



# **CMS@Caltech**

## **Computing Review**

### **October 5<sup>th</sup>, 2010**

**Michael Thomas**  
**Dorian Kcira**



# Storage Resources



**3 clusters, each running Hadoop for local storage**

Cluster	Hadoop (TB)	NFS (TB)	Total (TB)	Used
Tier2	527 <sup>1</sup>	0	527	80%
T3-higgs	7.62 <sup>2</sup>	22	29.62	85%
T3-susy	21.62	41	62.62	80%

<sup>1</sup>+60TB pending

<sup>2</sup>+30TB pending



# CPU resources



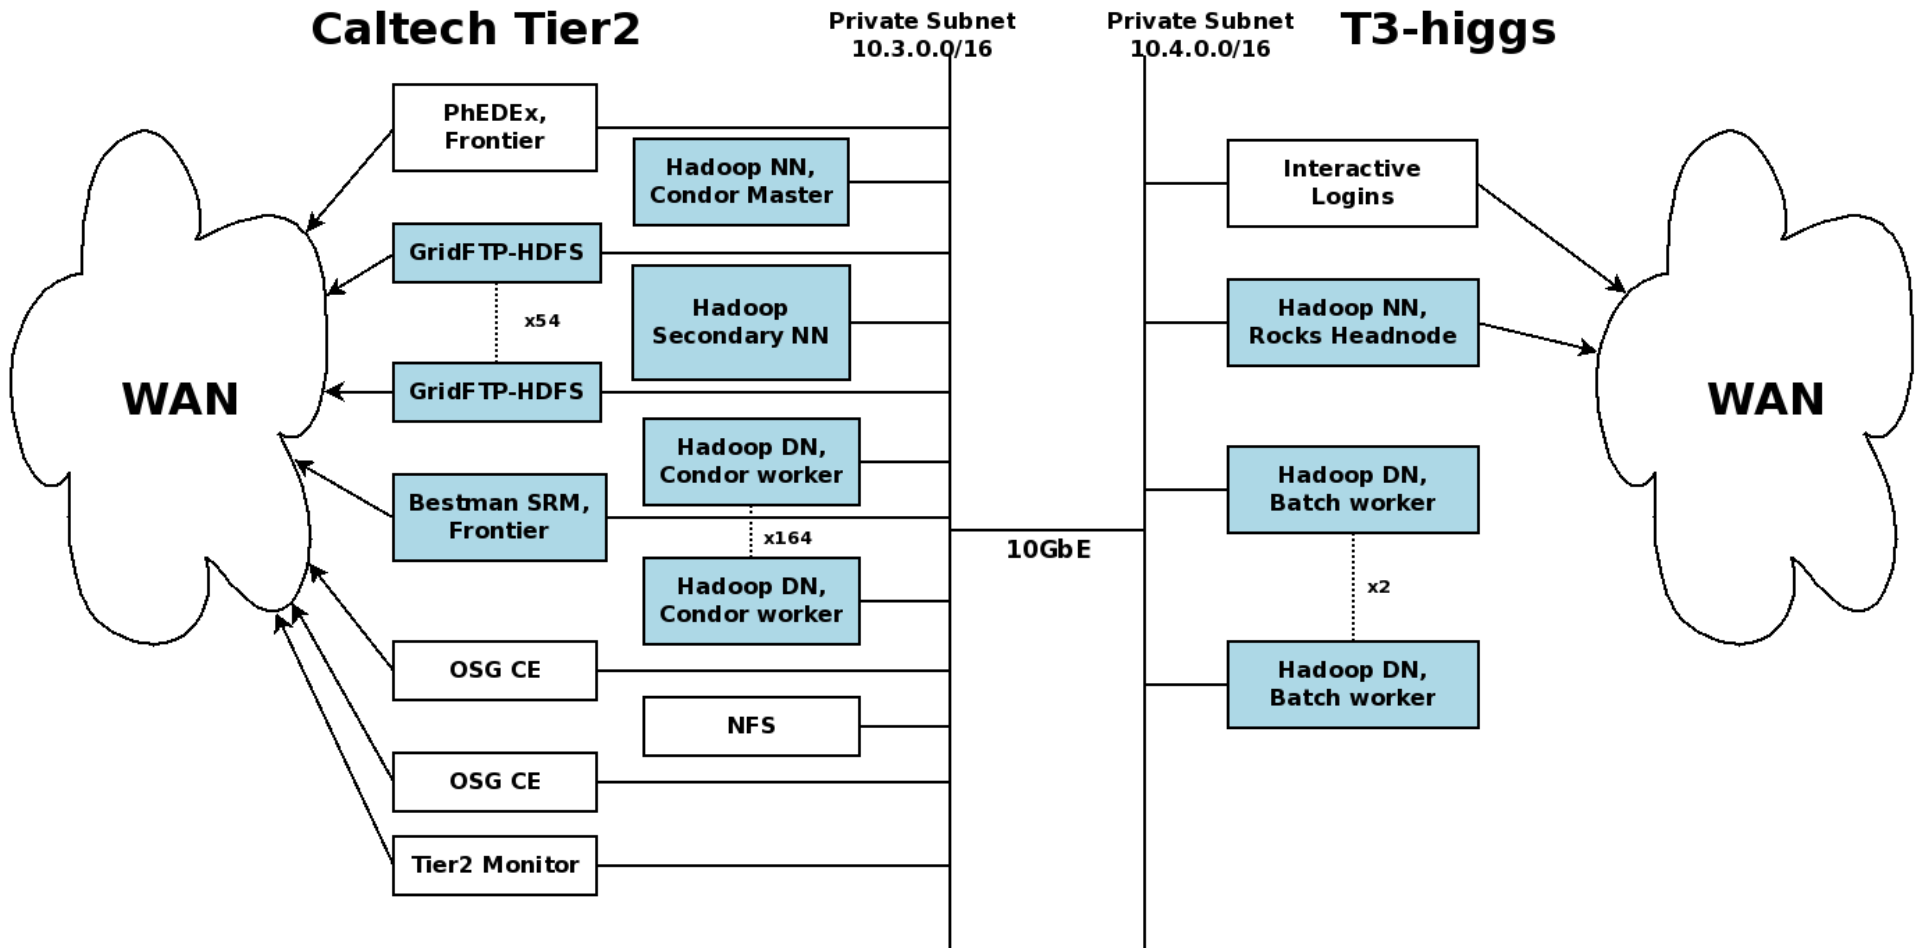
Cluster	Batch slots
Tier2	962 <sup>1</sup>
T3-higgs	8 <sup>2</sup>
T3-susy	112

