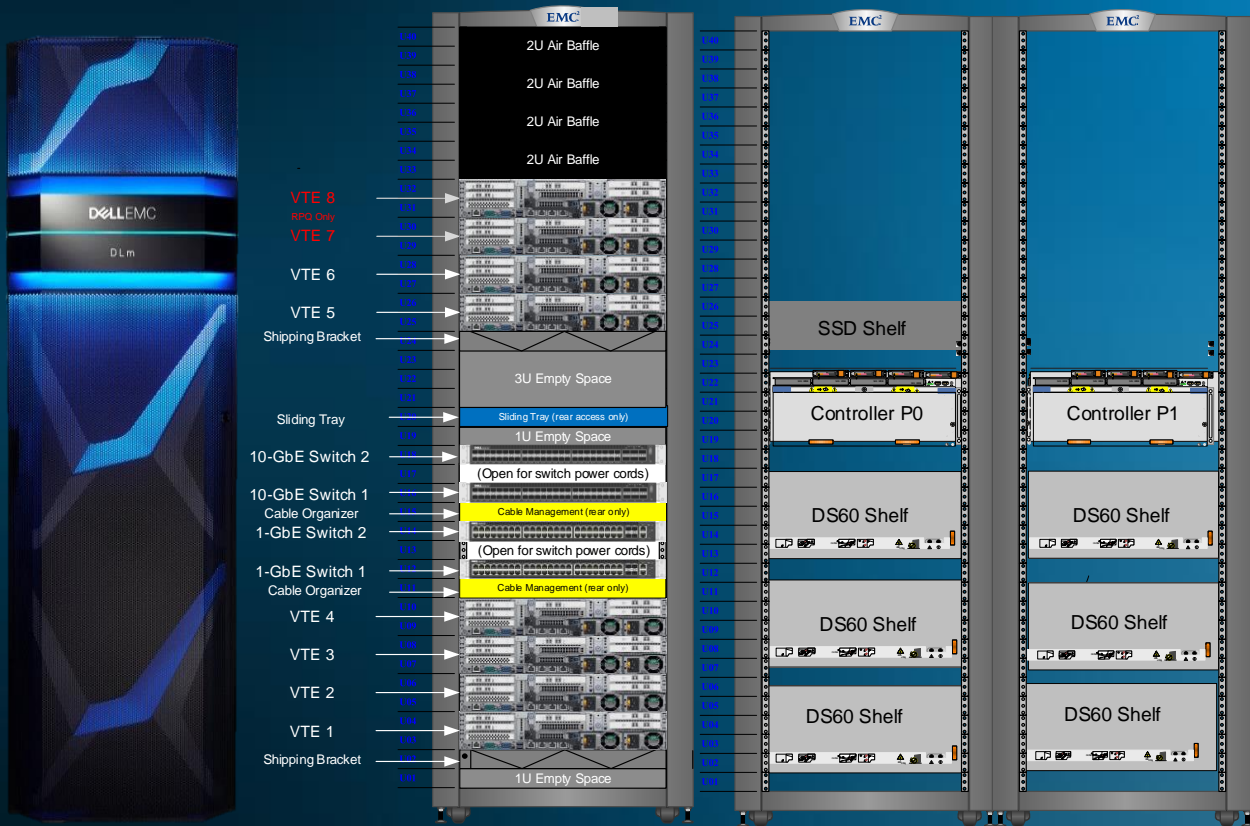
- 1 - 2 VTEs
 - Up to 4 FICONS per VTE
Total 8 FICONS
 - Up to 4 FICON Licenses per VTE
 - 512 – 1024 Tape Drive Support
 - Up to 3 GB/Sec

DD6300/DD6800

- Minimum useable capacity: 34 TB
 - 272 TB @ 6:1 Deduplication
- Maximum useable capacity: 420 TB
 - 3200 TB @ 8:1 Deduplication



DLm 8500 / DD9800 Scale Out



DLm8500

1 - 6 VTEs

- 1 – 24 FICON Channels
- 512 – 3072 Tape Drive Support
- ~750MBytes/sec single FICON
- ~1200MBytes/sec 32 FICONs

DD6800/DD9300/DD9800

- High Availability Configuration
- DD6800 Minimum useable capacity: 94TB
- 752TB @ 8:1 Deduplication
- DD9800 Maximum capacity: 1PB
- 8000TB @ 8:1 Deduplication

