

# Caltech Site Report

**Michael Thomas  
Dorian Kcira  
Azher Mughal  
Harvey Newman  
Julian Bunn  
Badar Ahmed**



**California Institute  
of Technology**

**USCMS Tier 2 Workshop**

**Fermilab, March 8<sup>th</sup> 2010**



# CACR @ Caltech



# Caltech Tier2 Team


Julian Bunn

Azher Mughal

Dorian Kcira

Not Pictured:  
Harvey Newman  
Badar Ahmed

Michael Thomas





# Batch Hardware



- 8450 HS06
- 962 cores
- Mix of 3GHz Woodcrest, 2.33GHz Clovertown, 2.5GHz Harpertown, and 2.5GHz Nehalem CPUs
- All servers are Supermicro
- Most recent purchases have been Twin, Twin<sup>2</sup> servers





# Management Hardware



- 2 CEs, 8 cores + 8GB memory each
- 1 node for phedex + frontier
- 1 node for bestman + frontier
- 1 node for monitoring services
  - MonALISA
  - RSV
  - Jobview
- 2 nodes for ITB testing (CE + SE)
- 1 node for NFS (\$OSG\_APP, \$OSG\_DATA)
- 1 node for NN and Condor negotiator/collector
- 1 node for SNN (also has rsync backup of NFS)



# Storage Hardware



- Raw space: 747 TB. Usable space: at least 50% but variable depending on replication
- Migrated from dCache to Hadoop last year
  - Much lower operational load
- Storage has been increasing by putting large disks in new worker node purchases and replacing smaller  $\leq 500$ GB drives
  - All new disks will be 2TB or larger
- Disks used as individual partitions, or as RAID-0. The SE provides our redundancy
- Worker nodes are also HDFS datanodes
- Assorted dedicated datanodes: 2 x Sun x4500, 4 x 2U whitebox, 4 x 4U whitebox, 1 x 3U JBOD
- 4 x 10GbE gridftp, 8 x 1GbE gridftp
  - Preparing tests with additional compute+gridftp nodes

