



LLNL Site Presentation

Pam Hamilton

June 18, 2002





Overview

- **Site configuration**
- **Current hardware configurations**
 - > OCF System
 - > SCF System
- **Archival storage performance trends**
- **Achieving bandwidth goals**
- **Issues/concerns**
- **Future Plans**



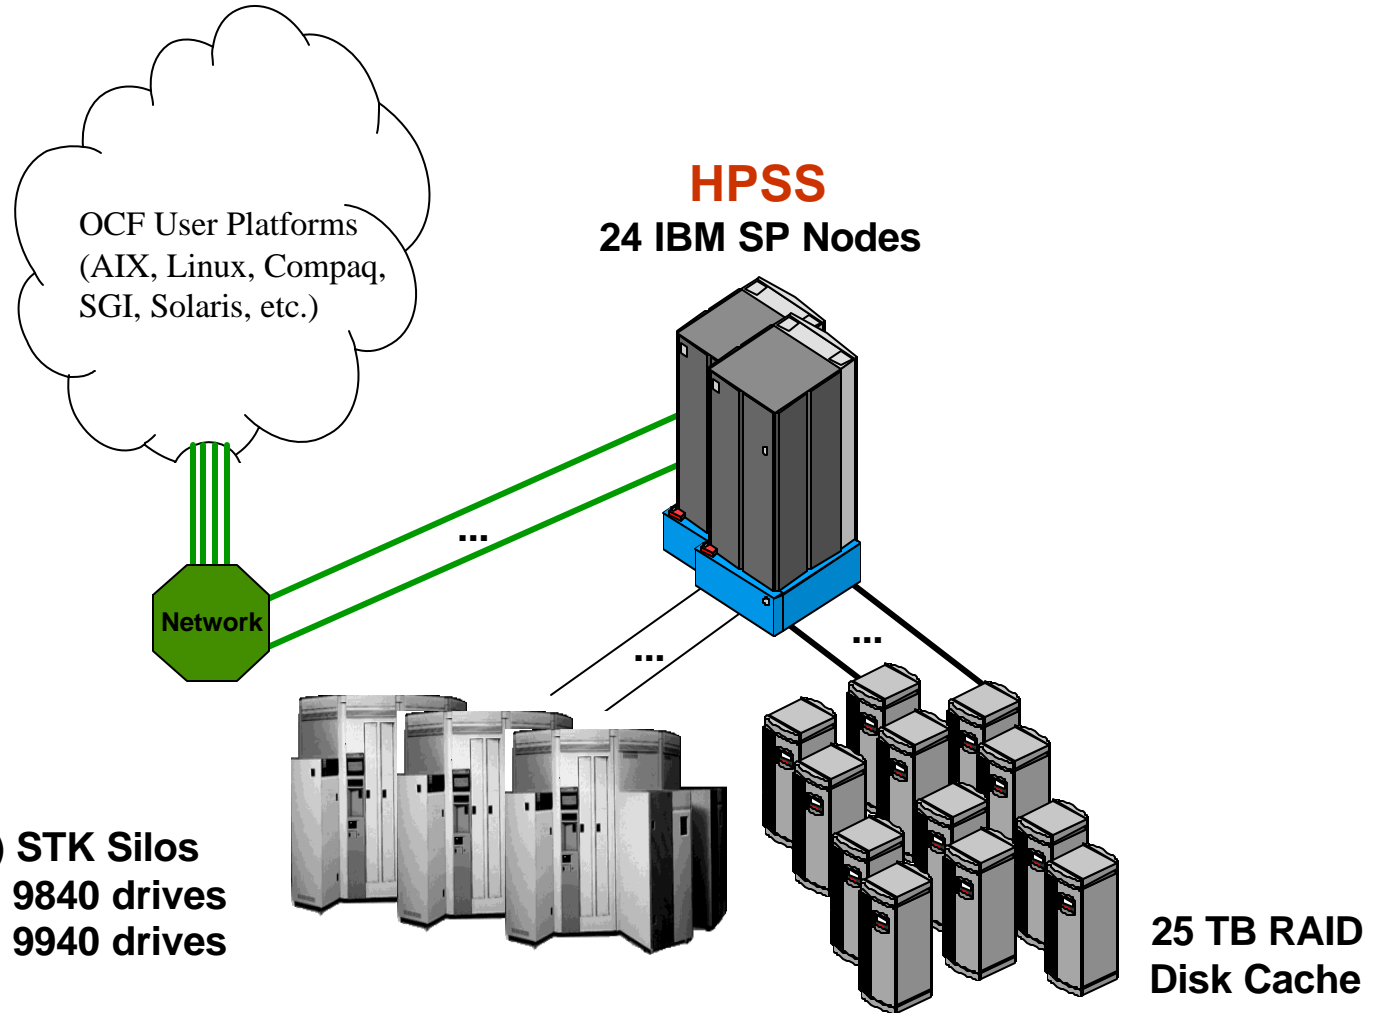
Site configuration

- **The LC operates classified (SCF) and unclassified (OCF) HPSS systems.**
 - > Both centers run HPSS R4.3
 - > Both systems are legacy systems converted from UniTree.
 - > SCF and OCF systems are kept as similar (and vanilla) 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.**



