

Physical Tape Replacement with DLm – a Financial Institution's Journey to Savings

Carl Swanson
Advisory Mainframe Systems Engineer
DELL EMC

November 2018
Session BC



Legacy Virtual Tape System

- Virtual tape system relied on backend physical tape
 - Small disk cache (buffer), large tape library
- Helped drive up the utilization (capacity) of tape cartridges
- Highly dependent on physical tape drives to store data
 - Little to no reduction in the number of physical tape drives occurred
- Application groups were not accountable for how long data was kept on tape
 - They started keeping data longer thinking the virtual tape system would handle everything
- Therefore things evolved to look like this.....

Physical Tape consumed a lot of Floor Space!



Minimum of 2500 SQ feet of floor space plus additional rack space not shown
At each data center

What were the issues?

- Classic Virtual Tape system that was highly dependent of physical tape
 - Small disk buffer, large tape library
 - Not enough physical tape drives in the system
 - 90% or more of the data was on physical tape, not in the buffer, causing performance issues
- Initially system did improve tape cartridge utilization
- System did not scale
 - Expect for floor space, which it consumed very well!

Modern Virtual Tape system requirements (new times new requirements)

- All tape data must be replicated
- All data must be encrypted during replication
 - Later requirement: All data at rest must be Encrypted
 - Where possible, use the same methods as in use by the distributed systems encryption
- Disaster Recovery RPO / RTO is to be less than 8 hours
 - Previous IEBTRUCK method was measured in days not hours
 - AKA: CTAM Chevy Truck Access Method
- Elimination of all physical tape
- Systems need to be scalable
 - Capacity
 - Performance
 - Replication

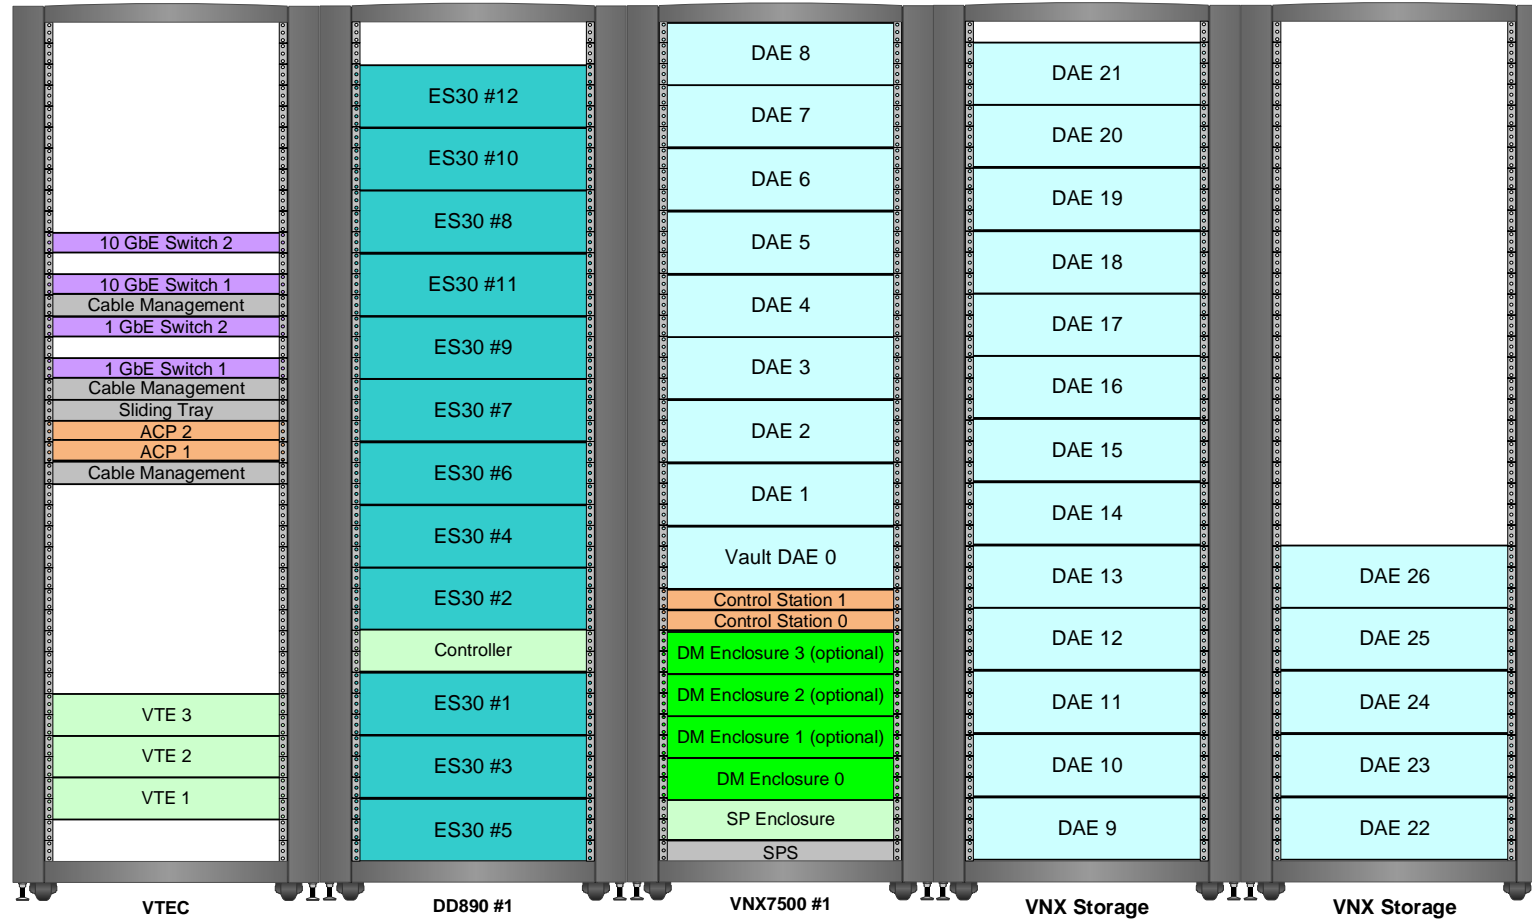
DLM6000 selected to be new z/OS Virtual tape system