Comprehensive PowerProtect DD Portfolio



**BACKUP
INGEST**
(with DD Boost)

Virtual Edition

Up to 11.2TB/hr
for 96TB

**LOGICAL
CAPACITY**
(with Cloud Tier)

Up to 14.8PB
for 96TB

**USABLE
CAPACITY**
(with Cloud Tier)

Up to 288TB
for 96TB

DD3300

Up to 7.0TB/hr

Up to 4.8PB

Up to 96TB

DD6900

Up to 33TB/hr

Up to 56.1PB

Up to 864TB

DD9400

Up to 57TB/hr

Up to 149.8PB

Up to 2.3PB

DD9900

Up to 94TB/hr

Up to 211PB

Up to 3.25PB

Logical capacity based on up to 50x deduplication (DD3300) and up to 65x deduplication (DD6900, DD9400, DD9900) based on additional hardware-assisted data compression of up to 30%. Actual capacity and throughput depends on application workload, deduplication and other settings.

Typical Disaster Recovery testing with DLm



Read-only mounts

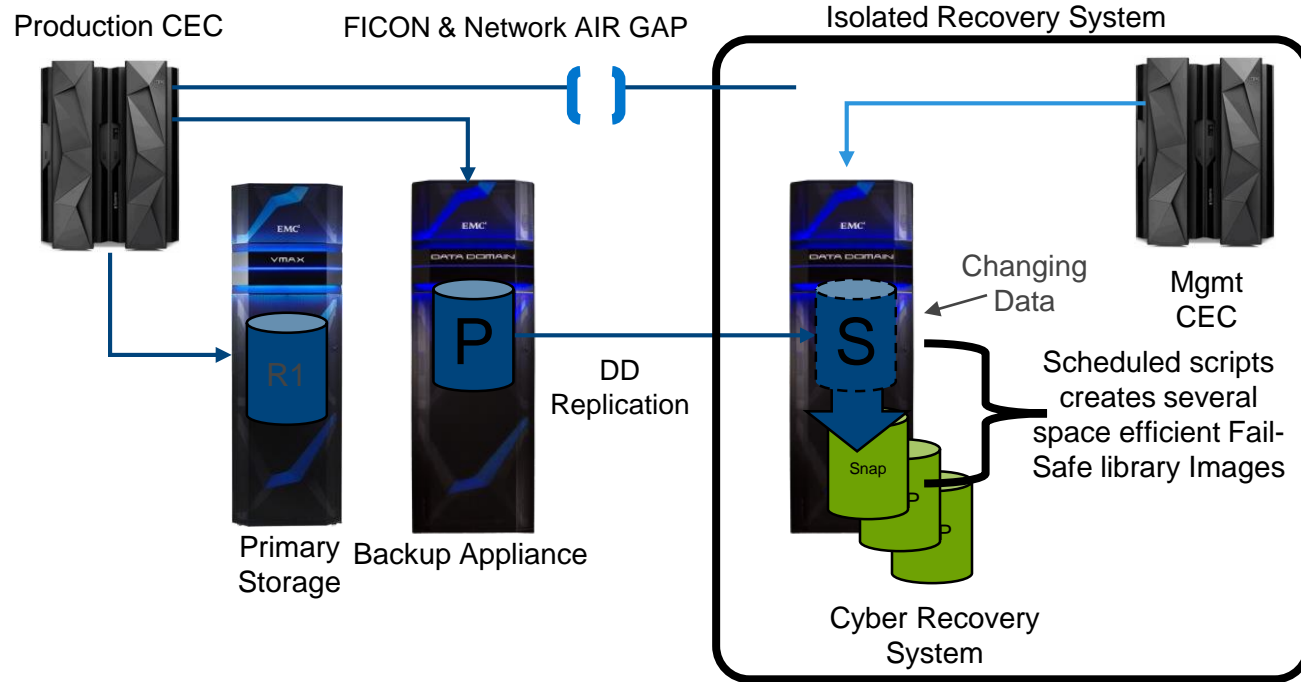
- Disk arrays allow instant “read-only” copies
- Confirm that tapes can be mounted and all required data can be accessed
- No incremental storage capacity required

Snapshots

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

Remote replication is uninterrupted during testing

DLM CYBER RECOVERY SOLUTION COMPONENTS



- P2P dedicated Ethernet replication ports
- All TCPIP Sockets Closed replication ports except Replicator
- Separate Dedicated Replication Network
- Coordinate TMC Catalog Backup with Restore Points
- Instant Access
- Retention Periods / WORM



