




Fermi Large Area Telescope

Collaboration Computing

Richard Dubois
SLAC National Accelerator Laboratory
richard@slac.stanford.edu





The Computing Model

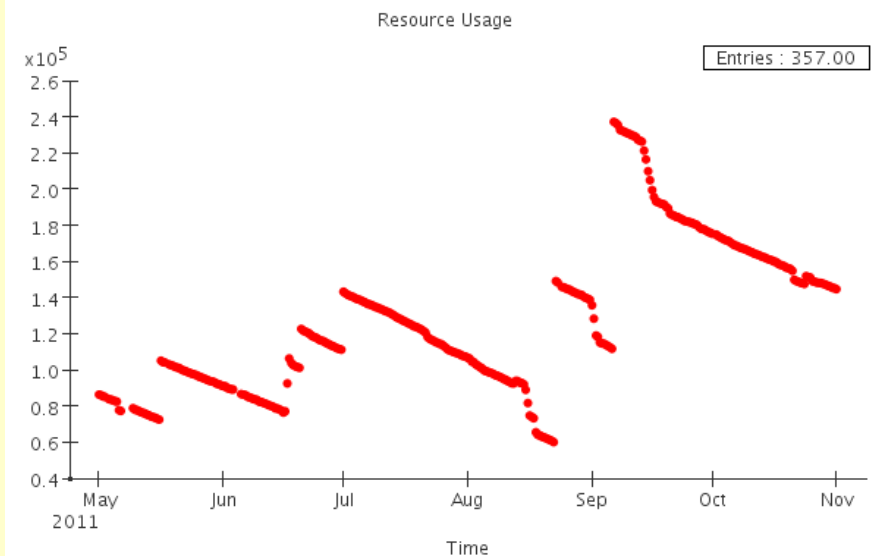
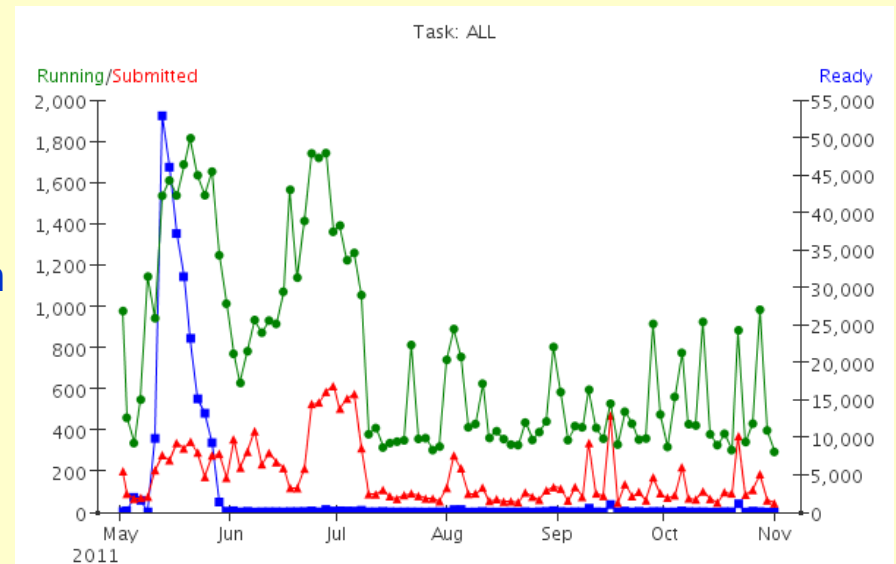
- **“Core” computing**
 - **Direct IFC support**
 - **Level 1 data processing, reprocessing & ASP at SLAC**
 - **All latest data on disk + copy of all versions on tape**
- **Simulations for Collaboration analyses**
 - **Done at SLAC, Lyon and some at INFN GRID**
 - **In support of science analyses: Pass7, Pass8, CRE etc**
- **Big Picture Plan**
 - **Still deciding on Pass7 reprocessing: ~3 months + 1 PB files – would be imminent**
 - **MC for Pass8 development: bounded by big HE run last year**
 - **Pass8 reprocessing: early 2013**




