


Lights-Out J90s with DLT Robot

*Randolph Gary Smith, The University of Texas at Austin,
Academic Computing & Instructional Technology Services,
High Performance Computing Facility*

Background

- Downsizing: CHPC -> HPCF (9/94)
 - \$3.5M -> \$1.0M Annual Budget, \$2.0M Transition Capital
 - 61 -> 12 FTEs (4 systems, 4 appl/US/viz, 4 operations)
 - Y-MP8 -> 2-J90s (4x4, 26 GB/s memory bandwidth), 9/95
 - DMF & Interactive Services (J916/5-1024) - UTCSR/NASA
 - Batch & Dedicated Services (J916/16-4096) - UT (HiPPI, no SSD)
 - 3.5 TB DMF-managed filesystem (66GB), 2.4M files (1M single-sector files), 3-5 of 10-20 GB/day, 33,000 CARTs
- SCSI Peripherals (VERITAS SEISMIC), UT spares
 - 24-Seagate Elite-9 (DD6S)-1/4 TB finally affordable
 - 6-STK 4890-primary MSP, 1/4x shelf space, 1/8x if 8mm, operator mounting, contingency plan-3-4280-A2s (loaned)
 - 4-Exabyte 2505-2ndary MSP experiment (low duty cycle)
 - 4-STK 9914-ancient data

Histogram of DMF-managed Files



Quantum DLT 4000 Experiment Drivers

- Continuing STK 4890 (Twin Peaks) delays
- Computer Technology Review articles (Aug-Dec 95)
 - very high capacity - 20 GB w/o compression, 3.5 TB~120 DLT cartridges, small volume

- moderate performance - 1.5 MB/s, 3x Exabyte, LOCATE
- capable HSM duty cycle - 500,000 head passes vs. 1,500 8mm due to low-tension linear recording vs. helical scan
- reliability excellent - tape-industry-unique multilevel error detection, parity checks, 64-bit CRCs, 16-bit EDCs
- low price & 2-year warranty

- Unsolicited communications regarding DLT (OIL)
- CRAY PAN - J90 UNICOS DLT support 11/95
- Initial plan - replace Exabytes with DLTs (2nd MSP)
- VERITAS swaps UT 2-9914 for 2-DLT 4000 1/96

DLT 4000 Results, STK 9710 Robot Drivers

- After correct configuration (EPRIV), DLTs perform:
 - 32 GB per cartridge, 20 GB incompressible (tested locally)
 - 1.5-2.4 MB/s transfer rate
 - 120-second access time for typical single file recalls
 - multiple recalls (up to 25 files!) per mount observed

- Begin considering 12 manual DLTs for both primary & secondary MSP based on observed performance and cost of media (\$25K for 8TB DLT vs. \$40K 4TB 3490E + 4TB DLT) and maintenance, small volume

- DLT Panther (Timberwolf) announced COMDEX 95
- VERITAS flexibility swapping STK 4890s/STK 9710
- Conversations w/ Pat Hughes positive re: 9710/DLT
- Continuing budget pressure results in administration suggesting lights-out operation with 9710 - requires both primary & secondary copies in robot

DLT Robot Plan Evolves, Library Conversion Underway

- Negotiations regarding STK 9710/DLT
 - VERITAS - 10->6 **robot-ready** DLTs (OMX firmware), we believed earlier in project that we could move manual DLT 4000s into robot-6 DLTs in 3+3 configuration

- STK - ACSLS software negotiations (\$5K), maintenance
- CRAY - various technical & business assists from Pat Hughes, single console for both J90s -> ACSLS server
- Sun - X1062A FWD SCSI adapter for SPARC5 (\$800)

■ Conversion of 33,000-tape 3480 library initiated

■ Various problems related to new IOS-V software - required for DLT drivers, but 4280-A2s fail, CRAY loans UT 2-TD3s for duration of library conversion

■ Exabytes mediate operations load - 1/3 of 3.5 TB library resident on 8mm (secondary MSP), so 1/3 of (1/2 of 33,000) CART mounts avoided

Library Conversion Proceeds, Robot Installed

■ Had to plan for relocation from JJPRC to Main UT Campus due to budget pressures (could not move 33,000 CARTs!)

