



LLNL Site Presentation

Kim Cupps
June 5, 2001



Overview

- **Current hardware configurations**
 - > OCF System
 - > SCF System
- **Recent performance accomplishments**
- **Current COS configurations**
- **Drivers to changes in COS strategy**
- **Proposed COS configuration**
- **Challenges**
- **Issues/concerns**
- **Future Plans**



Site configuration

- **The LC operates classified (SCF) and unclassified (OCF) HPSS systems.**
 - > Both centers run HPSS R4.1.1.4
 - > Both systems are legacy systems converted from UniTree.
 - > SCF and OCF systems are kept as similar as possible.
- **Applications are quite varied, but most focus on scientific modeling for defense, energy and global climate applications.**
- **Accelerated Strategic Computing Initiative (ASCI) requirements and funding drive our HPSS development efforts.**

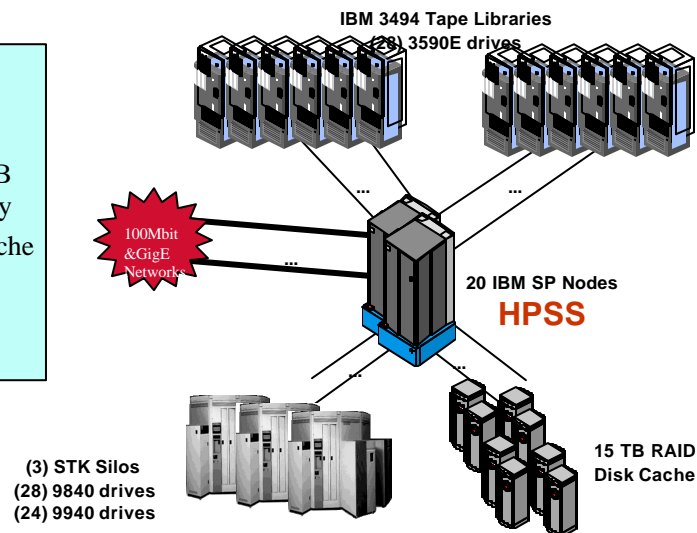
Kcc 6/5/01

3



OCF HPSS storage system

- ✱ 85 TB Stored
- ✱ 6 million files
- ✱ 50 GB-1.66 TB stored each day
- ✱ 15 TB disk cache
- ✱ 80 tape drives
- ✱ In production since 1/97




Kcc 6/5/01

4



SCF storage system

- ✧ 198 TB Stored
- ✧ 12 million files and directories
- ✧ 50 GB –4.82 TB stored per day
- ✧ 30 TB disk cache
- ✧ 104 tape drives



Kcc 6/5/01

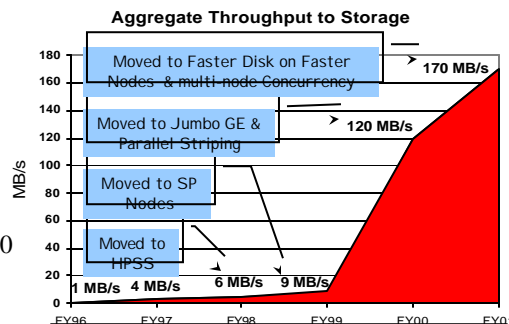
5



Archival storage performance

• Accomplishments

- > A 20x performance increase in 15 months
 - » Faster network
 - » Faster disks
 - » Attention to tuning
- > PSE Milepost demonstrated 170 MB/s aggregate throughput White-to-HPSS
- > Large single file transfer rates of up to 80 MB/s White-to-HPSS




At 170 MB/s, 2 TB of data moves to storage in less than 4 hours. A year and a half ago it took two and a half days to move that same data.

Kcc 6/5/01

6



SCF maximum pftp single file write performance by file size

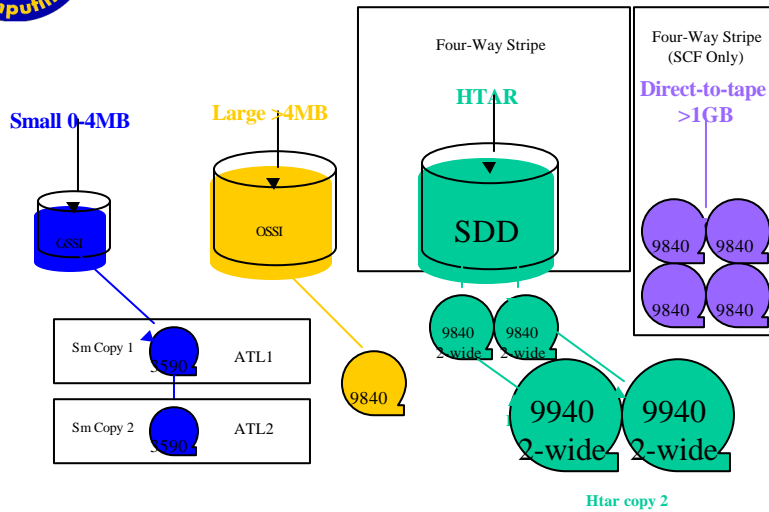


Kcc 6/5/01

7



Current COS configuration



Kcc 6/5/01

8




Drivers to changes in strategy

- **Want to minimize the risk of data loss without making two copies of everything**
 - > Last year we instituted dual copies for our small Class of Service (COS) - each file < 4MB
 - > Statistics showed we could do better by increasing the small COS to < 32MB
 - > Wanted to use 9940 drives and media appropriately
- **Increase performance for lots of small files**
 - > htar will be generally available in the next month
 - > htar files will be dual copied

