University of Mississippi Computing Farm

An Example of a Small Tier-3

David A. Sanders

University of Mississippi

OSG Site Administrators & CMS Tier 3 workshop
ACCRE, Vanderbilt University
August 10, 2010

Outline

- University of Mississippi High Energy Physics Group
- Goals and Requirements of a Small Tier-3
- Computing Farm
  - Infrastructure
  - Hardware
  - Software
- Monitoring
- Future Plans

UMHEP

The University of Mississippi High Energy Physics Group

Faculty

Lucien Cremaldi, Rob Kroeger, Breese Quinn, and Don Summers

Postdocs/Research Scientists

Terry Hart, Alex Melnitchouk, Lalith Perera, Rahmat Rahmat, and David Sanders

4 New Students

Experiments

Current:	BABAR, CMS, D0, MICE
Past:	E791

Goals and Requirements

Goals

- Setup a CMS Tier-3 Computing Farm that requires minimum operator intervention.
- Service our other VOs (D0 and LIGO).
- Allow our scientists to access the rest of the Grid.

Requirements

- Must be easy to setup and configure.
- Minimal Maintenance.
- Simple Monitoring.
- Easy for physicists to use.

Computing Farm

Status and Usage

- We have a 25 node (48 cores WN) Linux computing Farm used by D0, LIGO and CMS.
- We have installed the Open Science Grid (OSG) software and are operating smoothly on the GRID.
- We are members of the DOSAR (Distributed Organization for Scientific and Academic Research) group.
- Our small cluster is also used for PYTHIA level generation of black hole (BH) events and some Geant4 work, as well as other departmental work.

Computing Farm

Physical Layout

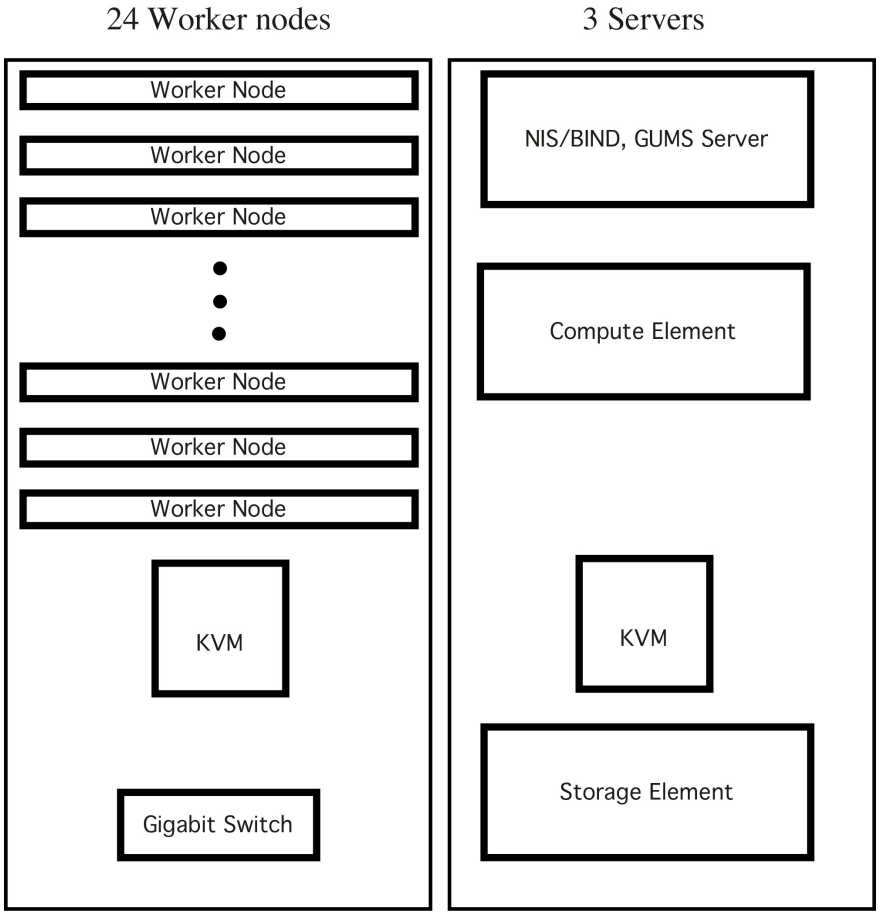

Infrastructure and Support

Infrastructure

- Computing Room with independent Air Conditioning.
- 18 KVA Building UPS for the Computing Farm.
- Gigabit Ethernet connection to Grid.

