

NERSC Reliability Data

NERSC - PDSI

Bill Kramer

Jason Hick

Akbar Mokhtarani

PDSI BOF, FAST08

Feb. 27, 2008

Production Systems Studied at NERSC

- **HPSS:** 2 High Performance Storage Systems
- **Seaborg:** IBM SP RS/6000, AIX, 416 nodes (380 compute)
- **Bassi:** IBM p575 POWER 5, AIX, 122 nodes
- **DaVinci:** SGI Altrix 350 (SGI PropPack 4 64-bit Linux)
- **Jacquard:** Opteron Cluster, Linux, 356 nodes
- **PDSF:** Networked distributed computing, Linux
- **NERSC Global File-system:** Shared file-system based on IBM's GPFS

Datasets

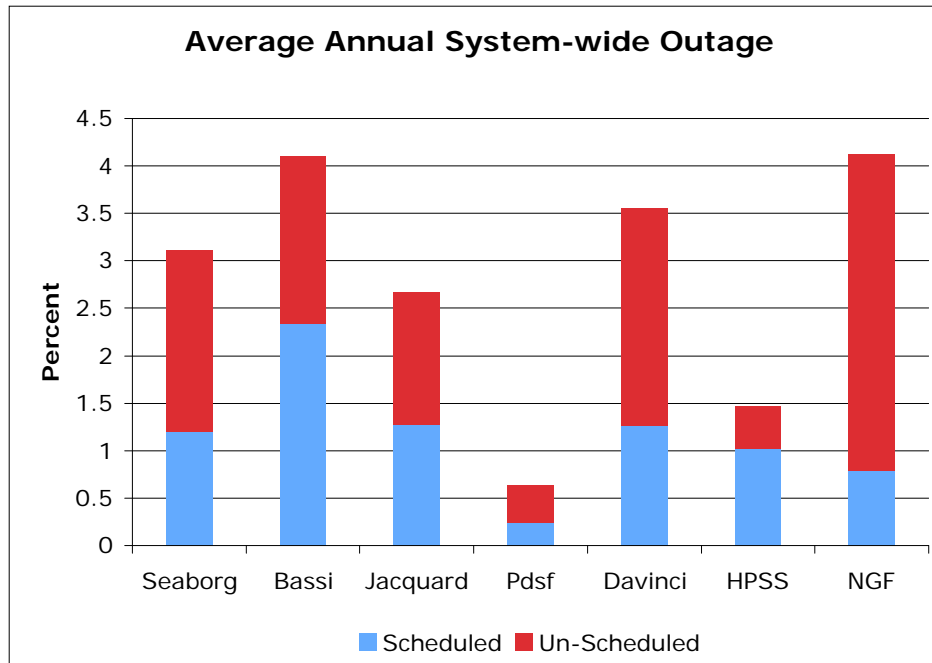
- Data were extracted from problem tracking database and paper records kept by the operations staff, and Vendor's repair records
- Coverage is from 2001 - 2006, in some cases a subset of that period
- Preliminary results on systems availability and component failure were presented at HEC FSIO Workshop last Aug.
- Have done a more detailed analysis trying to classify the underlying causes of outage and failure

- Produced statistics for the NERSC Global File-system (NGF) and uploaded to the CMU website. This is different from fsstats; used fsstas to cross check results on some smaller directory tree on NGF
- Made workload characterization of selected NERSC applications available. They were produced by IPM, a performance monitoring tool.
- Made trace data for selected applications
- Results from a number of I/O related studies done by other groups at NERSC were posted to the website.

