Backups and Cold Storage: Scaling And Cost Containment With Tape Storage At The PB Scale And Beyond Without the Disk Cloud. Tape Vs. Disk In the Real World.

Matthew Parker

Chief Technologist

Dimensional DBA, LLC

An Interesting Story…

- How did we get here…

Business Realities

- You must do more with less.
- You must be cost competitive compared to the Cloud, compared to prebuilt solutions or compared to what you could build.
- You must be cost competitive compared to emerging or alternate technologies.
- You must allow the business to be nimble and execute on opportunities instead of being the “No” in the process.

Choices and Detours

- Manage two teams.
- Fix the world.
- Only spend what has been budgeted ($0) if you really need to.
- Finance 10 year plan.
- Why not the cloud?

Scaling Vectors

- Servers.
- Network.
- Tape Drives.
- Tape Library.
- Backup software/Backup Policies.
- Human Processes.

Server (Scaling)

- Buses/Channels
- 2Gb Fiber Cards
- Single 1Gb NIC
- OS
- Memory
- Monitoring.

Network (Scaling)

- 1Gb NIC
- Second NIC – NIC Bonding.
- Oversubscription

Tape Drives (Scaling)

- Add more drives.
- Upgrade drives.

Tape Library Scaling (Scaling)

- Density of slots.
- Number of Drive Slots.
- Speed of Robotics.
- Size of Cap.

Backup SW/Backup Policies (Scaling)

- Backup Policies
  - How many different backup policies?
  - How many different tape pools?
  - Can you separate compressible and non-compressible data?

Human Processes (Scaling)

- Installation.
- Repair.
- Tape Loads/Unloads.
- Tape to/from offsite vendor.
- Tape to/from onsite storage location.

Story Continues…

- Robot arm breaks.
- Put single robot on support.
- Var bought part off of Ebay as HP had stopped manufacturing parts 2 years before.
- Further research on disk and tape.

Costing Vectors

- Tape Library/Tape Drives.
- Media.
- Servers.
- Networking.
- Data Center Floor Space.
- Power/Cooling/Clean Air.
- Backup SW.
- Support.
- Human Processes.

Tape Libraries/Tape Drives (Cost)

- HA Architecture.
- What single point of failure risks are acceptable.
- Examine Depreciation time frame for your Company.
- List price comparisons of vendors.

List Price Comparison of Vendors

	List Price Robotics	List Price LTO4 Drive	List Price 16 LTO4 Drives	List Price LTO5 Drive	List Price 16 LTO5 Drives	Year 1 Monthly Gold Support Robotics	Year 2+ Monthly Gold Support Robotics	Monthly Gold Support LTO4 Tape Drive	Monthly Gold Support LTO5 Tape Drive	Slots	Drives	Media Price Point	LTO4 Cost GB/Month	LTO5 Cost GB/Month
IBM TS3500	$774,964	$22,800	$79	$23,940	$83	$2,388	$3,885	$178	$187	10,293	16	$50	0.0066	0.0035
Oracle SL8500	$850,832	$22,500	$86	$26,000	$100	$8,508	$8,508	$225	$260	10,088	16	$50	0.0073	0.0038
SprectraLogic T950	$1,175,405	$16,785	$80	$17,960	$85	$0	$5,598	$118	$118	9,880	16	$80	0.0090	0.0048
SprectraLogic T-Finity	$1,819,189	$16,785	$92	$17,960	$98	$0	$6,609	$118	$118	10,000	16	$80	0.0104	0.0055
Quantum i6000	$958,617	$20,528	$117	$24,028	$137	$5,175	$5,175	$125	$125	4,800	16	$50	0.0108	0.0057
Qualstar	$331,399	$12,350	$51	$17,890	$75	$299	$746	$40	$40	3,126	16	$50	0.0080	0.0042

Media (Cost)

- Examine costing over time.
- Know when to get into the market.
- Simple Rule: When less than 2x of previous version.
- Complex Rule: When tape price for equivalent size is 1/3 the price of disk.
- Examine replacement time frame.

