



LIGO Data Analysis System and Simulation Updates

8th Meeting of the LIGO PAC

Caltech

Pasadena, California

2 May 2000

Albert Lazzarini

LIGO Laboratory Caltech



These talks

- Pending LDAS HW procurement planning
- LDAS SW Status (Kent Blackburn)
- LDAS Data Archive Status (Stuart Anderson)
- LDAS Database Access (Peter Shawhan)
- End-to-End Simulation Environment (Hiro Yamamoto)



LIGO Data Analysis System Hardware Procurement Overview

- LDAS hardware consists of:
 - » PCs for parallel computation
 - » Unix(Sun/linux) workstations for users (on private LDAS LANs)
 - » RAID (HW & SW) systems for data caching, staging, storage
 - » Data Servers (Sun 60, 450, 420)
 - » Hierarchical Storage Management system (HSM) for main archive
 - » Smaller robotic tape archives for sites
 - » Networking switches
- Amount budgeted from LIGO Project Construction for this procurement: \$5M.



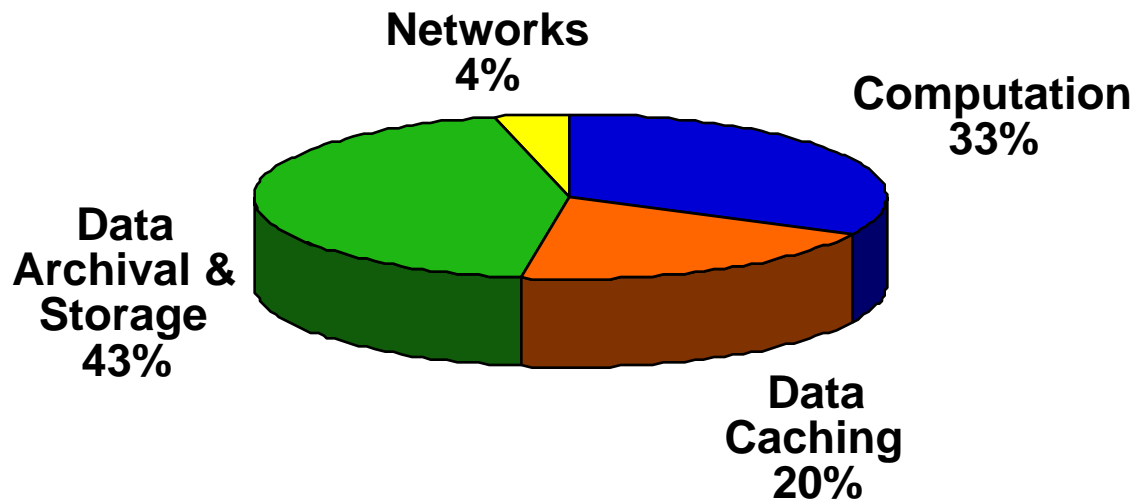
LIGO Data Analysis System Procurement Strategy

- Consider only commercial products (“catalog items”)
 - » Multiple vendor quotes for identical (or similar) configurations
 - » Select best value (reliability, warranty, cost, service, ...)
- Maintenance
 - » Bundle multi-year service contracts with purchase of major (unique, costly) equipment -- saves operations costs
 - ATM, Ethernet switches
 - Servers
 - RAID
 - » “Self insure” for commodity units (PCs) through:
 - Redundancy, spares
 - Prepaid multi-year service/replacement contracts




LIGO Data Analysis System Hardware Procurement Overview

LDAS Procurement By Function



LDAS Procurement By Site





Networking Technology

- LDAS uses heterogeneous mixture of networking technologies:
 - » Ethernet (100BT and 1000BT) for compute clusters, LANs
 - » ATM (OC3, OC12) for connection to internet, databases, server-server, etc. for high BW, point-to-point high priority connections
 - LAN Emulation (LANE)
 - Large (16 x OC12) switches connect databases & servers to compute clusters, local users
 - » Fibre channel for disk systems (where needed)



CPU Technology

Servers, Workstations, Compute Clusters

- Servers used for relational database and to distribute framed time series data to users
 - » DBMS
 - » Network file systems
 - » LDAS software modules
- e.g.:
 - » Unix/linux for users
 - » Linux (Intel) for compute clusters
 - » Enterprise class or equivalent servers (SMP)



