

HPSS Upgrade Project

Alf Wachsmann
SLAC National Accelerator Laboratory

alfw@slac.stanford.edu



Current Mass Storage

- SLAC has 6 StorageTek Powderhorn silos
 - End of service by end of 2010
- HPSS on unsupported version (v. 5.1 on Solaris)
- No disk cache

- BaBar already ejected ~4100 old tapes
 - Hopefully ~3500 more tapes to come
 - No need to migrate this data!

Old vs. New System

Old silos:

- 6 9310 Powderhorn silos with 6000 slots each
- 40 9940 drives (6 TSM, 2 AFS, 32 HPSS)
- 6 9840 drives
- 25000 9940 tapes in HPSS
- ~1000 9940 tapes in TSM
- ~1000 9840 tapes in AFS backup



New silos:

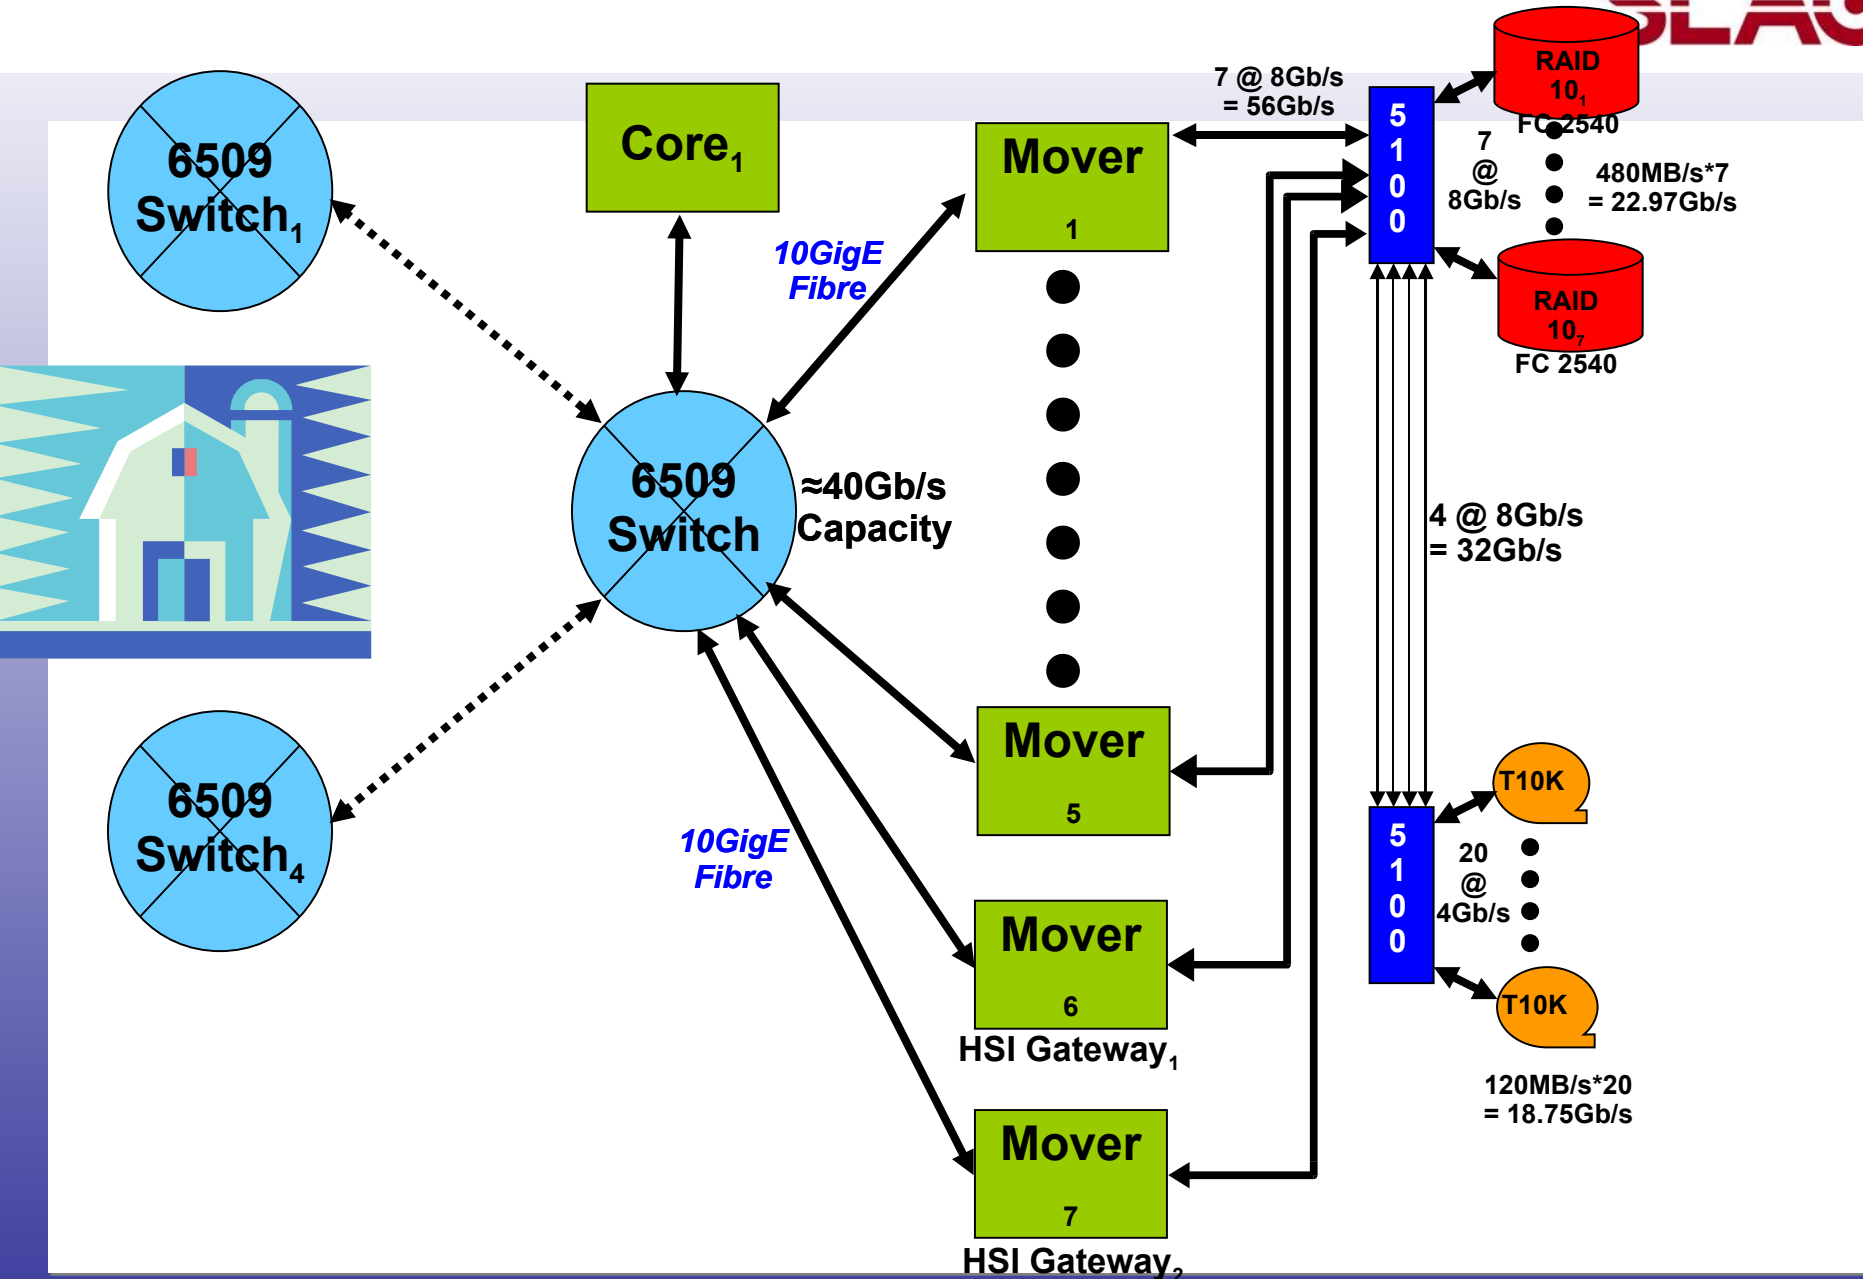
- 2 Sun/STK SL8500 with ~6000 slots each
- Pass-through port
- 28 T10K tape drives (more coming)
- 720 T10K tapes (many more coming)



