LIGO Data Analysis System
Hardware Needs, Considerations

Presented to SUN Microsystems

31 January 2000

Caltech

Pasadena, CA 91125

LIGO

Data Storage Architecture

(→ dataflow)

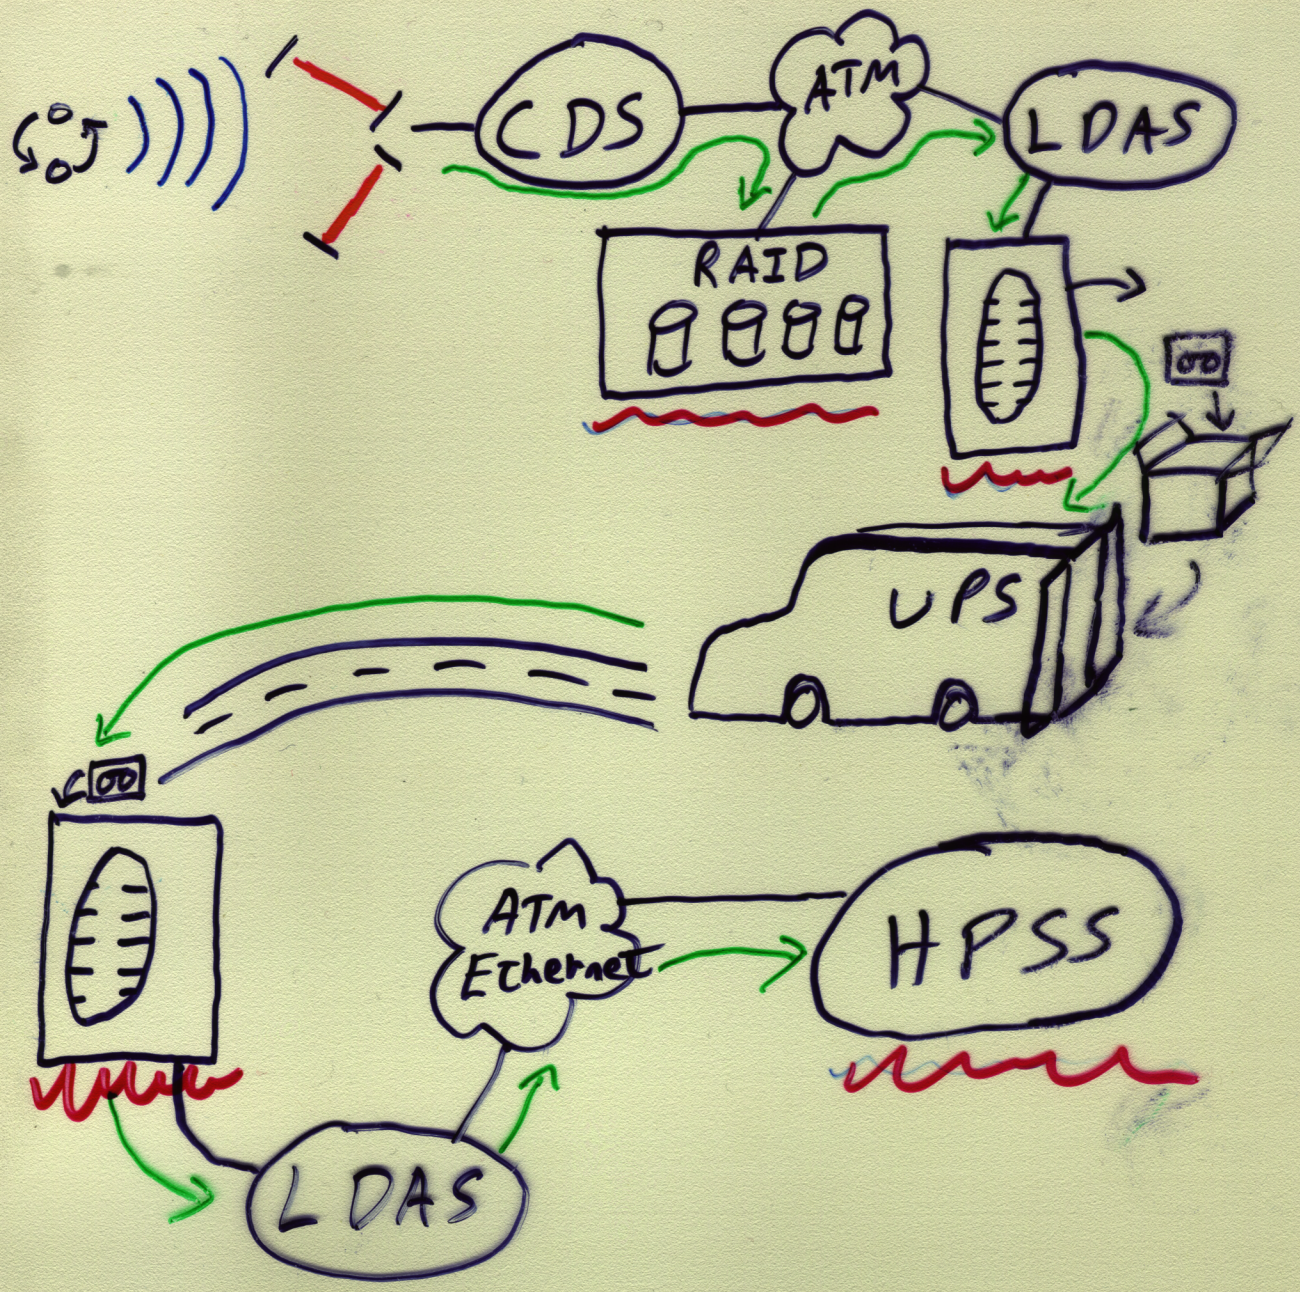

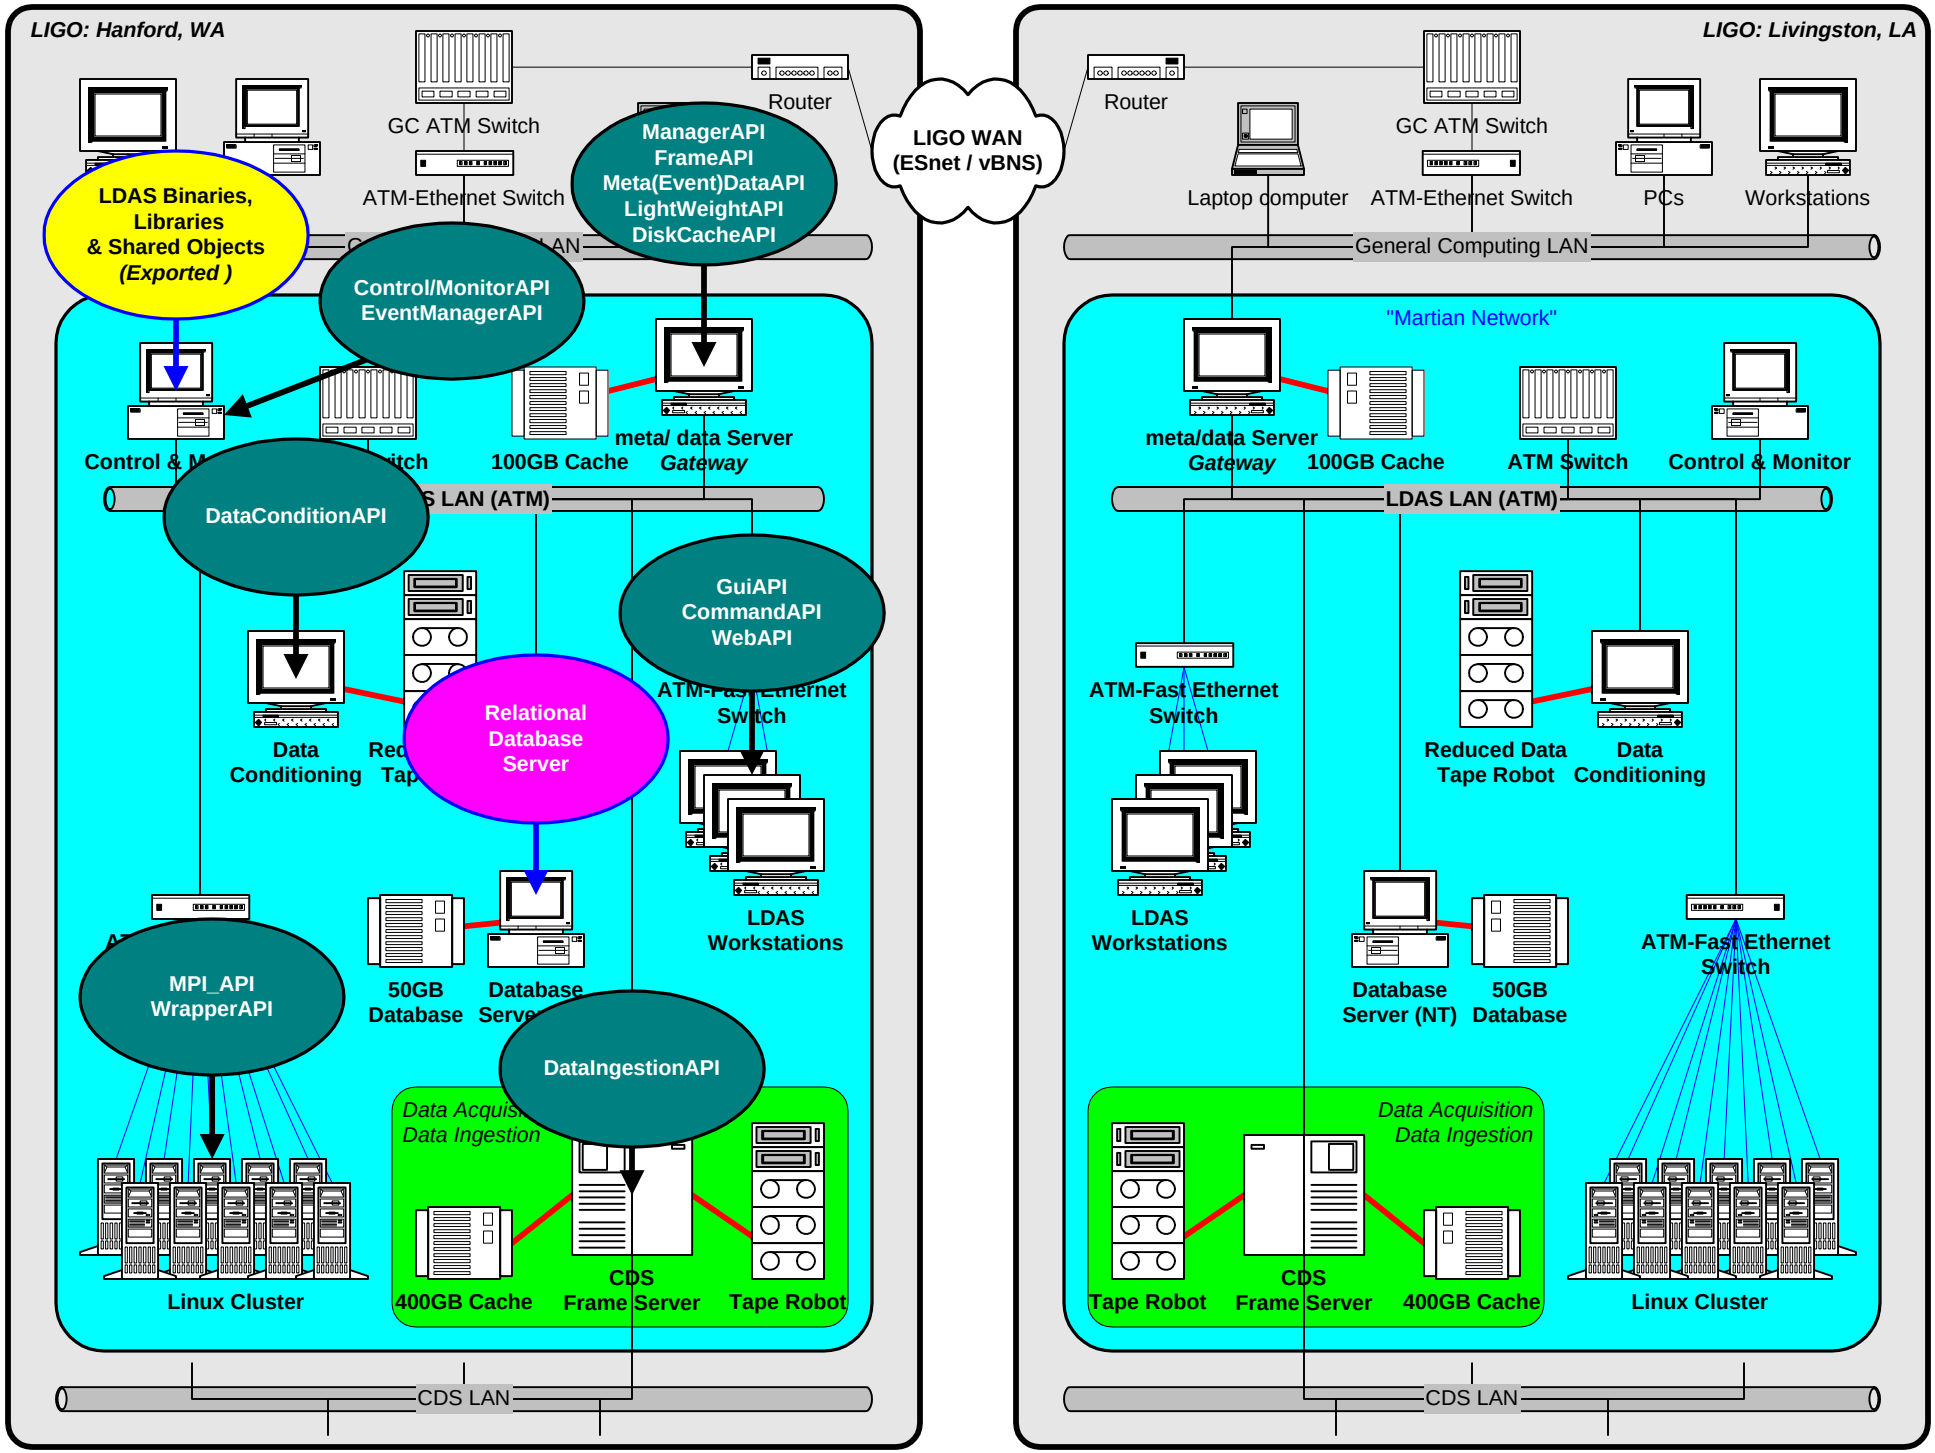
LIGO: Hanford, WA
LIGO: Livingston, LA
LIGO WAN
(ESnet / vBNS)
Router
GC ATM Switch
ATM-Ethernet Switch
Laptop computer
PCs
Workstations
General Computing LAN
LDAS Binaries,
Libraries
& Shared Objects
(Exported)
ManagerAPI
FrameAPI
Meta(Event)DataAPI
LightWeightAPI
DiskCacheAPI
Control/MonitorAPI
EventManagerAPI
DataConditionAPI
GuiAPI
CommandAPI
WebAPI
Relational
Database
Server
MPI_API
WrapperAPI
DataIngestionAPI
"Martian Network"
meta/data Server
Gateway
100GB Cache
ATM Switch
Control & Monitor
LDAS LAN (ATM)
ATM-Fast Ethernet
Switch
Reduced Data
Tape Robot
Data
Conditioning
LDAS
Workstations
Database
Server (NT)
50GB
Database
Data Acquisition
Data Ingestion
Linux Cluster
400GB Cache
CDS
Frame Server
Tape Robot
CDS LAN