CPU Technology

Servers, Workstations, Compute Clusters

- Sun & Linux machines (Intel):

Component	Qty	Comments
Compute clusters 1GHz CPU; 512 MB RAM 1x72GB IDE; 100 Mbps	336	linux/Intel commodity PCs for parallelized beowulf MPI analysis
Beowulf server 4x1GHz CPU; 4GB RAM 2x72GB IDE; 1000 Mbps	5	master nodes
Data Conditioning PCs 4x(Xeon CPUs); 512MB RAM 1x72GB IDE; 100 Mbps	8	For regression analysis, calibration, pre-processing of data
Control & Monitoring PCs 1 CPU; 256 MB 2x18 GB IDE; 100 Mbps	5	control room monitors; log file servers; software mirrors
Intel Workstations 1 CPU; 512 MB 1xGB IDE; 100 BT	16	Commodity PCs; for work within LDAS LAN
Sun Workstations 1 CPU; 512 RAM 1x36GB SCSI; 100 Mbps	16	For work within LDAS LAN

LIGO-G000120-00-E



Mass Storage Technology

- Comparative scale of LIGO I needs

Program	Top Level Requirements
LIGO I	<ul style="list-style-type: none">100 TB total storage0.2 TB per day growth<1 TB per day transfer10⁶ files5 concurrent transfers @ 1-10 MB/sAvailability > 95 % (up-time)
ASCI	<ul style="list-style-type: none">5 PB total storage10 TB per day growth50 TB per day transfer10⁹ files4 concurrent transfers @ 2 GB/sAvailability ?
LIGO II	<ul style="list-style-type: none">750 TB - 1 PB total storageOther requirements TBDAvailability > 95%

Major impetus for HPSS development





Mass Storage Technology

- Hierarchical Storage Management
- Large volume robot at Caltech
 - » 100 TB - 200 TB after 2 - 3 years
 - » Target: HPSS (likely on Sun)
 - » Dedicated HSM disk cache (1 TB)
 - » Data movers(servers) & tape controllers
- Smaller volume robots at each of two sites
 - » 20 TB max,
 - » Non-HPSS (initially for LIGO I)
 - » Same tape drive technology as at Caltech
- Miscellaneous small robots (~ 1 TB max)
 - » User access to reduced datasets
 - » AIT-2 (Sony, Cybernetics)
 - » 30 x 50 GB tapes



Mass Storage Technology

- **HPSS System for LIGO I (IBM baseline*):**

Component	Qty	Comments
SP2 rack with backplane switch	1	
4-way PCI nodes	5	Tape head servers, HPSS metadata server
High performance gateway node (HPGN)	1	High BW service to disk cache, users, PC cluster, etc.
High BW tape drives (STK 9840 or later, or equivalent)	8	High BW, high volume tape cassettes for archive
Misc. disk storage	~200GB	HPSS metadata
SSA RAID disk cache	1 TB	Intermediate disk cache for R/W tape transfers

****Baseline configuration shown is for IBM HW for budgeting purposes; attractive second option is to install equivalent Sun configuration***



Disk systems

- RAID systems (file servers, e.g. NetApp 720)
 - » 1500GB across 4 installations
 - » High BW access:
 - SCSI or fibre channel
 - Hostless, network-attached
 - » Metadata servers
- Disk cache for frame data
 - » 2850GB across 4 installations
 - » RAID 0+1, plus servers




Disk systems

- Disk cache systems:

Use	Qty (GB)	Comment
Framed data at observatories	1000/500	LHO/LLO
Metadata (DBMS)	750/150/300/300	CIT/MIT/LHO/LLO
LDAS caches	1250/250/500/250	CIT/MIT/LHO/LLO




LDAS Procurement Timeline



LIGO-G000120-00-E



LDAS Expenditure Profile



LIGO-G000120-00-E

LSC Meeting 2000.03.16-18

LIGO Laboratory at Caltech



LIGO Data Analysis System Summary

- This is the final major procurement for the LIGO Construction Project
- Procurement Plan follows from the requirements of LIGO I data analysis
- Procurement will ramp up during 2Q & 3Q 2000
 - » Mounting pressure/need from the detector characterization activities to archive, analyze data
 - » Procurements will be phased to match demand
 - » Select the most cost effective options at each phase
- Design is extensible
 - » Allows for expansion, upgrades to support LIGO II when needed