SLAC OHEP Review

Facilities & Funding for Scientific Computing

Randal Melen

Office of the CIO

February 9, 2011

SLAC NATIONAL ACCELERATOR LABORATORY

Key Messages

- SLAC has an experienced team for efficient operation of major scientific computing systems
  - Sufficient space, power and cooling with good expansion capability in the existing building over the next 4 years
  - Provides additional services in database, cluster design and recommended vendors
  - Partnering with Stanford campus to build a shared research computing data center beyond this point
- SLAC has developed a financially attractive hybrid model for funding scientific computing operations
  - Indirect funds support a base level (80%), while direct costs are passed on to participants in proportion to usage (20%)

SLAC National Accelerator Laboratory

SLAC Computer Center

Building 50 Computing Center

- 4 floors (1st and 2nd floor computer room space and phone switch, 3rd floor offices, 4th floor HVAC and one power substation)
- Second power substation behind building 50
- Some air-side economizer cooling; building uses chilled water, 8 “ pipes, central cooling tower for Computer Room Air Conditioners

Building 50 Computing Center

- Built in 1975 for IBM mainframes, 12“ raised floor
- 1st and 2nd floor about 20,000 sq. ft. of computer room space
  - 3000 sq. ft. available for expansion
  - 8 water-cooled Rittal racks for 344 hottest systems
  - Adding 10 water-cooled racks (up to 430 rack units) late summer for head room

Building 50 Computing Center

- Current maximum server power load is 1.7 MW with expansion capacity of ~1 MW
- Current cooling load is about 411 tons plus airside economizing
- Recent installation of diesel generator for core functions – 500 kW standby, 10 sec. max cutover, ~24 hour (690 gal.) capacity

External Capacity

- Two Sun Modular Datacenter units (Project Blackbox) behind Computer Center with separate chiller, hosting 518 Linux batch nodes
- Off-site Forsythe Hall (SU Campus) hosting 190 Linux batch nodes directly on SLAC network

500kW Diesel Generator

Two Sun Modular Datacenters

Computer Hardware

- Managing about 1,800 Linux batch systems (about 11,000 cores) with 43 public interactive load-balanced “head” systems, part of about 3,100 total managed UNIX servers and desktops
- About 6.31 PB of server disk space (plus 2 PB for LCLS)
- About 4 PB of robotic tape data managed with HPSS (8PB remaining capacity)

Dual SL8500 Mass Storage

Operating System Support

- Support for Red Hat Linux – RHEL 4 and RHEL 5 and a few SL 3 build systems
- System administration tools used across servers and desktops to manage configurations and apply patches
- Use of Red Hat satellite server
- Support contract plus TAM from Red Hat, weekly phone meetings

Batch Workload Management

- All batch capacity managed by Platform LSF
- Dedicated cores/queues for KIPAC (122 hosts with 885 cores) and some BaBar work (230 hosts with 1120 cores)
- Shared queues (about 830 hosts from PPA, 5200 cores) available to everyone using a fairshare algorithm
- A few Windows, MacOS and Solaris batch servers under LSF for batch builds