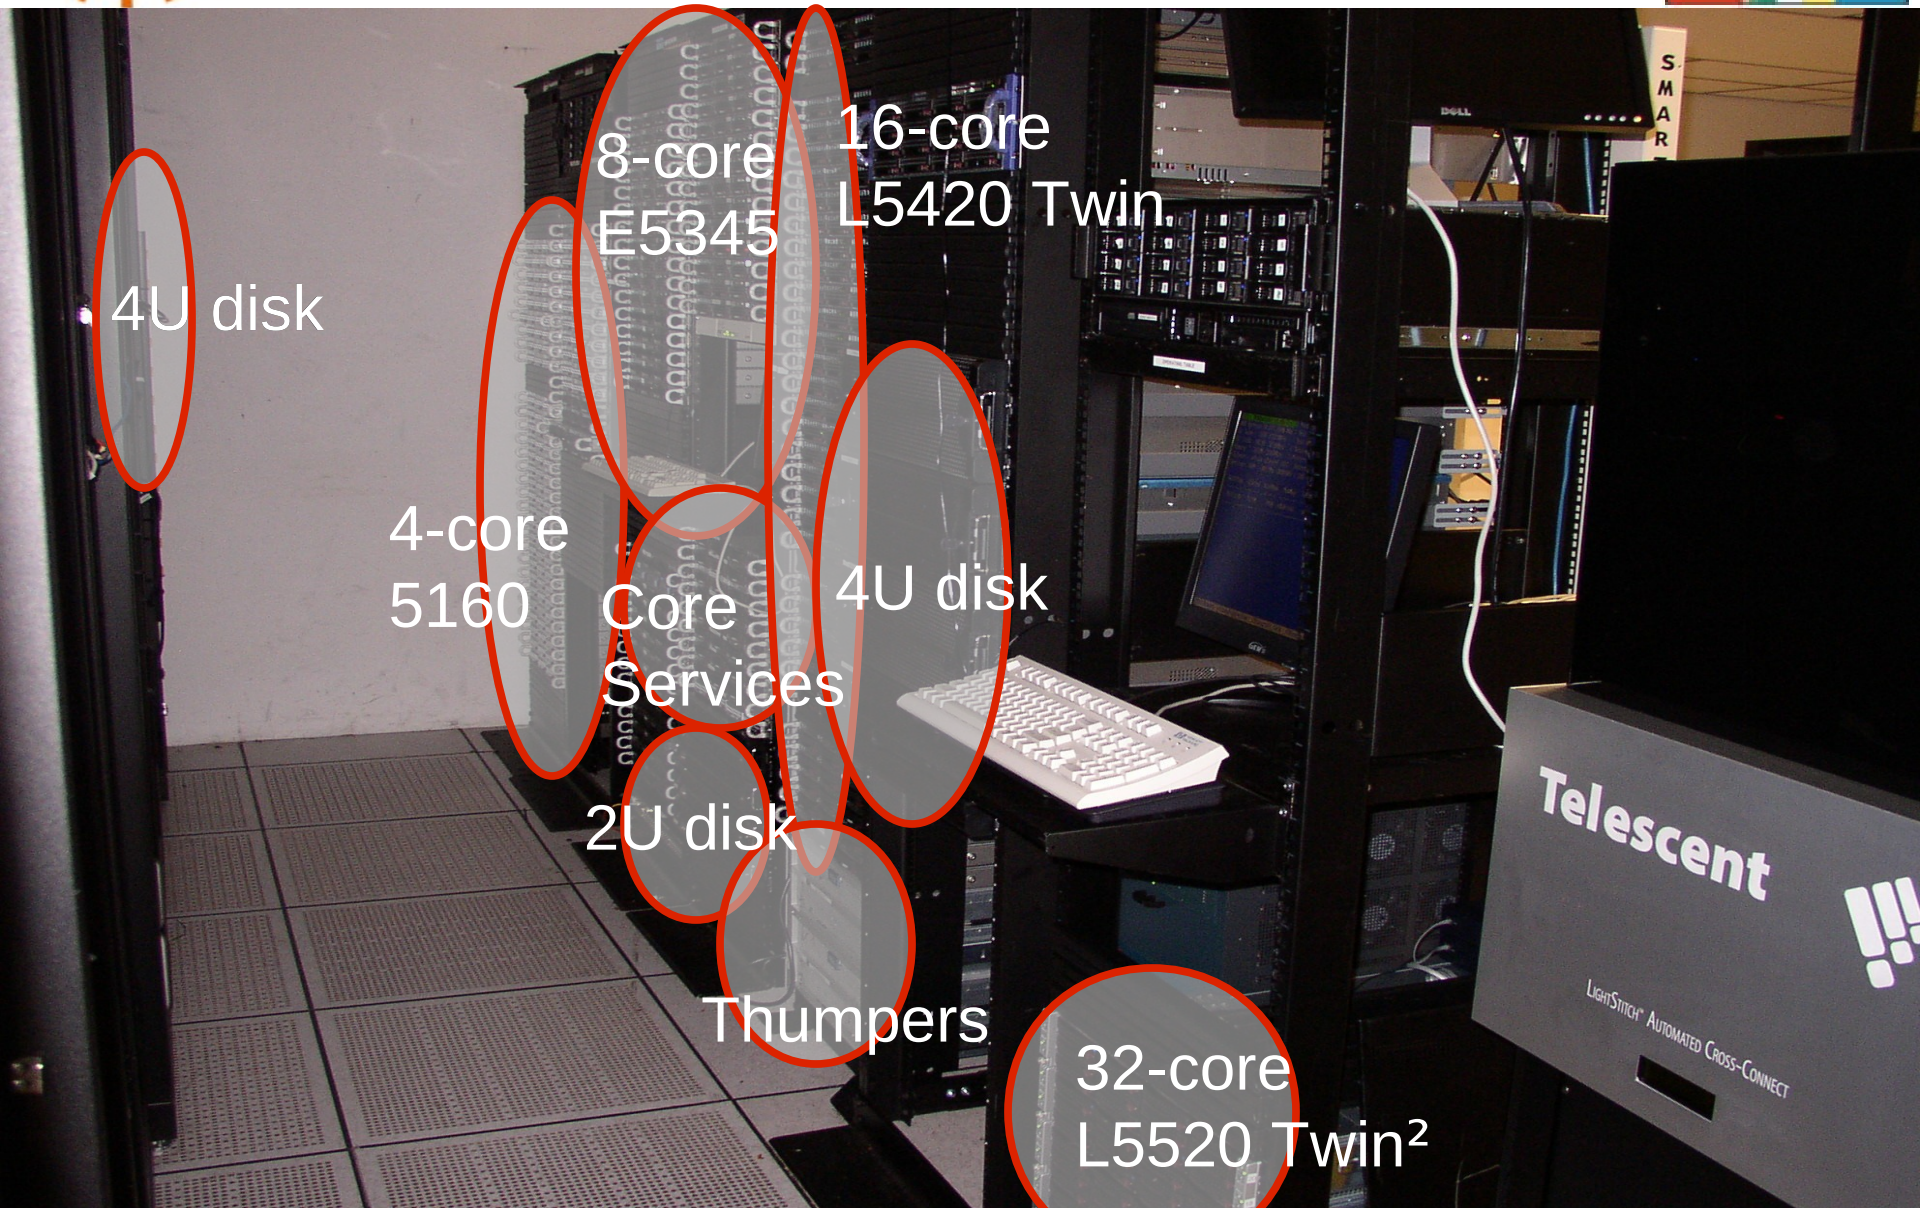


# Caltech Tier2 @ CACR





# Caltech Tier2 @ CACR



4U disk

8-core  
E5345

16-core  
L5420 Twin

4-core  
5160

Core  
Services

4U disk

2U disk

Thumpers

32-core  
L5520 Twin<sup>2</sup>


Telescent

LightSwitch<sup>®</sup> AUTOMATED CROSS-CONNECT





# Caltech Tier2 @ CACR





# Caltech Tier2 @ CACR





# Network Hardware




- 6509 chassis provides main trunk for private subnet
- Cisco 3560e
- Foundry x448 (2)
- Force10 Tera E600
- Force10 E600
- Force10 s50n (2)
- Force10 s2410
- SMC 8708

