

Migration of Large Tape Archives

Henry Newman

CEO/CTO

Instrumental, Inc.

hsn@instrumental.com

“The first 90% of a project takes 90% of the time. The last 10% takes the other 90% of the time.”

***Richard Templar,
The Rules of Management***

- Backup
 - Backups have two distinct purposes
 - The primary purpose is to recover data as a reaction to data loss, be it by data deletion or corrupted data
 - The secondary purpose of backups is to recover data from a historical period of time within the constraints of a user-defined data retention policy, typically configured within a backup application for how long copies of data are required
- Digital Archive
 - Storage of a collection of digital data
 - Archive data is data which does not change much or at all and must be kept long periods of time
 - Most often today digital archive data is managed by HSM systems

Topics to be Covered

- Need for Migration Planning and budgeting
- Issues
- Vendor Developments
 - Data Integrity
 - Archive Migration
- Planning Issues
- Questions and Comments

Why are archives growing fast



HIGH PERFORMANCE INNOVATION

- Businesses and organizations collecting more and more data
 - Monitoring is needed to ensure the integrity of the data
 - Analysis is needed to ensure SLAs are met
- Greater accountability for data storage and management
 - Regulatory requirements
 - Difficulties finding the data that is needed
- Support of duplicate environments
 - Multiple copies for data
 - business continuity
 - Latency
 - DR

Why Migration Planning


- Rapid archive growth
 - Archives are growing faster than tape density
- Meeting operational requirements
 - Ensure the integrity of the data being migrated
 - Ensure service level agreements (SLAs) continue to be met
- Architectural considerations and cost
 - Hardware: servers, disks, tapes
 - Software: HSM, file system
 - I/O connectivity: PCIe buses, switches

Rapid Archive Growth (cont.)



HIGH PERFORMANCE INNOVATION

Projected Data Growth (ZB) ¹



¹ The Digital Pileup;
http://www.nytimes.com/2011/03/13/opinion/13podolny.html?_r=2

- Channel Error Rates
 - Fibre channel originally developed at 25 MB/sec and it is now 32 times faster at 800 MB/sec and is half duplex (now 64x)
 - Channel error rate was 1 in $10E12$ bits corrected to 1 in $10E21$
 - No one expected SANs with many 1000s of switch ports
 - IDE channel originally was .625 MB/sec and it is now 480 times faster at 300 MB/sec
 - The channel error rate was again 1 in $10E12$ bits corrected to 1 in $10E17$
 - Channels are corrected with Reed Solomon encoding
 - Uncorrected error rate of 1 in $10E17$ to 1 in $10E19$ bits

Hard Error Rates



HIGH PERFORMANCE INNOVATION

Device	Hard error rate in Bits	PB Equivalent	MB/sec rate per 100 devices at full rate in seconds	Errors per year for 100 devices at full rate	Amount of Data Lost in bytes
Disk Consumer SATA	10E14	0.11	7800.00	2,063	1056479
Disk Enterprise SATA	10E15	1.11	8400.00	222	113775
Disk FC/SAS	10E16	11.10	12500.00	33	16931
Tape LTO	10E17	111.02	24000.00	6	6
Tape Enterprise	10E19	11102.23	36000.00	0	0

- Error rate for disk is number of bits read before the failure of a sector
 - 512 bytes today 4096 soon
- Error rate for tape is number of bits read before the failure a bit
 - Very different than disk



*The secret of getting ahead is getting started.
The secret of getting started is breaking
your complex overwhelming tasks into
small manageable task, and then starting
on the first one.*

Mark Twain

Meeting Requirements – Silent Corruption

UDBE R	Sustain Transfer Rate Per Second for a Year						
	0.5 GB/sec	1 GB/sec	10 GB/sec	100 GB/sec	1 TB/sec	10 TB/sec	100 TB/sec
10E28	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10E27	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10E26	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10E25	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10E24	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10E23	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10E22	0.0	0.0	0.0	0.0	0.0	0.0	0.3
10E21	0.0	0.0	0.0	0.0	0.0	0.3	2.7
10E20	0.0	0.0	0.0	0.0	0.3	2.7	27.1
10E19	0.0	0.0	0.0	0.3	2.7	27.1	270.9
10E18	0.0	0.0	0.3	2.7	27.1	270.9	2708.9
10E17	0.1	0.3	2.7	27.1	270.9	2708.9	27089.2
10E16	1.4	2.7	27.1	270.9	2708.9	27089.2	270892.2
10E15	13.5	27.1	270.9	2708.9	27089.2	270892.2	2708921.8

