LIGO I Data Analysis System

Overview

NSF Annual Review of LIGO
9 - 11 May 2000

Albert Lazzarini
Caltech
Pasadena, CA 91125

LIGO Data Analysis System
LDAS Overview

- Key LDAS requirements for LIGO I
  1. Analysis pipelines for (dedicated) astrophysical searches requiring parallelized computation
     - Analyze data at *same rate* they are produced
       - Requires 30 -100 GFLOPS/interferometer
     - Two stages: individual interferometers (on-site) + multiple interferometers (off-site)
  2. Data archive for LIGO I science run
     - ~100 TB raw time series data (binary data in frames)
     - Few TB reduced data -- selected channels, data QA
     - 1 TB Metadata -- frame summaries, triggers from on-line diagnostics, events from pipeline analyses, calibrations, etc.
  3. Data distribution - reduced datasets
     - pftp, scp, tapes

LIGO Data Analysis System
LDAS Overview

- LDAS: Software & hardware environment for LIGO I

1. *On-site* analysis pipeline at each of 2 LIGO Observatories -- individual interferometer analysis (Production only)

2. O*ff-site* analysis pipeline at Caltech to perform coherent multiple-interferometer analysis --- *likely includes international partners* (Production + Development)

3. Archive for LIGO I data at Caltech with open access to collaboration members

4. Local compute cluster, small archive at MIT (Development)

LIGO Data Analysis System
LDAS Overview

- LDAS Design Process
  - » LDAS White Paper (1997), issued jointly with visiting LSC collaborator (B. Allen)
  - » LIGO Laboratory design reviews (reqmnts:1998; prelim: 1999; final: 2000 - *pending NSF procurement approval, first MDC*)
    - – Review panels *include* collaboration members outside lab.
  - » LSC White Paper (1999), jointly authored by Lab members & collaborators
    - – Discusses analysis plans, data usage model, long term plans (preliminary)
- Two LDAS Layers
  - » Infrastructure (data management, movement, pipeline control) built principally by LIGO Laboratory, with LSC participation (UTB, PSU, ANU)
  - » Scientific analysis software developed as LSC activity -- includes LIGO Laboratory scientific participation

LIGO Data Analysis System LDAS Overview

- LDAS implementation
  - » C++ for compiled code; Tcl/Tk command/steering language for controlling analysis flow, processes, user interactions
    - – 20 major software modules (with multiple layers - C++, Tcl)
      - 10 modules *complete* to date (use mock data challenges to test integration)
      - 6 modules *in progress* (5% - 50% complete, work in progress with LSC)
      - 4 modules to be started before end of 2Q2000
    - – 460k lines currently in CVS repository.
      - 120k lines from LIGO; 340k lines of external sources adopted, modified, maintained by LIGO.
  - » Target O/S: Unix -- Solaris(Sun), linux(intel)
  - » 1Q2000 : began work of integrating analysis software into environment
    - – Work with LSC collaborators from SW development groups
    - – Concentrate on a series of mock data challenges aimed at validating integration & performance.

LIGO Data Analysis System
LSC White Paper Overview

- LSC White Paper draft issued Fall 1999 identifies:
  - » Principal astrophysical searches which will be performed
    - – Burst, chirped signals, CW signals, broadband stochastic signals, unmodeled sources
  - » Techniques, analysis methods
    - – Detector cross-correlation (time/frequency domains) for non deterministic/unmodeled sources
    - – Template based optimal (Wiener) filtering for sources with deterministic, calculable phase evolution
  - » Data products, data usage model (still evolving)
    - – Time series data -- available as frames
    - – Metadata -- available as XML

LIGO Data Analysis System Data Archive Overview

- LIGO Laboratory Archive located on campus at Caltech
  - » Shared space, Institute infrastructure at Center for Advanced Computing Research (CACR)
  - » Resources available to LIGO/LSC through MOU with CACR
    - – Use of robotic silo owned by CACR/Caltech
      - 2000 tape slots (50 - 100 TB, depending on cassette type)
    - – HPSS site license for operating archive
    - – High bandwidth access to internet2/Calren2
    - – Shared personnel with expertise in HPSS/parallel computing
  - » LIGO owned, controlled & administered:
    - – Data movers (servers) + tape drives for Hierarchical Storage Management system (HPSS)
    - – Disk cache systems for data distribution
    - – Parallel linux compute clusters (beowulf)

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

LIGO Data Analysis System Hardware Procurement Overview

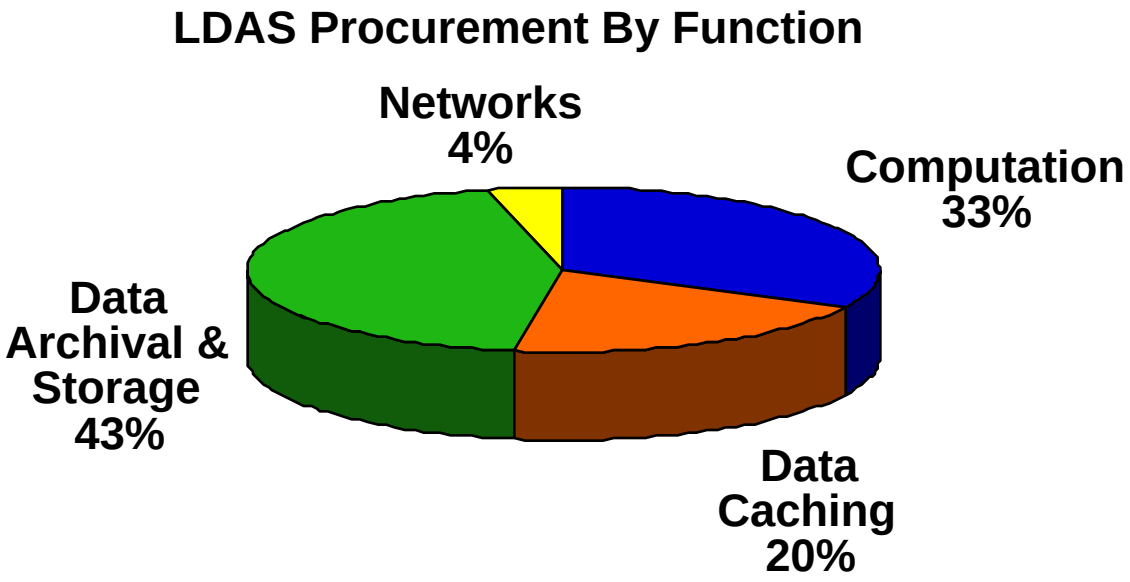

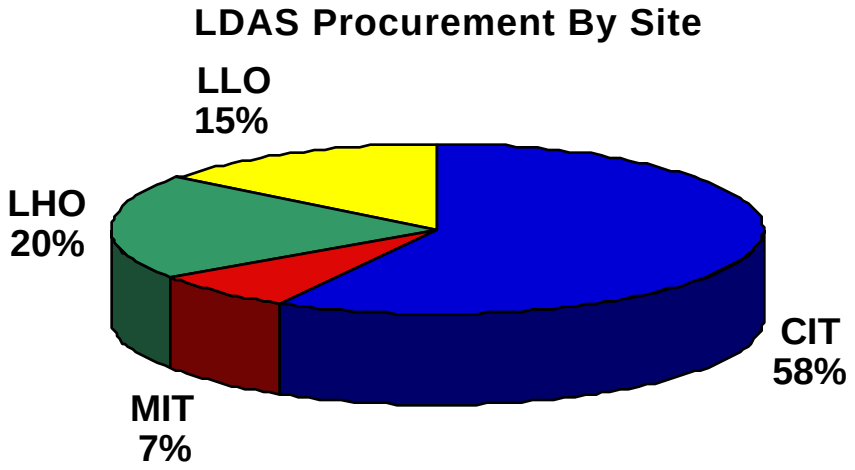

LIGO I Data Analysis System

Procurement Overview

NSF Annual Review of LIGO
9 - 11 May 2000

Albert Lazzarini
Caltech
Pasadena, CA 91125

LIGO Data Analysis System Procurement Plan

- LDAS Procurement Plan follows from the requirements of LIGO data analysis
- The basis for the cost estimate is an example system representing available market options
- Procurement will be phased to match demand
- We will identify the most cost effective options for each phase of the procurement

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, warrantee, 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