- First SL8500 was unused for almost one year
- Not much progress with software migration either
- Formal project was started in November 2008:
 - Project Manager: Alf Wachsmann
 - Technical Leader: Lance Nakata
 - System Architect: Andy Hanushevsky
 - IBM expert: Michael Gleicher
 - Andrew May
- Weekly conference call

- Windows (NetBackup) and Unix (TSM) backup systems were moved to SL8500 in March – independent of HPSS migration
 - Unix AFS backup (butc) still to be moved
 - Probably move to new software (TiBS?)

HPSS Network Architecture with San3P and 10Gb Fibre



- Some servers+disks were ordered with first SL8500
- Was unused for almost one year
- New additional hardware is about to be ordered:
 - 2 Sun X4250 + Sun STK2540 as core servers
 - 14 Sun X4150 + Sun STK2540 as movers
 - 12 x 15K 300GB SAS drives (3.6 TB raw)
 - Total of ~50 TB raw disk cache
 - 16 T10K tape drives for HPSS
 - 2 Brocade 5100 FC switches with 4 x 8 Gb/s trunk lines
 - 1 Cisco 6509 with 10 Gb/s ports; 4 x 10 Gb/s uplinks

Logical “building block”:

- 1 mover machine
- 1 RAID array
- Up to 4 tape drives



Sun X4250




Brocade 5100



Sun X4150



Sun STK2540



Sun T10K drive

- Migrate from HPSS version 5.1 on Solaris machines to HPSS version 6.2 on Linux (64 bit RHEL4)
 - Turn off HPSS core machine
 - Copy HPSS' DB2 DB dump to an AIX machine
 - Consistency and health check
 - Convert it to other endian DB2
 - Move new DB2 dump to new Linux core server
 - Do all the configuration of new HPSS (this took the longest)
 - Start the core services
- Took 2 full days (May 4+5) to do
 - Open conference call during entire time
 - Micheal Gleicher did the main work with local help from Lance Nakata and Andrew May

- Keep the two HPSS instances:
 - A small instance for `mstore`
 - A large instance for BaBar etc.
- Use San3P
 - Allows tape drives and disk arrays to talk directly over FC
 - No server involved in data movement
- Use HSI for data migration (**no HPSS repack!**)
 - HSI and HTAR were written by Michael Gleicher
 - Free for HPSS sites
 - If wanted/needed: Michael provides commercial support
- Migrate all data with HPSS version 6.2
- Move to HPSS version 7.2 once migration is done

- Smaller HPSS instance was converted May 5
- `mstore` is now writing data into new instance
 - One Solaris mover with HPSS 6.2 connected to old Powderhorn for reading old data (data migration?!)
- Gaining first real experience and performance data with new hard- and software (no numbers yet)
- Larger instance depends on some electrical work
 - Endless amount of paper work
 - Hopefully approved in July
 - Work projected to be done August 2, 2009
- Switch to HPSS v. 6.2 end of August
- Start of BaBar data migration; will take 1(?) year