Yale High Energy Physics Tier 3 Experience

Rochelle Lauer
Yale University Physics
rochelle.lauer@yale.edu

Yale Computer Facility History

- Started $>$ 40 Years Ago
  - There Were No Other Departmental Computers On Campus
  - Our Requirements Are Often Ahead of What Yale Provides
  - Our Requirements Exceed What Yale Provides
  - The Architecture Evolved and Is Evolving
- Shared Facility
  - High Energy Particle Physics
  - Astrophysics
  - Nuclear Physics

Tier3 Decisions

- Integrate the T3g into the existing facility
  - Use the existing servers for NFS access to data
  - Use the existing storage subsystem this year.
    The existing storage subsystem can provide at least 20 TB of storage for ATLAS in this fiscal year.
  - Purchase a compatible "compute farm"
- Create an Independent T3g
  - Purchase of compute nodes, servers and storage
- Use Yale's High Performance Computing (HPC)Facility
  - Cost and storage capabilities were undetermined

$\Rightarrow$ The decision was to integrate into the existing facility

Configuration Before Adding Tier3

- Servers
- Storage
- Linux Desktops
- Linux Blade Compute Farm
- Network
- Platform LSF

⇒ **Tier3 Ready**

- All infrastructure is in place
- Only compute nodes are required

The Existing Infrastructure

Servers – 2 Node HP-UX Serviceguard Cluster

⇒ NFS

- External Logins, Web Server, Samba ...

Why HP-UX ?

- Stable conservative supported OS
  - With stable conservative supported Cluster Software
- Stable supported volume manager/filesystem (VxVM/VxFS)

  ⇒ Support for very large filesystems. Needed by Astrophysics several years before LHC Requirements.

Serviceguard

- HP-UX/Serviceguard Is An Established Product
  - Still Being Enhanced
  - Supported
    $\Rightarrow$Upward Compatibility Is Provided
- Failover Model
  - Load sharing when all nodes are up
  - Failover to a good node if primary node is down
- Specialized software (serviceguard packages) for clustered NFS (Samba, Apache)

Storage
Fibre Based Storage Area Network (SAN)

- Controller(s) + Disk Enclosures in a rack
- Easy Incremental Expansion
- 2 racks with SCSI Based controllers
  - SCSI + SATA Disks
  - SATA disk enclosure translates SATA to SCSI
- Connected to Servers Via Fibre Based Storage Area Network(SAN)

New Storage Subsystem
Already in place With room for ATLAS

- Dual Controllers
- Fibre Backplane
- Fibre + FATA Disks

⇒ "Virtualized" storage

- FATA disks are currently 1 TB
- Two fully configured subsystems fit in 1 Rack (192 disks)

Linux

- ∼40 SL4 and SL5 Desktops

- Blade Compute Farm (pre Tier3)
  - 20 single core, dual processor blades
  - SL4, SL5
  - Astrophysics processing + Occasional Monte-Carlo and Analysis

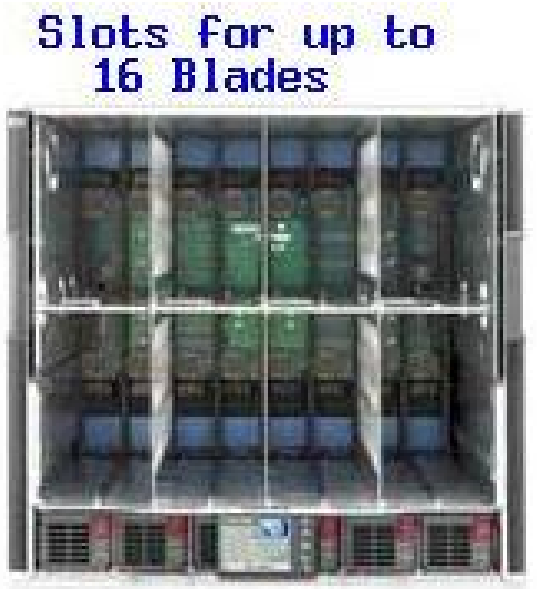

Network

Local Area NetworkLAN

- Managed Locally
- 1Gb and 10Gb Connections
  - 1Gb on desktops and blades
    10Gb between switches

$\Rightarrow$ Single 1GB uplink from blade enclosure(s).