- Initial DLM6000 had both VNX and Data Domain Storage
 - VNX used for WORM data initially and z/VM data
 - VNX also used for MOBIUS, SARS and CBROAM data
 - Data Domain used for everything else
- All tape data replicated to DR site
 - DLM starts replicating data at open instead of waiting for close improving RPOs
- System needed to be scalable
 - Initial implementation allowed for doubling of performance
 - Capacity on each the VNX and Data Domain could be more than doubled
 - Network bandwidth reduced substantially and available replication ports less than 50% used

DLM6000 selected to be new z/OS Virtual tape system

- Deduplication was key to making things work
- Without deduplication bandwidth costs would eliminate any savings
- Performance needed to improve, waiting on physical tapes was impacting batch processing
 - The performance improvement was noticeable for all applications, especially in batch processing, were all jobs using the DLM were running measurably faster
 - Batch window reduced by 3 hours daily on average

Step 1: Eliminate two of the TS7740 systems



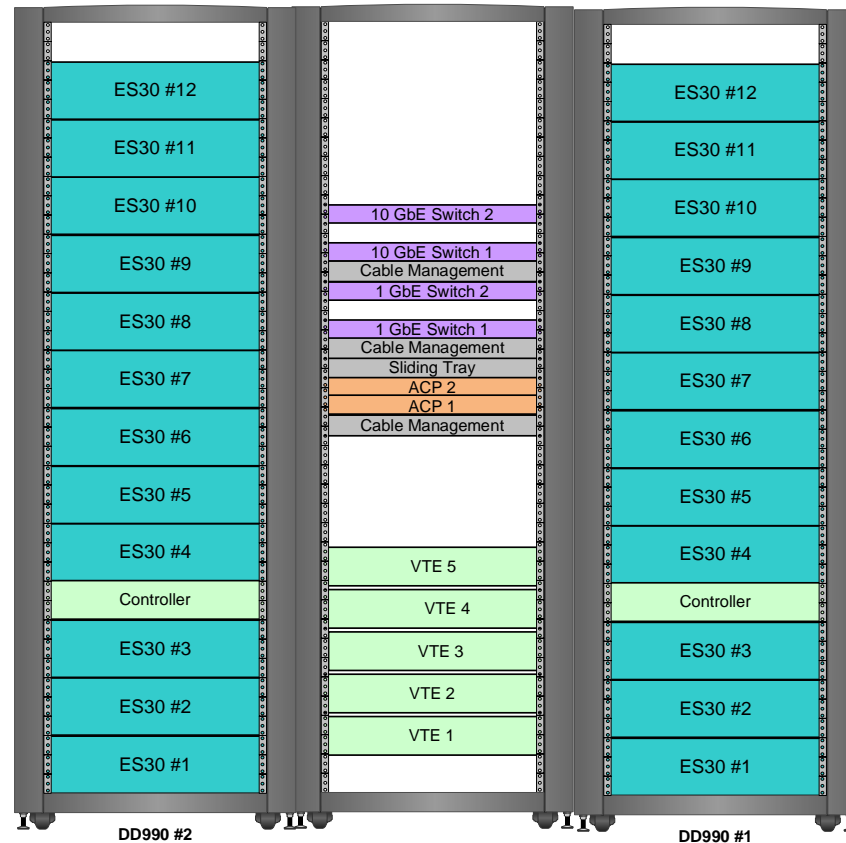
DLm6000

20 SQ Feet of Floor Space:

VNX for WORM and Specific applications Data

Domain for everything else

Step 2: Eliminate remaining two TS7740 systems to DLm8100 DD only system



DLm8100
12 SQ feet of Floor Space
All Data Domain

DLM8100 installed to replace ageing DLM6000

A change to all Data Domain occurs as part of this implementation

Why eliminate VNX?

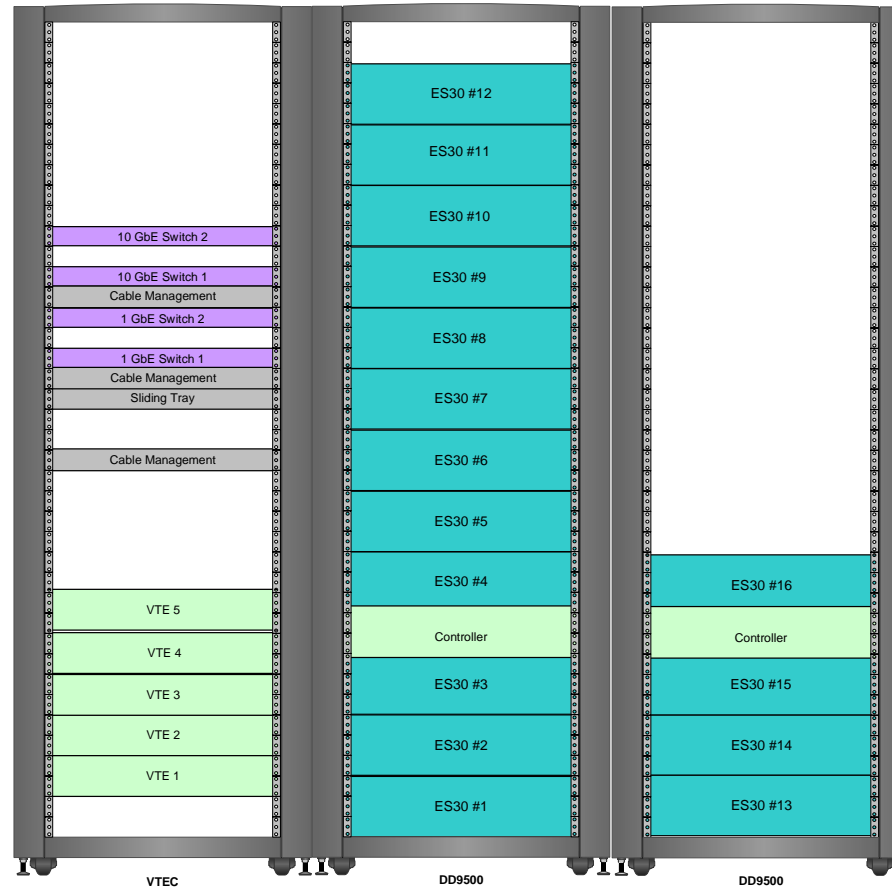
- Replication bandwidth was the key, VNX required substantially more bandwidth to move less data
- Data Domain had proven its capabilities with each generation getting faster and larger
- The worry about how well Data Domain could rehydrate data had long since passed
 - No complaints about performance from any application groups moved onto Data Domain
- WORM was now available on Data Domain
- Single type of backend storage made disaster recovery testing simpler

