



WLCG Update



6th May, 2008


HEPiX – Spring 2008
CERN

Ian Bird
LCG Project Leader





LCG Organisation – Phase 2





WLCG Collaboration



- The Collaboration
 - 4 LHC experiments
 - ~130 computing centres
 - 12 large centres (Tier-0, Tier-1)
 - 56 federations of smaller "Tier-2" centres – 121 sites
 - Growing to ~40 countries
 - Grids: EGEE, OSG, Nordugrid
- Technical Design Reports
 - WLCG, 4 Experiments: June 2005
- Memorandum of Understanding
 - Agreed in October 2005
- Resources
 - 5-year forward look

MoU Signing Status

Tier 1 – all have now signed

Tier 2:

Australia	Netherlands	<u>Still to sign:</u>
Belgium	Norway *	Austria
Canada *	Pakistan	Brazil (under discussion)
China	Poland	
Czech Rep. *	Portugal	
Denmark	Romania	* Recent additions
Estonia	Russia	
Finland	Slovenia	
France	Spain	
Germany (*)	Sweden *	
Hungary *	Switzerland	
Italy	Taipei	
India	Turkey *	
Israel	UK	
Japan	Ukraine	
JINR	USA	
Korea		




LCG Service Hierarchy



Tier-0: the accelerator centre

- Data acquisition & initial processing
- Long-term data curation
- Distribution of data → Tier-1 centres



Tier-1: “online” to the data acquisition process → high availability

- Managed Mass Storage –
→ grid-enabled data service
- Data-heavy analysis
- National, regional support

Tier-2: ~130 centres in ~35 countries

- Simulation
- End-user analysis – batch and interactive







Recent grid use

- Across all grid infrastructures
- Preparation for, and execution of CCRC'08 phase 1
 - Move of simulations to Tier 2s

CPU Usage Jan-Feb 2008




Tier 2 federation use





Recent grid activity



- WLCG ran ~ 44 M jobs in 2007 – workload has continued to increase
 - 29M in 2008 – now at ~ >300k jobs/day
- Distribution of work across Tier0/Tier1/Tier 2 really illustrates the importance of the grid system
 - Tier 2 contribution is around 50%; > 85% is external to CERN

- These workloads (reported across all WLCG centres) are at the level anticipated for 2008 data taking




Combined Computing Readiness Challenge - CCRC'08



- Objective was to show that we can run together (4 experiments, all sites) at 2008 production scale:
 - All functions, from DAQ $\Rightarrow$ Tier 0 $\Leftrightarrow$ Tier 1s $\Leftrightarrow$ Tier 2s
- Two challenge phases were foreseen:
 1. **Feb:** not all 2008 resources in place – still adapting to new versions of some services (e.g. SRM) & experiment s/w
 2. **May:** all 2008 resources in place – full 2008 workload, all aspects of experiments' production chains
- Agreed on specific targets and metrics – helped integrate different aspects of the service
 - ❑ Explicit “**scaling factors**” set by the experiments for each functional block (e.g. data rates, # jobs, etc.)
 - ❑ Targets for “**critical services**” defined by experiments – essential for production, with analysis of impact of service degradation / interruption
 - ❑ WLCG “**MoU targets**” – services to be provided by sites, target availability, time to intervene / resolve problems ...




LHC OPN





Data transfer


- Data distribution from CERN to Tier-1 sites
 - The target rate was achieved in 2006 under test conditions
 - Autumn 2007 & CCRC'08 under more realistic experiment testing, reaching & sustaining target rate with ATLAS and CMS active





Castor performance - Tier 0

Migration
policy mod.



t0export: in from WN, out to tape

Recovered


Ran out of
tapes

■ CMS:

- Aggregate rates in/out of castor of 3-4 GB/s
- Sustained rate to tape 1.3 GB/s with peaks > 2 GB/s

■ May:

- Need to see this with all experiments





Resource ramp up for 2008




- CPU:
 - Most Tier 1 sites will have full 2008 pledges in place for 1 May
 - Total of 36725 KSi2K.
 - Largest missing is +2500 at NL-T1 due Nov.
- Disk and tape
 - Many sites will catch up later in the year as need expands:
 - 2008 disk requirements are 23 PB with 15.5 PB expected by 1 May
 - 2008 tape requirements are 24 PB with 15 PB expected by 1 May.
- May run of CCRC'08 at 55% only requires +1PB of disk and +1.5PB of tape (mostly reusable) so should have no resource problems.
- Full status of resource installation will be reported at C-RRB in April.
- Many sites had problems with procurement process/ vendor delivery/ faulty equipment
 - These issues must be taken into account in future – the process is long, but yearly deadlines are important