# Caltech Tier2 Network

## Public Network



## Private Network





# Software

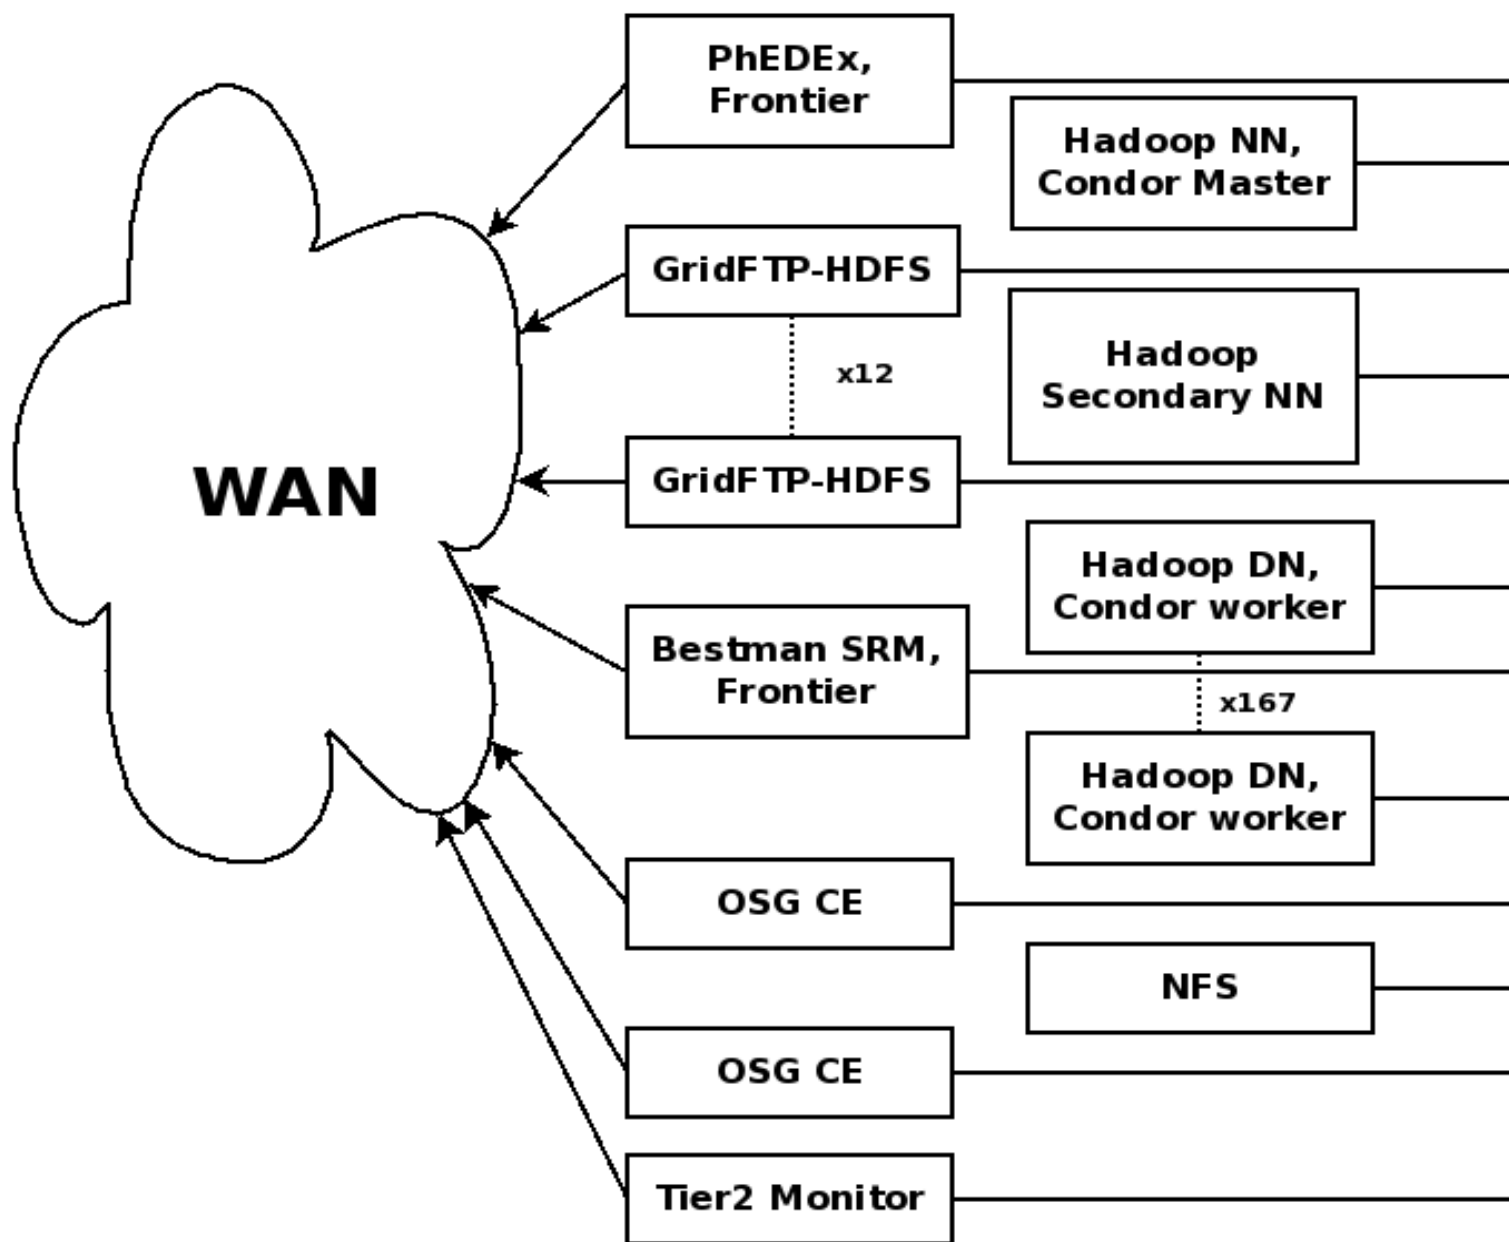


- OSG 1.2.6 (two CEs)
- Hadoop 0.19.2
- PhEDEx 3.2.10
- Frontier 4.0rc8
- Bestman 2.2.1.3.10
- GUMS 1.3.14
- Condor 7.4.1
- Rocks 5.1 with CentOS 5.4 x86\_64
- Custom 2.6.30 Ultralight kernel with cubic TCP<sup>13</sup>

