

IN2P3 Computing Centre

LCG Tier1 Status And Plans

Fabio Hernandez

LCG Tier-1 Technical Leader

IN2P3 Computing Centre

fabio@in2p3.fr

LCG-GDB Lyon, March 16th, 2005



Contents

- Deployment status
- Site usage by LHC experiments
- Involvement in grid operations
- Work in progress
- Planned work for 2005
- Summary
- Questions

Deployment status

- **Context**
 - Using a shared batch facility, both for grid and non grid jobs
 - ♦ scheduling is based on each VO's CPU share, regardless of the submission method (i.e. grid or local)
 - We support the 4 LHC experiments
 - ♦ although their shares are not equal
- **Computing Elements**
 - RH7.3: 50 dual-processor worker nodes
 - ♦ We intent to phase this CE out by the end of march
 - SL3.0.3: 300 dual-processor worker nodes
 - ♦ All the Linux-based batch facility is usable through LCG middleware
 - ♦ The 50 WNs currently running RH7.3 will be added to this CE
 - Both managed by BQS (our local LRMS)
 - ♦ Operations people (including on-call team) monitor the grid and the non-grid production

Deployment status (cont.)

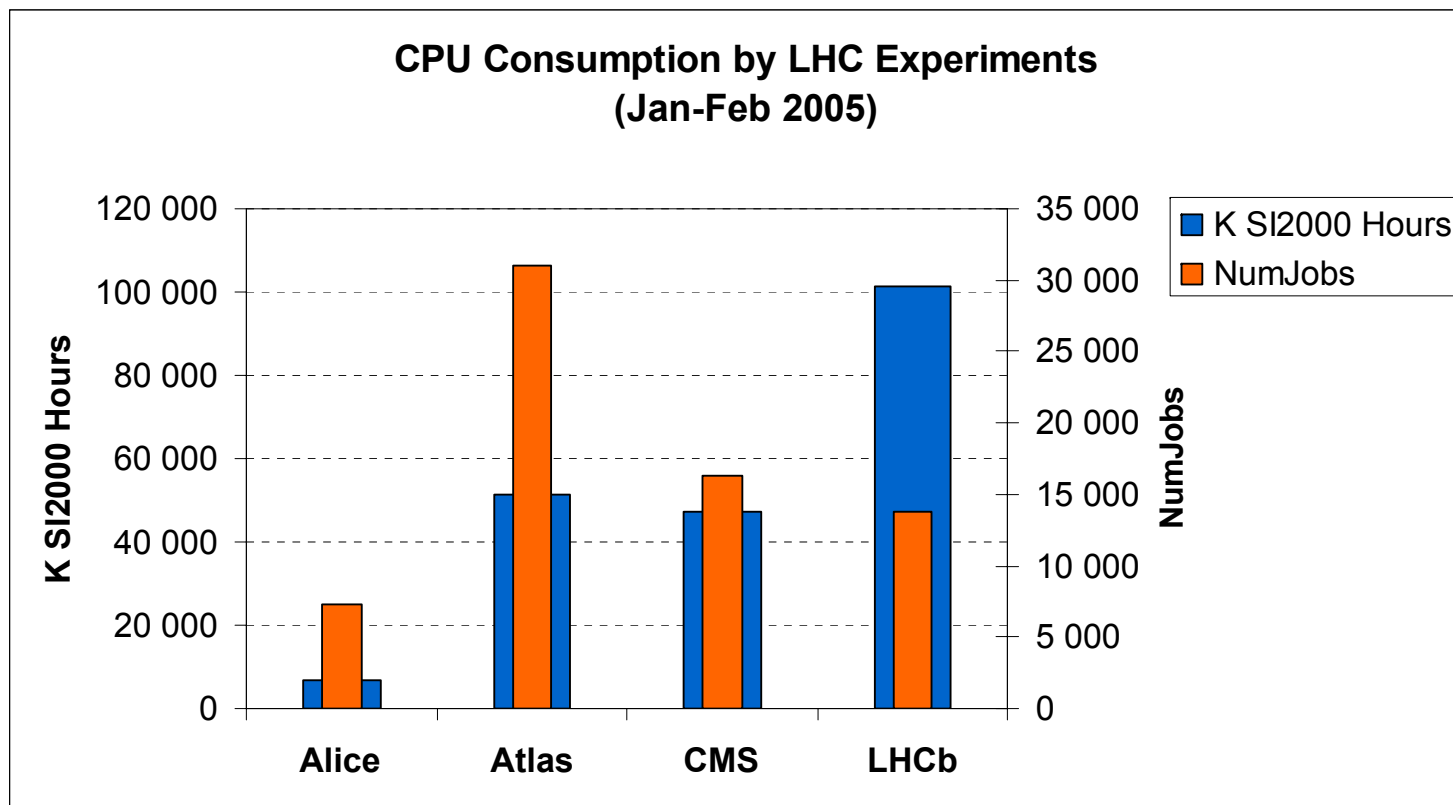
- **Storage Elements**
 - HPSS-backed SE (gridFTP server interfaced with HPSS)
 - Classic SE (~2.5 TB)
 - Import/export storage element (access to HPSS and semi-permanent storage)
 - ◆ used by CMS, but not declared in the site information system
- **Non-LHC supported VOs**
 - dzero, virgo, biomed, esr, egeode, cdf (being deployed)
- **Core grid services for non LHC VOs**
 - VO management service for biomed and egeode
 - Replica management service for biomed

Deployment status (cont.)