Results


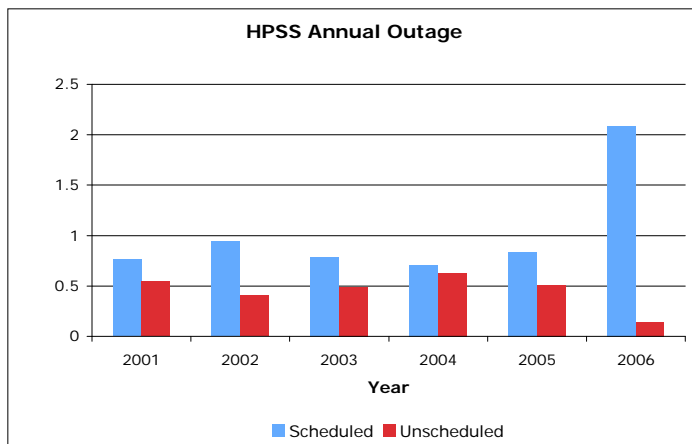


- Overall systems availability is 96% - 99%
- Seaborg and HPSS have comprehensive data for the 6 year period and show availability of 97% - 98.5% (scheduled and unscheduled outage)
- Disk drives failure rate for Seaborg show rates consistent with “aging” and “infant mortality”, average of 1.2%
- Tape drives failure for HPSS show the same pattern, average rate (~19%) - higher than manufacturer stated 3% MTBF

Average Annual Outage


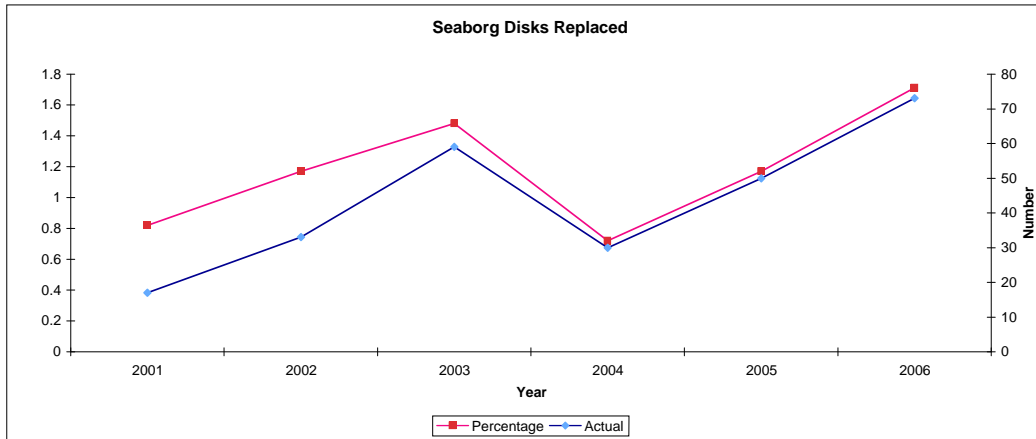


Data since: Seaborg(2001), Bassi(Dec. 2005), Jacquard(July 2005), DaVanci(Sept. 2005), HPSS(2003), NGF(Oct. 2005), PDSF(2001)

