Caltech Site Report
Mar. 3 2009

Michael Thomas

Facility Photos

Harvey Newman
Julian Bunn
Azher Mughal
Dorian Kcira
Michael Thomas

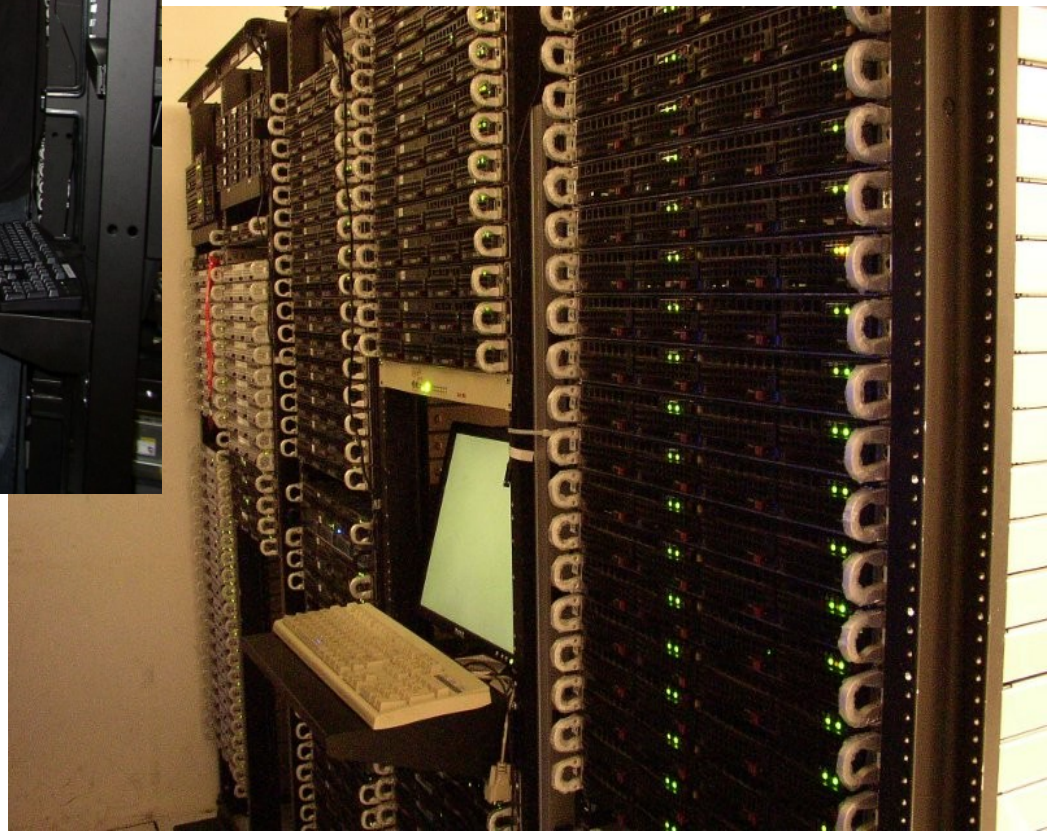

Current Hardware Status

Compute nodes:

- 26 x Opteron 275, 1TB dCache pool
- 35 x Intel 3.0GHz dual-core
- 24 x Intel 2.33GHz quad-core
- 6 x Intel 2.5GHz quad-core
- 458 batch slots, 1120 kSI2k
- +512 batch slots, 1075 kSI2k “RSN”

Storage nodes:

- Pools on all worker nodes (software raid-0)
- Pool sizes vary from 1TB to 3TB
- +2 Sun x4500 Thumpers, 44TB each (Solaris + ZFS)
- +4 2U whitebox disk servers, 8TB each
- +5 4U whitebox disk servers, 16TB each
- 276TB usable space, +208TB “RSN”

Current software status

Most nodes running 64-bit CentOS 4.7 (Rocks 4.2.1)

