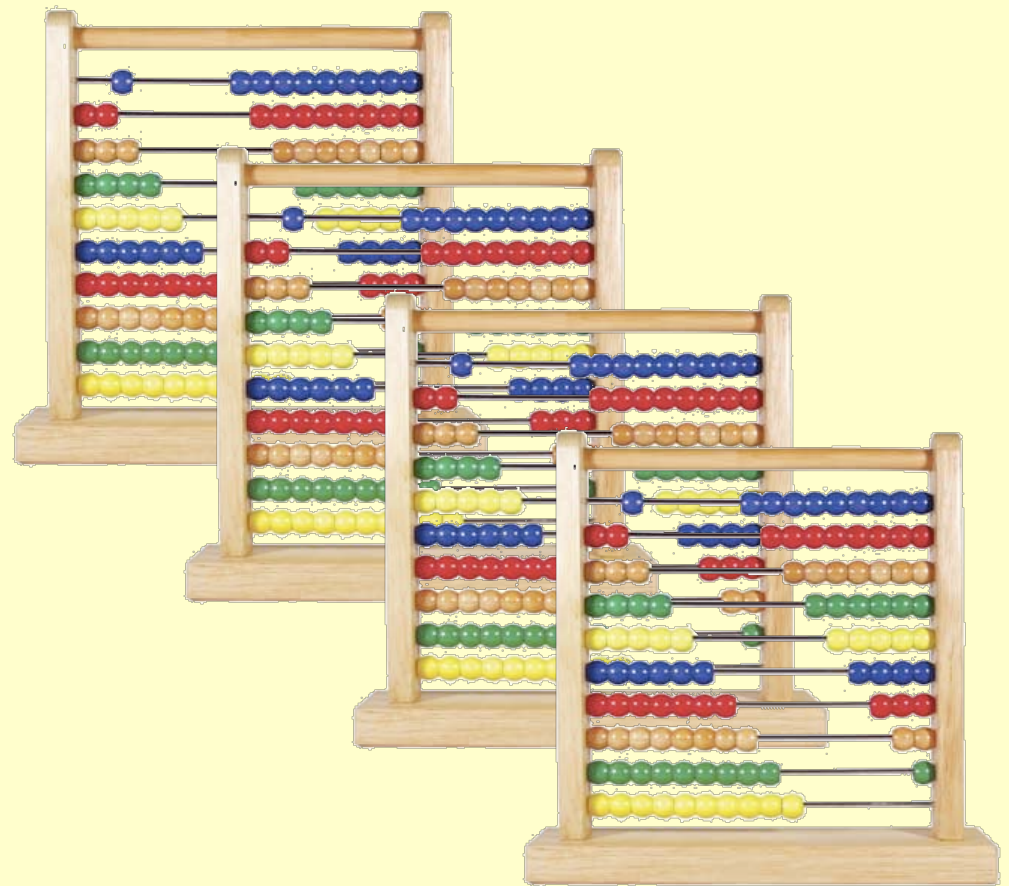




# Fermi Large Area Telescope:

## Collaboration Computing

**Richard Dubois**  
SLAC National Accelerator Laboratory  
[richard@slac.stanford.edu](mailto:richard@slac.stanford.edu)






# Current State

---

- **Resources**
  - **SLAC**
    - 1200 cores + 650 TB disk + 350 TB tape – 55 TB disk free
    - 128 TB additional disk ordered underway.
    - 6 TB/wk use by L1 Processing
  - **Lyon**
    - 600 cores
- **Level 1 processing + ASP**
  - ~300 cores to turn around downlink promptly – recon + monitoring etc
  - ASP not noticeable
  - L1 will exhaust disk in 2 months (burn rate ~25 TB/month)
- **Simulations have been mostly targeted to Cosmic Ray Electron simulations and to validating code builds**
  - Continuing with ghost tracks/Pass7 development
  - Encourage MC tasks to be run at Lyon
- **Reprocessings –**
  - A reprocessing of the first year data for the public release
  - Reprocessing underway now for next round of event cuts (“Pass7.2”)
  - We anticipate a full reprocessing (from raw data) in the coming year for Pass8, which should better handle the ghost tracks




# 6 Months in CPU Cycles





# CPU history


CPU-use by Month and Application (as of 12:00pm OCT 23, 2009)





