



WLCG: The Experiments' View

Graeme Stewart

(with a lot of help from colleagues in experiment computing teams; particular thanks to Roberto Santinelli, Patricia Mendez Lorenzo, Josep Flix and Stephane Jézéquel)

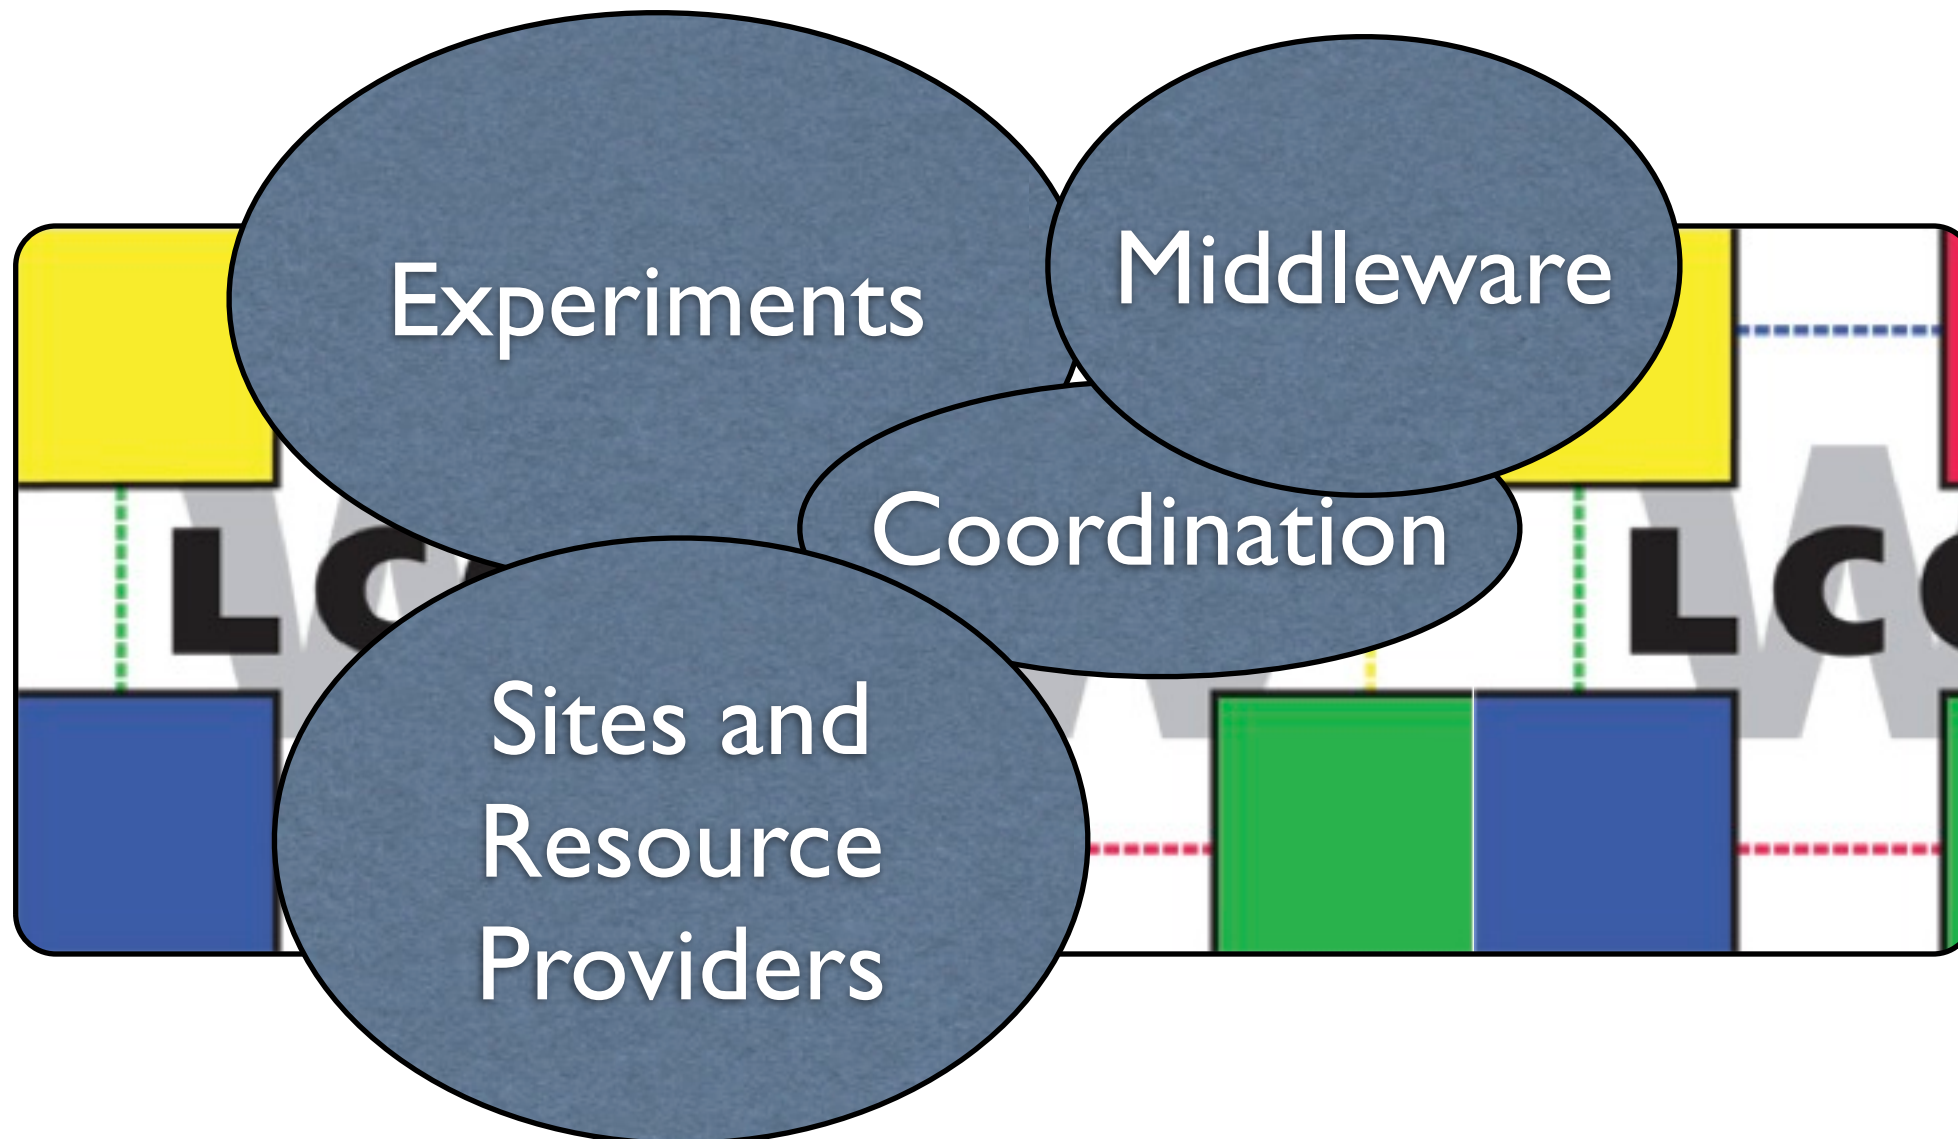


Overview

- I'll try to:
 - Describe the interactions between the experiments and sites
 - Showing how WLCG helps
 - Consider the impact of LHC data on operations
 - Bring up some specific issues experiments have

WLCG

“The Worldwide LHC Computing Grid (WLCG) is a global collaboration of more than 170 computing centres in 34 countries. The mission of the WLCG project is to build and maintain a data storage and analysis infrastructure for the entire high energy physics community that will use the Large Hadron Collider at CERN.”