Resource Usage of Late

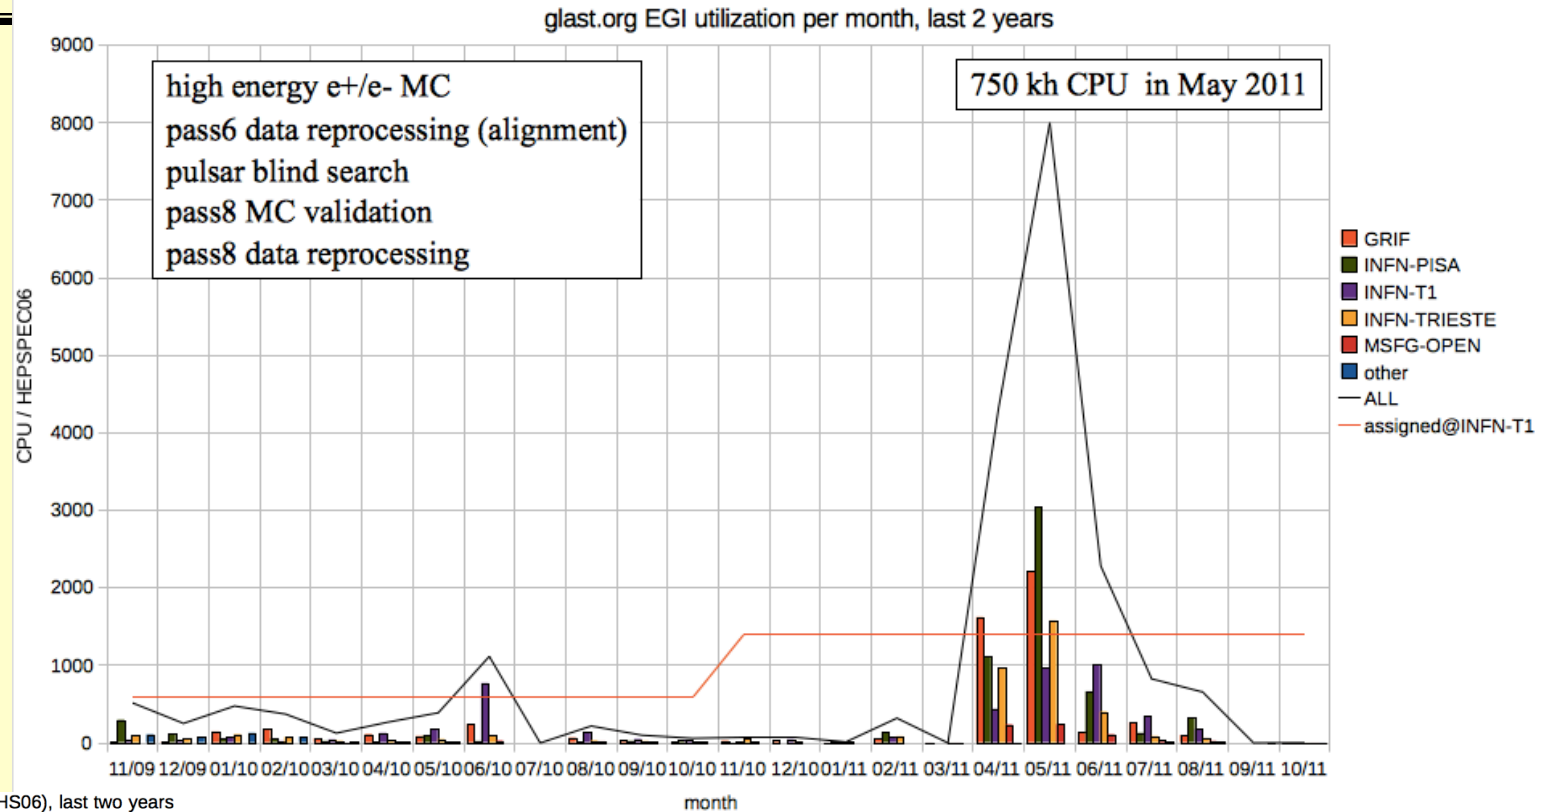
- **Not much simulation after massive runs of first half of 2011**
 - Episodic for IRFs, Pass8 etc
 - Run by Stephan Zimmer (Stockholm)
 - Upgraded Lyon to GridEngine batch system
 - New computing professional at Montpellier to work the pipeline