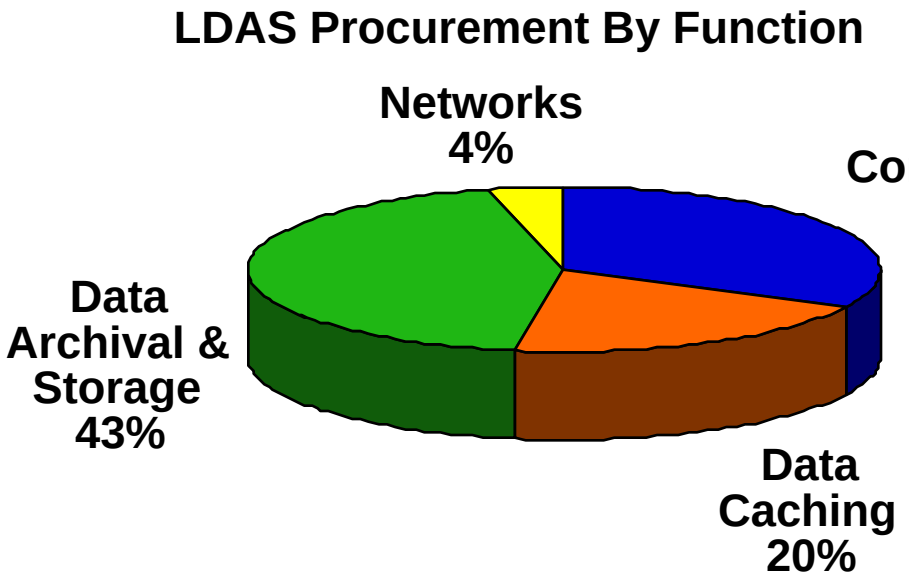

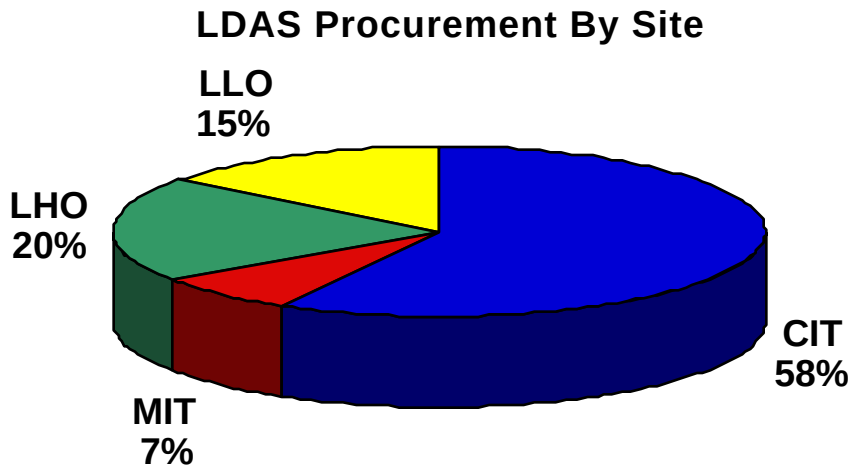

Networking Technology

- 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)

Networking Technology

- Networks:

Component	Quantity	Comments
ATM Switches	1ea ASX 1000[1] 1ea BX200[1] 5ea 2810[1,2]	LANE LANE workstations
Ethernet Switches	8ea 1000BT/100BT	compute clusters

[1]FORE products

[2] Ethernet/ATM

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

- Data Servers:

Component	Quantity	Comments
2 CPU 512+ MB 3x36GB SCSI	8	DB servers, Data servers
4 CPU 1GB RAM 4x36GB SCSI	4	Data servers

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

Mass Storage Technology

- Comparative scale of LIGO I needs

Program	Top Level Requirements
LIGO I	• 100 TB total storage • 0.2 TB per day growth • <1 TB per day transfer • 10^6 files • 5 concurrent transfers @ 1-10 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%

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 casettes 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

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 & 3 tape heads	2	Tape units compatible with HPSS units; placed at sites