■ 3.5 TB pumped thru 2 manual DLTs, running 24 hours/day, 7 days/week for 3 months beginning March 96, 1 TB secondary copies accomplished

- No DLT media failures
- One DLT drive failure (mechanical positioning - rare)

■ 3-week delay in conversion process due to world-wide shortage of CompacTape IV cartridges 5/96

■ STK 9710 installed July 16, 1996 - 1st CRAY site

Problems & Solutions

■ Forward Space Block, Drive Not Ready

- drive unloads unexpectedly during sequence: LOCATE, SPACE -1, SPACE +1, READ (or WRITE)
- above sequence determined observing IOS trace (x44440)
- only robot-ready DLT 4000s exhibit failure, moving robot-ready DLT 4000 outside robot -> OK (stknet not involved thus no serial I/O 9710), still no explanation from Quantum
- Solved with installation of UNICOS 9.0.2 & IOS-V 2.1, difference is state of Immediate Bit on LOCATE command

■ ERRNO 202 - CRAY solution being tested (special MSP workaround developed locally)

■ 1 drive failure, 1 stuck DLTR failure, expansion door

■ RO manual drives complain WRITE FILEMARKS 0

■ SCSI bus reset results in directory rebuild (cascade)

■ Loss of power results in directory rebuild, 2-hr wait

Performance

■ Typical file recall times (total, what user sees)

- Aug 95 - CART-recalled (<50MB) files ~ 420 sec
- Sep 96 - DLTR-recalled (<50MB) files ~ 196 sec

■ load & cache directory-45s, position-136s max, transfer->1.5MB/s, rewind-136s max, unload-17s (always rewrite directory, even RO!), 1700 ft

■ Wrote+Read 2,500,000 MB per month, 7-8/96

■ 2.5 TB intra-robot secondary copies, <371 GB/day

■ No avrproc processing for DLTR (rewind->delay)

■ Current Operations include 2 sets of 2 MSPs

- backups, precompressed, cyclic deletions - 25 GB/DLT
- normal production migrations - 36 GB/DLT

■ BMX 3480s ~ 300KB/s/drive (single controller), SCSI 3480s 2MB/s/drive, DLTRs 1.5-2.6 MB/s/drive

Futures

■ Add 4 DLT 4000s -> faster recalls (drive queuing)

■ DLT 7000s preferable

- 18 TB -> 31 TB (assuming current LZ compression ratio)
- 1.5 MB/s/drive -> 5 MB/s/drive (50MB/s aggregate native)
- No media change (CompacTape IV, 35GB native->56GB)
- 2+2+2+2+2 configuration given >1.6x, FWD SCSI devices

■ STK 9710/DLT support on GigaRing/T3E (sn6507)

■ NPACI - SDSC/UT Partnership -> 24-hour/day operator coverage -> effectively doubles robot capacity since secondary copy can be outside robot

■ Could get by with prime shift operator coverage IF user waits when secondary copy mount necessary

Summary

■ Neither DLT nor robotics planned but took advantage of opportunities as they arose, luck (all STK swap)

■ Many factors necessary for success

- prior experience with large tape library conversions
- VERITAS, STK, CRAY, UT communication & cooperation
 - Quantum microcode (OMX-60), 9710 microcode (1.0.0.3)
 - Sun SPARC5 (Solaris 2.4/ACSL 5.1.1), UNICOS 9.0.2/IOS-V 2.1
- accepting reasonable risks, working through inevitable problems of being first 9710 DLT robot production facility

- Excellent value-18 TB robot+media for <\$170K, \$1.2K/mo
- Users see faster recalls-primary driver along w/ cost

Reference Addresses

- <http://www.quantum.com/products/menus/tape.html>

- <http://www.quantum.com/support/faq/dltfaq.html>
- http://www.storitek.com/StorageTek/9710/9710_frames.html
- gary@hpcf.cc.utexas.edu (Gary Smith)
- markh@vsl.com (Mark Hopkins)