Cost Effective



RPO = Hours to 1 Day*



RTO = Hours to Days**

- Can co-exist with and enhance an existing BC/DR Solution
- Tape Data is periodically copied in to CRS environment
- File System Snapshot multiple space efficient Fail-Safe Library images
- Preferred that Management CEC is "hardwired" to CRS Solution
- Un-addressable Snapshot Copies allow multiple restore points
- Management CEC can be used used to periodically validate data

* Adjustable based on desired security window

** Depends on # of Volumes, Data etc.


CYBER RECOVERY EFFICIENCY

Filesys Compression

From: 2019-10-23 06:00 To: 2019-10-30 06:00

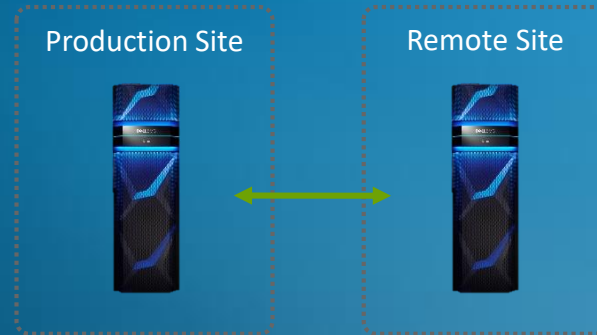
	Pre-Comp (GiB)	Post-Comp (GiB)	Global-Comp Factor	Local-Comp Factor	Total-Comp Factor (Reduction %)
-----	-----	-----	-----	-----	-----
Currently Used:*	30609529.7	643832.0	-	-	47.5x (97.9)
Written:					
Last 7 days	19662362.7	57548.9	100.4x	3.4x	341.7x (99.7)
Last 24 hrs	2461299.6	6592.0	107.4x	3.5x	373.4x (99.7)
-----	-----	-----	-----	-----	-----

Data Domain - Guaranteed Replication



- Three Modes of Operation


- **Guaranteed Replication** – Job Termination Waits Unit Tape is Replicated
- **RUN** – Final Tape Replication At Job Termination – NO Wait
- **Sync Mode** – Replication and Wait On Tape Sync Command
- Active Replication Cycles During Job Execution Reduce Final Replication Delays
- Most Data Is replicated Prior to job Termination



Replacing Physical tape:

Cloud Object Storage and DLm5.1 Long-Term Retention

Benefits of DLm8500 + ECS For Long-Term Retention



1. Reduce costs

- **Keep Only The “Working Set” Of Tape Datasets On Primary Tier Of Storage**
- **Eliminate** Physical Tape Media
- **Repurpose** all or a portion of an existing **ECS** environment