# 6 Months in Disk

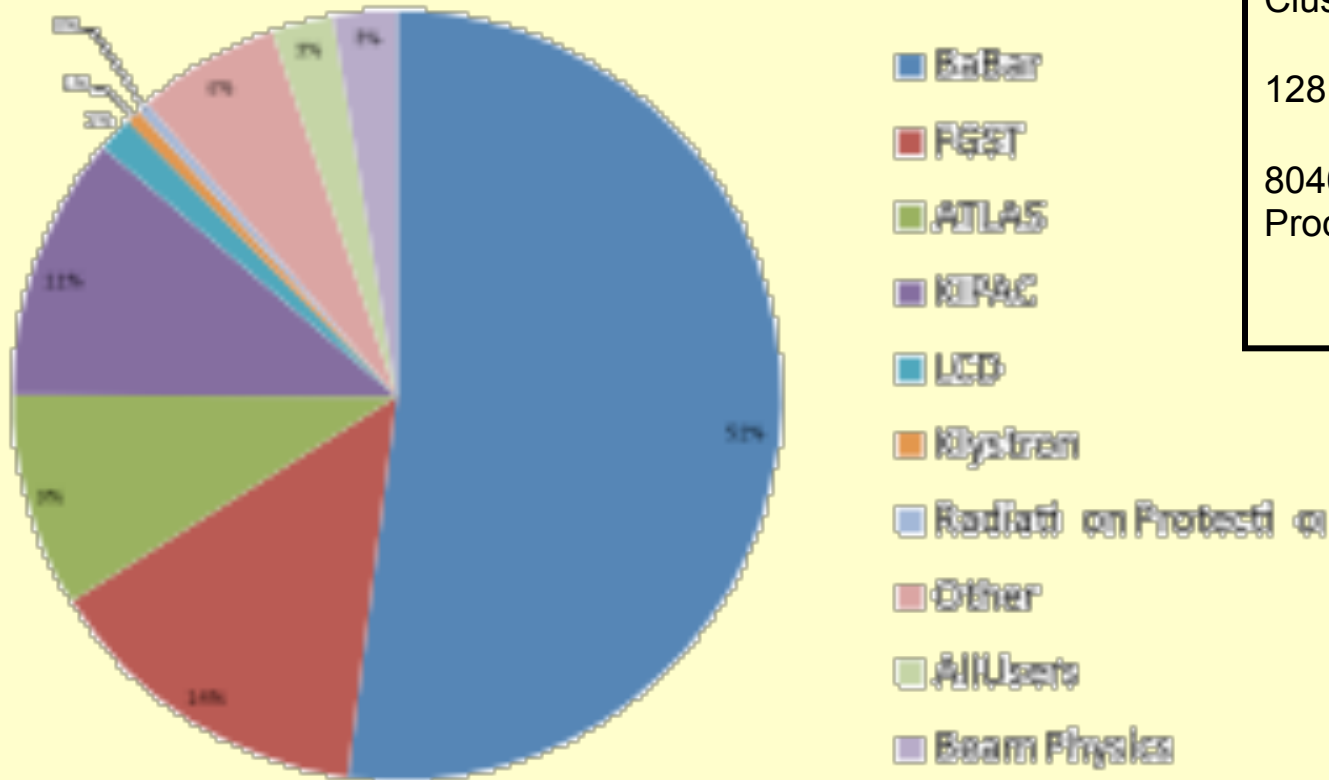
Available xrootd Space (GB)





# SLAC HEP Computing Facilities

**SLAC HEP CPU 9000 Cores**  
Average core age 2.6 years



72 Cores: SMP (SGI Altix)

768 Cores: Infiniband Cluster

128 Cores: Myrinet Cluster


8040 Cores: HEP Data Processing

**BABAR is moving 3000 cores into the general queues on Nov 16**



# SLAC HEP Computing Facilities

SLAC HEP Disk Space      2052 Terabytes  
Average TB age 2.3 years



34 TB: AFS  
47 TB: NFS-Network Appliance  
700 TB: NFS  
1200 TB: xrootd  
74 TB: Lustre, Objectivity ...



# Where Did the '09 Money Go?

---

- **2 shots of disk**
  - **275 TB in March**
  - **128 TB in Nov/Dec (in process now - \$125k)**
- **100 TB tape for backups**
  - **New 1 TB tapes**
  - **200 TB on order now (included in \$125k order)**
- **400 cores added in SLAC farm**
- **Miscellaneous little stuff**
  - **Cables, etc**
- **\$180k left in kitty from 2009**



# Towards an Ongoing Computing Model

---

- **Price of tape has dropped by x ~4!**
  - **Can now afford to keep a copy of everything on tape**
  - **And the latest version of data always on disk**
- **It looks like we wisely stalled long enough for hardware prices to drop enough to avoid deleting any of our processed data files**
- **New model is to copy all data to tape right away. When reprocessing, we delete the previous processed files from disk – but they stay on tape if needed.**
- **We expect no more than one full reprocessing from raw data per year. But (at least for a while!), we would reprocess all previous data each time.**



# Computing Model Projection for 2010


---

- **L1 eats ~350 TB/yr**
  - **Disk**
    - 350 TB (\$260k)
  - **Tape**
    - 350 TB tape for previous data
    - 350 TB for 2010 data
    - 700 TB for complete reprocessing in 2010
    - Total 1.4 PB tape (\$175k)
- **Allocate ~64 TB/yr for MC (\$47k)**
- **Add 400 cores @ SLAC (\$200k)**
  - **Prepare for full recon with slower reconstruction code**
- **Replace 5 older linux servers; 1 fileserver (\$49k)**
- **Note that \$180k carryover from 2009 will be applied here also**



# Computing Resource Projections

LAT Computing Resource Projection



## Needs in '10:

- 350 TB disk: \$260k
- 1400 TB tape: \$175k
- 400 cores: \$200k
- Sundry servers: \$50k

**Total: \$685k**  
(less \$180k carryover)

## Needs in '11:

- 350 TB disk: \$260k
- 1400 TB tape: \$175k
- Sundry servers: \$50k

**Total: \$485k**

2009 pricing: \$0.75k/TB disk; \$0.125k/TB tape; \$0.5k/core



# The Lyon Story

---

- **Pipeline 2 running well at Lyon**
  - **Negotiated initial allocation with Dominique Boutigny (Lyon CC Director) for 600 cores in 2010**
    - **MoU agreed to: doubling of cpu capacity annually**
  - **Proving very useful to the collaboration!**
    - **Just ran off simulations for Pass7.2 IRFs**
    - **More Cosmic Ray Electrons simulations**
  - **They plan to hire a support person for astro groups usage**
- **A PS on Italy**
  - **Our simulations can now run on the Grid in Italy**
  - **Work has started in Pisa to interface our pipeline to the Italian Grid**



# Take-away Messages

---

- We now understand L1, ASP and RSP
- Still need to understand full reprocessing loads
  - Full reprocessing for Year 1 release
    - Improved reconstruction and event classification ( but slower!)
    - Implement the new disk/tape model

Computing Model not settled yet but getting there!