Communications

- Experiments use a range of channels to communicate with sites
 - from basic setup information
 - CIC Portal
 - Twiki pages
 - to ticketing systems
 - GGUS
 - to dedicated operational meetings
 - WLCG and Experiment

- Some information for all VOs
- Most complete for LHCb, sparse for ALICE
- Additional sources of information reflecting complexity of setup
- And the fact that we use multiple grids

LHCb

| | | |
|--|-------------------|---|
| Number of Users | 130 | ? |
| Job CPU limit | 4500 min/1000SI2K | ? |
| Job Wall Clock Time Limit | 0 mn | ? |
| Job Scratch Space (GB) | 10 GB | ? |
| Per-SE Storage Space (GB) | 0 GB | ? |
| RAM per CPU core (MB) on 32 bits server | 750 MB | ? |
| RAM per CPU core (MB) on 64 bits server | 1000 MB | ? |
| Required virtual memory (MB) on 32 bits server | 1500 MB | ? |
| Required virtual memory (MB) on 64 bits server | 2000 MB | ? |
| Minimum number of job slots needed | 0 | ? |

Other Requirements ?:

ATLAS

Software installation requirements:

- an experimental software area (shared filesystem) with at least 250 GB free and reserved for ATLAS;
- the full compiler suite (c/c++ and fortran) should be installed in the WNS, including all the compat-gcc-32* and the SL_libg2c.a_change packages in SL4-like nodes;
- the recommended version of the compilers is 3.4.6;
- the f2c and libgfortran libraries (in both i386 and x86_64 versions, in case of x86_64 systems) are also required to run the software;
- other libraries required are:
 - libpopt.so.0
 - libblas.so
- other applications required are: uuencode, uudecode, bc, curl;
- high priority in the batch system for the atlasgsm user;
- for nodes running at 64 bits, a copy of python compiled at 32 bits is also needed to use the 32 bits python bindings in the middleware. See <https://twiki.cern.ch/twiki/bin/view/Atlas/RPMcompatSLC4> for more details;
- for SL5 nodes please refer to <https://twiki.cern.ch/twiki/bin/view/Atlas/RPMCompatSLC5> and <https://twiki.cern.ch/twiki/bin/view/Atlas/SL5Migration>.

Site Oriented Documentation

- Used by ALICE, ATLAS and CMS to discuss aspects of site setups
- Reflects some of the complexity of LHC experiment workflows and specific components used by each experiment

SL5Migration

- ↳ [Introduction](#)
- ↳ [SL4 Kits on SL5](#)
 - ↳ [Base Install](#)
 - ↳ [YUM Metapackages](#)
 - ↳ [Compatibility libraries](#)
 - ↳ [ATLAS Software Infrastructure Team \(SIT\) list](#)
 - ↳ [LCG Applications Metapackage](#)
 - ↳ [Kickstart Files](#)
 - ↳ [US ATLAS Kickstart File](#)
 - ↳ [SELinux](#)
 - ↳ [Software Areas](#)
 - ↳ [EGEE](#)
 - ↳ [SL4 and SL5 Resources](#)
 - ↳ [Installable Kits](#)
- ↳ [Virtual Memory](#)
- ↳ [Post Migration](#)
- ↳ [Other Resources](#)

Introduction

This topic describes and collates knowledge necessary for ATLAS grid site

The screenshot shows the AliEn website interface. At the top, there is a dark blue header with the 'AliEn @GRID' logo. Below the header is a navigation bar with links for 'About', 'Development', and 'MonALISA ALICE Grid Monitor'. The main content area is divided into two columns. The left column, titled 'AliEn', contains a list of links: Home, User distribution, Download, Documentation, Howto, User, Site (highlighted), Exclusive SEs, Install LCG VO Box, Install AliEn Site, Install Xrootd, VO Admin, Developer, FAQ, Publications, and Virtual Organizations. The right column, titled 'AliEn Site Howto', contains a list of links: Exclusive SEs, Install LCG VO Box (with sub-links for Debug, Manage Proxies, Configure LDAP, and 64 Bit VO Box), Install AliEn Site, and Install Xrootd.