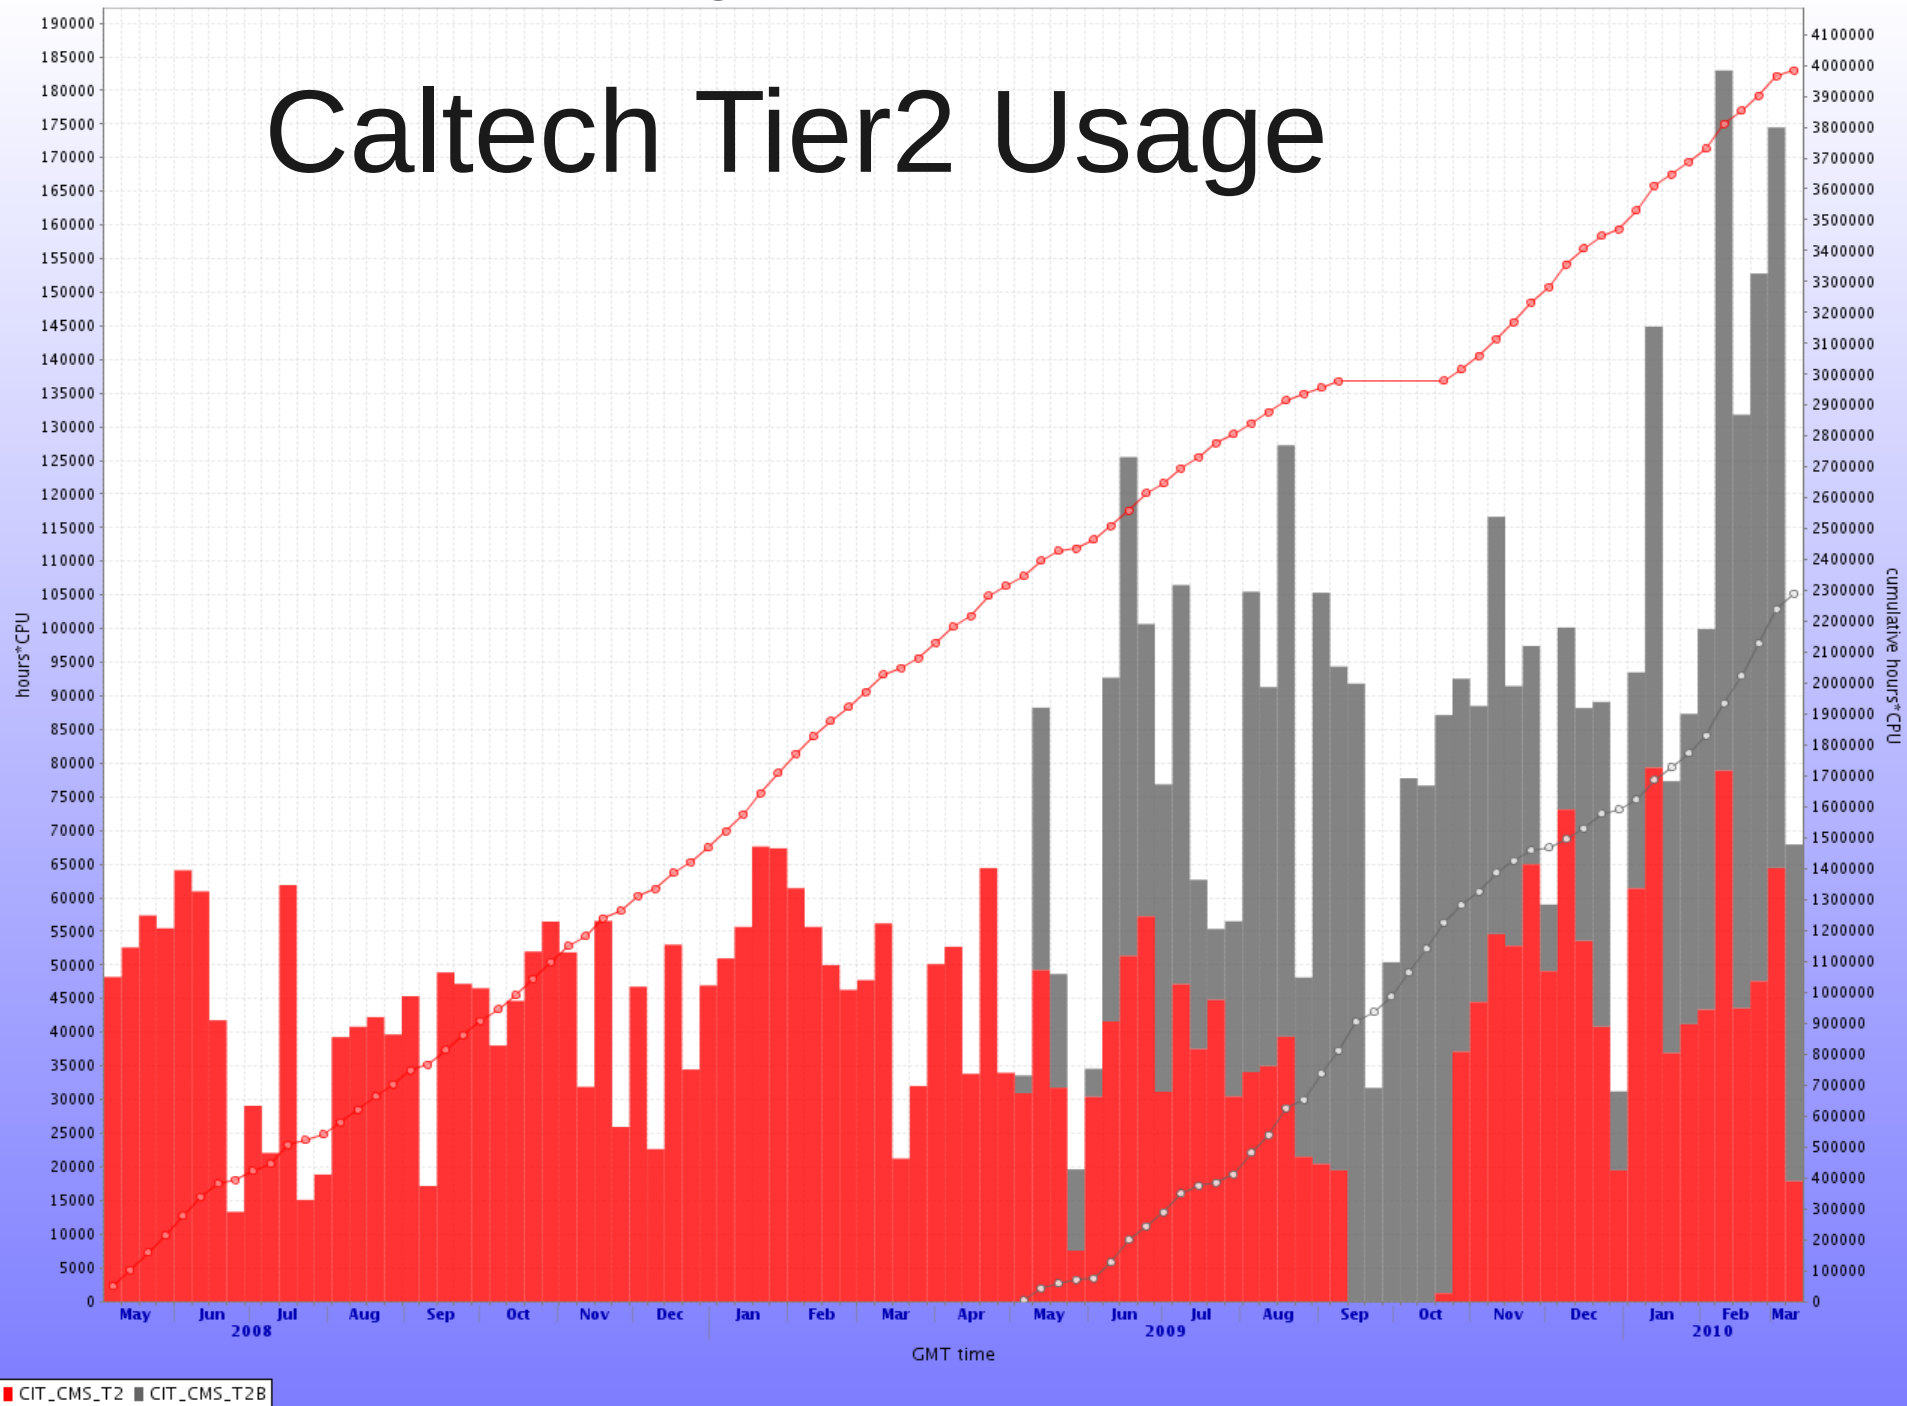


# Caltech Tier2 Services

Private Subnet  
10.3.0.0/16




# Caltech Tier2 Usage



Total No of finished jobs at each site

# Caltech Tier2 Usage







# Future plans



- FY10 milestones already met
- Continue to increase raw storage capacity to maximize the amount of 2 x replicated space
  - Replacing all 500GB drives gives us ~400TB usable replicated space, replacing all 750GB drives gives us ~545TB usable replicated space
- Will continue to evaluate new hardware offerings
- Expand local user analysis resources
  - Streamlined read-only access to Tier2 storage
- Focus on WAN transfer efficiency
  - Change spikes to plateaus
- Hadoop update to 0.20
- HA GUMS
- HA Namenode