- Rocks headnode, pnfs server still 32-bit

OSG 1.0.0

Gums 1.2

PhEDEX 3.1.3

Frontier 4.0rc6

Dcache 1.8.0-12

Public IPs on (almost) all worker nodes

Performance Plots

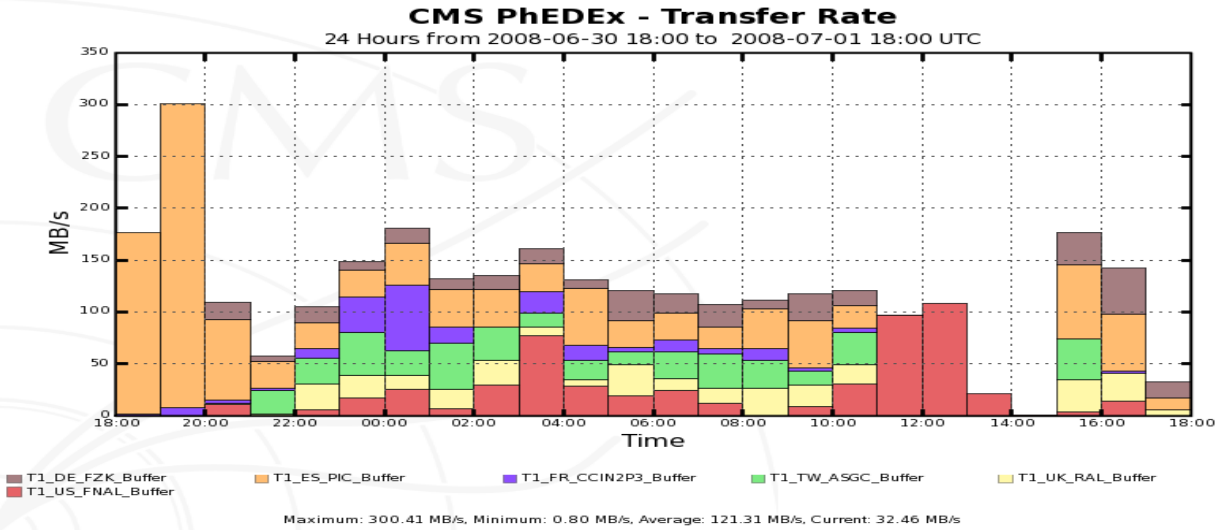

PIC @ 300MB/s

FNAL @ 300MB/s

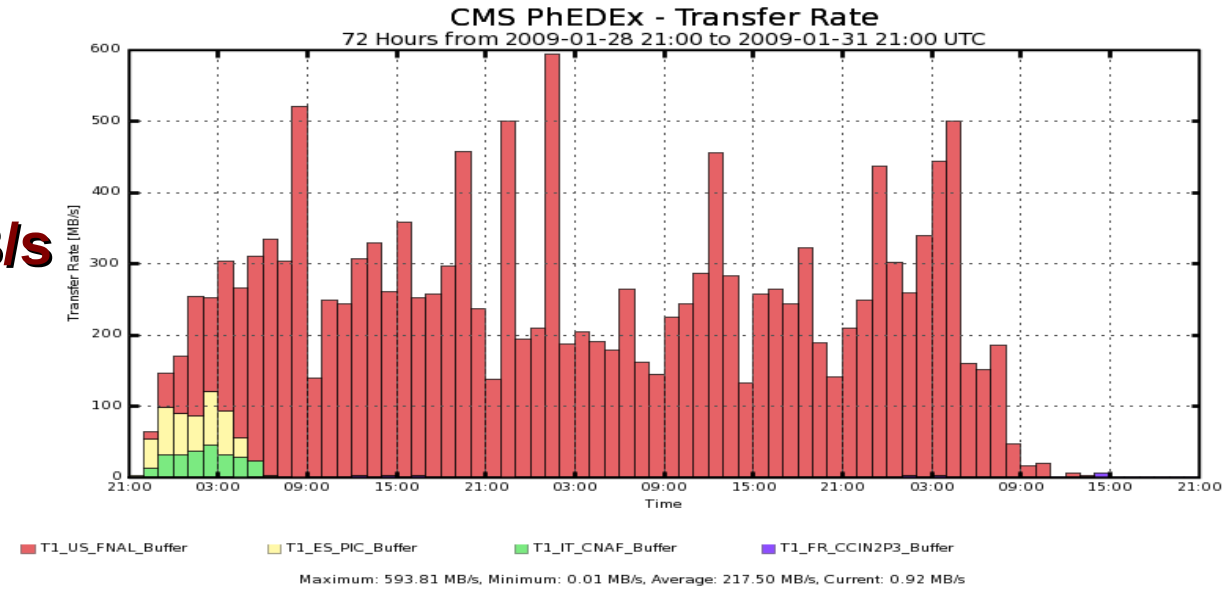

Plans for '09

Replace older servers

- Opterons go to user analysis cluster
- 5+ year-old Xeons in user analysis cluster go to e-waste
- Purchase 13 new Supermicro Twin to replace # cores
  - +26TB due to larger disks
- Purchase additional disk storage on-demand
  - Individual disks vs. additional servers