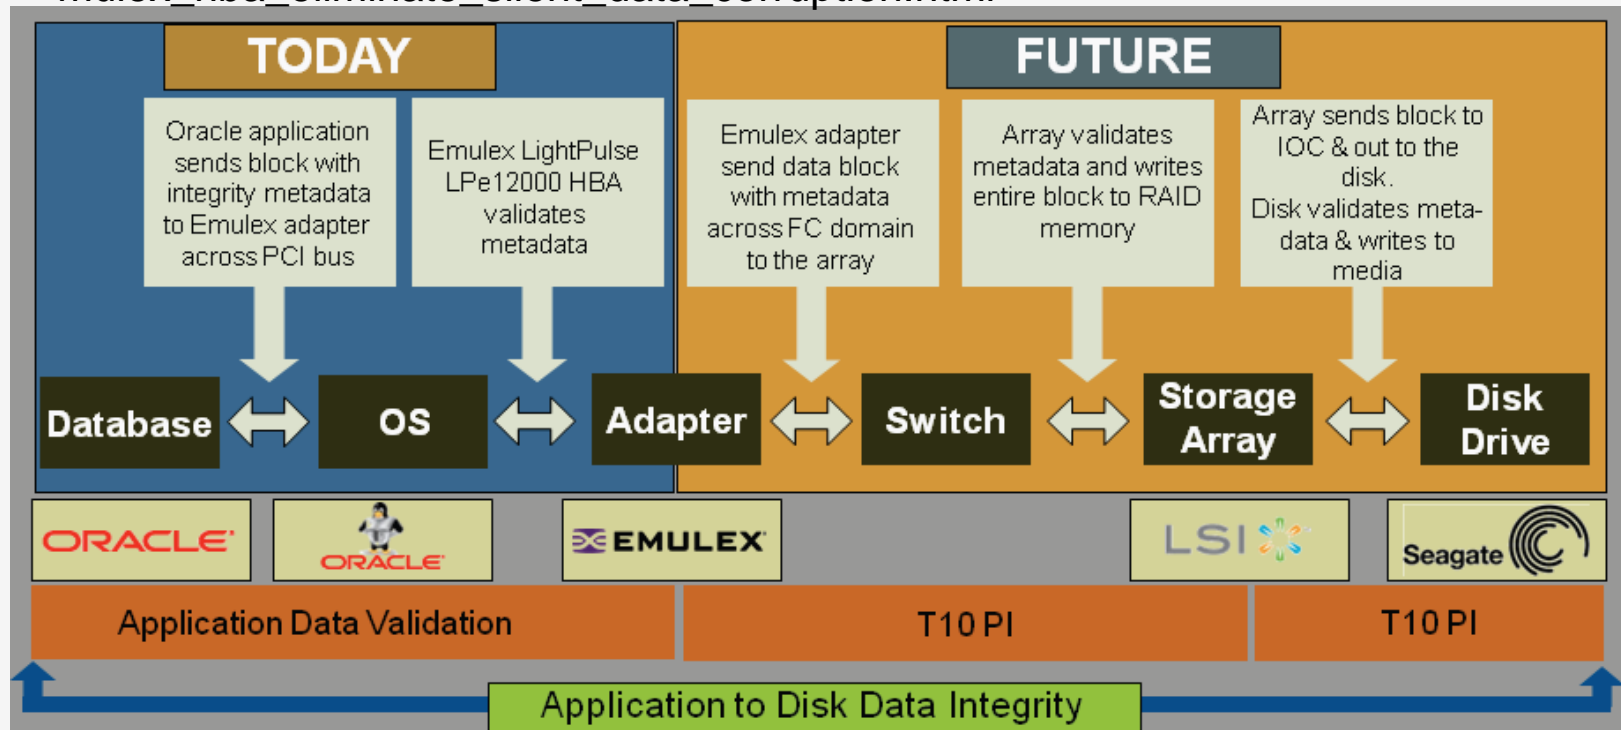
- These annual failure rates are for a perfect world where the channels are operating at the specified rate of 1×10^{-12} and corrected to $1 \times 10^{-17}/10^{-19}$.
- What happens when the world is not perfect?
 - Trust me the world is not perfect.

