LIGO

Data Analysis Systems

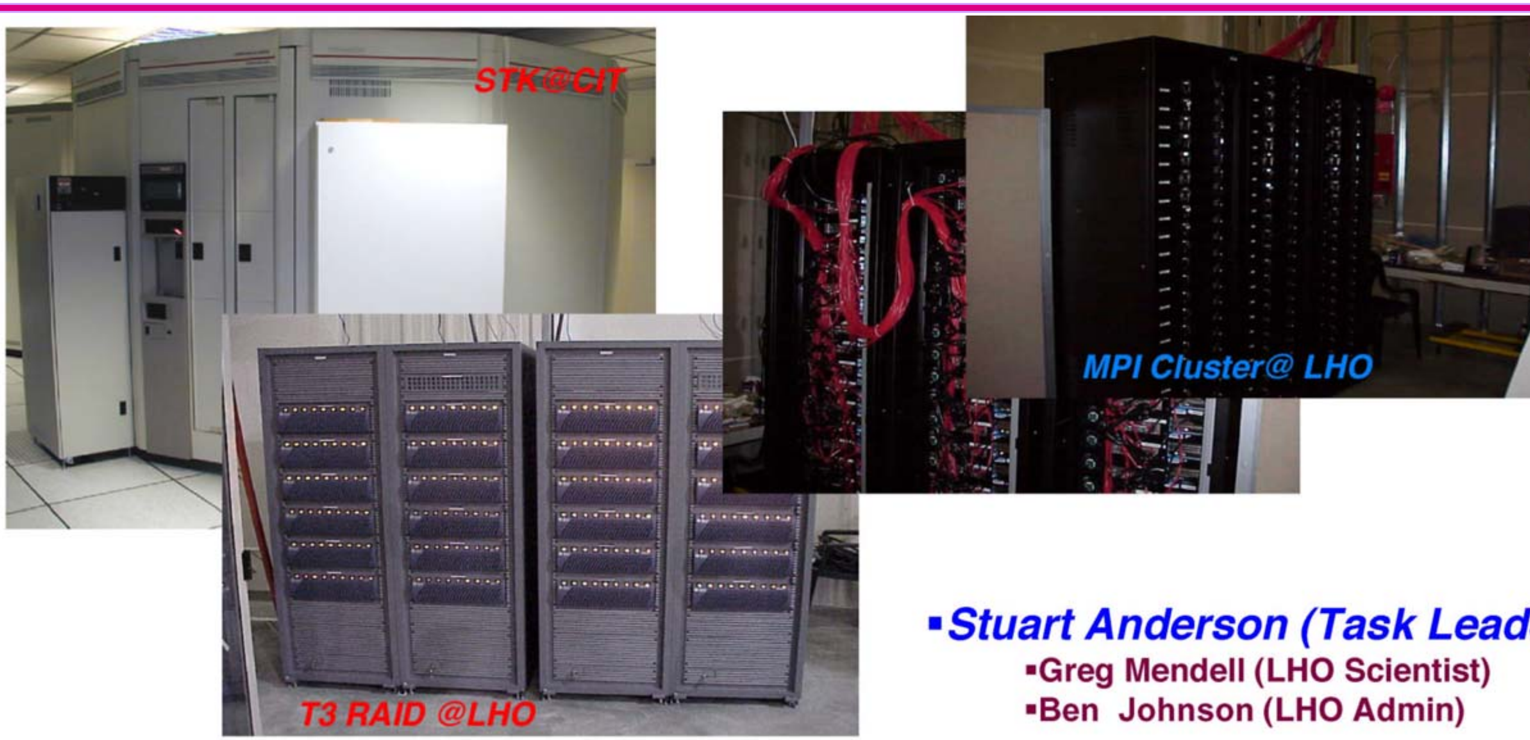

- ***Stuart Anderson (Task Leader)***
  - **Greg Mendell (LHO Scientist)**
  - **Ben Johnson (LHO Admin)**

  - **Igor Yakushin (LLO Scientist)**
  - **Shannon Roddy (LLO Admin)**

  - **Dan Kozak (CIT Admin)**
  - **Al Wilson (CIT Admin)**

  Keith Bayer (MIT Admin)

LIGO-G030291-00-C

LDAS Instances

- 4 Production systems
  - » 2 on-site: observatory archive and real-time processing (LHO/LLO)
  - » 2 off-site: central archive and multi-IFO processing (CIT/MIT)
- 1 Development system
  - » Run daily builds of LDAS from CVS
- 1 Test system
  - » Pre-release testing
  - » Integration testing of LSC/LIGO software

LDAS Hardware Update

- Final LDAS Construction items have been delivered:
  - 420 2.66GHz/2GB dual-Xeon nodes.
    - 70 nodes/IFO and 112 nodes at MIT.
  - Copper Gigabit Ethernet switches for Beowulf clusters.
  - Tape robotics with 4 STK tape drives for each Observatory(140TB).
  - 12TB disk cache and 6 STK tape drives for central archive (1.2PB).
- All servers upgraded to Gigabit Ethernet.
- Shared QFS filesystem operational as data interface between LDAS and CDAS for E7, S1 and S2.
- SAM-QFS mass storage system used for archiving S1 and S2 data.

LDAS S1(S3) Configuration

"Install full size compute clusters and mass storage systems."

	FC (TB)	IDE (TB)	CPU (GHz)	Tape (TB)
LHO	10	2(4)	139(763)	2(140)
LLO	5	2(4)	107(391)	2(140)
MIT	1	2(8)	45(244)	
CIT	3(12)	18	34(1150)	90(500)
DEV	1	2	25(42)	2(5)
TEST	1		8(64)	
BOX(x5)		0.1	5	
TOTAL	21(30)	26(36+85)	383(2679)	96(785)

Post-construction LDAS hardware

- Grow Archive system out to 1.2PB with subsequent Science runs.
- Grow IDE RAID to hold frequently accessed data.
  - S1(full) 13 TB / S2(RDS) 7TB
  - Per annum 270 TB
    - Disk may be cheap enough to hold all of this on-line if sufficient facility resources can be obtained from Caltech (sq ft, kW, kTon).
- Upgrade servers
  - Linux DataconditioningAPI servers to 64-bit (4x2GHz/16GB).
  - Solaris frame data servers to 8x1.2GHz/16GB.
  - Database servers to faster CPU to meet higher than expected event rates (TBD)
    - E7 and S1 ran at 5Hz average rates but >50Hz may be needed.

Central Archive Selection Criteria (order of importance)

- HPSS advantages
  - Several years of experience
  - Free at Caltech
  - 62TB 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)

LIGO

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 62 TB/1.8 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) *<u>$0.400/GB (2002)</u>* "free" (2003)	$0.046/GB (2001) *<u>$0.400/GB (2002)</u>* "free" (2003)

Estimate that LDAS integration with SAM-QFS is 1 man-week and that HPSS is 1 man-year.

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

SAM-QFS Selection

Selected SAM-QFS over HPSS for the following main 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 allows LDAS (and others) 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 LIGO data using ANY computer system that supports the FC tape drives holding 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.

- S2 (660k files/48TB) was successfully archived to SAM-QFS only and post-S2 all trend and reduced data where also switched from HPSS to SAM-QFS.