- Room for Tier3 uplink(s) on switches
  Either 1GB or 10GB connections

Wide Area Network (External Network)

- 1GB Connection to Yale Network

LSF

- Simple batch was needed for system management
- Simple batch was needed for long jobs and for the blade cluster
- A supported multi-platform solution was required (Tru64, Linux, HP-UX)
- LSF was being used at Fermi and BNL

- LSF is pricey (but not anymore)
  - Various license types
  - ATLAS discount
- LSF is overkill (we use a very small subset of features)
- LSF is serving us well

Tier3 Blade Testing

With the existing infrastructure

- We purchased an AMD dual processor blade for testing. Six cores/processor, total of 12 cores per blade.
- HP provided an Intel dual processor blade for testing. Four cores/processor, total of 8 cores per blade.

- We obtained technical details and quote from Dell
- We obtained quotes from HP

Blade Tests (Summary)

- ROOT and Athena jobs
- 12(8) parallel jobs with data on local disk
  - Performance degrades to unusable with 12/8 jobs
  - SAS disks a bit better than SATA
- 12(8) parallel jobs using NFS
  - Better (and acceptable) performance
    15 - 20% degradation
  - How will it scale to 80+ jobs ?

The Tier3 Purchase Configuration

The Compute Farm

- Blade Infrastructure
  - One Blade Enclosure
  - Management modules (redundant)
  - Power infrastructure (redundant)
  - Internal networking (redundant)
- Blades (10 blades, 80/120 cores)
  - 10Gb network connections ?
- UPS

- This configuration is expandable to 16 blades with the one enclosure.
- And it is expandable to at least 48 blades with the purchase of the 2 Blade enclosures.

HP or Dell

- Very similar blade architecture
  - Except for 10Gb blade network capability
- Pricing was very similar (with DELL ATLAS discounts)
- HP is more compatible with our existing infrastructure

With HP

AMD or Intel ?

- Both have 10Gb uplink
- Intel blades have 10Gb local ethernet
  - And 10Gb internal switch option

HP Purchase

AARA Funds

- **Compute Farm**
  10 Intel 8 core Blades (E5530)
  – Choice based on 10Gb network capability
- **Blade infrastructure**
  One blade enclosure, power shelf and redundant power, redundant management modules, redundant internal network switches, UPS

Computer Facility Funds

- **Disk Enclosure and disks for existing SAN**
- **HP-UX Server to add to the Serviceguard cluster with 10Gb network interface card**
- **10Gb network interface card for existing server**

Implementation

- Up and running in 2 days after hardware delivery
- Kickstart network install (SL 5.4)
- Add `tier3_long`, `tier3_short` queues to LSF
- Queue access only to those in `atlas` group
- Simple fairshare scheduling

  `tier3_short` has a slightly higher priority
- Limit on the number of jobs per host
- Software (Athena, globus, etc.,) was already installed

$\Rightarrow$ The Yale Tier3 is being used as we speak !

Local Issues

- What happens to performance as the number of parallel jobs increases ?
- Is this simple queue scheduling adequate ?
- How to manage disk space ?
  - Reliability versus ease of use
    Multiple data areas versus one very large (>40TB) data area
    * HP-UX can handle very large filesystems but... recovery from failures is questionable.
  - Users do not delete unwanted files.

After the Hardware is Up and Running...What Next ?

Tier 3 Issues

Rochelle's Random Tier3 Issue List

- Understanding the ATLAS software environment
- Keeping software up to date
  Athena, globus (glite-UI), pacman, dq2, etc.
  - No defined consistent methodology (that I know of)
  - So far ... reactive system management
- Local requirements conflict with "required" software
- Unstable software environment
  $\Rightarrow$ A consistent environment for users is necessary
- Changes should be upward compatible
- More than one method to do the same thing
  (e.g., xrootd, storage element ...)

$\Rightarrow$ 64 bit machines/32 bit software

This Group Can Help

Rochelle's Atlas Tier3 Group (Random) Wish List

- Minimize the complexity and amount of software T3g sites are required to install
- Keep the software environment stable
- Provide documentation for remote system managers
  - Twiki is not sufficient (my opinion)
- Help with keeping ATLAS software up to date
- Provide input to the developers

Thank You