Batch Fairshares

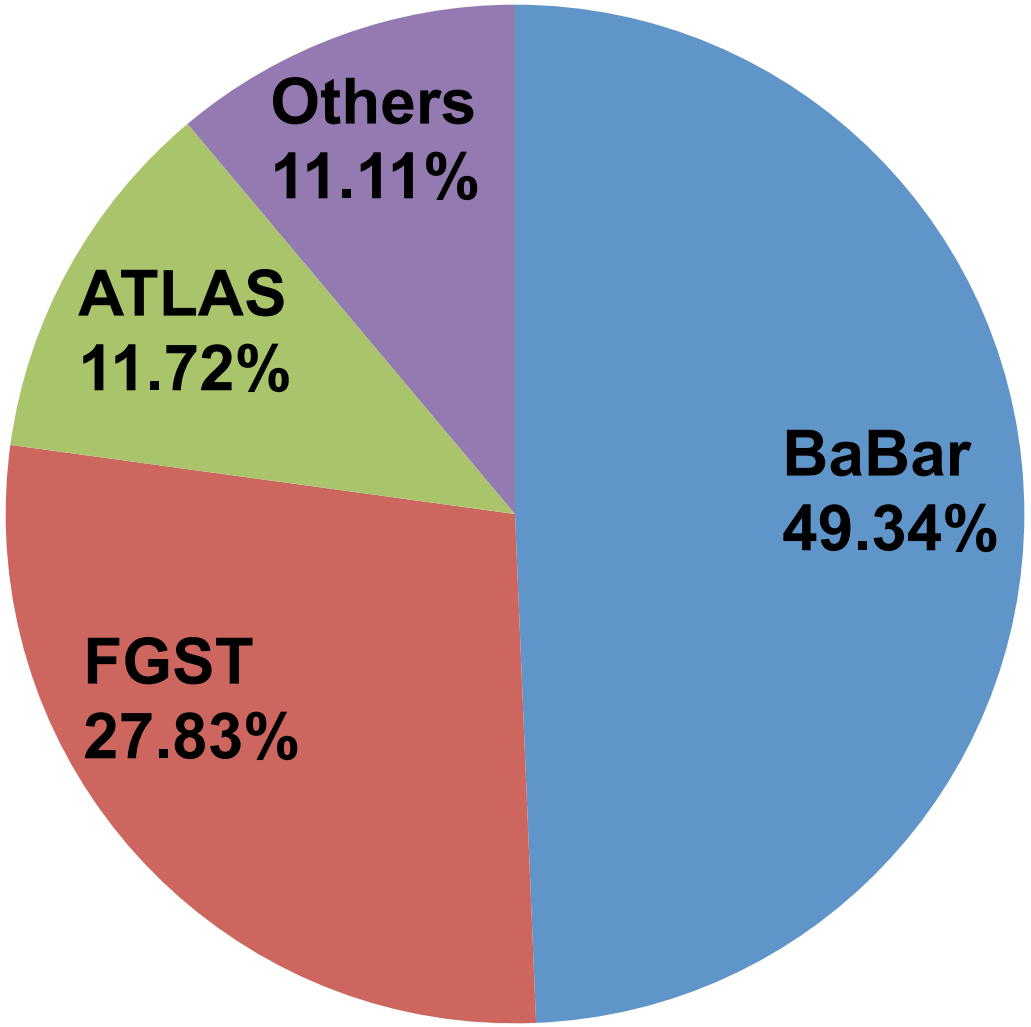

Example: Batch Workload Management

- Shared queues for Fermi GST production
  - Download of satellite data every ~4 hours
  - Current jobs are suspended, ~800 preempting FGST jobs are launched
  - When FGST job completes (~1 hour), suspended job resumed
- No need for a dedicated FGST data processing cluster
- Able to optimize use of resources, including short-term peaks and special needs

Proposed Stanford Research Computing Facility

- Early planning stage for a shared, Stanford-funded data center
  - To be housed on SLAC leased land, and therefore close to Stanford campus
  - About $40M project phase I, Stanford Board of Trustees approval expected late 2011
  - Occupancy likely ~36 months after approval
  - SLAC as a major tenant

Proposed Stanford Research Computing Facility

	Present Building 50	Initial Stanford Research Computing Facility (2014)	Ultimate SRCF capacity
SLAC	2.7 MW	1 MW, 65 racks	3 MW
Stanford	n/a	2.5 MW, 135 racks	7 MW

Proposed Stanford Research Computing Facility

- 16,000 sq. ft. computer space with receiving areas and staging areas
- Full capacity will have UPS and diesel generators for resiliency (e.g., for power line maintenance)
- No current building design done yet

Funding Scientific Computing

- SLAC has evolved from HEP landlord model to a multi-program financial model in 2010
- Created a Scientific Computing Task Force in spring 2010 to address implications for scientific computing
  - Representatives from all science directorates
  - Charged with devising governance for SLAC scientific computing and proposing a competitive funding model

