

Large Tape Users Group



Sun StorageTek Tape Hardware Migration Experiences

Jason Hick

9940B & T10KA to T10KB


Serial 9310s to TCP/IP SL8500

Large Tape Users Group



Replacing Old Hardware

- **8ea 9310s containing 40,000+ cartridges (91% occupied)**
 - 16,130 x 9940B tape cartridges
 - 17,700 x 9840A tape cartridges
 - 6,859 x T10KA tape cartridges
 - 34 x 9940B tape drives (3-way & 1-way)
 - 18 x T10KA tape drives (3-way & 1-way)
- **Some reasons Sun provided us:**
 - 9940B availability of replacement parts/drives (several weeks on 2-3 occasions)
 - T10KA sensitive to unavailability of drives and couldn't repair fast enough
 - 9310s going end-of-life
- **Some temporary solutions for above:**
 - Buy local spare parts cache
 - Support for 9310s from 3rd party

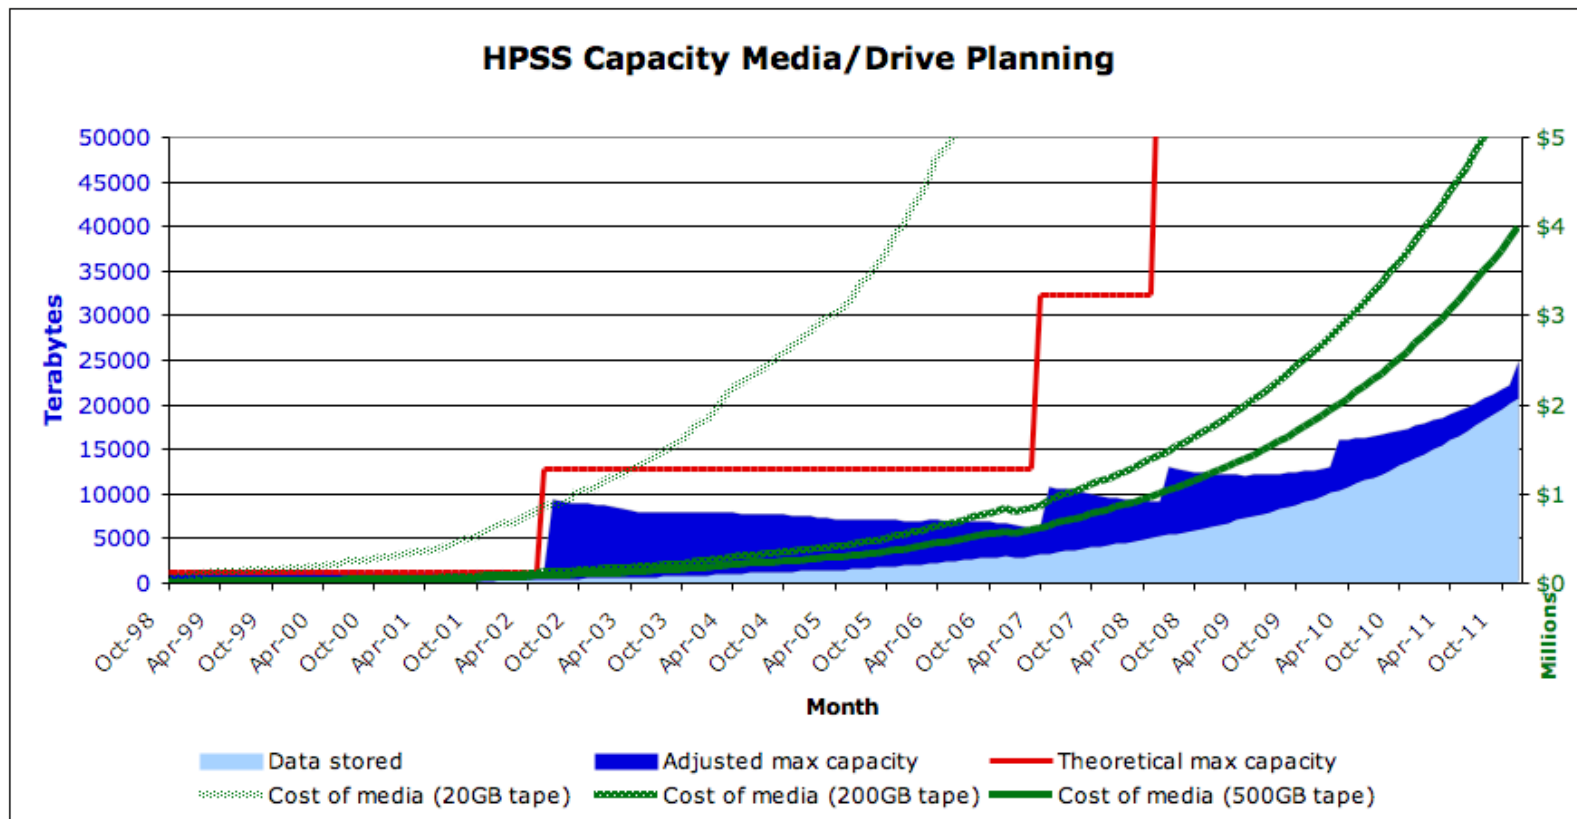


Large Tape Users Group



Replacing Old Hardware (cont'd)

- **Our own reasons for needing new hardware:**
 - 50% new data each year, more and more 9940 tapes (media budget)
 - 9840A and 9940B use 73% of our library slots (running out of free slots)



Large Tape Users Group



Purchasing the T10KBs

- Selecting the right number of tape drives for HPSS is determined by:
 1. Performance requirements - width of tape stripe
 2. Number of migration streams required to keep up with data ingest
 3. If direct to tape, average number of concurrent drives in use
 4. Frequency of tape read requests
 5. Repacks (deletions and troubled media)
 6. Target date to have old data off previous media (number of streams reading old and writing new)
- Didn't have SL8500s, got quote for T10KB in 9310s
- T10KB not supported in 9310s
- Our local Sun sales team provided us with a **great deal** on SL8500s with the tape drives
- 3ea SL8500s containing
 - 40 x T10KB drives (2-way & 1-way)