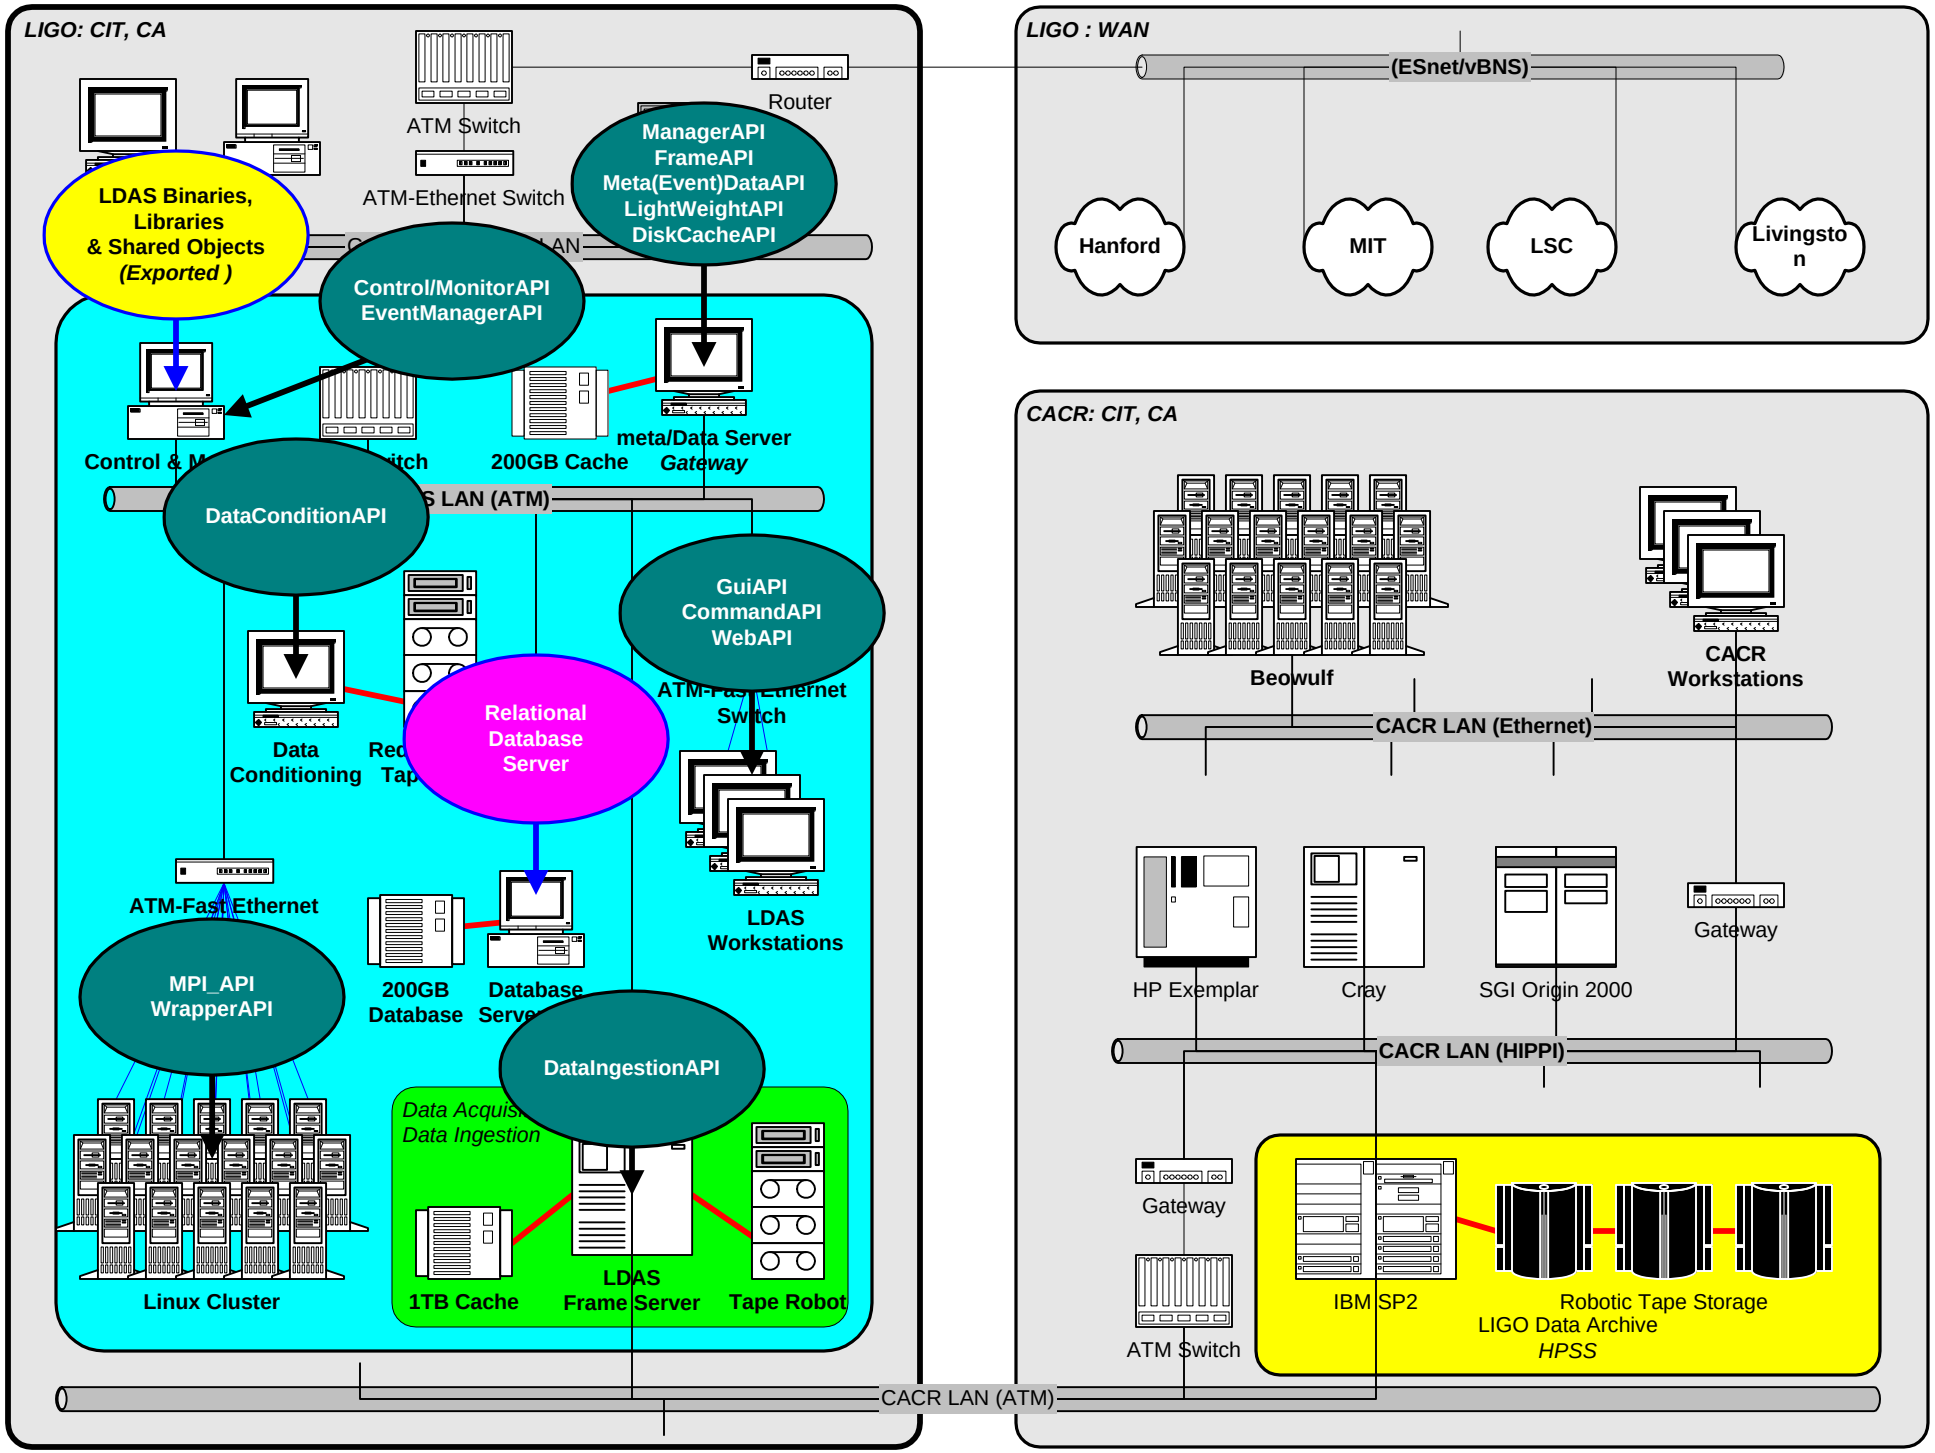
LIGO: CIT, CA
ATM Switch
Router
ATM-Ethernet Switch
ManagerAPI
FrameAPI
Meta(Event)DataAPI
LightWeightAPI
DiskCacheAPI
LDAS Binaries,
Libraries
& Shared Objects
(Exported)
Control/MonitorAPI
EventManagerAPI
200GB Cache
meta/Data Server
Gateway
DataConditionAPI
GuiAPI
CommandAPI
WebAPI
Data
Conditioning
Relational
Database
Server
ATM-Fast Ethernet
Switch
ATM-Fast Ethernet
LDAS
Workstations
MPI_API
WrapperAPI
200GB
Database
DataIngestionAPI
Data Ingestion
Linux Cluster
1TB Cache
LDAS
Frame Server
Tape Robot
CACR LAN (ATM)
LIGO : WAN
(ESnet/vBNS)
Hanford
MIT
LSC
Livingsto
n
CACR: CIT, CA
Beowulf
CACR
Workstations
CACR LAN (Ethernet)
HP Exemplar
Cray
SGI Origin 2000
Gateway
CACR LAN (HIPPI)
Gateway
ATM Switch
IBM SP2
Robotic Tape Storage
LIGO Data Archive
HPSS