Scientific Computing Governance

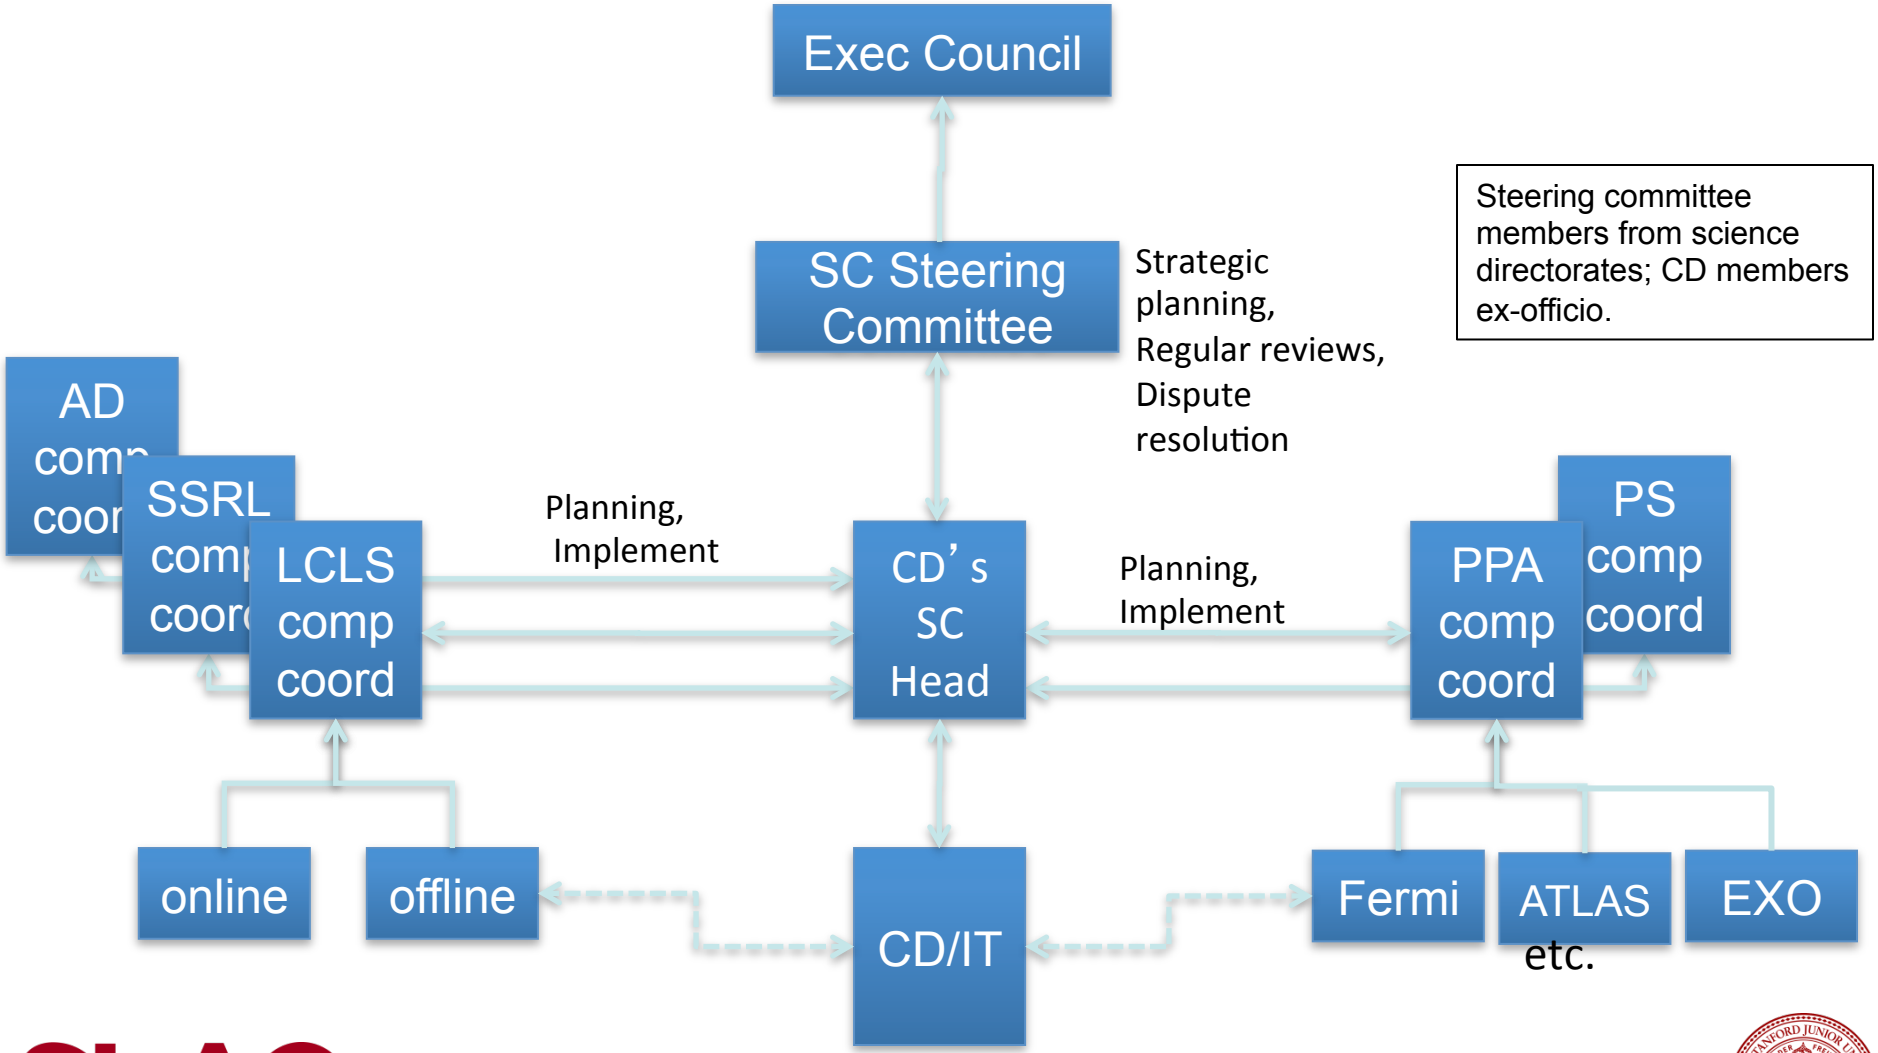

Funding Scientific Computing

- The present scientific computing capacity requires support from 15 FTEs
  - 4 FTEs worth of shared scientific computing services to be capable to minimum support
  - 11 FTEs worth of more specific support
- SLAC senior management agreed to fund the 4 FTE for shared support and the first 5.5 FTEs of the 11 FTEs from indirect funds

Funding Scientific Computing

- Remaining 5.5 FTEs direct charged to users using a recharge algorithm based on:
  - # of batch servers, clusters, file servers, and special servers
  - Some direct charges for special support (e.g., after hours support, premium response time, GPU support)
- Maintenance and other M&S costs also assumed by indirect funding

Funding Scientific Computing

- Net impact: 80% of operations costs indirect funded, 20% direct charged to users
- Using inventory database to track owners, sizes, and services
- Expect billing snapshots every 6 months
- Will evolve over several years to encourage sharing resources across groups, encourage lower power/cooling use

Summary

- SLAC has an experienced team for efficient operation of major scientific computing systems
  - Sufficient space, power and cooling with good expansion capability in the existing building over the next 4 years
  - Provides additional services in database, cluster design and recommended vendors
  - Partnering with Stanford campus to build a shared research computing data center beyond this point
- SLAC has developed a financially attractive hybrid model for funding scientific computing operations
  - Indirect funds support a base level (80%), while direct costs are passed on to participants in proportion to usage (20%)