Large Tape Users Group



T10KB Deployment

- Ensure you are at least at microcode 1.40.208
- Sun provided ODM for AIX hosts had messed up device names (case #70765646, bugtrac #6820161):
 - Previous ODM had T10000A
 - Expected T10000A and T10000B in new ODM
 - What it had:
 - ◆ T10000X - both A & B drives map to this which makes it impossible to differentiate A & B drives on the host.
 - ◆ 9840C, but no 9840D entry... our 9840D drives showed up as 9840C's, so the device response was wrong or a D & C drive have the same type string?
- Testing showed 1TB (base 10) of uncompressible data on tape
- No failures with read/write & checksum across all our drives
- Uncompressible data showed 100-120 MB/s
- Under HPSS 6.2 application (migration) to 2-way stripe, see 180-200MB/s, formerly saw 200-250MB/s with 3-way T10KA stripe
- 1 drive failed in first week of usage (HBD "logic" card bad), replaced same day from regional parts cache

Large Tape Users Group



Purchasing the SL8500s

- **Redundancy eliminating single point-of-failure is key in our environment**
 - 2N power (1 power supply per robot, 1 power supply per 4 drives) & SL8500 #1 needs 2ea Auxilliary racks
 - Dual TCPIP (SL8500 #1 still single point-of-failure), want new multi-TCPIP option
 - Single HBZ (dual wasn't available) is single point-of-failure
 - Webcams almost useless
 - Do NOT ask for a pen and stylus
 - Handbots do fail (8ea even if don't need exch/hr)
 - Do fully build out (5 expansion)
 - Saving money ideas
 - ◆ License for 5,000 slots
 - ◆ Consider GOLD 24x7 support 4hr response, significant cost reduction and don't get 2hr response anyways
 - ◆ Trade-in credit on 9310s (we couldn't do)
 - Fire suppression strongly consider Sun certified vendor
 - ◆ as opposed to 4 months of disputes and \$\$ for special re-certification