- **Level 1 Processing**
 - Steady at 10 CPU-yrs/month
 - 90k CPU-hrs






GRID in Italy



glast.org EGI utilization per site (HS06), last two years



- Manually run by Pisa folks
- No connection to existing infrastructure
- Files manually put in data catalogue
- Good incentive to extend the pipeline to GRID



2012 Hardware Forecast

- **Assume Pass8 full reprocessing**
 - **Disk pricing uncertain at present: assuming \$0.5k/TB**
 - **360 TB each disk and tape (\$0.12k/TB) for L1**
 - **50 TB disk for MC**
 - **Replace remaining 32 TB of borrowed storage**
 - **2 PB tape for Pass 8 reprocess**
 - **Total ~\$500k**

- **Start replacing aging servers**
 - **10 (of 25) standalone servers: \$50k**
 - **Replace 100 TB of >5 yr-old file servers: \$50k**
 - **Additional user disk: \$10-15k**

- **Total ~\$625k - less \$175k not spent in 2011 and carried forward**



The Way Forward on Storage

- **We've bought our last storage from Oracle**
- **SLAC SCCS now has recommendations for vendors:**
 - **Dell for Workgroup size**
 - **NetApps (bought from LSI) for mid-range/xrootd (us)**
 - **Data Direct Networks for large (>PB)**
- **Of course, it is never dull**
 - **Floods in Thailand threaten the hard drive market, so we are advised to buy early**
 - **Will trigger ~300 TB order asap!**



Progress on Computing Model

- **Major work areas**
 - **Transparent access to files on tape or disk**
 - **done**
 - **Better file transfer from Lyon to SLAC**
 - **xrootd proxy system for MC simulations**
 - **done**
 - **Could be useful for sharing reprocessing loads**
- **We are not buying CPUs for peak load anymore**
 - **Counting on leveraging current cores for spikes in processing and drawing on shared clusters**
 - **Full reprocessings will get more and more time-consuming**
 - **Pass8 will not be faster than Pass6/7!**