Servers (Cost)

- How many tape drives can a server support from a performance perspective?
- How much network bandwidth can the server support from a performance perspective?
- How many PCIe slots does the server support and of what type?
- How much power/cooling does the server require?

Networking (Cost)

- Cost in rack to the server.
- Cost rack to Data Center Core.

Data Center (Cost)

- Tiles Consumed by Tape Library/Backup servers.
- Standard Cooling space tiles.
- Power consumed by Tape Library/Tape drives/Backup Server Racks.
- Cooling requirements for Tape Library/Tape drives/Backup Server Racks.
- Air Filtering requirements for Tape Library/Tape drives.

Backup SW (Cost)

- Proprietary vs Open Source.
  - Which SW/Database Adapters Available?
  - Which Tape Libraries Supported?
  - Proprietary normally All.
  - Open Source not many but maybe enough.
- Open Source with support is normally less than 12% of proprietary price.
- With larger Tape Libraries Open Source doesn't work.
- Must custom write part or negotiate large scale license for proprietary vendor.

Support (Cost)

- Can you skip first year since it is under warranty?
- Can you self maintain the system?

Human Processes

- Tape Load: 30 minutes to load 10K tapes.
- Tape Unload and offsite: Iron Mountain/Recall.
- Cross siting backups through network and not removing tapes from library.
- Repair. Choices between what is repaired and what still requires support.

Design The System.

- Tape Drives Direct Attach
- LTO-3: 4 tape drives utilizing two 2 ports.
- 8 servers, 4 per rack.
- Network: one 10Gb Ethernet 1-port card per server fiber attached into a Cisco 7600 in each rack with two 10Gb links to the Core Routers of the data center.
- SL8500, 32 LTO3 tape drives with the full 10K slots. (One in each data center.)

Compare to Disk.

- Unknown cloud structure to compare against.
- No costing models produced by finance to assist in the process.
- There was published pricing as to what it would sell for.
- I could physically go look at the racks and determine what their costs were.
- Read latest papers and books on disk. (Mechatronics)
- Create price target model.

Negotiations and Legal.

- Negotiate a price and compare to pricing model.
- Then negotiate some more.
- Build 10 year plan with new system costs and overlay future disk technology along with current estimated S3 costs. (10x cheaper).
- Haggle some more with finance.
- Systems continue to die, arms and tape drives.

Finally Success.

- Purchase tape library 11/2005.
- 2 more months to get installed.
- 1 month set back on some parts.
- Put into production.
- Time to rest…

The Story Continues: Alternate Technology

- Data Domain – Dedupe.
- FalconStore - VTL.
- COPAN Green Power.
- Rumblings from the Cloud.
- Direction to move to the cloud – too expensive – stay on tape.
- Repeat, repeat, repeat.

2010 New Challenge

- Competition between tape and an internally designed disk system for Cold Storage in the cloud.
- How to sell Cold Storage in the cloud for ½ cent per GB per month.

Lots of Research and Discussions.

- Meetings with IBM Tuscon RAS team.
- Meetings with Fuji.
- Analysis with IBM on increasing data center temperatures to meet server temperatures while maintaining sufficient reliability.
- Analysis of mini racks on top of tape libraries for media servers.

Example System

- Tape Drives Direct Attach
- LTO-4/LTO-5: 6 tape drives per server utilizing one 4-port card and one 2 port card.
- Network: one 10Gb Ethernet 1-port card connected to in rack switch with equal bandwidth to the Data Center Core.
- Use 18u racks with blade chassis on top of robot.

Example System

- IBM TS3500 with minimum 10K slots.
- 51 Tape drives at peak during media conversion.
- Examination of relative merits of MTBF and Estimated Life Time.
- Extra tape drives for media verification.

The Future Is Bright

- The story goes on….