Support

- One person (Me) ~10% FTE.

  (Plus lots of help from THIS group and DOSAR.)

Hardware

- Main Server (OSG CE and Condor):
  Dual-single-core 1.6 GHz AMD Opteron, 4 GB RAM
- Clients (OSG WN): (48 core total)
  24 Dual-single-core 2 GHz AMD Opteron, 2 GB RAM
- Storage (OSG SE):
  8 TB RAID5 Disk Array (24 X 300 GB SATA disks)
- NIS/BIND and Local /home Server:
  Dual-single-core 1.6 GHz AMD Opteron, 4 GB RAM
- Interactive Node:
  Quad-core 2.66 GHz Intel Q9400, 4 GB RAM

Software

Basic

- Scientific Linux 5.4
- Condor 7.2.4 (Installed stand-alone)
- Use NSF and Autofs to mount remote volumes
- OSG 1.2.6 on CE (1.2.3 on WN, 1.2.8 on SE)
- Currently use grid-mapfile for authentication
- CMS software (CMSSW) is automatically installed
- On the Storage Element (SE) we have installed BeStMan v2.2.1.3.8 and PheDEx v 3.3.1
- Monitoring: MonaLisa, Ganglia, Gratia, RSV and SAM(?)

Software

Local Customizations/Features (quirks)

- No SELinux – Caused problems with Networking and NFS
- No IPtables – Using a strict hosts.allow and hosts.deny
- Use NSF and Autofs to mount remote volumes
- Condor 7.2.4 installed stand-alone with CE and WNs sharing single common file area
- OSG-WNs also share common install but with local OSGTMP
- Uses “condor_userprio” to set priorities for local users and VOs
- Farm is on its own subnet
- Use a BitSight temperature sensor via snmp to protect from AC failure