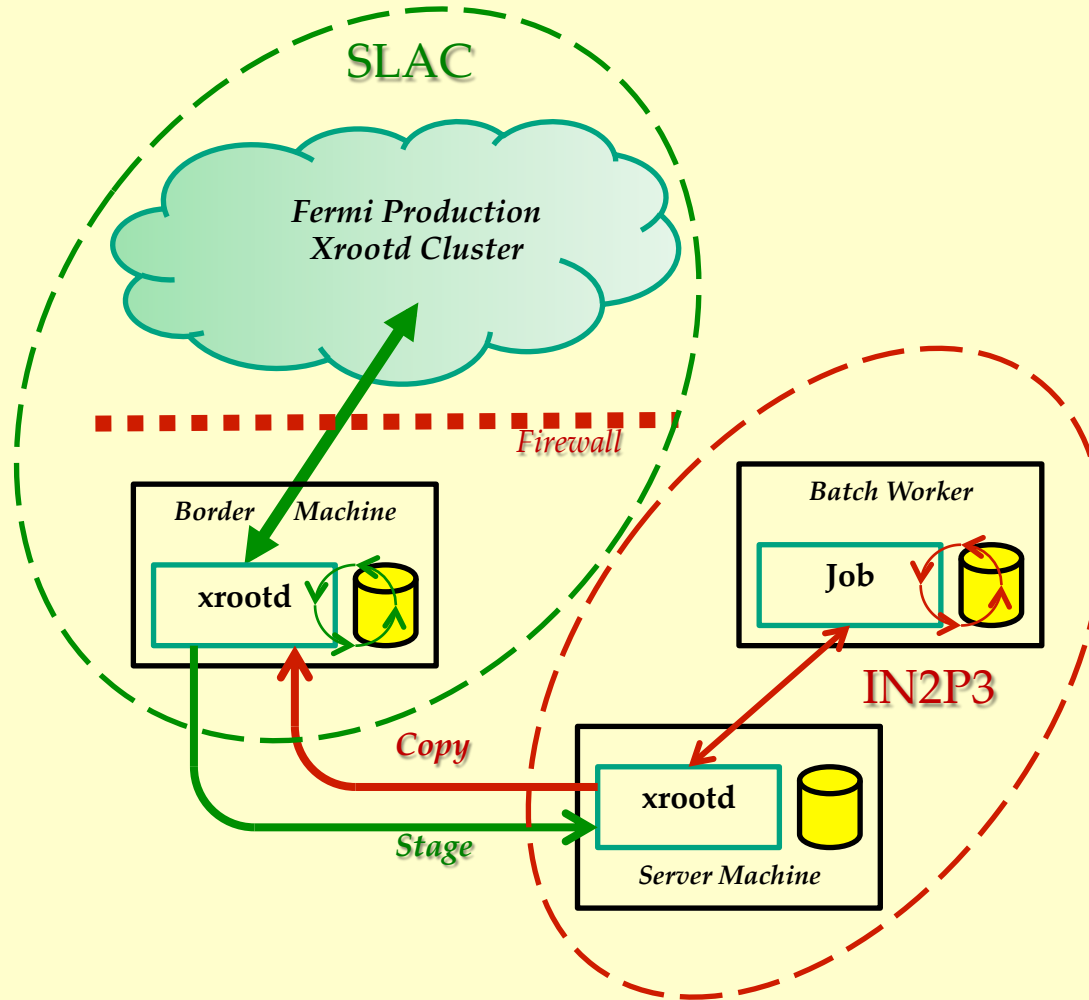


New Initiatives - 1

- Investigate data reprocessing at Lyon
 - 1200 cores at Lyon is a noticeable fraction of Fermi's total now
 - Would generate sizable amounts of recon files
 - 0.6 PB if done by the end of this year
 - Would take months to transfer back to SLAC
 - Summary files are much smaller and can be transferred back asap – that is all we need to go public
 - We have a design with xrootd to handle the I/O
 - Need a disk buffer at Lyon for the recon files to let us dribble back the recon files at our “leisure”



Concept for Reprocessing





New Initiatives - 2

- **Extending the pipeline to the GRID**
 - **Growing resources in Italy, Sweden, France on the GRID**
 - **Manual use in Italy is happening**
 - **Potential synergy with CTA/Montpellier**
 - **Investigating interfacing to LHC DIRAC project**
 - **Take care of GRID middleware**
 - **Will want to understand how to inter-operate between EU (EGEE) and US (OSG) grids**
 - **Will start up discussions with OSG imminently**



New Initiatives - Virtualization

- **We see the writing on the wall for ongoing OS support**
 - **A lot of work involved in keeping up with new OS versions: external libraries, compiler changes etc**
- **Virtualization is in our future**
 - **Able to freeze out to a particular OS version and run it in a VM in a modern machine**
 - **Starting the investigations now to get ahead of the game before no people are left**
 - **Good connections to BABAR and FNAL expertise**
- **We're doing the requirements gathering now:**
 - **When would such a system be needed?**
 - **To what extent must the VMs be isolated? Can they be integrated with SLAC resources such as the batch system?**
 - **How much of our processing would be run in VM's? Will this change with time?**



Summary

- **Disk issues are hopefully behind us**
- **Need \$626k for CY12 hardware**
 - **Pass7 and Pass8**
 - **Replacing old hardware**
- **Looking ahead to need for peak cycles:**
 - **Data reprocessing at Lyon**
 - **Extend the pipeline to the GRID**
- **Looking ahead to maintenance mode**
 - **Researching virtualization**