
HPSS User Forum - Site Report

Bill Weeks

Stanford Linear Accelerator Center

Produced under contract DE-AC03-76SF00515 between Stanford University and the Department of Energy



Bill Weeks - HPSS User Forum, June 5-7, 2001

1

BaBar & The B-Factory

- Use big-bang energies to create B meson particles
 - ◆ Look at collision decay products
 - ◆ Answer the question “where did all the anti-matter go?”
- 800 physicists collaborating from >80 sites in 10 countries
 - ◆ USA, Canada, China, France, Germany, Italy, Norway, Russia, UK, Taiwan
- The experiment produces large quantities of data
 - ◆ 200 - 400 TBs/year for 10 years
 - ◆ **Data stored as objects using Objectivity**
- Heavy computational load
 - ◆ ~15,000+ SpecInt95's needed



Bill Weeks - HPSS User Forum, June 5-7, 2001

2

HPSS Milestones

- BaBar HPSS system upgraded to 4.1.1.4
- Converted from 9840 to 9940 tape drives
 - ◆ 20 9940 FC tape drives
 - ◆ Brocade SilkWorm 2800 switch
 - ◆ Movers are Sun Netra 1405 running Solaris 7
 - ◆ 2 CPU, 1GB RAM, Gbit ethernet, 4 9940s
- Installed Public HPSS 4.2 Solaris system
 - ◆ Non-BaBar experimental data
 - ◆ Replacement for old staging system
 - ◆ Not in production



Public HPSS Software Levels

<u>Date</u>	<u>Machine</u>	<u>OS</u>	<u>DCE</u>	<u>Encina</u>	<u>HPSS</u>
Solaris core server					
May 01	420R	7	3.1	4.3.0.3	4.2
Solaris disk mover (160Gb RAID disk)					
May 01	420R	7	3.1	4.3.0.3	4.2
Solaris tape mover (4 9840 tape drives)					
May 01	220R	7	3.1	4.3.0.3	4.2



BaBar HPSS Software Levels

<u>Date</u>	<u>Machine</u>	<u>OS</u>	<u>DCE</u>	<u>Encina</u>	<u>HPSS</u>
AIX core server					
Jul 98	F40	4.1.5	2.1	2.1	3.2
Sep 99	F50	4.2.1	2.1	2.5	4.1
Jan 01	F50	4.3.3	2.2.0.8	4.2.0.14	4.1.1.4

- BaBar operation makes major software upgrades difficult
- Need to move to supported software versions




BaBar HPSS Software Levels (continued)

<u>Date</u>	<u>Machine</u>	<u>OS</u>	<u>DCE</u>	<u>Encina</u>	<u>HPSS</u>
Solaris 9840 mover (4 9840 SCSI tape drives)					
Jul 98	E250	2.5.1	1.1	2.5	3.2
Sep 99	E250	2.6	2.0	TX4.2	4.1
Jan 01	E250	2.6	2.0	4.2.0.14	4.1.1.4
Solaris 9940 mover (4 9940 FC tape drives)					
Mar 01	Netra1405	7	2.0	4.2.0.14	4.1.1.4




HPSS Bytes Written per Day



Bill Weeks - HPSS User Forum, June 5-7, 2001

7


HPSS Bytes Processed per Day



Bill Weeks - HPSS User Forum, June 5-7, 2001

8


HPSS Total Space Used



Bill Weeks - HPSS User Forum, June 5-7, 2001

9


HPSS Total Files



Bill Weeks - HPSS User Forum, June 5-7, 2001

10

HPSS Average File Size



Bill Weeks - HPSS User Forum, June 5-7, 2001

11

9840 Tape Statistics

- 290 TBs data written to 7,087 tapes
- Only one permanent read error causing data loss
- 112 write errors causing premature EOM
- 13 tape drives replaced



Bill Weeks - HPSS User Forum, June 5-7, 2001

12

9940 Tape Statistics

- 36 TBs written to 287 tapes
- Tape positioning about 40 seconds slower than 9840
- Maximum observed aggregate data transfer rate of ~200MB/Sec
- Peak individual file transfer rate of 30MB/Sec
 - ◆ Data compression
 - ◆ Fibre Channel



9840 vs 9940 Load/Unload Comparison

	9940	9940	9940	9840
microcode	29.201	28.223	28.220	
mount/load	19.00	26.21	24.06	11.41
open	1.51	1.77	1.61	1.73
close	0.00	0.04	0.03	0.03
unload	15.19	15.62	13.48	5.59
dismount	5.11	5.34	5.75	5.77

Time in seconds

Based on 100 iterations, no tape positioning on open



HPSS Reliability

■ Service outages since January 2001

- ◆ 5 Separate outages, total down time 38 hours
- ◆ 3 Unscheduled outages
 - ✦ Network problems (12 hours)
 - ✦ Power outage (12 hours)
 - ✦ SFS crashes (13 hours)
- ◆ 2 Scheduled outages
 - ✦ 9940 conversion (1 hour)
 - ✦ SFS parameter changes (.5 hour)

■ System available 98.9% of real time



HPSS Problems

■ SFS Invalid OFD

- ◆ Symptoms
 - ✦ Sfs_SetOfdIdleTimeout error: ENC-sfs-0059: Invalid OFD
 - ✦ "Bad Data Transfer. (error = -5, moved = 0)" in pftp_client
 - ✦ High number of transactions active reported by watchtran (> 2400)
- ◆ Solution
 - ✦ Changed idle timeout value from 60 to 181 seconds for SFS
 - ✦ Changed application to reduce/eliminate mkdir operations
 - ✦ Modified hpss_GetOfd to call Sfs_SetOfdIdleTimeout



HPSS Problems (continued)

■ PVR tape mounting problems

- ◆ Drive error limit doesn't work
 - ◆ 40 consecutive mounts of different tapes failed on drive
 - ◆ Drive not disabled
 - ◆ Problem reported. IBM response: working as designed
- ◆ Mount errors cause tapes to be marked EOM
 - ◆ Over 150 tapes marked EOM because of communication problems with ACSLS
 - ◆ Tapes had to be repacked and/or reclaimed
 - ◆ Change request submitted



HPSS Wish List

- Automatic system startup
 - ◆ Proper sequencing of DCE, Encina, hpsd, SSM, etc
- Improvement in handling large RAID disks
 - ◆ Segment size constraints
 - ◆ 16384 clusters per virtual volume
- PVR Error messages could be improved
 - ◆ PVR0166 STK gave unexpected return code
 - ◆ Include action, volume, drive
- Conversion from AIX to Solaris core server



Summary

- HPSS is reliable
- HPSS performance is adequate for current usage pattern
- Software upgrades will always be a problem