- Archives are growing faster than tape density or performance
 - Need more and more tapes, requiring more and larger tape libraries
 - Data reliability not keeping up with tape and archive capacity growth
 - Little change in disk and tape reliability even though capacities have grown more than an order of magnitude
 - Channel errors are a major concern as noted previously
 - Need end-to-end checksum capability to ensure data integrity
 - T10 PI disk and tape checksums standards are being implemented by some vendors to address this issue
 - LTO density is about 2x but performance increase is about 20%
 - This is not going to change soon if at all

- Depending on the archive software there are potentially 4 types of hardware that need to be considered when migrating from old tape drives to new tape drives:
 - Tape drives both old ones and new ones
 - Switch ports for the drives
 - RAID storage
 - Servers
- The number of both old tape drives needed depends on a number of two factors:
 - How fast you want to complete the migration
 - What the user load current is and how well you want and need to satisfy it during the migration
- Do not forget data integrity!
 - Data validation takes resources

Vendor Developments (cont.)

- Oracle Data Integrity Extensions and their interface to other vendor products
 - http://blogs.oracle.com/linux/2010/12/oracle_unbreakable_enterprise_kernel_and_emulex_hba_eliminate_silent_data_corruption.html



- Tape could be included in this picture given new ANSI tape checksums
 - Currently lacking support from software vendors

***“We will either find a way, or
make one.”***

Hannibal

- **Two types of migration**
 - **Migration to new tape drives and storage**
 - **Most HSM packages have internal tools**
 - **Migration to new HSM**
 - **Quick overview of 2 HSMs with conversion utilities**
 - **IBM HPSS**
 - » **Provides migration as a service**
 - » **Has migrated customers with Oracle SAM-FS QFS and LEGATO DiskXtender UNIX/Linux (DXUL) HSMs**
 - » **Formats metadata of legacy system into HPSS format and imports it into HPSS**
 - » **Legacy tape format is added as a read-only format**
 - **StorNext HSM**
 - » **Provides Archive Conversion Utility (ACU) to migration Oracle SAM-FS QFS**
 - » **Legacy file system structure is imported into StorNext File System allowing access to legacy files**
 - » **New files are stored in StorNext format and legacy files can be gradually moved, but this requires moving files to disk**

- Understand and measure your user load and SLAs
 - User load in many sites changes from day time to night time
 - Time to migration a single tape might be as much as 2 hours
 - If your user load changes at 7AM you might need to stop high speed migration at 5AM
- More hardware will be needed
 - Servers
 - Fibre channel switches
 - Disk storage and bandwidth
 - And of course Tape Drives

- More CPUs might be needed to validate checksums
- More PCIe buses or faster buses given the increased I/O bandwidth
- More memory bandwidth given the increased checksum load and increased I/O load
 - Remember memory bandwidth per core is going down not up

- More switch ports will be needed given increased number of tapes in the increase needed in I/O bandwidth
 - Balancing low and high density switch blades needs to be considered
- These ports might need to be updated and upgraded switch
 - E.g. 4 Gbit to 8 Gbit

- For most HSMs all the data must be read in from tape and written back to disk and read from disk and written to tape
- This requires both:
 - I/O bandwidth
 - Storage space
- Some HSM have options for tape to tape but checksums are no validated
 - Migration tape to tape is not a good idea as data is not validated

- Clearly more tape drives will be needed:
 - New drives for data to be migrated to
 - More old drives to meet requirement as drives will in use for migration

“Knowledge is Power”

Sir Francis Bacon

- You must understand the SLAs and user requirements to set expectations
 - This is one of the 1st steps in migration planning
- You must plan on purchasing more than just tapes and tape drives to meet these requirements
 - New requirements might dictate new hardware
- You must carefully plan on how much data you can migrate at a time based on bandwidth being used and technology in the path
 - Hardware, HSM software, operational scenarios, and user demands
- It is not easy

***For more detailed
information see my May
2011 column at:***

<http://www.enterprisestorageforum.com/continuity/features/article.php/3933731/Tape-Migration-Ignore-It-at-Your-Peril.htm>

Questions and Comments

Thanks for listening!