

Indiana University's HPSS – A Site Report

Anurag Shankar <ashankar@indiana.edu>

Manager, Distributed Storage Services
University Information Technology Services
Indiana University

2003 HPSS Users' Forum
June 10-12, 2003, Ashville, NC

http://storage.iu.edu/presentations/iu_rept_huf03.ppt

Outline

- Statistics
- HPSS Use at IU
- HPSS Hardware/Software/Access
- Unique Features
- Issues/Concerns/Problems
- Future
- Conclusions

Statistics

- Users : 1,000
 - Data stored : 112 TB
 - Files : 5,685,265
 - Bulk data transfer : up to 2TB per day
 - Data growth rate : +40-60 TB per year
 - Data storage pattern : > 98% writes
- (Detailed stats at <http://www.indiana.edu/~dssg/stats.html>)

Statistics...

- Uptime : 99%
- Downtime details :
 - 60 hours (regularly scheduled)
 - 20 hours (unscheduled)
- # of problems : 135

HPSS Use at IU

- Research – biggest users are from astronomy, physics, chemistry, geology, etc. but use is growing by users in theater, fine arts, business, library IS, history, anthropology, etc.
- Life sciences – genomics, proteomics, radiology, etc., at the IU School of Medicine.
- Administration – student records, scanned documents, HR data, etc.
- Digital libraries –images, audio, video, scanned books, etc.
- General –workstation backups, etc.


HPSS Hardware

- Core HPSS servers run on a wide Silver IBM SP node at Bloomington.
- Tape/disk movers and DFS servers are on 10xIBM SP Silver nodes at Bloomington, and on an IBM H70 and 3xp640s in Indianapolis.
- Two IBM 3494 and two STK9310 tape libraries distributed between Bloomington and Indianapolis.
- 9xIBM 3590E, 6xSTK9840, & 12xSTK9940 tape drives.
- IBM SSA and UltraSCSI RAID disks.

HPSS Software

- Running HPSS 4.3.
- Running AIX 4.3.3, DCE 3.1 clients, DFS 3.1 servers, DCE 2.2 core servers.
- Using Samba, Netatalk, Apache, s-FTP as gateways, to provide access to the masses.

Distributed HPSS between Bloomington and Indianapolis



HPSS Access

HPSS access to users provided via:

- ftp
- pftp, kerberized pftp
- hsi (v2.6.2)
- DFS
- Samba/Netatalk/s-FTP
- Apache (WWW)

Activities since HUF '02

- Upgrade from HPSS-4.1.1.4 to HPSS-4.3.
- Addition of two new tape movers at IUPUI.
- Addition of six new 9940A drives at IUPUI.
- Migration of 30TB of data from 9840 to 9940 tapes.
- Upgrade to gig-E networking on the SP nodes.
- Migration of 30TB from 9840 to 9940A tapes.
- Institution of four-way, striped disk and tape COS using T9940A drives. Obtained linear scaling in performance.
- Development of HPSS API modules for Samba and Apache.

Unique/Noteworthy Features

- DFS
- Remote tape and disk movers
- Disaster tolerance via two tape copies and metadata backups at two geographically distributed sites
- HPSS-enabled Samba and Apache (demo <http://storage.iu.edu/DOOQS/>)
- Four-way, striped disk/tape, accessed via multinoded pftp

Issues/Concerns

- Handling small files
- Replacing DFS functionality
- Serving Windows users
- Lack of 24x7 support
- High, after-hour support cost

Problems

- Frequent HDM crashes
- SSM crashes/freeze ups
- Disk caches filling up

Future

- Upgrade to HPSS 4.5 this summer
- T9940B/fiber channel technology insertion
- Implement Linux-based, large disk caches to handle backups
- Experiment with cheap, IDE disks
- Move off the SP for core services/tape movers
- Prepare for life without DCE/DFS
- Extend HPSS to the masses via Samba-HPSS, Apache-HPSS

Conclusions

- HPSS has been quite successful at IU.
- Use has been growing, both in terms of data volume and diversity of users.
- The ability to consolidate data often more important to users than performance!
- Distributed, file system access is essential for HPSS to be appealing to the masses.

For more info, visit <http://storage.iu.edu/>.