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

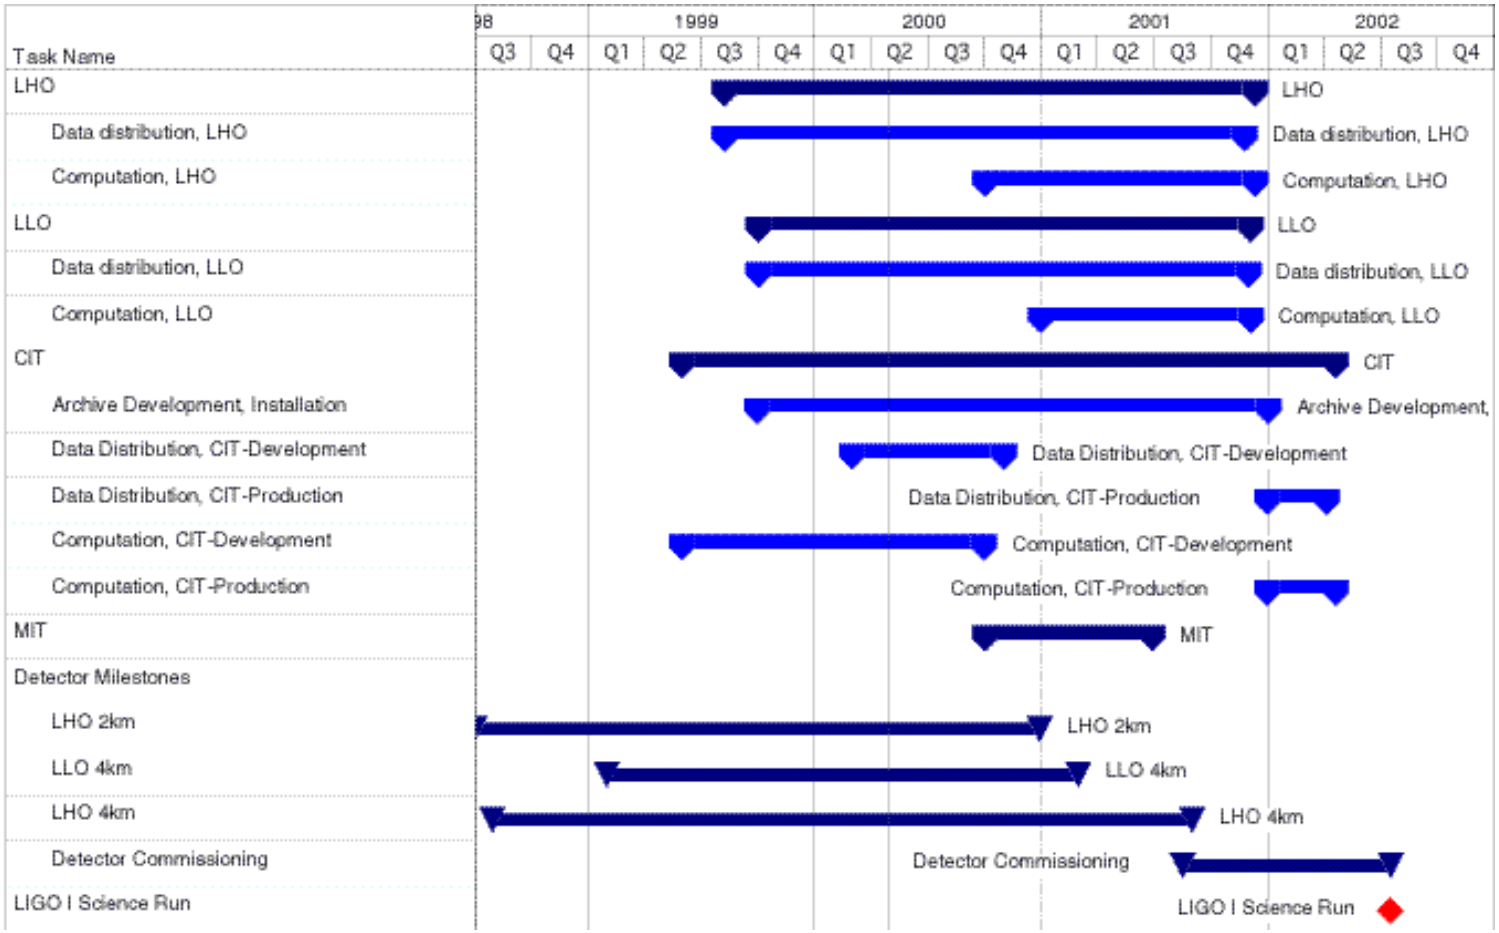

LDAS Expenditure Profile

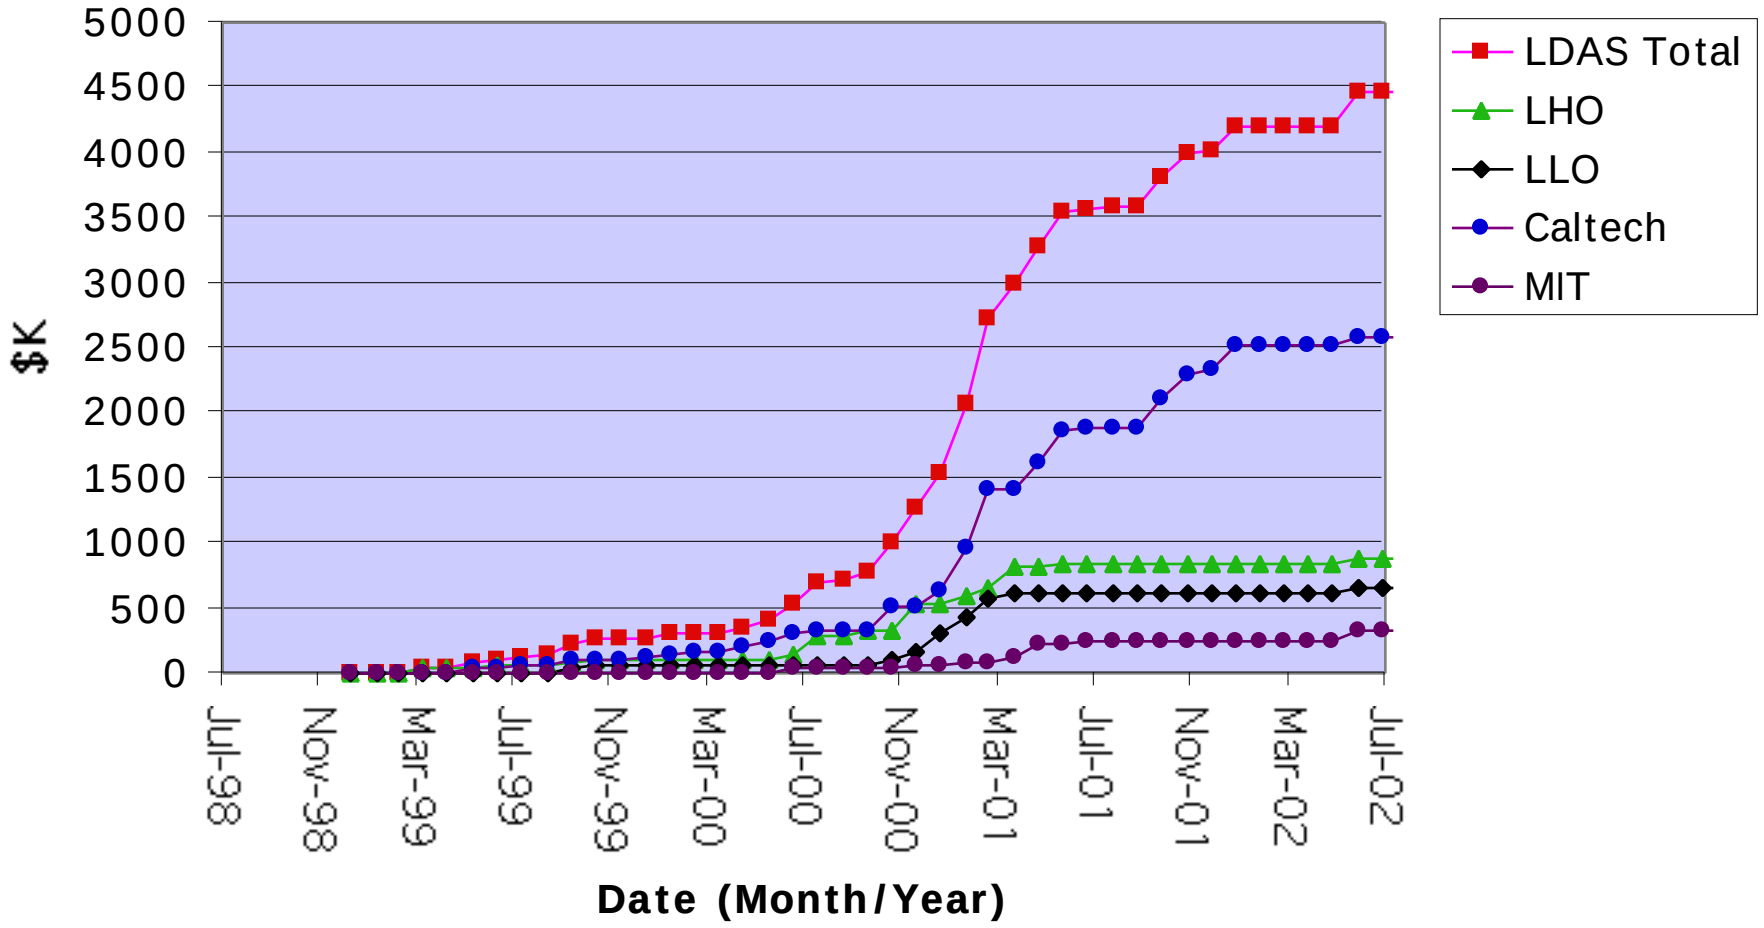

LIGO Data Analysis System Summary

- LDAS 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
- LDAS design is extensible
  - » Allows for expansion, upgrades to support LIGO II when needed