OCF HPSS storage system

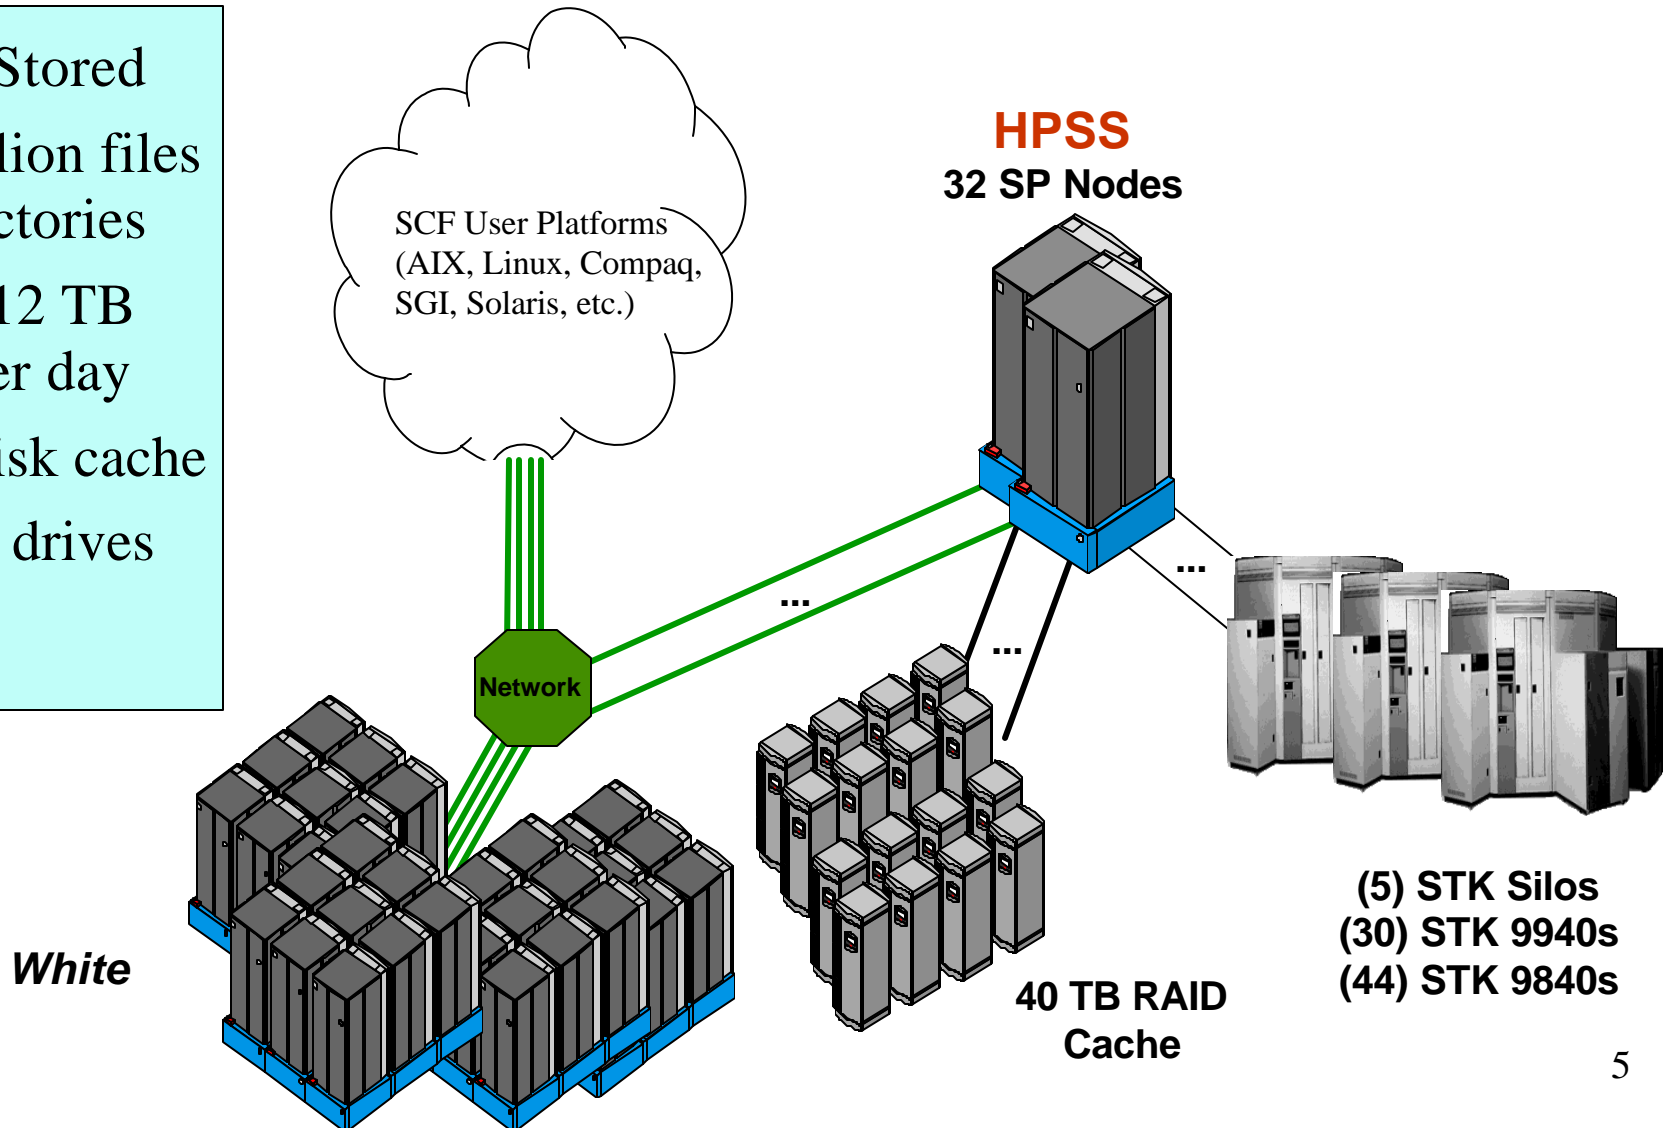
- ❖ 180 TB Stored
- ❖ 8.8 million files
- ❖ 50 GB-2.6 TB stored each day
- ❖ 25 TB disk cache
- ❖ 60 tape drives
- ❖ In production since 1/97





SCF HPSS storage system

- ❖ 500 TB Stored
- ❖ 15.7 million files and directories
- ❖ 50 GB -12 TB stored per day
- ❖ 40 TB disk cache
- ❖ 100 tape drives