CACR facility upgrade

- Expand floor space by removing unused offices/storage space
- Add new cooling equipment
- Increase power capacity
- ...may be delayed due to campus cost reductions

Software Plans for '09

Update cluster to CentOS 5.2 (Rocks 5.1)

- 64-bit on all nodes
- Enable selinux (permissive)
- read-only $OSG-APP on worker nodes
- Fix raid-on-reinstall bug
- Update to Condor 7.2
- Rename workers to reflect physical location
- Fewer public IPs
- Beefier hardware for service nodes
- Pioneer SL5 support
- Backup $OSG-APP nfs server via rsync
- Shorter $OSG_WN_TMP path for madgraph
- Kernel netdump to NFS server
- Update to dCache 1.9.2?

Other developments

FDT/dCache

- dcap support added to FDT
- ~4Gbps between FNAL -> CIT

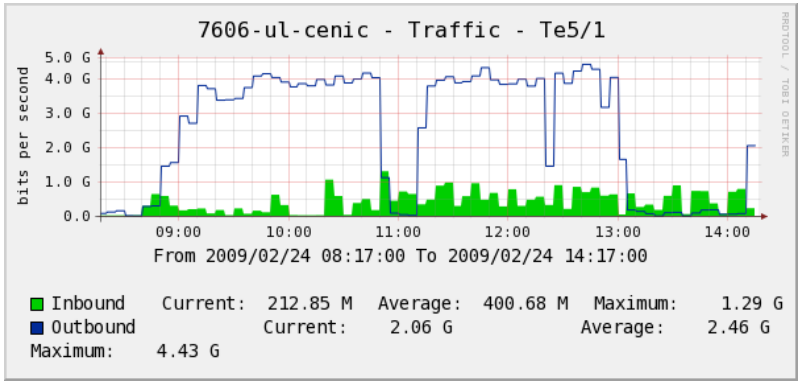

DBS Query Tool

- Help physicists discover new data globally and at specific sites

Frog Service

- Web service for generating event display
- CMSSW converts .root to .viz, frog converts .viz to .png

How we do it

Lots of monitoring

- Nagios
- MonALISA
- logwatch

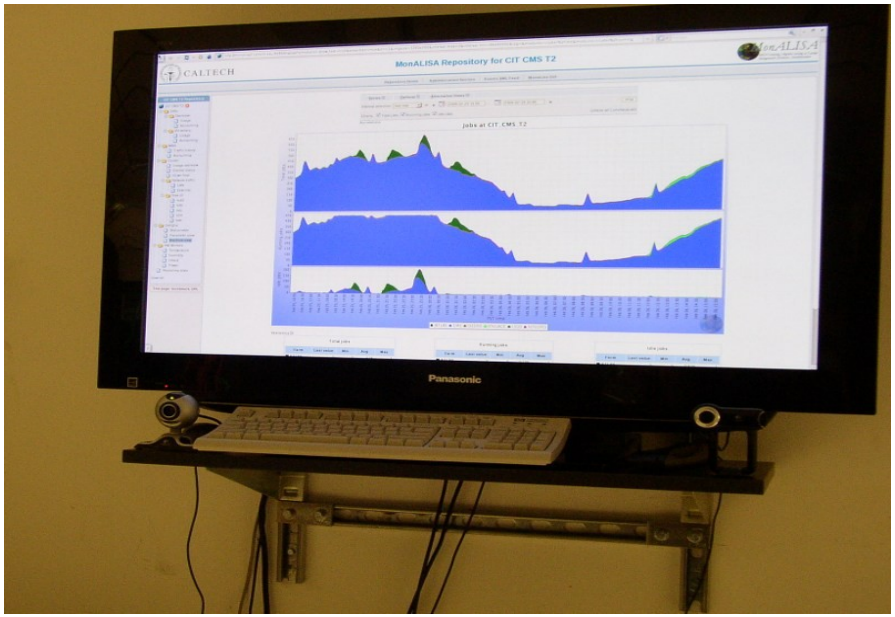

Keep things simple

- No interactive Tier2 use
- Homogeneous architecture
- Only CMS supported (other VOs allowed opportunistic use)

Keep current

- Software updates
- ITB participation

Proactively decommission old hardware

Shenanigans

http://pcbunn.cacr.caltech.edu/Tier2_Shenanigans_2.wmv

http://ultralight.caltech.edu/~dkcira/CaltechT2_Feb2009/