Supplementary Information

Software Block Diagram

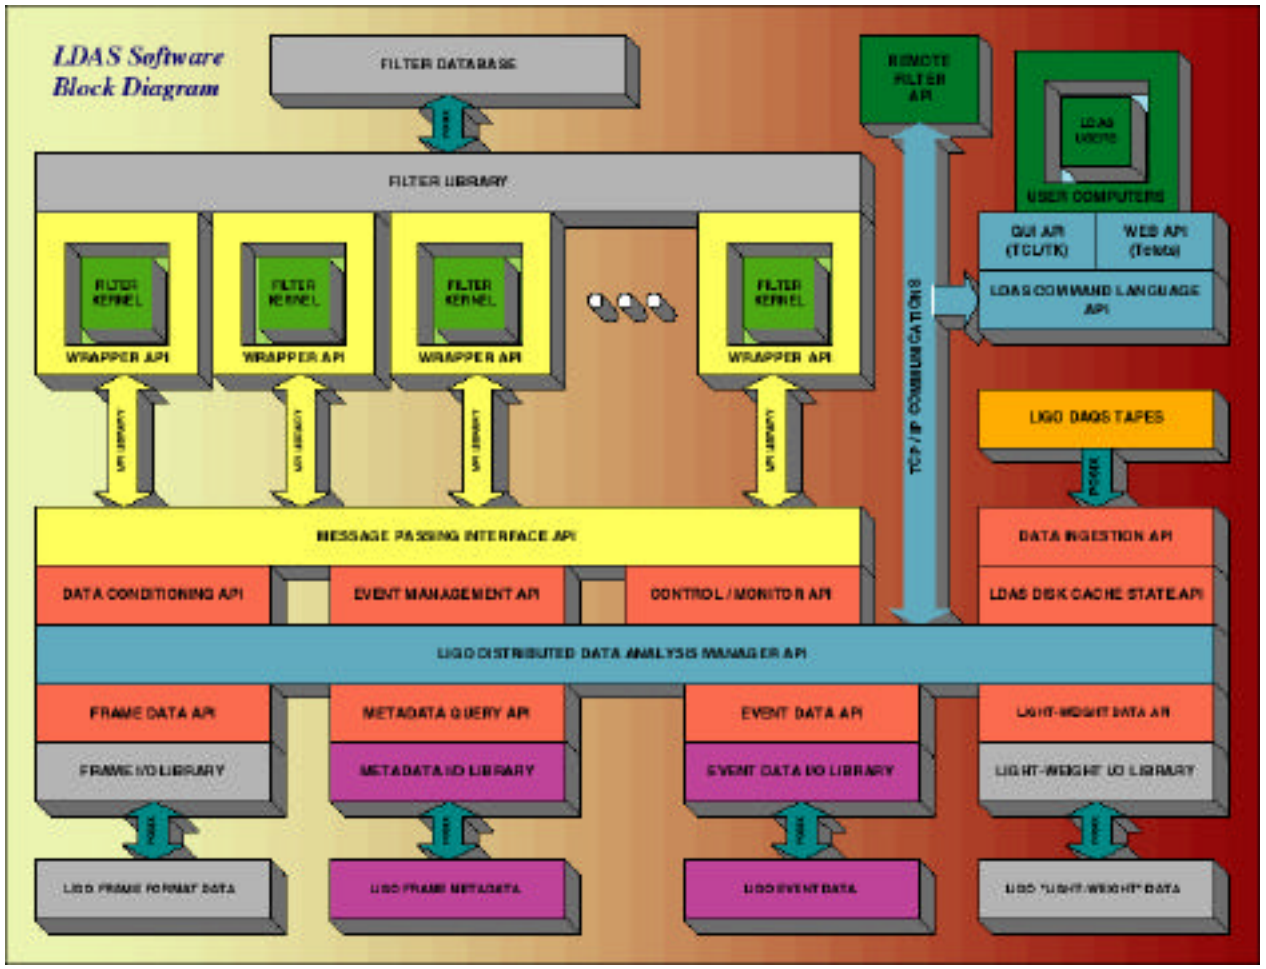

LDAS Hardware installation at Sites

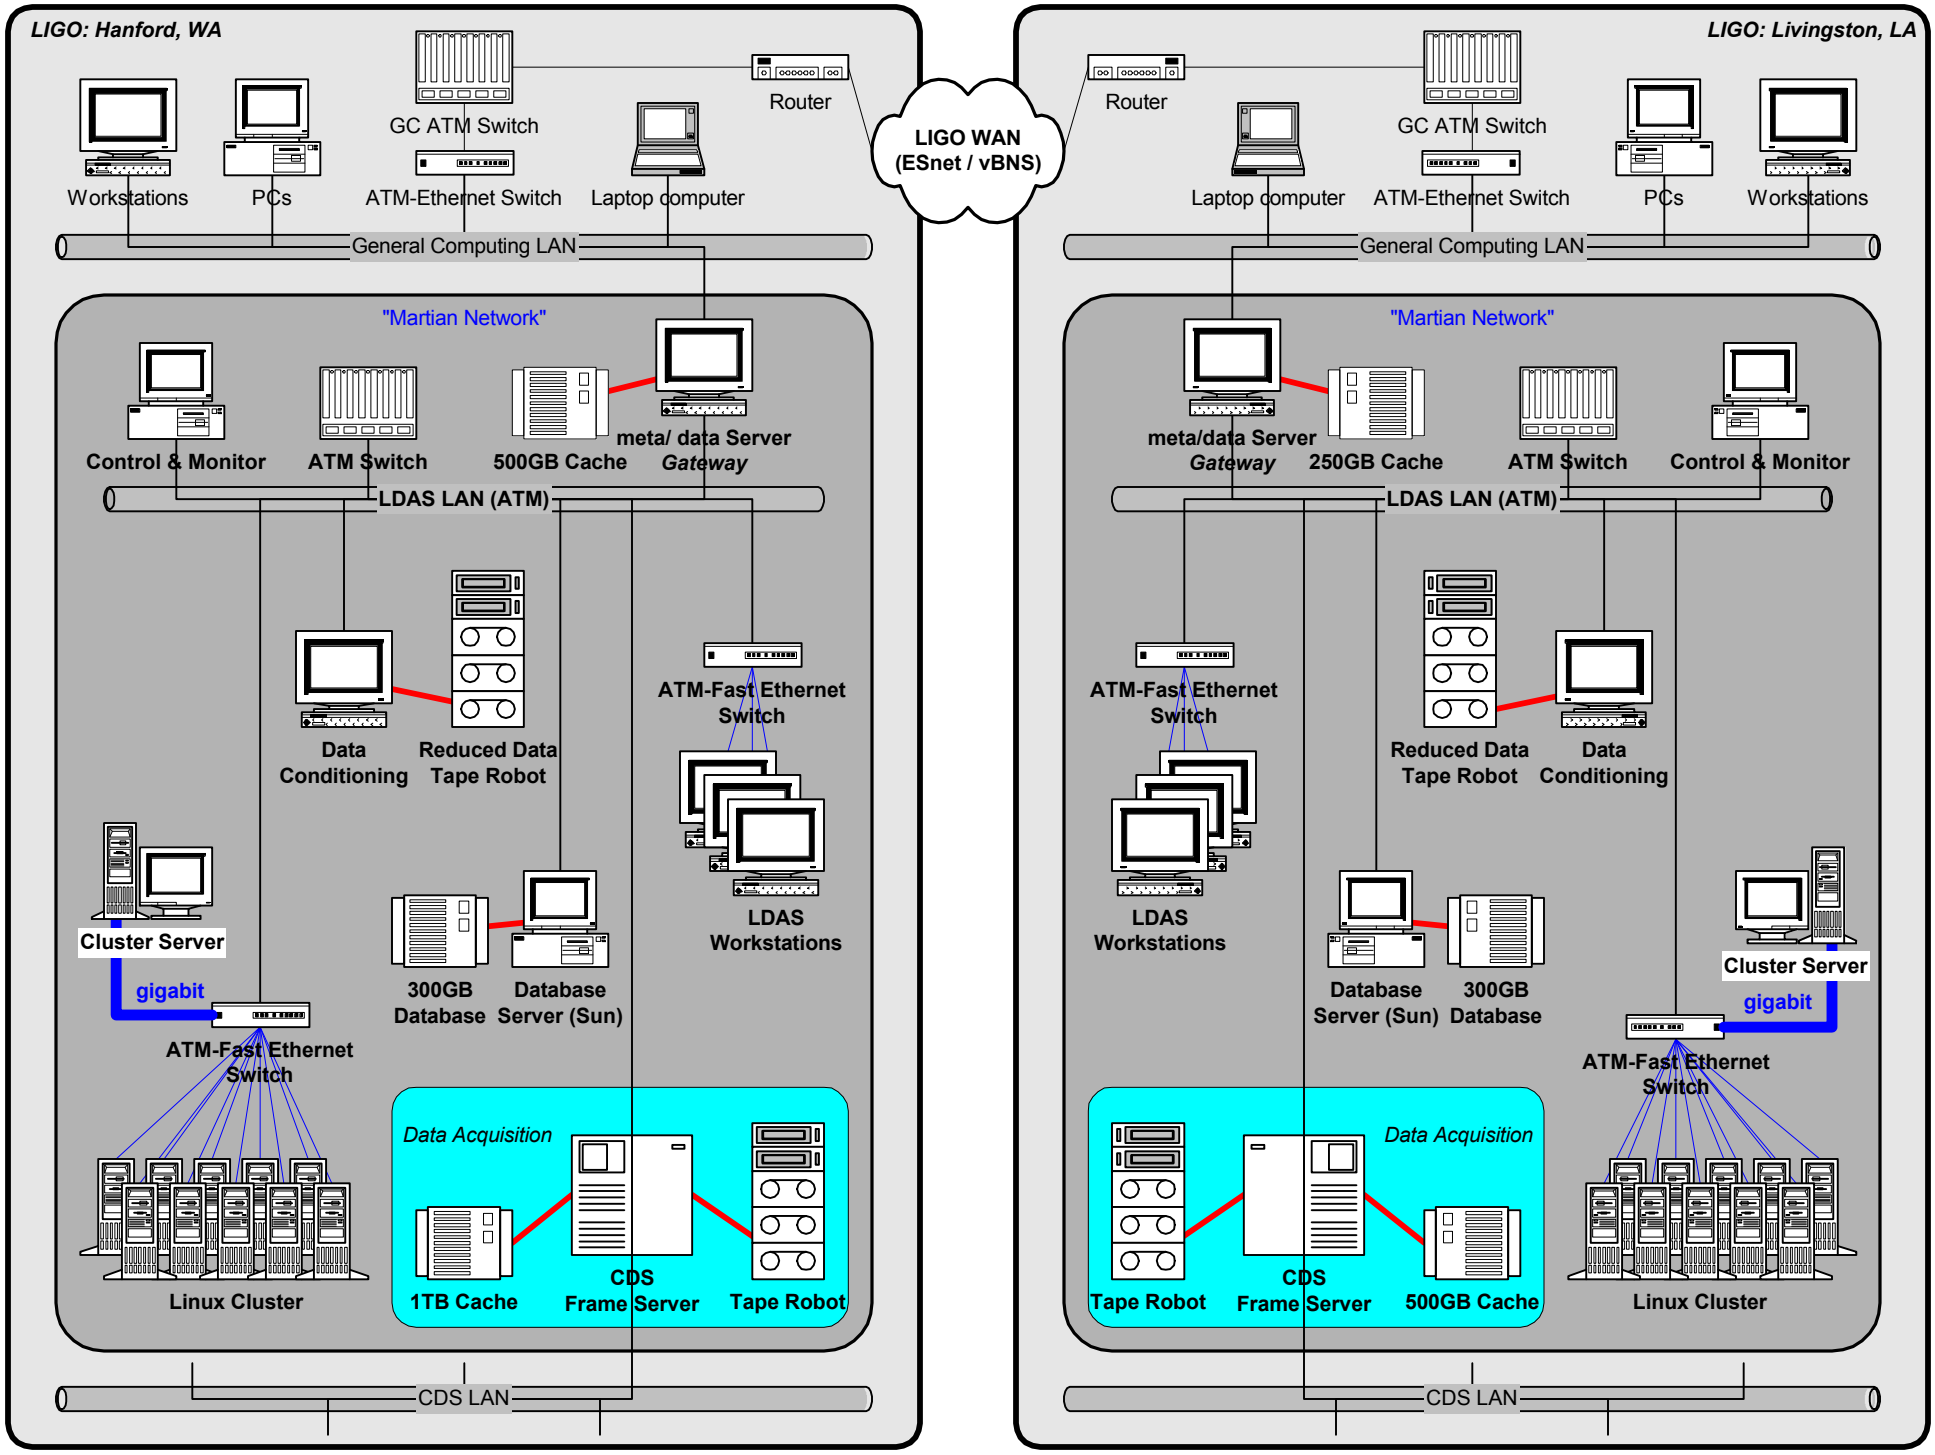

Mapping of LDAS APIs Onto Site Hardware Installations

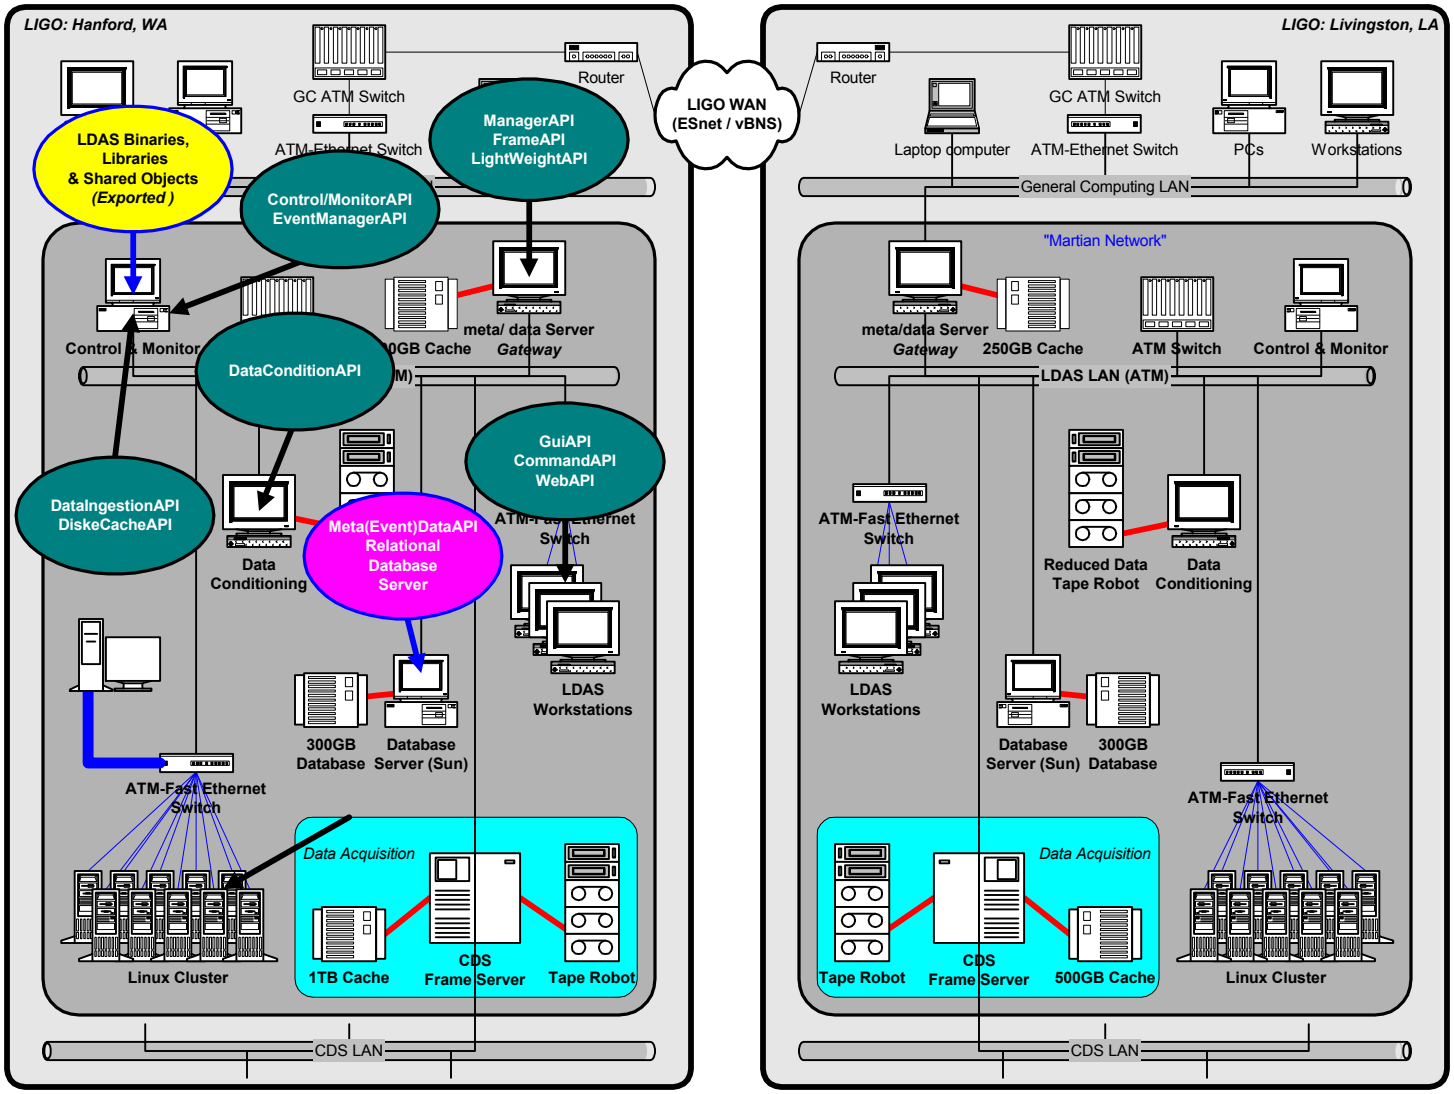