Seaborg and HPSS Annual Outage

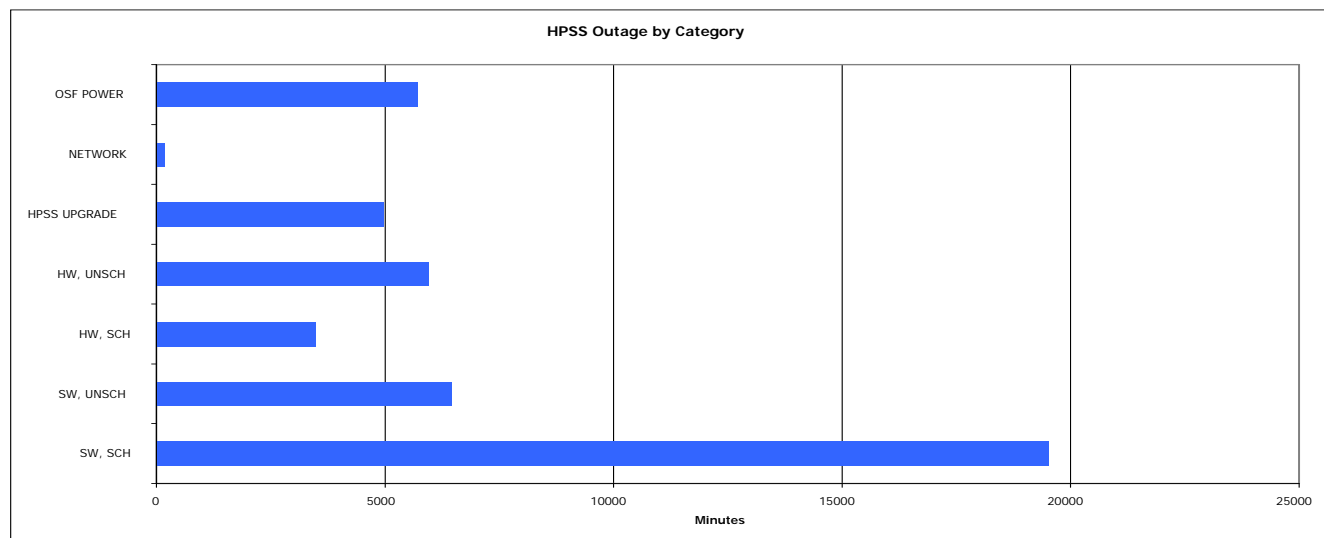



Disk and Tape Drives Failure



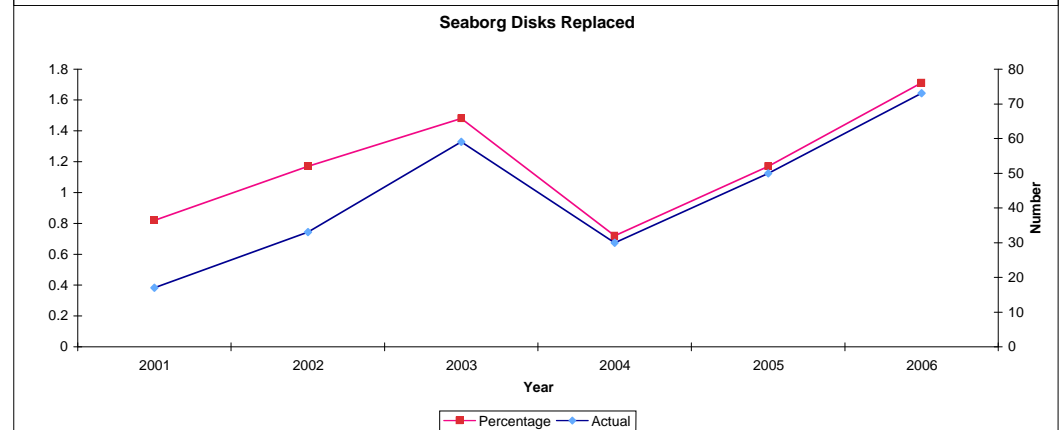
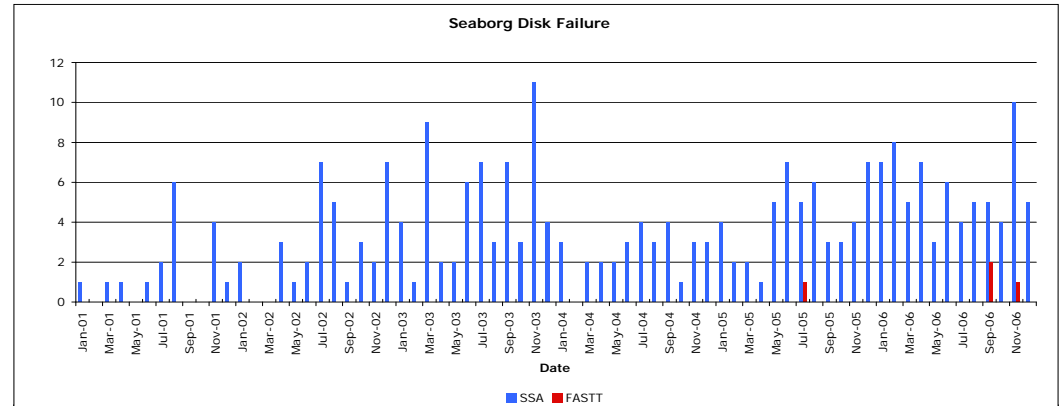

Extra Slides

Outage Classifications






Seaborg Data

- IBM SP RS/6000, AIX 5.2
- 416 nodes; 380 compute nodes
- 4280 disk drives (4160 SSA, 120 Fibre Channel)
- Large disks failure in 2003 can be attributed to “aging” of older drives and “infant mortality” of newer disks





HPSS Data


- Two HPSS systems available at NERSC
- Eight tape silos with 100 tape drives attached
- Tape drives seem to show the same failure pattern as seaborg's disk drives, "aging" and "infant mortality".



NGF stats



Seaborg Outages



HPSS Outages