- Integration of local tools/environment
 - Resource management system
 - ◆ Jobmanager and CE information provider
 - Collection of accounting data
 - ◆ Extracted from BQS accounting data base, formatted and sent to GOC
 - AFS for experiment software installation
 - ◆ Same software shared by all worker nodes
 - HPSS as storage element back-end
 - ◆ Modifications to gridFTP server
 - Integration with the local installation tool
 - ◆ for installing the worker nodes
- Every middleware release are deployed in our integration testbed before going to production
- Preparing deployment of LCG 2.3.1
 - Certification authorities related modifications are already deployed in production

Resource Usage for 2005

- LHC usage in Jan-Feb 2005: 18% of computing centre's capacity



Involvement in LCG/EGEE operations

- Developing and operating the portal for Core Infrastructure Centres (CIC)
 - <http://cic.in2p3.fr>
 - Definition of operational procedures and tools for supporting them
 - ◆ Active participation to the weekly GDA meeting
 - ◆ Developing tools for collecting information on the available grid services and providing a view according to several target audiences
 - End-user, VO, ROC, CIC on duty, site administrator, ...
 - ◆ Working with Piotr for improving the test service
 - Tools for integrating the tests results and triggering tickets through GGUS
 - This is being used in production from this week on (we are CIC on duty this week)

LCG/EGEE operations (cont.)

- Core Infrastructure Centre (cont.)
 - Actively contributing to the CIC-on-duty activity
 - ◆ with CERN, CNAF, RAL
 - ◆ This gives us the opportunity to use the operational procedures/tools and improve them
- Coordination of EGEE/LCG operations in France
 - Including 2 LCG Tier-2 candidates
 - Support for 10 more EGEE sites
- Hosting the EGEE all ROC managers meeting on march 17th and 18th

LCG/EGEE operations (cont.)

Core Infrastructure Center (CIC) Portal
 Communication portal between Central Operations to Users, VOs and Sites

HOME | VO MB VIEWS | VO VIEWS | RC VIEWS | ROC VIEWS | **CIC VIEWS**

- CIC VIEWS -

News
 Contacts
 On Duty Dashboard
 On Duty Procedures
 Operation Metrics
 Resources Infos
 Services Infos
 VO Infos
 Problem Tracking
 Publish info

Operational Procedures

► View **Current Operational Procedures Draft** (written by Markus Schultz and Piotr Nyczyk)

Daily Tasks

This section is an extract of the above document (Appendix A). It describes the regular task an operator has to carry out and how the operation team should prepare for the task.

Activities

The operations team is expected to watch a few things continuously. Here are the most important ones:

- Watch the LCG-ROLLOUT mailing list and react to operational problems mentioned there. The list can have a high trac and not every thread is related to operation. However, the operations team is expected to skim through the messages on the list. Check the list every hour.
- Keep a browser open with the GIS Monitor and look for sites that are dropping out of the information system. If a site is spotted a look at the history of published values will give an indication whether the services are down, or have just some transitional problem. Contact the site if needed. You should have a look at this page at least every 2 hours. Don't forget to have a look at the service page.
- Once a day go through the GIS Monitor reports and look for inconsistencies or missing services.
- Keep a browser open with the 'Live Job Monitor' and look for suspicious behavior. This can be sets of Jobs failing, large number of jobs queued on one site etc.. Try to understand the situation. Take a short glance at this display every 30 minutes.
- At around 10:30 go through the list of the VSite Tests Reports* and compare what you see there with the tasks in the Savannah tracking system. Do the followup as described in the escalation procedures. If you receive some feedback you should learn the facts for the specific sites.
- In the morning you should have a look at the 'Certificate Lifetime' monitor and trigger an alarm if you spot sites with host certificates valid for less than a week.

Work in progress

- Evaluating dCache + SRM + HPSS
- Contributing to service challenges
- Improving tools for local integration and validation of middleware releases
- Improving/adapting tools for the site grid operations
 - Including tools for standard operations and on-call people

Planned work for 2005

- Upgrade the site's infrastructure
 - Power & cooling: work started on 2004 and are expected to last until 2006
- Upgrade the network link to CERN to 10 Gbps by September-October 2005
- Acquisition of storage hardware
 - 150 TB
 - Tape drives and servers
- Involvement in LCG-3D project
 - Currently remotely attending the meetings
 - Active work to be started by the end of 2Q2005
 - ♦ Currently working on migration of DB infrastructure to a cluster of Oracle 10g

Planned work for 2005 (cont.)

- Expressed strong interest in taking part in the pre-production testbed
 - We need to deploy early to have the opportunity of doing our local integration
 - We can provide feedback on the “adaptability” of gLite to (“exotic”) site-specific constraints

Summary

- Significant contribution to the operations of LCG/EGEE
 - A lot of work needs to be done
- Expecting major changes with gLite
 - Both at the site level and for the whole grid operations

Questions

