



Jan Strube

Rutherford Appleton Laboratory

SID ON THE GRID EXPERIENCE AND OUTLOOK



LOI and PostLOI

- 73 million events processed for the LOI and IDAG follow-up questions
 - + beam backgrounds
- Generation (whizard)
 - Done by SLAC centrally
- Simulation (SLiC)
- Reconstruction (org.lcsim, MarlinReco)
- User Analysis (AIDA, ROOT)
- SLAC started in September '08
- RAL joined in December '08



SLAC Setup

FTP (wget, curl)

bsub/qsub

NFS

Huge advantage in accessibility of log files

POSIX access

Limited by CPU availability



RAL Setup (steps)

- Grid-enabled Storage (Castor)
- Get grid certificate
- Sign up with (DESY, not Fermilab) ILC Virtual Organisation (VO)
 - Not more than one VO per certificate, though!
- Glite-wms-job-submit or edg-job-submit ?
 - Google is not always your friend
 - GridPP pages are a much better entry point (now)
- File Checksums
 - not something I worried about on NFS or FTP
- Limited by disk throughput
 - (Only used RAL storage)



Learning curve

Initial solution

- Manual submission
 - Manual bookkeeping
 - Log files 'hard to parse' at best
- Indiscriminate submission of jobs
- Installation of software in each sandbox
- **84.7 % efficiency**

Current solution

- Ganga
 - Better bookkeeping
 - Better logfiles
- Whitelisting of trusted sites
- SidSoft in site-wide area
- **99.8 % efficiency**



Software installation


- SLiC v2r5p3_geant4-v9r1p2
 - Fairly compiler-agnostic, just called SimDist
- org.lcsim 1.5
 - LHC needs old compilers
 - org.lcsim needs at least Java 1.5
 - Ship a Java runtime
 - Drop in the libraries and run
- MarlinReco: ILC SOFT v01-041.6
 - Validated by E. Devetak for SiD
 - Flavor tagging tuned by T. Lastovicka





Production Stats

25 different data sets

Production Step	# events
Simulation (SLiC)	42,611,000
Reconstruction	31,991,000
Flavor Tagging	20,446,000



CPU times (for the jobs for which I have this info)



Jan Strube
 Rutherford Appleton Laboratory




Tape inventory

sample	stdhep	SLiC	reco	MarlinReco
SUSY @ 500 GeV	5147	5147	5147	0
sbottom	990	990	926	926
SM @ 250 GeV	936	936	7878	7878
SM @ 500 GeV	575	575	7182	7182
BhaBha sample	6317	6317	6296	0
two-photon backgrounds	0	26454		
top (174 GeV)	0	0	1109	1109
top (174 GeV) w/ bg	0	1034	999	999
top(173.5 GeV)	0	0	1072	1072
Higgs samples (recoil, triple H coupling)	1118	1118	1382	1280
SUM	15083	42571	31991	20446

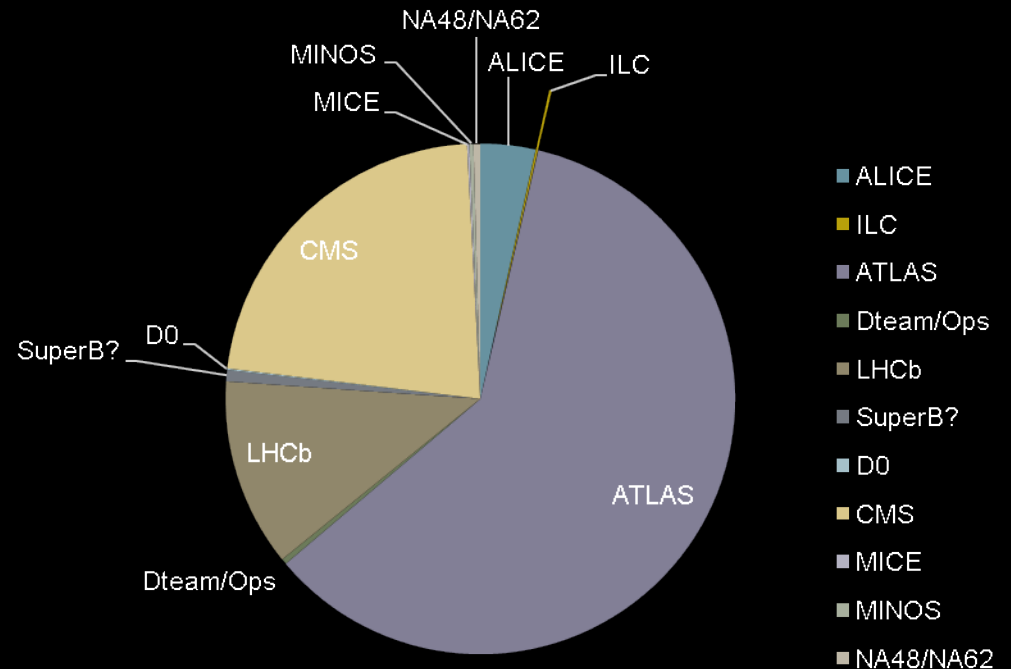


RAL Tier1 schedule 2010

CPU allocation



Disk / Tape allocation





Current Status of 'Middleware'

- 'sidsoft' directory for software installation
 - Currently contains only LOI versions
 - Submission scripts have benefited greatly from LOI experience
 - Stability (executable and copy operation)
 - List of trusted sites
- Efficiency: 99.4% in May (postLOI)
- 2009 avg.: ATLAS 87 %, CMS 77 %



Current and Future Developments

- File Catalog
 - Consistency between files at SLAC / in the RAL Tier1 castor
 - Production steps / metadata
- Production / Bookkeeping tools
 - Make sure the queues stay filled
 - Automatically, if possible...
- Grid liaison
 - Talking to people in person proved very useful during LOI production
 - Keep track of software versions / validations of external packages, publication of software tags at designated sites
 - Coordinate w/ ILD
- File size requirements
 - Castor likes larger files (~5 GB)
- Exploit concurrency



Summary

- LHC Computing Grid played an essential role in the SiD LOI effort
- LOI production was completed
 - Lauded in IDAG report
 - Without prior experience
 - Helped by LHC misfortune
- The LCG is a vast resource
 - We'd be wise to use it
 - But should use it wisely



THANK YOU



BACKUP




Imperial College

- New switches - this should enable us to integrate 300 TB of storage seamlessly into the cluster.
- dCache - Lots of small issues:
Hardware tweaks/headnode transferred to more powerful machine.
Suspended files not automatically retried.
- dCache has been the main cause of our downtimes (hotly followed by College turning off the power on Saturdays: no weekend coverage of Tier 2s)



Tier 1 outages affecting ILC in past year

- 30 Scheduled 'At Risk' events
 - Time: 37 days, 8 hours
- 18 Unscheduled 'At Risk' events
 - Time: 18 days, 14 hours
- 23 Scheduled 'Outage' events
 - Time: 38 days, 21 hours
- 39 Unscheduled 'Outage' events
 - Time: 37 days, 8 hours





The others

- ATLAS, CMS:
 - Infinite resources
- MINOS: Already working on ganga-based submission script
 - 50 Million events
- MICE: N / A
- H1: low-scale production
- BaBar: Using resources (CPU, Disk), but not middleware





Data Rates

- 1 event + several underlying events every 25 ns
- 100s of particles in every event
- Event size: few MB
- 1 event every 200 ns
- 10s of particles in every event
- Event size: ~50 kb



1 GB / s

LHC



1MB / s

ILC