Large Tape Users Group



ACSL S Servers and LMUs

■ ACSLS

- Sun E250 (end-of-life), Solaris 8, ACSLS 7.1
 - ◆ CPU utilization was 70% with application that queries it frequently via `cmd_proc`
- Replaced with T2000 (Sparc), Solaris 10, ACSLS 7.3
 - ◆ Serial port + ethernet for SL8500s
 - ◆ CPU utilization is now 10% with same application
- Migration went smoothly except we discovered that hostname of server is retained in ACSLS database
- **cmd_proc hangs, must restart ACSLS, no resolution Sun Case #70901966**

■ LMU

- Serially attached to 9310s but need new DB25 male (LMU) to DB9 female (T2000)
- Sun LMU serial cables not available, **generic cables work**
- Sun serial cards for T2000 not available
- TCPIP cards far and few between and expensive
- Jerry-rigged 3 connectors/cables to Sun serial port and it worked
- Found and installed TCPIP card and LMU upgrade to support (price reasonable finally), pricing through Sun Support

Large Tape Users Group



Deploying SL8500s

- Build them in 24 hrs and few problems (1 handbot, T10KB tape drive or two replaced early on)
- Private ethernet for libraries recommended, **LLNL reports this is fixed**
- SLC is slow to use, **LLNL reports reboot SLC is temporary workaround**
- Rail segregation requires more drives per library than 9310
- Timings on elevator and passthrough are acceptable
 - 1m30s passthru 3 libraries
 - 30s elevator top to bottom
- Drive installation/movement is a snap



Large Tape Users Group



Summary of lessons learned

- 9840Ds have been solid, no problems
- T10KBs weren't as resilient. Had 2-4 issues mandating drive replacement w/in first few weeks of use.
- Support for both is great so far. No parts availability issues.
- SL8500 handbots break, you'll want 8
- Serial-attached LMU to ACSLS is difficult with new ACSLS server hardware
- DO NOT get the pen & stylus for SL8500
- SL8500 fire suppression, use Sun vendor if possible (i.e. library recertifications are non-trivial anymore!)
- Need Sun to continue providing 2x tape capacity and performance improvements (secondary for us) every 2 yrs!
- Need new generations of tape to have media reuse. If not, this removes my #1 justification to sticking with Sun STK tape drives.

Large Tape Users Group



Future stuff

- **Would I do it again?**
 - Tape drives, yes, necessary.
 - Libraries, no thank you. Please make SL8500 SOL (EOSL) 2040.
- **Migration of data from 9940B & T10KA to T10KB**
 - When not altering the stripe width on tape (1 9940B cartridge to 1 T10KB cartridge) use HPSS tech insertion to redefine all 9940B resources to T10KB. All new data goes to T10KB. Then as many streams of repack as can handle (1 9940B, 1 T10KB) to move 9940B data. This takes hours to plan/complete.
 - Redefining stripe width from 3-way to 2-way required new HPSS class of service. All new data to 2-way. Need multiple threads of COS change. This will take years (2-3) to complete with 24x7 effort and 2-3 streams.
- **Migration of data from 9840A to 9840D**
 - Using HPSS repack, as many streams as we have drives, will take 1-2 years of 24x7 work to complete.
- **Weighing 3rd party support for 9310s and old drives past EOSL 12/2010 to purchasing drive trays and placing old hardware in SL8500s (move drives/media)**
 - Trade-in for 9310s reducing, pay for removal eventually
 - Drive-tray/installation costs
 - Drive to media locality in SL8500s
 - Slots in SL8500 near full

Large Tape Users Group



Questions?