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*

CPU Technology
Servers, Workstations, Compute Clusters

- Servers for large relational databases and to provide data to users
  - *Compilation*
  - *Simulations*
  - *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

- Servers:

Component	Number	Comments
2 CPU 512+ MB 72+GB SCSI	5	DB servers, File servers
4 CPU 1+GB RAM 128+GB SCSI	2	Site-based DB, file servers

CPU Technology
Servers, Workstations, Compute Clusters

- Linux machines (Intel):

Component	Qty	Comments
Compute cluster, data preprocessing 1–2 CPU 256–512 MB 2x18 GB IDE 100 BT	224	linux/Intel commodity PCs
Control & Monitoring PCs 1 CPU 256 MB 2x18 GB IDE + 4x18 GB SCSI 100 BT	4	linux/Intel Commodity PCs
Workstations 1 CPU 512 MB 1 x GB IDE or SCSI 100 BT	14	linux/Intel Commodity PCs

CPU Technology
Servers, Workstations, Compute Clusters

- Unix (SUN or equivalent) machines:

Component	Qty	Comments
Workstations **1 CPU** **512 MB** **1 x GB SCSI** **100 BT**	**14**	**Ultra 10 or equivalent & higher**

Mass Storage Technology

- Comparative scale of LIGO I needs

Program	Top Level Requirements
LIGO I	• 200 TB total storage • 1 TB per day growth • 4 TB per day transfer • 10^6 files • 5 concurrent transfers @ 1-15 MB/s • Availability > 95 % (up-time)
ASCI	• 5 PB total storage • 10 TB per day growth • 50 TB per day transfer • 10^9 files • 4 concurrent transfers @ 2 GB/s • Availability ?
LIGO II	• 750 TB - 1 PB total storage • Other requirements TBD • Availability > 95%

Mass Storage Technology

- Large volume robot at Caltech
  - *100 TB - 200 TB after 2 - 3 years*
  - *e.g.:, HPSS*
  - *Associated disk cache (! - 2 TB)*
  - *Servers, tape controllers*
- Smaller volume robots at each of two sites
  - *`10 - 20 TB max,*
  - *Non-HPSS*
  - *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

- ~100 TB HPSS System for LIGO I:

Component	Qty	Comments
SP2 rack with backplane switch	1	Mainframe
4-way PCI nodes	5	Tape head servers, HPSS metadat 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)	4	High BW, high volume tape casettes for archive
Misc. disk storage for HPSS metadata	~200 GB	
SSA RAID disk cache	1 TB	Intermediate disk cache for R/W tape transfers

Mass Storage Technology

- Mid and small size tape robot systems:

Component	Qty	Comments
~ 1TB AIT-2 robot, 30 tapes, 2 heads	4	For user-based generation and readback of reduced data tapes.
~ 20TB robot 2 or 3 tape heads	2	Tape units to be compatible with HPSS

Disk systems

- RAID systems for on-site data storage (acquisition)
  - *500+GB*
  - *SCSI or fibre channel*
  - *Hostless, network-attached*
- Disk cache for HPSS
  - *1 TB*
  - *JBOD or RAID*

Disk systems

- Disk cache systems:

Use	Qty (GB)	Comment
Framed data	200/400/200	Excludes HPSS cache at Caltech 3 sites
Metadata (DBMS)	300/100/100/100	4 sites
Template storage	360/240/120	3 sites