Savannah

- In general used by experiments to track issues which are often internal
- Used more widely by CMS to raise site issues to site squads
- And by ATLAS to manage site storage (lost files to SE migrations)
- Better issue tracker and easier to involve multiple parties



The screenshot shows a Savannah bug tracker interface. At the top left is a bomb icon. The main heading is "DQ2-DDM Operations - Bugs: bug #69409, Change SE name for UKI-LT2-IC-HEP". Below this is a navigation bar with links: Group, Main, Download, Docs, Support, Mailing Lists, Source Code, Bugs (highlighted), Tasks, and News. Below the navigation bar, it says "You are [manager](#) for this tracker." At the bottom, there is a red-bordered box containing the text "bug #69409: Change SE name for UKI-LT2-IC-HEP".

GGUS



- Global Grid User Support
- Used more for experiments to ticket sites
 - But sites can contact experiments this way
- Ticket priorities seem to have different interpretations:
 - For LHCb it is the priority of the issue at that site; for ATLAS it is the overall priority of the issue to the experiment
- Standard channel for error reporting
 - CMS use a savannah-ggus bridging tool

Ticket Types

- Team tickets have helped experiments use GGUS in a coherent way a lot
- Ticket can be opened, modified, verified by any member of the experiment operations team
- Alarm tickets used for critical issues at Tier-0 and Tier-I sites
 - Ensure that on-calls receive an alert message about an important and immediate issue
 - But ALARM ticket chain has failed a number of times

▶ Open **TEAM** ticket ▶ Open **ALARM** ticket

Submit ticket

User information

| | | | |
|---------------------|--|---------|--|
| Name | Graeme Stewart | E-Mail | <input type="text" value="graeme.andrew.stewart@cern.ch"/> |
| Notification mode ? | <input checked="" type="radio"/> on solution <input type="radio"/> on every change <input type="radio"/> never | CC to ? | <input type="text"/> |

WLCG MoU

- A rather formal document between sites and WLCG
- Does set out the maximum delay in ‘responding to operational problems’
 - e.g., 12 hours for an interruption to data taking capacity at a T1 during accelerator running
 - e.g., 2 hours for problems with a ‘end-user analysis facility’ at a T2
- In practice T1s and many T2s are largely perceived as offering a good service to experiments
 - ‘fix yourself as quickly as is practically possible’

WLCG Daily Operations Call

- Happens at 3pm (CERN) time, Monday - Friday
- <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsMeetings>
- Attended by experiments and TIs
 - Little direct T2 participation
- For discussing immediate problems and tracking progress towards their resolution
- Perceived by all experiments as very useful

WLCG Tier-1 Service Co-ordination

- Biweekly meetings between with experiments and Tier-1s in attendance
- Deals with longer term issues oriented at T1 services, e.g., FTS upgrades, database issues
- Has been very useful to track deployment progress

Grid Deployment Board

Grid Deployment Board

There are 4 more events in the *future*. [Show them.](#)

September 2010

 08 Sep [GDB](#)

August 2010

 11 Aug [GDB](#)
 10 Aug [Pre-GDB](#)

- Experiments attend each GDB and provide specific feedback every three months
- GDB follows up on longer term middleware deployment: e.g., CREAM CE
- Tends towards a review of operations in the previous months
- Some contentious issues raised:
 - LHCb - consistent publication of normalised CPU available
 - ATLAS - storage reliability during reprocessing campaigns