Monitoring

MonaLisa

Ganglia

CPUs Total: **48**
Hosts up: **24**
Hosts down: **0**

Avg Load (15, 5, 1m):
84%, 85%, 85%
Localtime:
2010-08-06 12:18

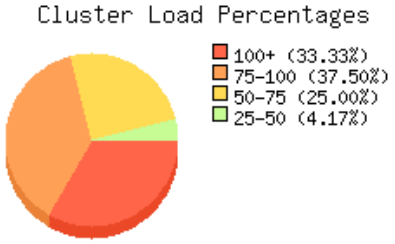

Monitoring

Gratia Accounting

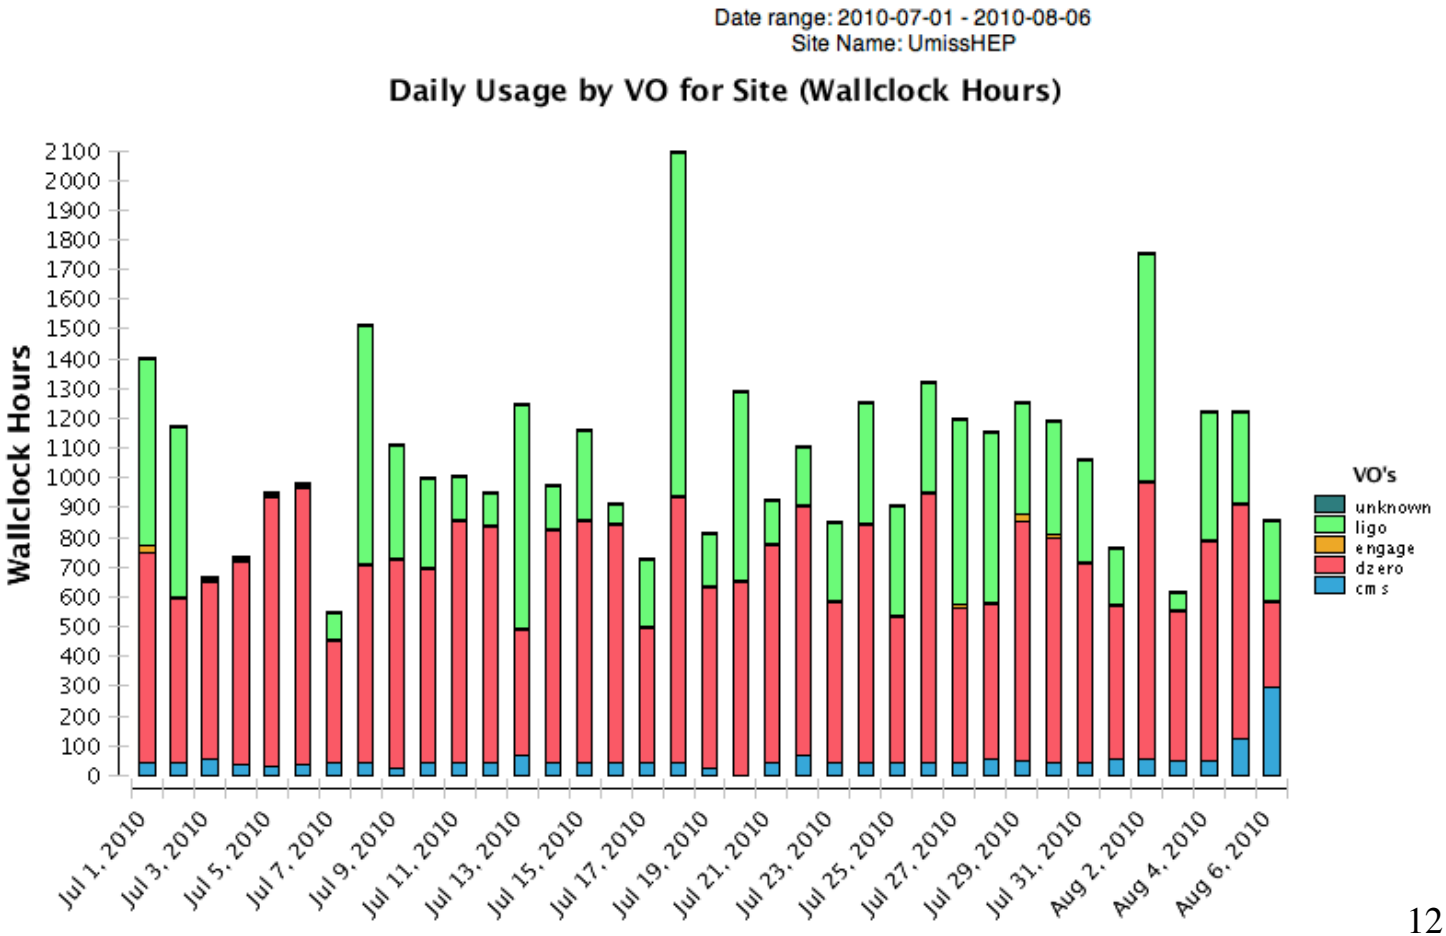

Short-Term Future Plans

- Complete the PheDEx installation and testing so we can become fully integrated into the LHC Computing Grid Project.
- Test new Interactive Node an OSG-Client. (CRAB and the gLite-UI ?)
- Setup a “CMS Center” for remote Computer Shifts.
- Add 2 to 4 more nodes that are now dual 6-core (or 12-core) CPUs, depending on their power consumption. This will double or triple our computing power.

Long-Term Future Plans

Upgrade our Computing Room infrastructure.

- This would mean doubling the capacity of both the Uninterruptible Power Supply and the Air Conditioning of our computer room.

Upgrade our 24 worker nodes with dual 6+ core CPUs. Upgrade the disks in our SE.

- Using 6-core CPUs we would increase our computing power by a factor of 6.
- Going to 2 TB Disks would give us 50 TB of Storage.

Summary

Hardware:

- 24 clients (48 cores) 2 GHz Opterons
- 8 TB RAID5 disk array storage
- 2 Servers (CE and SE)
- Gigabit switches/links locally.

Software:

- OSG working – v1.2.3+
- Condor – v7.2.4
- Scientific Linux – v5.4
- BeStMan and PheDEx on SE