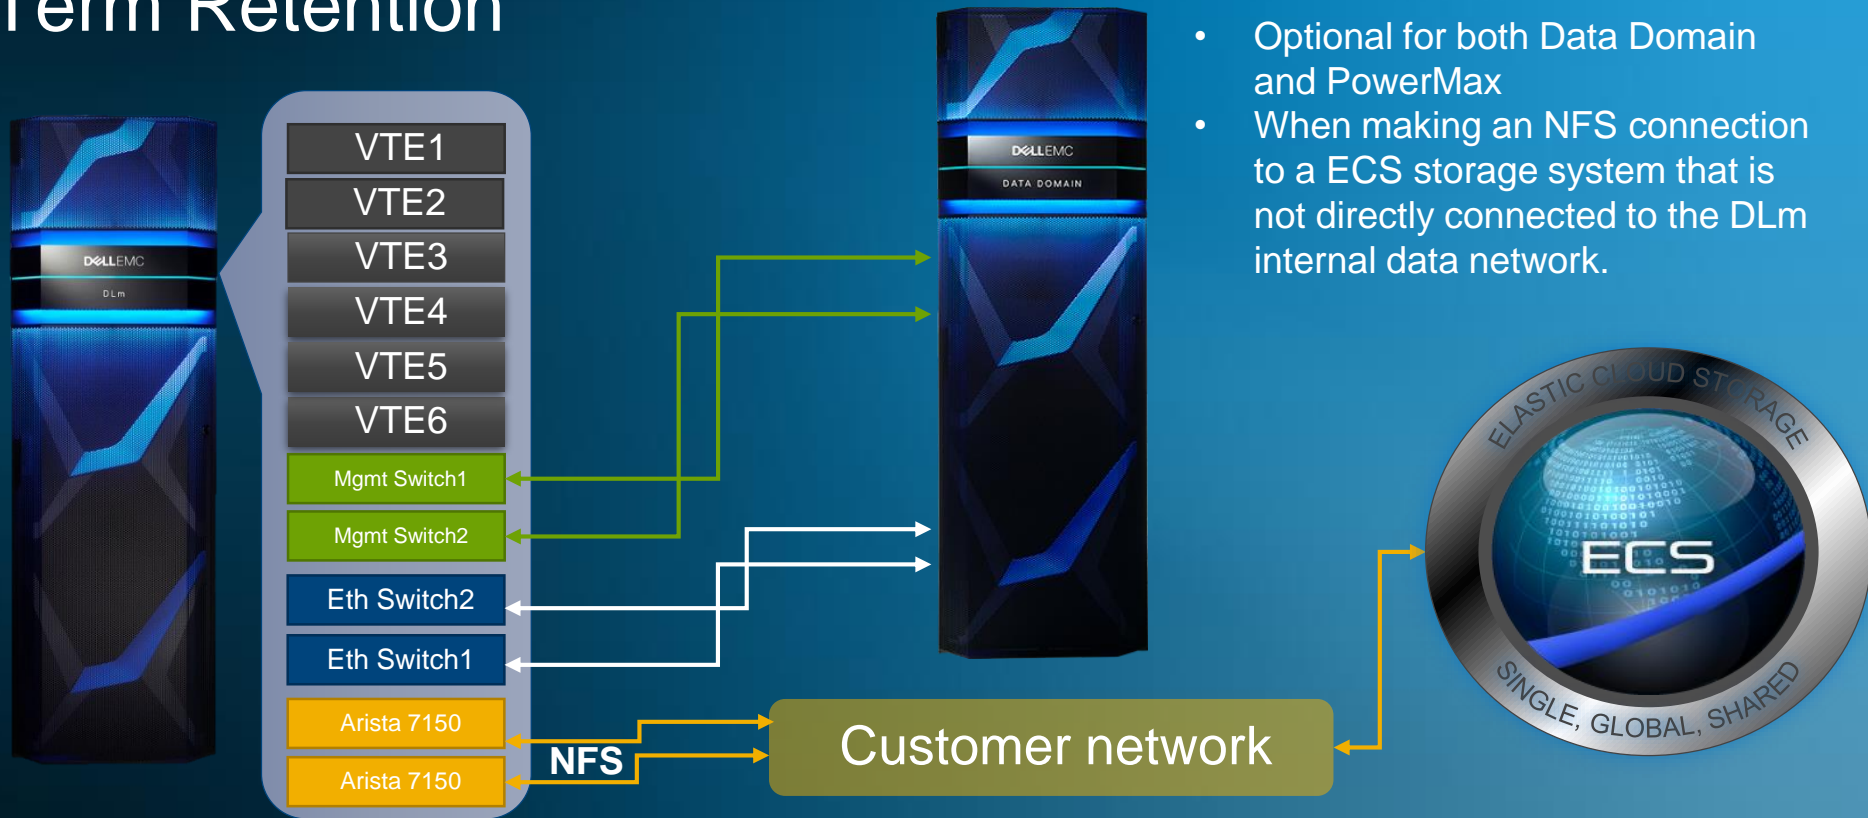
2. “Simplify & Quantify” DR

- Knowing what’s on ECS vs. DLm creates more accurate accountability

3. Create a “Safety Valve” To Accommodate Unexpected Increases In Tape Workload

4. Create “political capital” as mainframe participates in the organization’s cloud strategy.

Elastic Cloud Storage Connectivity (Object) for Long Term Retention



GDDR Tape: DLm DR & Test Automation



- Planned and unplanned outages
- Leverages Dell EMC GDDR Technology
- Automates DLm DR Test Setup and Tear Down
- Automates Switch Over / Failover / Failback
- Supporting DLm w/ VNX Data Domain & PowerMax

Summary of DLm 5.1 and ongoing advantages



Universal Data Consistency between disk & tape + tape Synchronous Copy via PowerMax 8000 SRDF/S & ability to r/w test 100% of data with no DR interruption



Extreme storage efficiency through deduplication combined with Superior and consistent performance



Additional enhancements include SNMP V3 network security, the ability to use a customer-supplied rack and 3-Phase power, configurable at installation



Lower Tape TCO, eliminate physical tape with DLm + Elastic Cloud Storage with simplified and expanded capabilities



Continuous, High Availability in a single frame to lower costs, via unique virtual tape engines & deduplication storage

D~~E~~LL EMC