LIGO Data Archive Selection

Stuart Anderson

Presentation to breakout session on LIGO computing

NSF Annual Review

24 October 2002

LIGO Data Archive Selection

(S2 data will exceed current disk capacity at Observatories)

- Current Status
- Hierarchical Storage Management (HSM) selection:
  HPSS vs SAM-QFS
  - Executive summary
  - Technical comparison
  - Validation tests
  - Cost comparison
  - Recommendation

Current Status

- Observatories
  - 5/10 TB Fibre Channel (FC) disk (LLO/LHO)
  - 2.2 TB Integrated Drive Electronics (IDE) disk
  - 1.5 TB tape robot
- Archive Center
  - 2 TB FC disk
  - 20 TB IDE disk
  - 1.2 PB tape silo
  - 54 TB HPSS frame archive
- Science runs
  - S1 13 TB
  - S2 47 TB
  - Per annum 270 TB

HPSS vs SAM-QFS Executive Summary (order of importance)

- HPSS advantages
  - Several years of experience
  - Free at Caltech
  - 54TB successfully stored
  - Scalability (raw data)
- SAM-QFS advantages
  - Simplicity (both use and administration)
  - License cost allows for use at observatories
    - Media import/export
  - Stability (asymptotic performance with increasing load vs. crash)
  - Metadata performance (x1000)
  - Reduced dependency on CACR
  - Disaster recovery (GNU TAR)
  - Single vendor solution (server, software and OEM storage)

HPSS vs SAM-QFS Technical Comparison

	HPSS	SAM-QFS
Topology	Network based 3rd party transfer	Single server (recent demo at 830MB/s)
Metadata	Nested transactional database (roll back changes)	Inode (1000x performance) Traditional backup
Tape format	Raw data only	GNU Tar (disaster recovery)
Software	AIX/Solaris DCE + Encina + …	Solaris Single package
User Interface	FTP (PFTP) hsi shell	POSIX filesystem (ls, emacs, …) QFS (already selected)
Data migration	Raw data copy (extra tape drives)	Physical media ingestion Metadata copy

HPSS vs SAM-QFS Validation

- HPSS
  - Archived 54 TB/1.3 M files of frame data over multiple years.
    - Very little retrieval due to difficulty of use, I.e., traditional backup.
      - Even though network bandwidth was larger than tape I/O, the E7 data replication to UWM was done via labor intensive tape shipping/ingestion.
- SAM-QFS
  - Archived all of S1 (13 TB/198 k files).
  - Retrieved every byte in 1 week with 2 tape drives.
    - Unattended weekend run at 27.6 MB/s.
  - Each file positively verified to have the correct MD5SUM from IDE-RAID system at 227 MB/s.
  - Retrieved 273 GB of early S1 data while archiving later data without any performance degradation, I.e., no tape thrashing.
  - S1 data replicated to UWM from QFS until UWM disk full.

HPSS vs SAM-QFS Cost Comparison

	Caltech	Observatories
HPSS	Covered by CACR MOU in exchange for 1 FTE (unlimited size)	$300k + $100k/yr + 1-2 FTE (per observatory)
SAM-QFS	$0.046/GB (2001) *$0.400/GB (2002)*	$0.046/GB (2001) *$0.400/GB (2002)*

Estimate that LDAS integration to SAM-QFS is free and that HPSS is 1 man-year.

Note: Tape ($0.4/GB), Disk ($4/GB)

HPSS vs SAM-QFS Recommendation

Select SAM-QFS over HPSS for each of the following sufficient reasons:

- SAM-QFS supports the import/export of original tapes.
  - HPSS fails for both technical and financial reasons.
- SAM-QFS will allow 1yr of automated data access at each Observatory.
- SAM-QFS should allow LDAS (and other?) direct access to deep archive.
- In my opinion, SUN will drop support for HPSS unless they win a large government contract leaving us stuck with IBM hardware and OS.
- When the next best thing comes along in a few years we will be able to migrate our data using ANY computer system that supports the FC tape drives with the data and is able to run GNU Tar.
- To do a directory listing of the current LIGO archive in HPSS takes more than 24hr, whereas in SAM-QFS it is extrapolated from the 1/7th size S1 dataset to be just 4min.