# FDT + HDFS



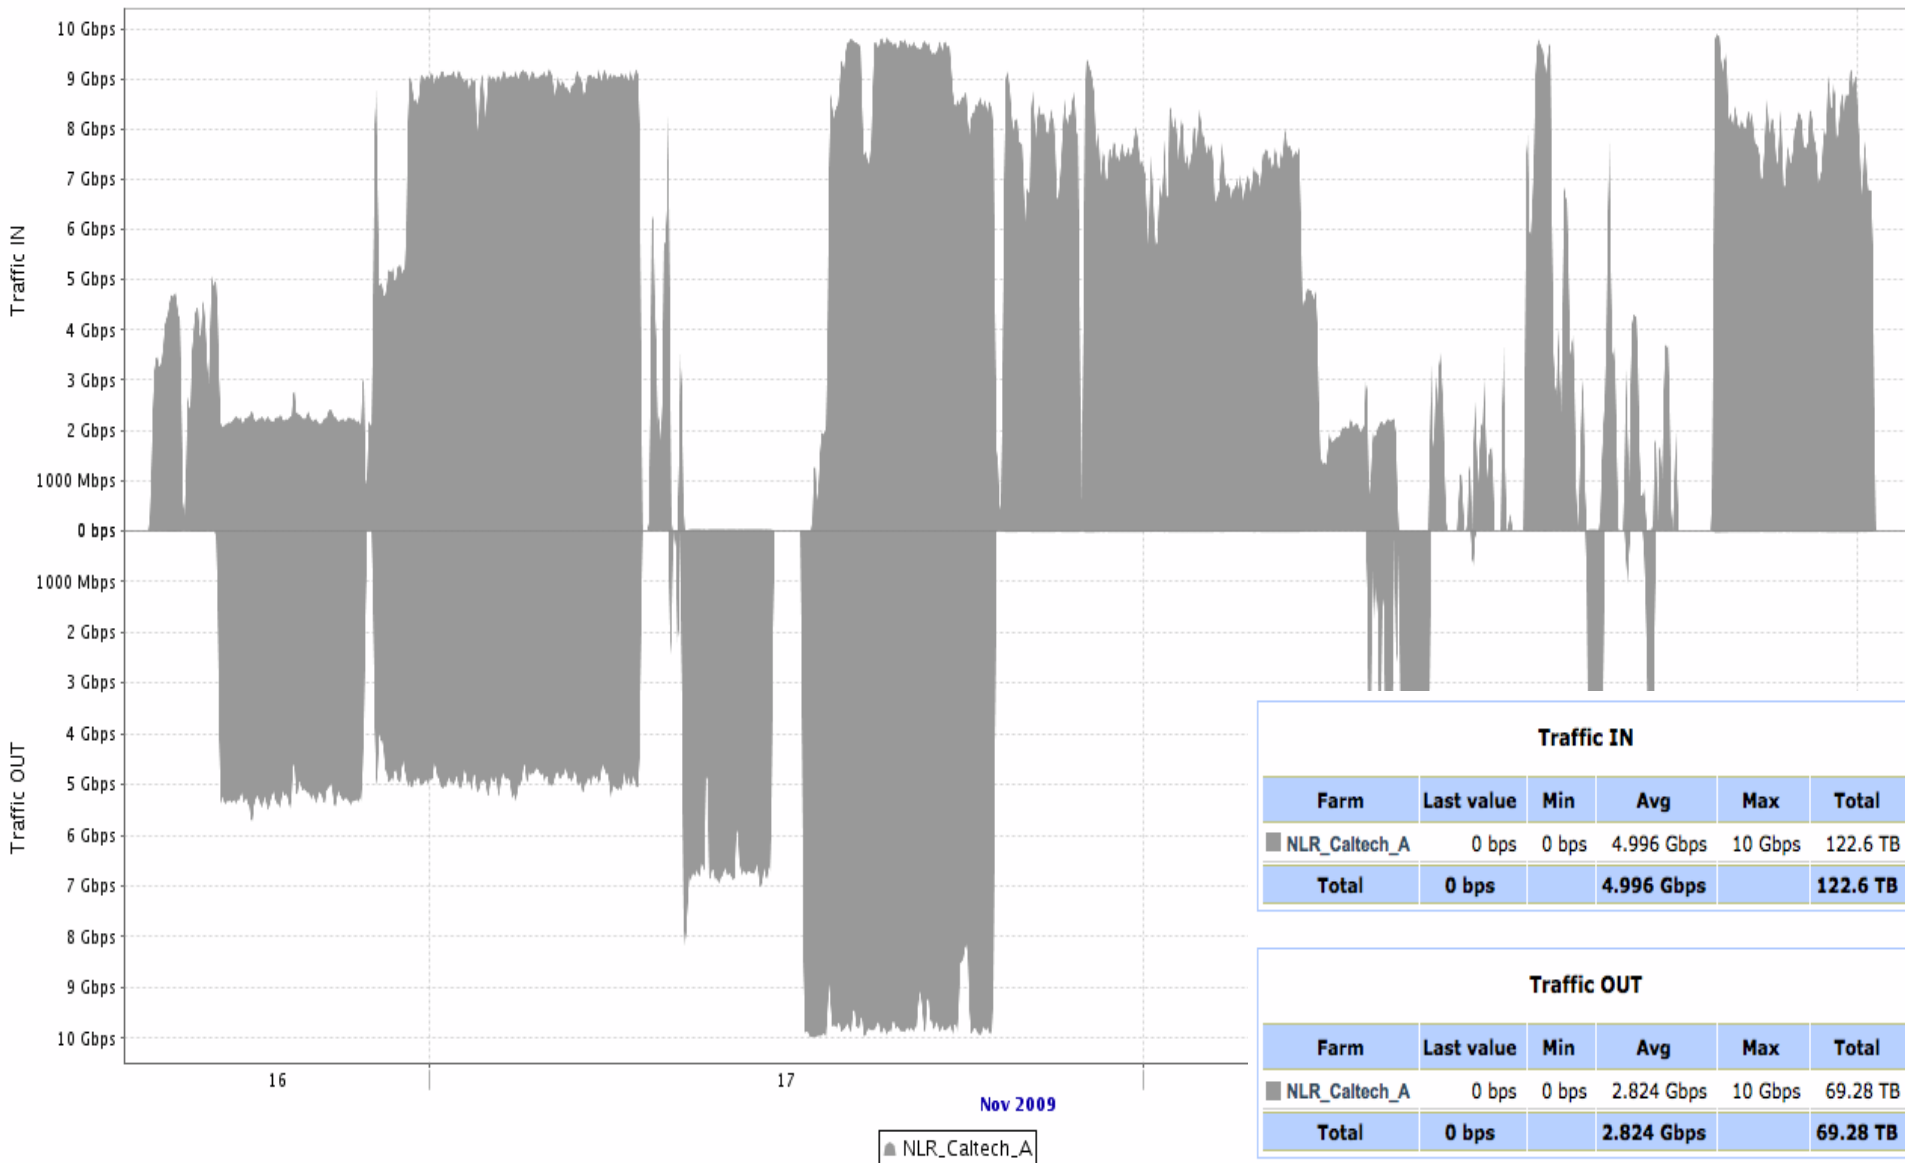
- Plugin to FDT to allow optimized read/write access to HDFS
- Initial work by Badar Ahmed is very promising
- Zdenek Maxa working to integrate FDT with PhEDEx



# SC09: FDT+HDFS



## WAN links





# Analysis {model, Ops}



- Problems with user jobs using excessive resources (memory, disk) reminds us of the need for local resource control
- Pilot jobs lose information on original submitter
  - Submitted request to have this information put back into the condor job classad



# Summary



- Caltech Tier2 in good shape
- Several improvements in the last months
- Ready for the 2010 LHC data taking