Dm8100 installed to replace ageing Dm6000

Migration considerations

- All data needed to be copied from the VNX to the Data Domains
- Initial configuration had the VNX in its own SMS Storage Groups and Data Domains in other Storage groups
 - No new mount scratch or mount privates would be permitted on VNX
 - All new tape mount directeds to the 3 Data Domain systems, each in their own Storage Group (concatenated)
- Application based utilities and CA-COPYCAT used to move the data

Step 3: Replace Original DLm6000 with DLM8100



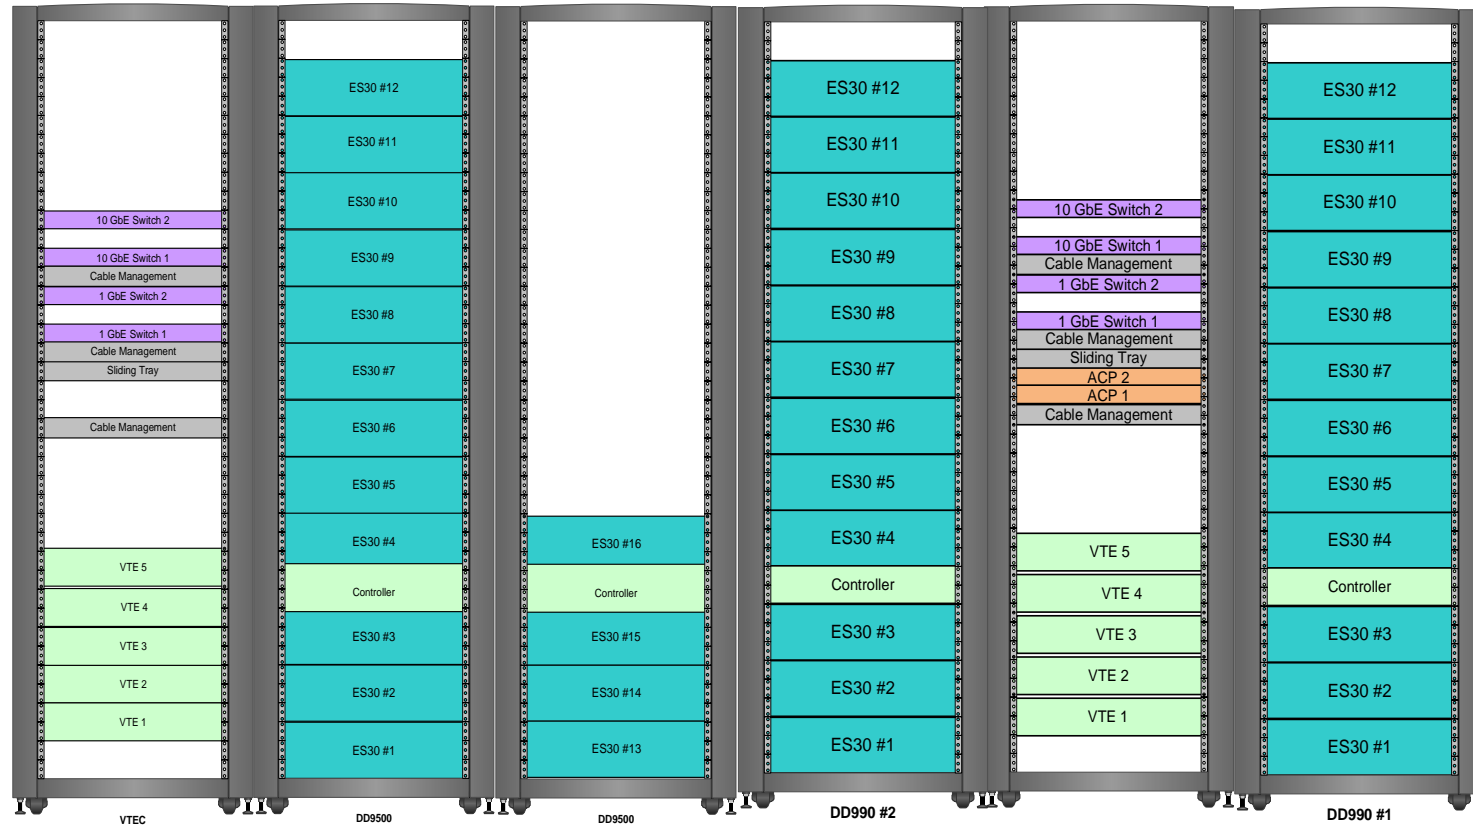
**DLM8100 12 SQ feet of Floor
Space
All Data Domain**

Before DLm:



Over 2400 Square Feet of floor space Freed!