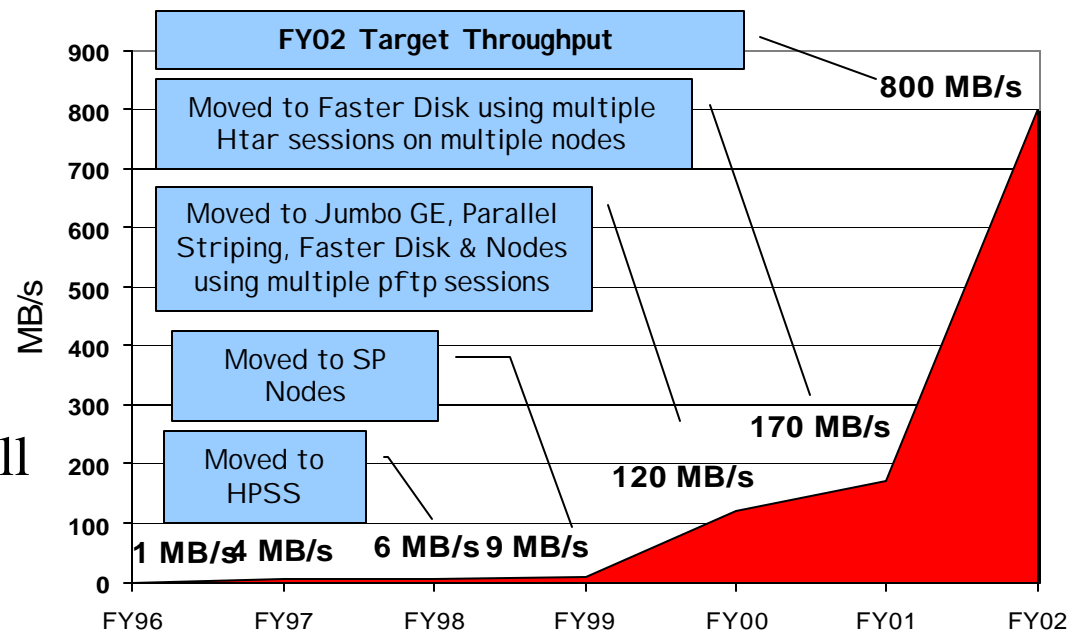


Archival storage performance

- **Accomplishments**

- > *A 28x performance increase in 24 months*
 - » Faster network
 - » Faster disks
 - » Attention to tuning
- > Delivery of Htar improved small file transfer performance by a factor of 50.
- > Demonstration of large single file transfer rates of up to 250 MB/s White-to-HPSS

Aggregate Throughput to Storage



At 250 MB/s, 2 TB of data moves to storage in 2.2 hours. Two years ago it took two and a half days to move that same data.



Achieving bandwidth goals

- **Moving to IBM pSeries 660 6M1 servers**
 - > 6-way RS64 IV 750 MHz Processors
 - > 3GB memory
 - > 4 RIO drawers w/9 GigE cards and 8 FC HBAs
- **Distributing disk across more nodes for more pipes**
- **Adding SAN FC Director Class Switches**
- **Moving to Cisco jumbo frame GigE for backend network**



LLNL issues and concerns

- **DCE replacement**
- **SAN integration**
- **PDATA push protocol improvement**
- **Inexpensive mover platforms (Linux)**
- **Disk allocation algorithm to improve performance to disk**
- **RAIT – mitigate issues of ever-increasing tape density**
- **Object manageability**
 - > Undelete
 - > File lifetimes



Future plans

- **Integrate (and relocate) more equipment**
- **Install HPSS 4.5**
- **Upgrade to AIX 5.1**
- **HPSS 5.1 conversion**
- **Explore Linux movers**
- **SANitize HPSS movers**
- **Continue performance tuning**
- **Satisfy ever-increasing bandwidth requirements**





Acknowledgements

- **The data storage group and many others make it happen**
 - > Kim Cupps, [Jim Daveler](#), [Dave Fisher](#), Keith Fitzgerald, Mark Gary, Mike Gleicher, [Pam Hamilton](#), Todd Heer, [Donna Mecozzi](#), [Jim Minton](#), [Debbie Morford](#), Alice Potts, Vickie Renbarger, George Richmond, Roger Skowlund, [Stan Solbeck](#), Jerry Shoopman, George Warda