<sup>1</sup>+192 pending

<sup>2</sup>+96 pending



# Cluster diagram

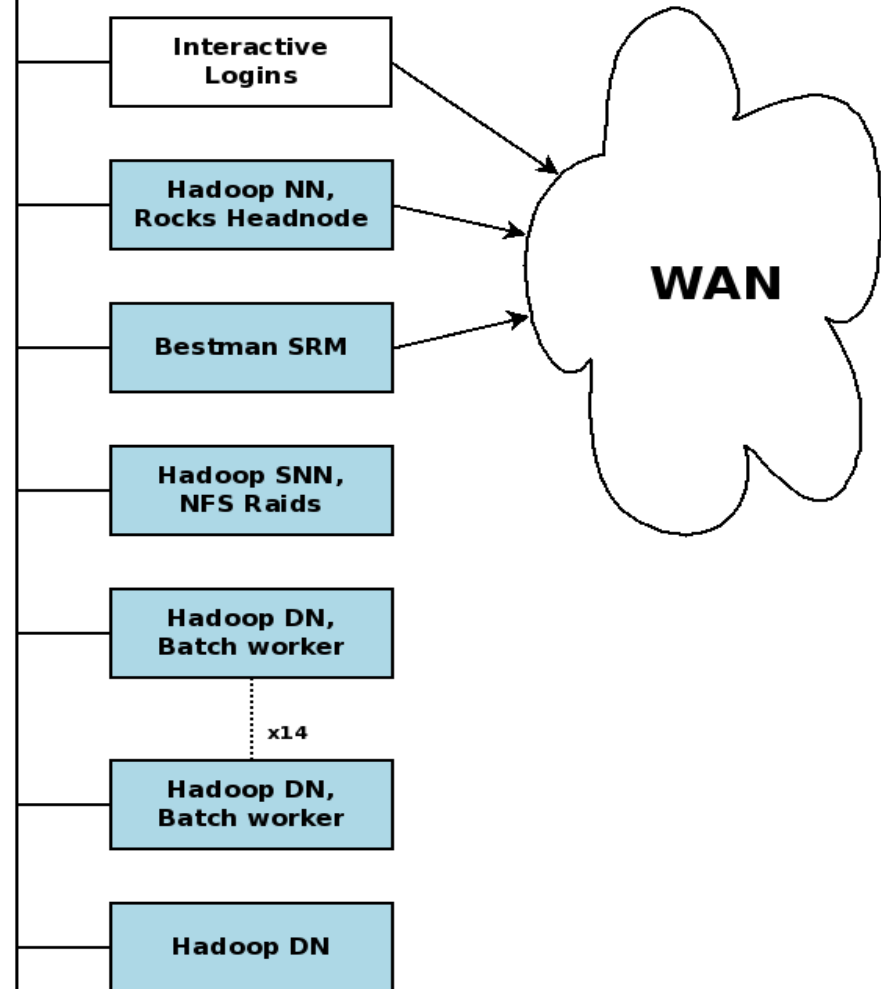




# t3-susy diagram


Private Subnet  
192.168.0.0/16

## T3-susy






# T3-susy storage by user





# T3-higgs storage by user

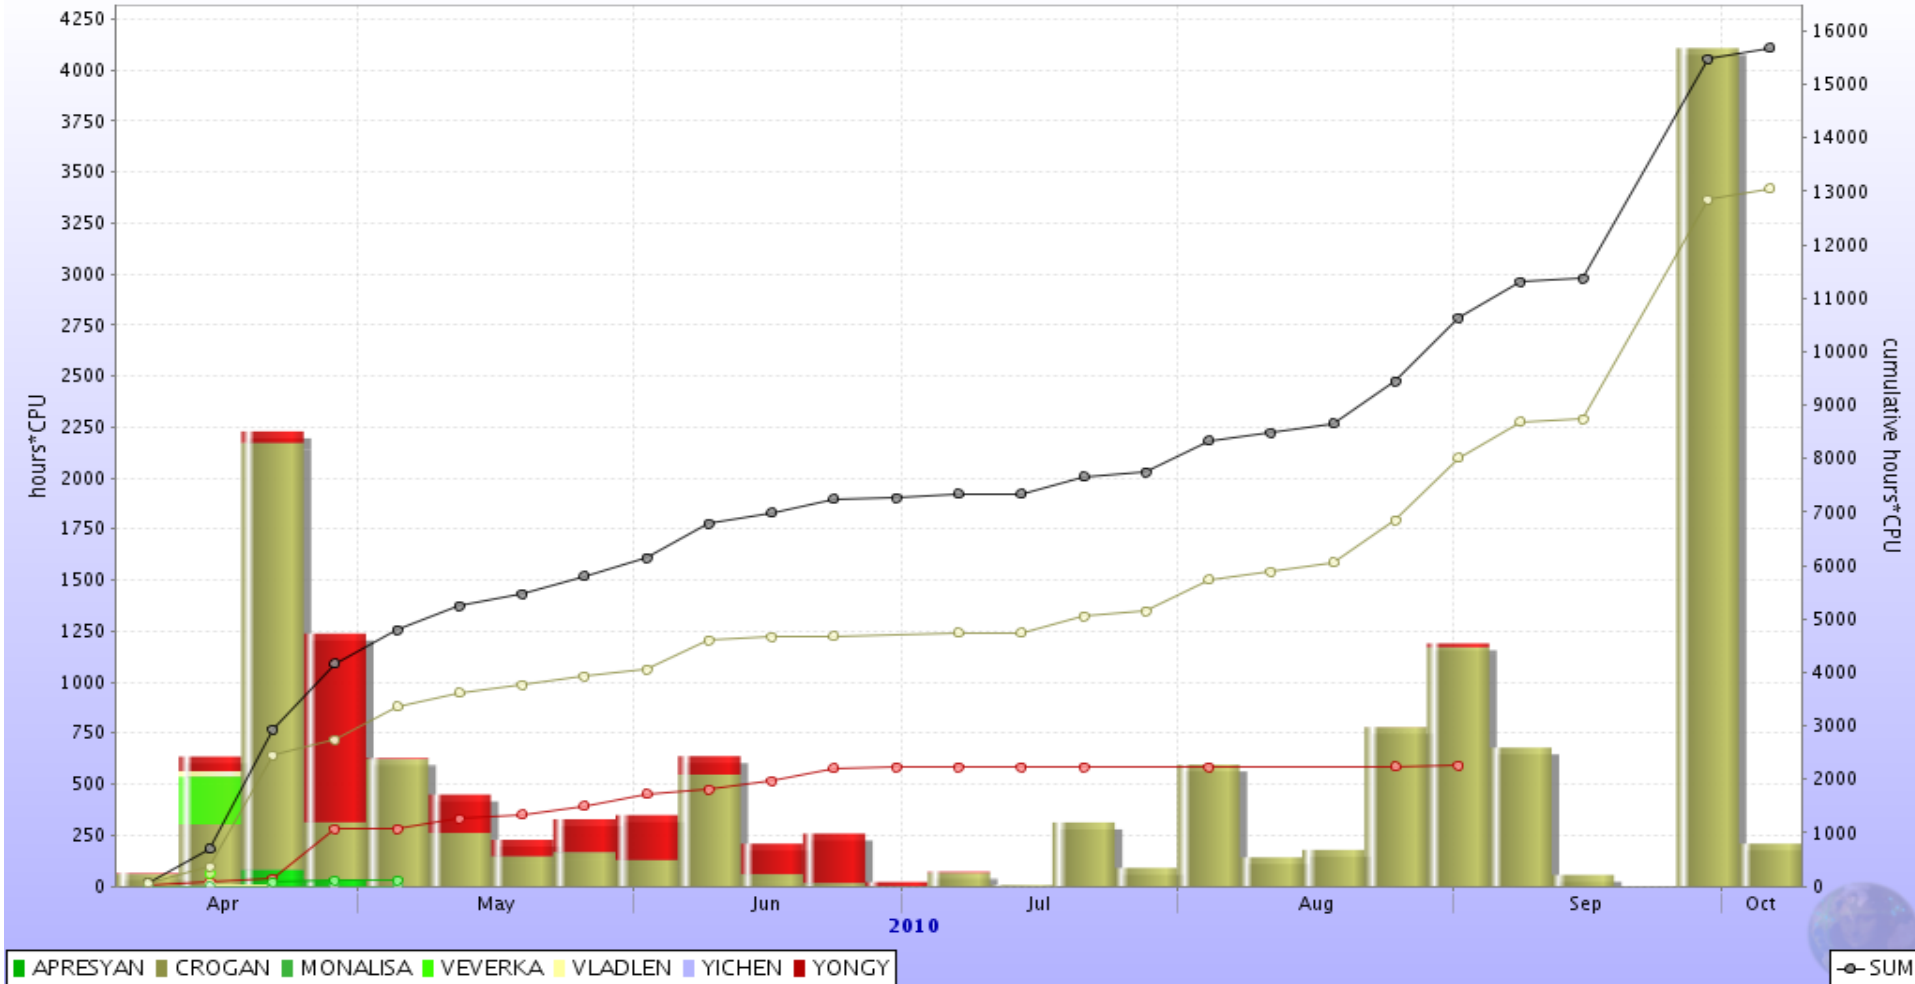




# t3-susy Batch system usage



### Integrated CPU time consumed per user



1 Day = 2668 CPU hours





# **t3-susy 6 mo. batch statistics**



- **Cumulative wall time used: 6 days (MonALISA) or 10 days (SGE)**
- **Total jobs run: 261738**
- **Distinct users: 6**
- **Average job length: 379s**
- **Average job cpu time: 134s**



# **t3-higgs 6 mo. Batch Statistics**



- **Cumulative wall time used: 18 days**
- **Total jobs run: 35265**
- **Distinct users: 1**
- **Average job length: 44s**
- **Average job cpu time: 18s**
- **Last job: July 14**



# Additional Services 1

- **CVS** - <https://cms.hep.caltech.edu/viewvc>
  - SVN? Do users need it? CVS also mainly used by MT/DK
- **Group page:** <http://hep.caltech.edu/cms>
- **Personal pages**
  - <http://hep.caltech.edu/~username>
  - On positron: \$HOME/public\_html
  - Or use Twiki
- **Twiki** <https://twiki.hep.caltech.edu>