(including rack space outside of libraries)



24 SQ feet of floor space
All Data Domain

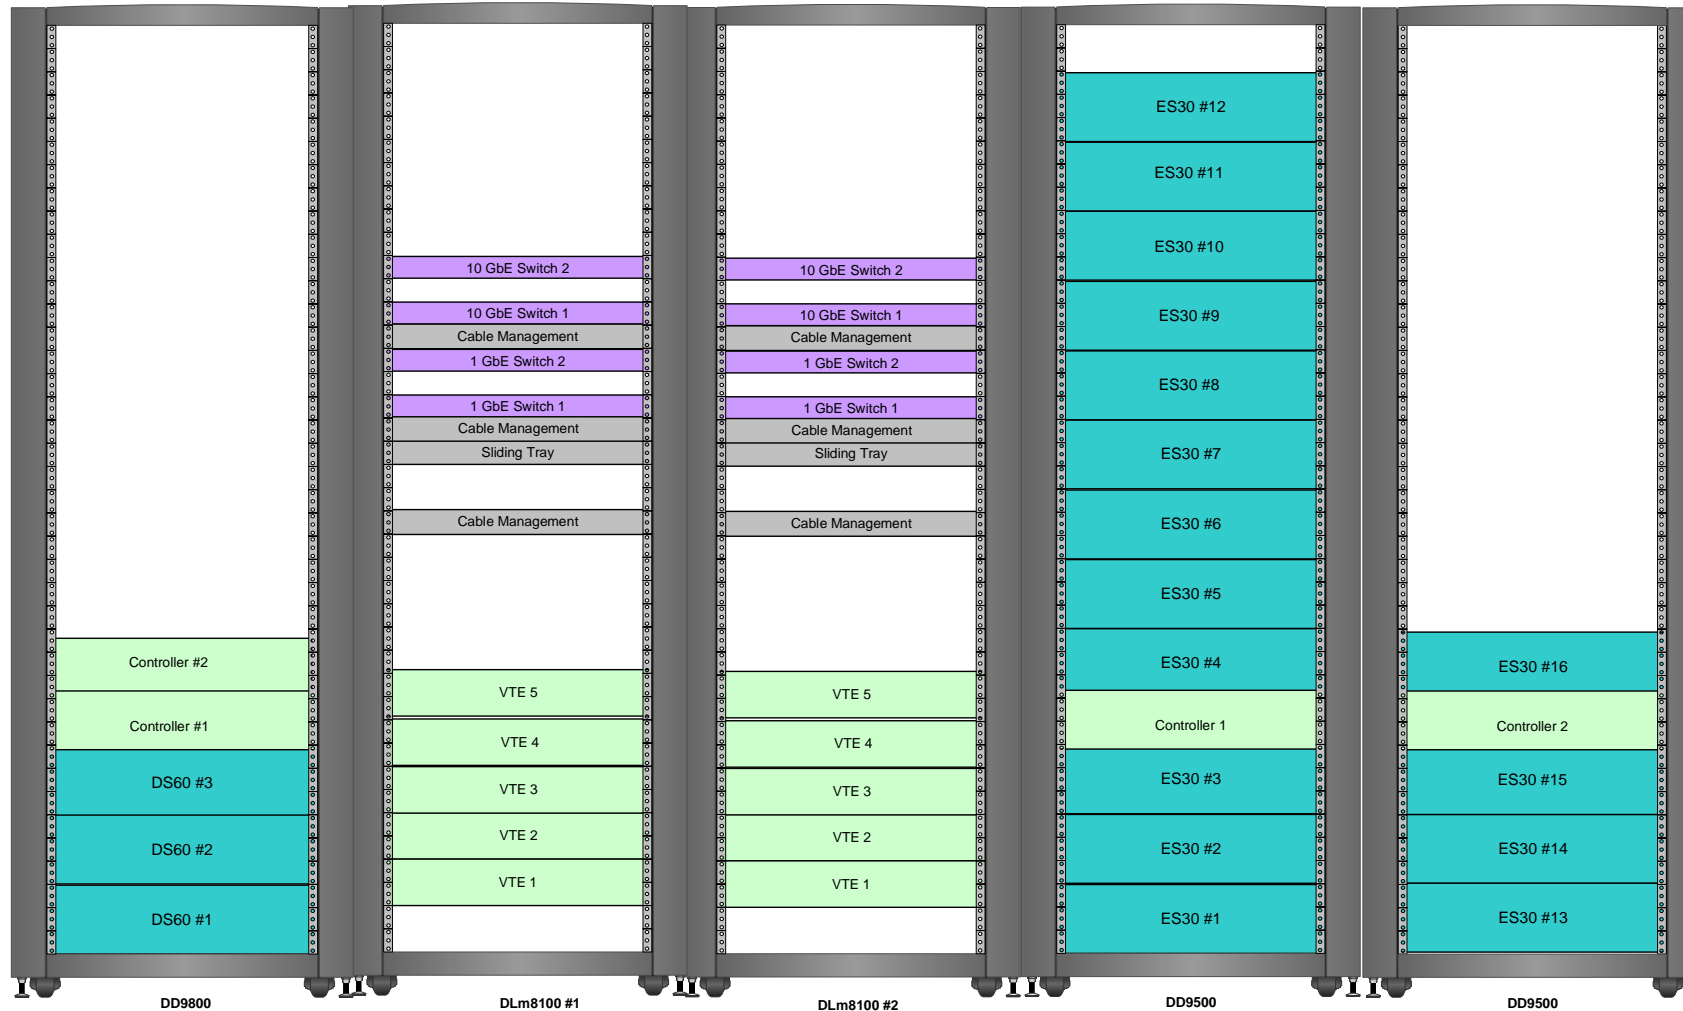
Where is this customer now?