LDAS Functionality Mapped Onto Site Hardware Installations

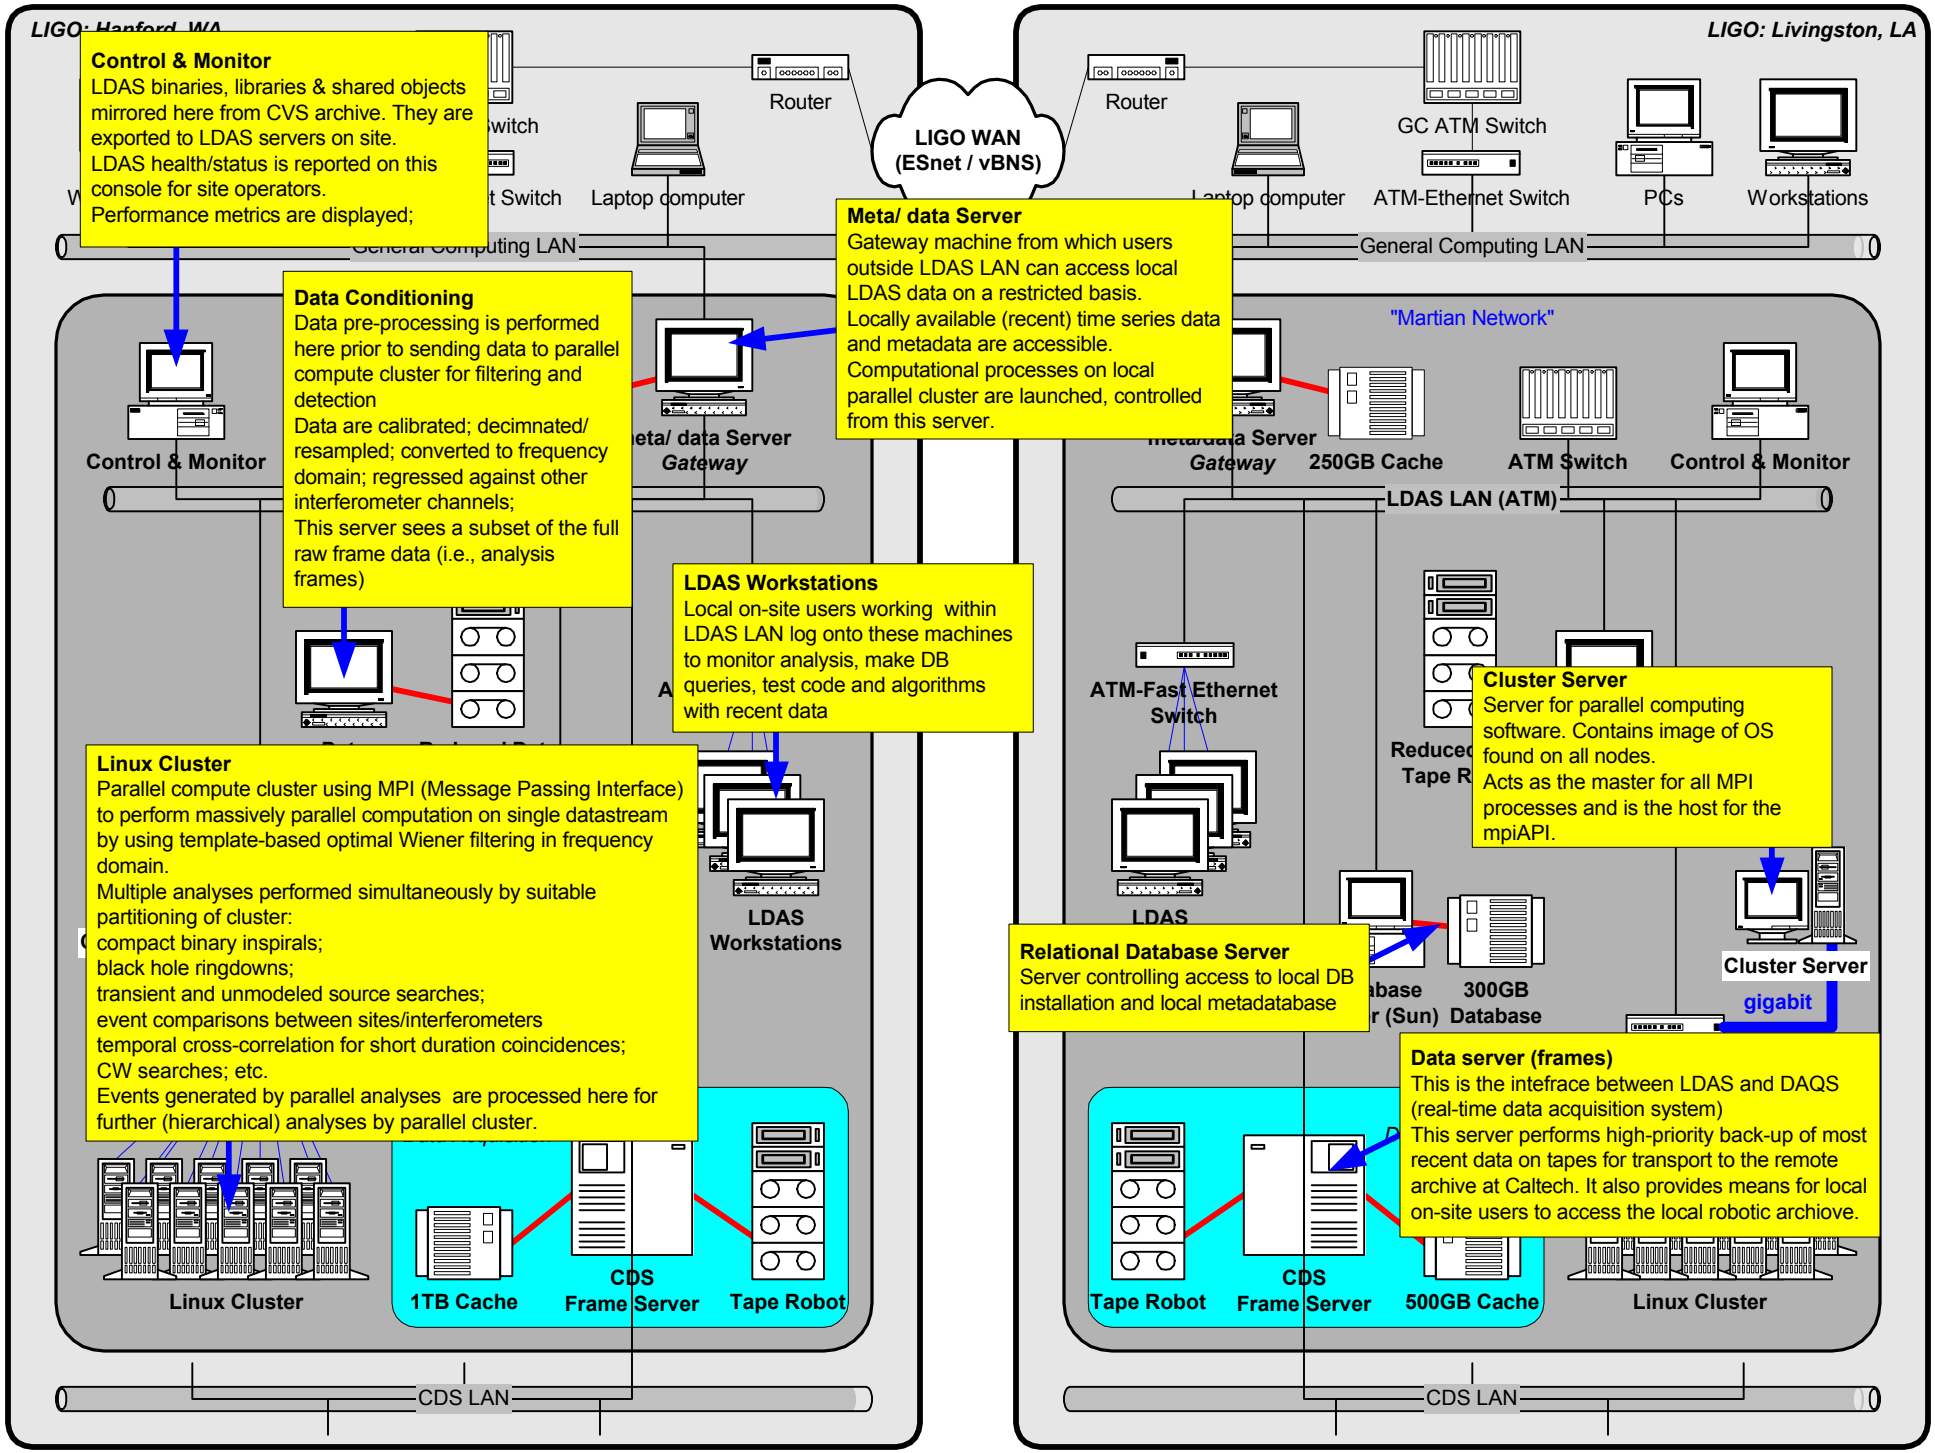

Caltech LDAS Hardware Installation

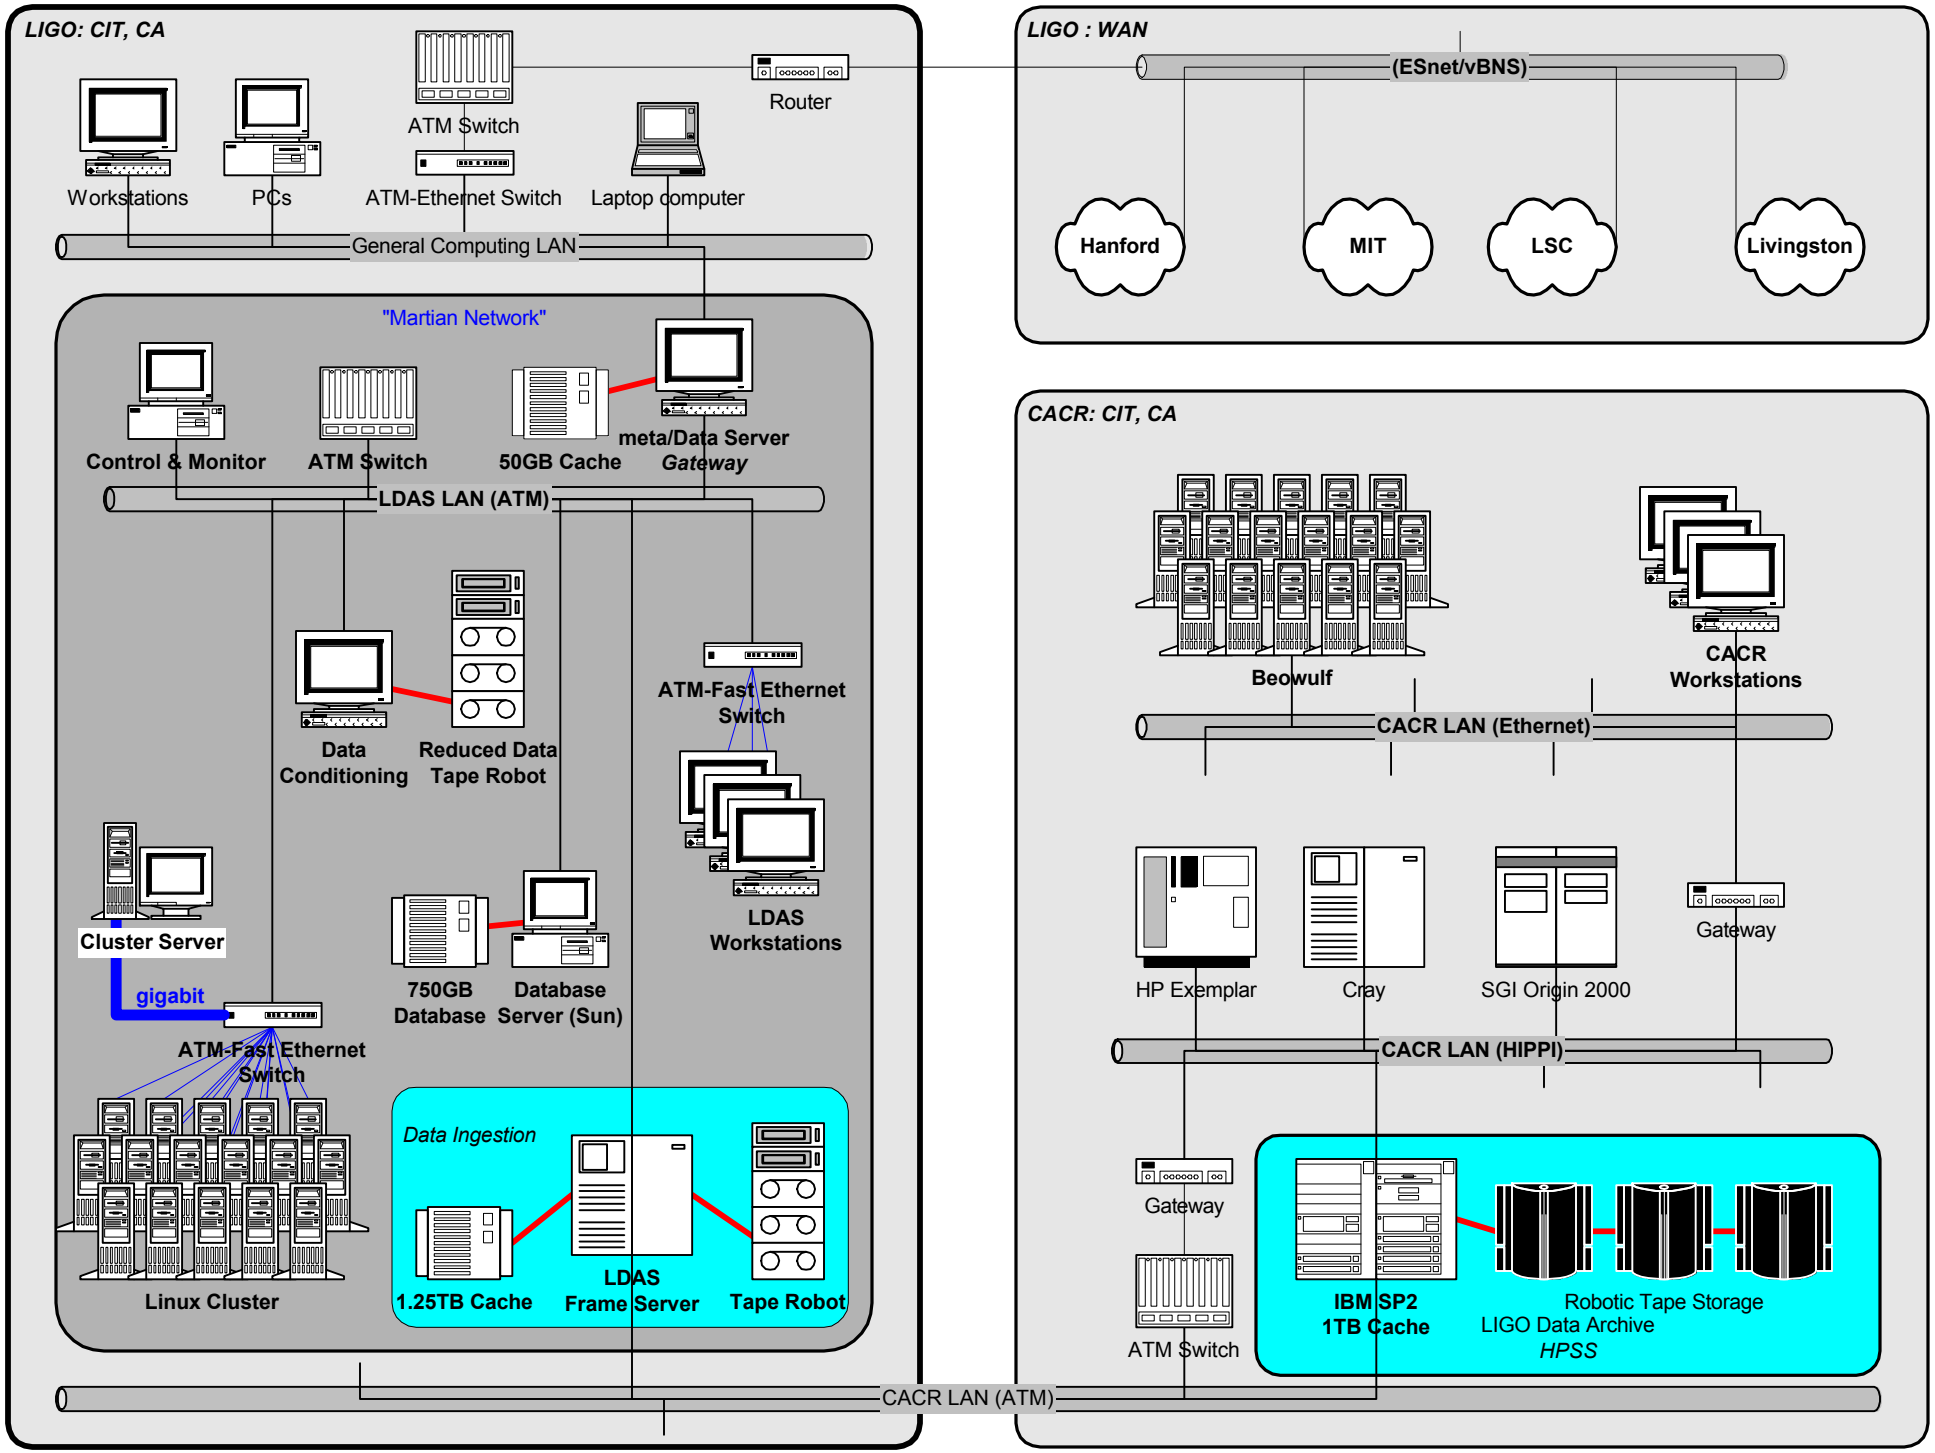

Mapping of LDAS APIs Onto Caltech Installation

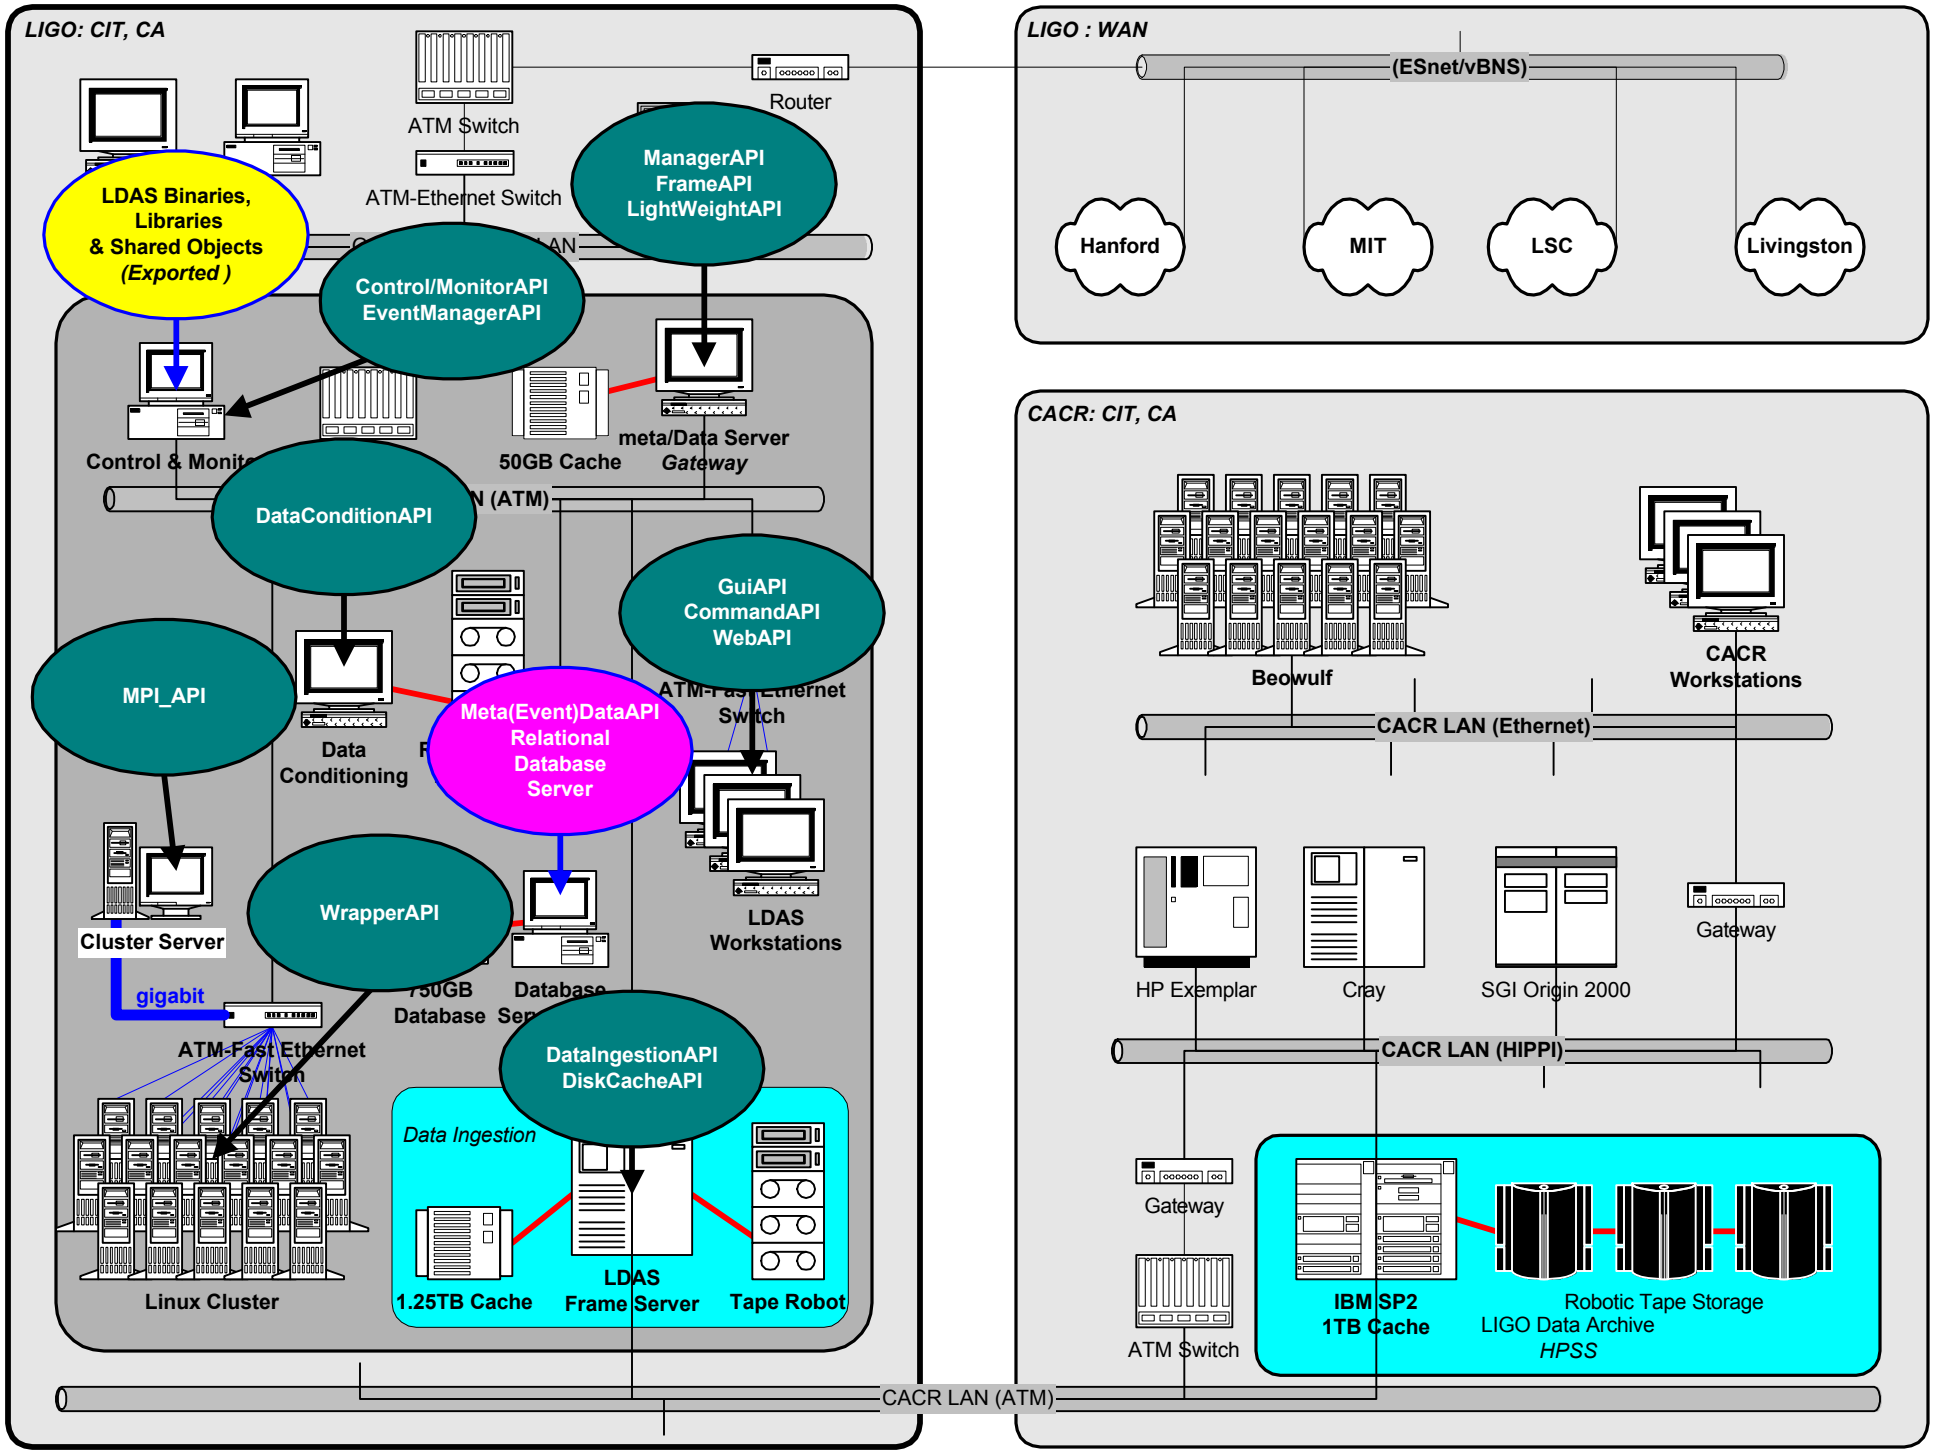

Acronyms used in this talk

Acronym	Description	Acronym	Description
API	Application Programming Interface	LLO	LIGO Livingston Laboratory
CACR	Center for Advanced Computing Research	LSC	LIGO Scientific Collaboration, also
DTD	Document Type Definition, dictionary of valid data for XML parsers	MDC	Mock Data Challenge
FFT	Fast Fourier Transform	MOU	Memorandum of Understanding
FTE	Full-Time Equivalent	MPI	Message Passing Interface, protocol for parallelized computation
GDS	Global Diagnostics System (for the detector)	ODBC	Open Database Convention
GEO	British-German Interferometer Program	OOP	Object oriented programming (design)
GRA	Graduate Research Assistant	PDR	Preliminary Design Review
IFO	Interferometer	TAMA	Japanese Interferometer Project
LDAS	LIGO Data Analysis System	Tcl/Tk	Tool Command Language, a scripting language
LHO	LIGO Hanford Observatory	XML	Extensible Markup Language
		VIRGO	French-Italian Interferometer Project