Kcc 6/5/01

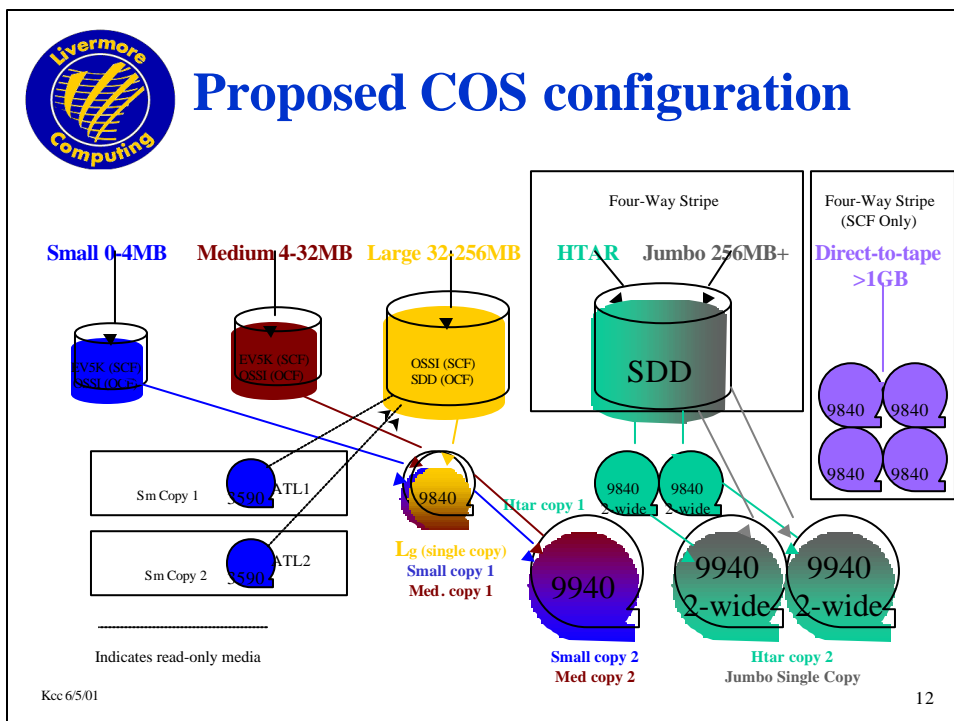
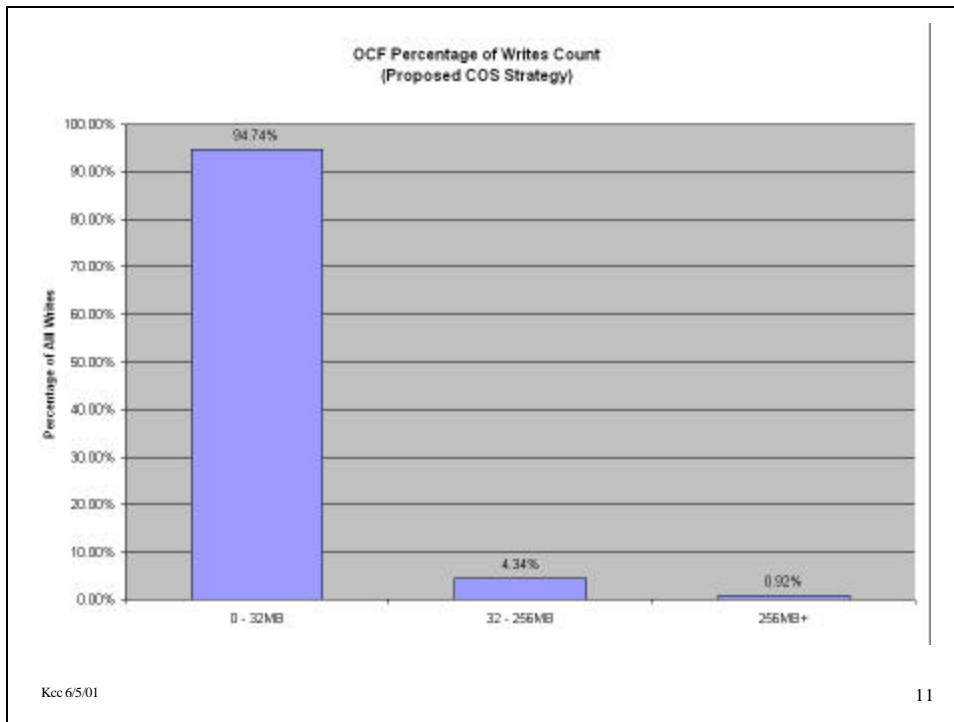
9

OCF Write Count vs. Write MBs
(Proposed COS Strategy)



Kcc 6/5/01

10





Operational challenges

- **Continuous integration of multiple software releases (HPSS, DCE, Encina, AIX)**
- **Integration and maintenance of an expanding hardware base**
- **Legacy data (continuous tape repack)**
- **Need to continue to provide > 2X throughput and capacity increases on a yearly basis in preparation for the 100 TFLOP platform**

Kcc 6/5/01

13



LLNL issues and concerns

- **PFTP data transfer at distance protocol improvement**
- **RAIT – direct-to-tape with performance and safety**
- **Object manageability**
 - > Quotas
 - > File lifetimes
 - > Undelete
- **File migration policy**

Kcc 6/5/01

14



Future Plans

- **Provide htar to all users**
- **Integrate more equipment**
- **Install HPSS 4.2**
- **Provide quota capability**
- **Provide delete tool capability**
- **Continue performance tuning**