- All Data is replicated
- Using all Data Domain instead of compression only systems reduces bandwidth required by more than 75% for the same RPO/RTO
 - On a Daily basis over 90% reduction
- Performance has been outstanding for all workloads
 - Data Domain systems have no issues handling peak workloads
 - On average more than 50TB of data is handled by the DLM's with Data Domain
 - 75% write 25% Read on average
 - With each generation the performance has improved substantially in the Data Domain family
 - The DD9xxx family is capable of handling workloads that require high performance both read and write without issue
- Having High Availability heads on Data Domain Systems makes consolidation practical
 - Allowing the customer to consider moving to a single system.
 - Prior to this the Data Domain was a single point of failure (for that reason we had more than one)

Where is this customer now?

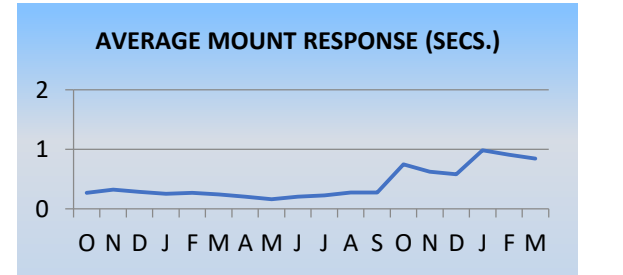
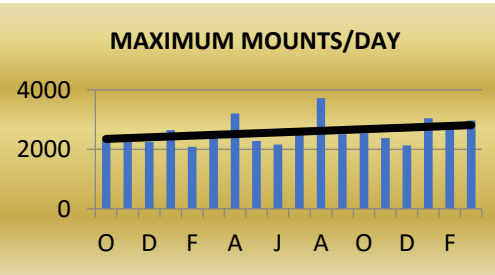
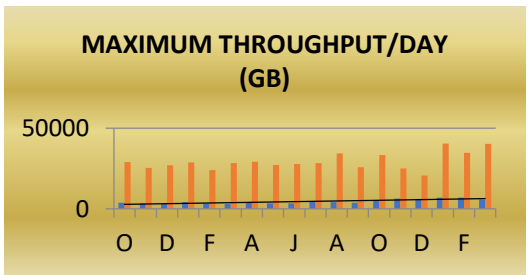
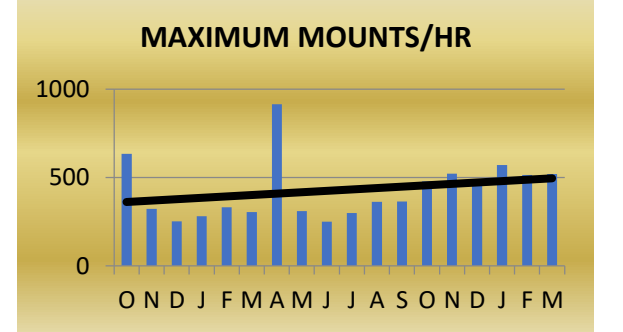
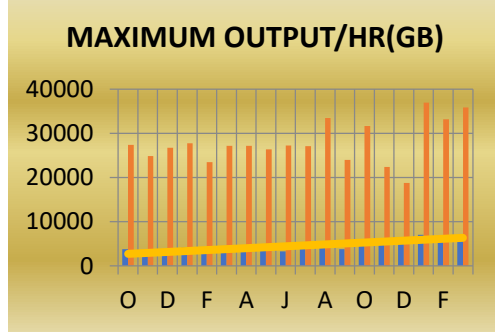
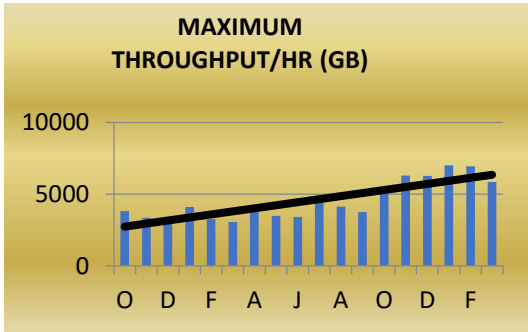
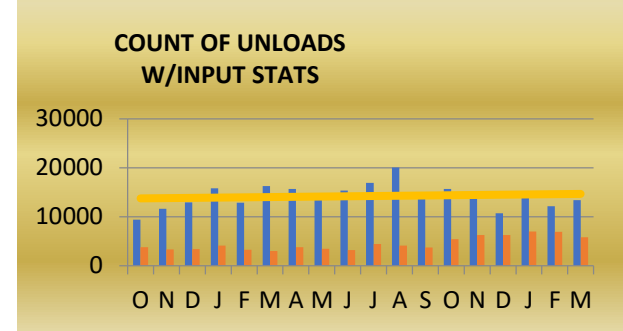
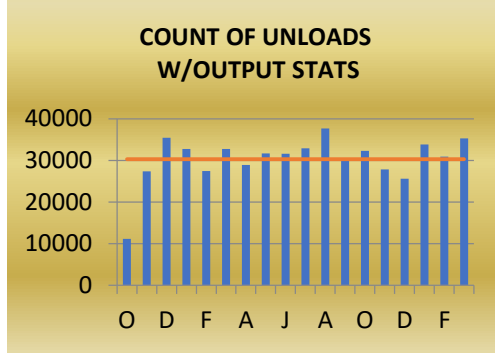
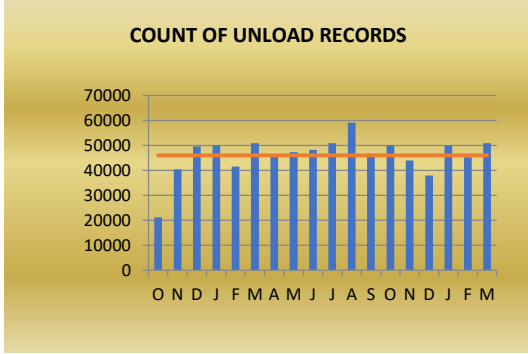
- Recent update replaced TWO DD990s with Single DD9800 further improving deduplication rates
- The ability to have z/VM and z/OS on the same system made things easy
- The ability to have WORM data on Data Domain eliminated the need for VNX storage
- System is highly scalable and supports all changing requirements
- Deduplication is the glue that makes everything work!