Experiment Meetings


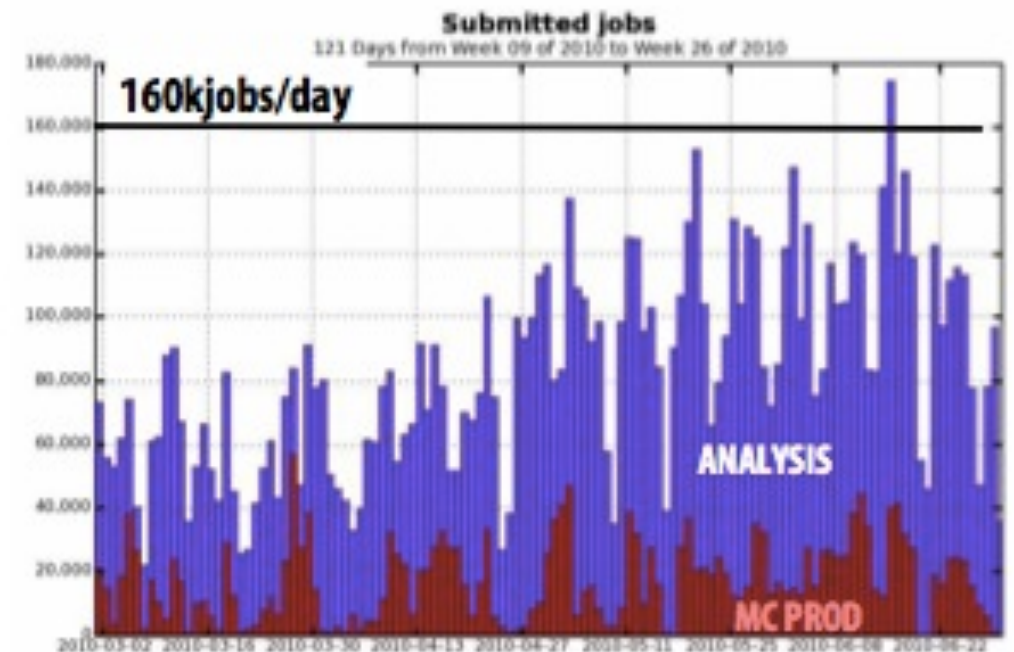
- All experiments hold dedicated computing operations meetings each week
 - See backup slides for details
 - Attendance from sites is usually good where there is a strong link between the site and the experiment
- Opportunity to discuss both immediate operational issues and experiment and site planning
 - e.g., major site upgrades and downtimes, new requests for site support or services

Experiment Organisation

- Experiments have teams of shifters and experts working in monitoring grid computing
- In the larger experiments these are subdivided:
 - e.g., CMS: Data Operations, Facilities Operations, Analysis Operations
- It's certainly very useful for sites to know how to contact experiments in case of problems
 - Daily meeting can be used
 - LHCb encourage GGUS tickets to the experiment
 - ATLAS, CMS and ALICE have operations email lists and local or regional site contacts


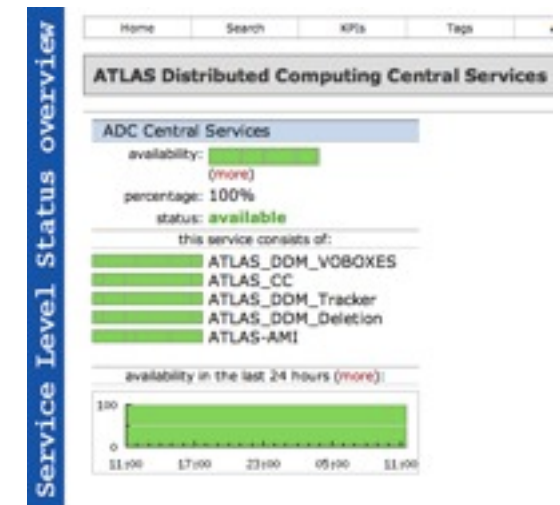
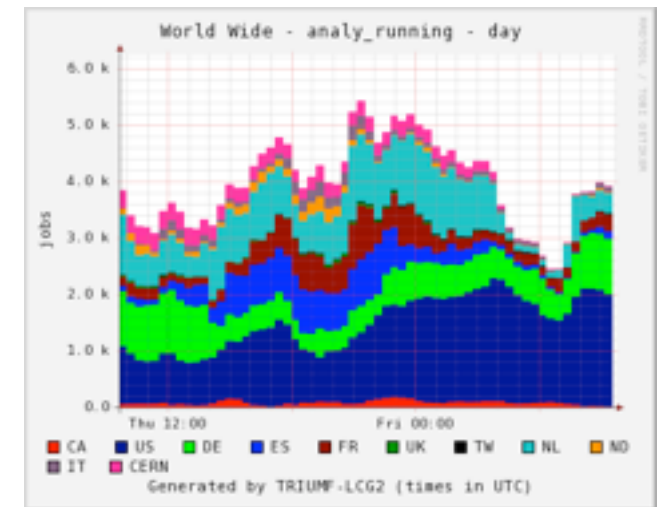
Challenge of LHC Data