  - Will migrate (part of the) information from the old gaewiki
- **LDAP for user management**
  - Most other services use it



## Additional Services 2



- **Email: HyperNews or web interface to existing archive from Mailman**
- **Meetings**
  - Indico, not yet done
- **CMSSW**
  - Vladimir used to do the installations (+MT/DK)
  - Yousi could take over (+ MT/DK)
  - First set up of installation area difficult, further installations easy.

**Note:** Sakai not used anymore. Link will disappear, parts of the info will be used for the group page and the new twiki. <https://cms.hep.caltech.edu/sakai>



## **Additional Services 3**

- **Analysis code**
  - **CITAnalyzer used previously but now each Caltech sub-group has own analysis code**
- **Use Twiki to collect information**
- **Keep code in CVS or in user areas that are automatically backed up. Do not keep important code on data disks or on Hadoop.**



# Additional Services 4



- **Desktop resources**
  - **Caltech**: mainly through HEP computing. Linux or Windows PCs.
  - **CERN**: need list of desktops with IPs and access to them. Prepare a Twiki page with these resources.



# Additional Services 5



- **Support email:**  
**[tier2admin@hep.caltech.edu](mailto:tier2admin@hep.caltech.edu)**
- **For issue tracking, we use internally**  
**[bugzilla.hep.caltech.edu](http://bugzilla.hep.caltech.edu). Can be made**  
**available to users if needed/useful.**



# Issues 1



- **T3 CPU is underutilized**
  - **Use T3 login nodes for running interactive jobs**
    - ♦ **t3-susy, t3-higgs: 8 CPUs each**
  - **Use the SGE batch system for non-interactive jobs**
    - ♦ **Priorities taken care by SGE. Simple default SGE priorities for now: resources and user based. Can be refined if needed: group based, different queues.**





## Issues 2



- **Storage is in short supply**
  - **Need to get organized**
  - **Automatic notifications for unused data in hadoop?**
  - **Twiki page with datasets and ntuples as well as code for running over them?**
  - **Strategy for scaling the T3 storage with increase in data size**



## Issues 3



- **AFS kernel module is currently not working**
- **Ganglia does not seem stable**
- **Need fixed times for routine maintenance**
  - 1 hour weekly time slot
  - 8 hour monthly time slot
- **Need closer collaboration with users**



# For further discussion

<http://twiki.hep.caltech.edu/twiki/bin/view/Main/Computing2010>