Existing Configuration as of 8/11/18



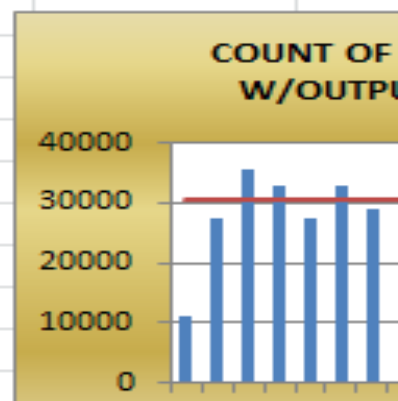
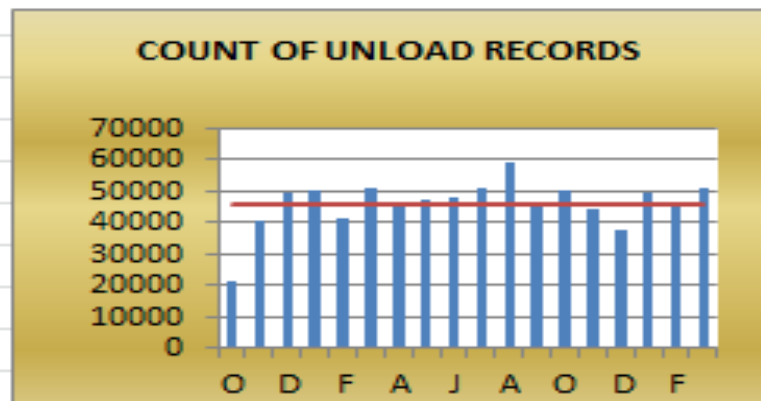
20 SQ feet of floor space
All Data Domain

System Overview “Summary Dashboard”



Sample Performance of single DLm DD9500

	A	B	
3			
4		1	
5	DLm8198	O	
6	Counts and Maximums in GiB	Oct-16	Nov
7	COUNT OF UNLOAD RECORDS(mounts)	21132	
8	AVG UNLOAD Records	46057	
9	COUNT OF UNLOADS W/OUTPUT STATS	11160	
10	Avg Unloads w/Output	30331	
11	COUNT OF UNLOADS W/INPUT STATS	9434	
12	MAXIMUM OUTPUT/HR(GB)	3810.44	
13	MAXIMUM OUTPUT/DAY (GB)	27388.86	249
14	MAXIMUM INPUT/HR (GB)	1208.32	7
15	MAXIMUM INPUT/DAY (GB)	3216.41	41
16	MAXIMUM THROUGHPUT/HR (GB)	3810.44	
17	MAXIMUM THROUGHPUT/DAY (GB)	28932.53	253
18	MAXIMUM MOUNTS/HR	634	
19	MAXIMUM MOUNTS/DAY	2323	
20	EFFECTIVE COMPRESSION RATIO	1 TO 1	1 T
21	AVERAGE MOUNT RESPONSE (SECS.)	0.262	



Interesting observations over time

Performance improves as more data is loaded into the system

- › With a deduplication system, the more duplication the less data that needs to be written, it points to the block instead of writing the block
- › Higher deduplication rates have no impact on system READ or WRITE performance

Primary Focus is always on the Overall De-Duplication rate

- › While important it is not everything (based on 04/11/18 data)
 - Over all deduplication rate 9.1:1
 - Translation 3.75PB stored using 412TB of storage
 - This number to stabilizes over time
(past 2 years the rate has been between 8.7:1 and 9.5:1)
 - More than 4x more effective than compression only (@2.5 compression from SMF21 records)
 - Last 24 hours
 - Deduplication rate of 19.4:1 translates into 94.9% reduction in bandwidth. Only new blocks are replicated

Availability

- › 99.996% of all systems measured from 2011 to date

Summary

DLM systems with Data Domain allowed this customer to exceed the RPO and RTO requirements

- › Normally within 15 minutes

Without Data Domain the network bandwidth would have increased by over 3x to meet the RPO /RTO requirements

DLM + Data Domain performance is unaffected by changing workloads or applications

- › Mobius, SARS and CBROAM recall performance outstanding
- › ADATABASE and DB2 are the primary databases
- › Not a DFHSM shop

Using the same data encryption methods for distributed systems and mainframe systems

keeps the auditors happy

- › D@RE for DASD and Disk (HDS and Dell EMC) plus network encryption
- › Using Data Domain Encryption for Open Systems and Mainframe Tape keeps thing simple (no Network Encryption required)

We want your feedback!

- Please submit your feedback online at
 - <http://conferences.gse.org.uk/2018/feedback/nn>
- Paper feedback forms are also available from the Chair person
- This session is BC