- LHC data was the critical test of the WLCG infrastructure
- In general things are thought to have gone well, with sites responding actively to problems and early analysis successful on the grid
- Real activity was often in excess of pre-tests



Monitoring

- Common infrastructure of SAM tests provides basic validation of sites
- Different experiment monitoring frameworks probe at deeper levels
- Pro-active sites monitor using the experiment frameworks
- It's probably rather difficult to get a coherent view across WLCG right now



| US Sites | Job Nodes | Jobs | Latest | Pilot Nodes | defined | assigned | waiting | activated | sent | running | holding | transferring | finished | failed | tot | trf | other |
|-------------|-----------|------|-------------|-------------|---------|----------|---------|-----------|------|---------|---------|--------------|----------|--------|-----|-----|-------|
| Site Name | 2051 | 961 | 04-09 09:33 | 3065 | 0 | 211 | 10 | 18252 | 1 | 9375 | 384 | 6761 / 1 | 15031 | 981 | 6% | 0% | 6% |
| AGLT2 | 285 | 43 | 04-09 09:33 | 478 | 0 | 0 | 0 | 4142 | 0 | 1900 | 73 | 3288 / 0 | 2210 | 43 | 2% | 0% | 1% |
| BNL-ATLAS_1 | 374 | 526 | 04-09 09:33 | 638 | 0 | 0 | 0 | 4728 | 0 | 2371 | 68 | 0 / 0 | 5533 | 526 | 9% | 0% | 9% |

Problem Areas

- Areas where there are non-self contained problems do to cause more difficulties
- Poor transfer rates between sites
 - T1-T1 transfers
 - Upload of data from T2s to T1
- Problems with authentication of a single user or a particular CA

Data Management

- It's still the case that data access remains the hard problem (reliable, robust, performant)
- Site storage stability and reliability are noted as being particularly problematic by ATLAS, CMS and LHCb
- Initiatives such as the Amsterdam Data Management Jamboree are welcome
- But effective management of the current system has to be done

Other Issues Raised

- Software areas at sites (LHCb)
- Interference between services provided to multiple experiments (CMS)

Conclusion

- WLCG does work as a community to provide working distributed computing to the LHC experiments
- Providing resources and service on this scale is not trivial
- The current system is labour intensive and relies on key experts in many areas
- The system will have to scale up (definitely) and become cheaper (probably)
- Can we provide a community which better serves the LHC's computing needs?

Some Items to Discuss

- MoU and site response to problems
- Data Management Stability
- Software areas
- Systems interference
- Consistent CPU publishing

Backup Slides

Experiment Computing Information

- ALICE: <http://alien2.cern.ch/>
- ATLAS: <https://twiki.cern.ch/twiki/bin/view/Atlas/AtlasDistributedComputing>
- CMS: <http://cms-computing.web.cern.ch/cms-computing/index.html>
- LHCb: <https://twiki.cern.ch/twiki/bin/view/LHCb/LHCbComputing#Operations>

WLCG Meetings

- Daily Operations Meeting: <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsMeetings>
- Tier-I Service Co-ordination: <http://indico.cern.ch/categoryDisplay.py?categId=2726>
- Grid Deployment Board: <http://indico.cern.ch/categoryDisplay.py?categId=31181>
- WLCG Workshops: <http://indico.cern.ch/categoryDisplay.py?categId=890>