Tier 0/Tier 1 Site reliability

- Target:
 - Sites 91% & 93% from December
 - 8 best: 93% and 95% from December
- See QR for full status



	Sep 07	Oct 07	Nov 07	Dec 07	Jan 08	Feb 08
All	89%	86%	92%	87%	89%	84%
8 best	93%	93%	95%	95%	95%	96%
Above target (+>90% target)	7 + 2	5 + 4	9 + 2	6 + 4	7 + 3	7 + 3





Tier 2 Reliabilities

- Reliabilities published regularly since October

Overall	Top 50%	Top 20%	Sites
76%	95%	99%	89→100

- In February 47 sites had > 90% reliability



- For the Tier 2 sites reporting:


Sites	Top 50%	Top 20%	Sites > 90%
%CPU	72%	40%	70%




Reliability...




- Site reliability/availability affects resource delivery and usability of a site
 - This has been ***the*** outstanding problem for a long time (slow improvement of reliability)
- Addressed through human oversight (Grid Operator on Duty) ...
 - Teams of experts on duty to flag problems and follow up with sites
 - In place since late 2004; effort intensive
 - Instrumental in stabilisation and gradual improvement of reliability
 - Unsustainable in the long term
- ... and though better monitoring ...
 - Monitoring tools (many!)
 - Aggregating information
 - Understanding (visualising) the data
 - Automation





GS

Messaging based archiving and reporting

CERN IT
Department


- Top-level visualization of distributed systems*
- Prototype deployed since EGEE'07

1st version Oct 2007, 2nd version Nov 2007, Link: <http://gridmap.cern.ch>

region

site

- size = #CPUs
- colour = status or availability



details, links to monitoring tool

- GridMaps visually correlate "importance" to status



Tier 2 Communication



- How are the Tier 2s kept informed, and does it work?
 - Flow of information from Management Board, - do Tier 2s read the minutes?
 - Is everyone engaged in the GDB (or even aware that they can be)?
- How can we structure the communication with the great number of Tier 2 sites, so that we can have a workable process to communicate problems and follow up (in both directions)??
- How can we aggregate Tier 2 status to report in LHCC/OB/RRB/CB etc?
 - Today it is extremely difficult to get an overview of Tier 2 status and problems
- Agreement in the CB is to use the Tier 1s to coordinate communication to / from the Tier 2s



Progress in EGEE-III



- EGEE-III now approved
 - Starts 1st May, 24 months duration (EGEE-II extended 1 month)
 - Objectives:
 - Support and expansion of production infrastructure
 - Preparation and planning for transition to EGI/NGI
 - Many WLCG partners benefit from EGEE funding, especially for grid operations: effective staffing level is 20-25% less
 - Many tools: accounting, reliability, operations management funded via EGEE
 - Important to plan on long term evolution of this
 - Funding for middleware development significantly reduced
 - Funding for specific application support (inc HEP) reduced
- Important for WLCG that we are able to rely on EGEE priority on operations, management, scalability, reliability



Operations evolution



- Anticipates ending EGEE-III with the ability to run operations with significantly less (50%?) effort
 - Moving responsibility for daily operations from COD to ROCs and sites
 - Automation of monitoring tools – generating alarms at sites directly
 - Site fabric monitors should incorporate local and grid level monitoring
 - Need to ensure all sites have adequate monitoring
 - Need to provide full set of monitoring for grid services
 - The manual oversight and tickets should be replaced by automation
 - Remove need for COD teams
 - Operations support should have streamlined paths to service managers
 - Eliminate need for TPMs for operations
 - Need to insist on full checklist (sensors, docs, etc) from middleware
 - This process is the subject of formal milestones in the project



Other operational aspects



- User support
 - Streamlining process for operations
 - Focus of TPM effort on real user “helpdesk” functions – TPMs now explicitly staffed by ROCs (were not in EGEE-II)
- Security
 - Operational security team also have explicit staffing from regions to ensure adequate coverage of issues
 - Focus on implementing best practices at sites
- Policy
 - Efforts should continue at the current level



Middleware ?



- In EGEE-III the focus on middleware is the support of the foundation services
 - These map almost directly to the services WLCG relies on
 - Should include addressing the issues with these services exposed with large scale production
 - Should also address still missing services (SCAS, glExec, etc)
 - Should also address the issues of portability, interoperability, manageability, scalability, etc.
 - Little effort is available for new developments
 - (NB tools like SAM, accounting, monitoring etc are part of Operations and not middleware)
- Integration/certification/testing
 - Becomes more distributed – more partners are involved
 - In principle should be closer to the developers
 - ETICS is used as the build system
 - Probably not enough effort to make major changes in process



EGEE Operations in 2010 ...



- Most operational responsibilities should have devolved to
 - The ROCs, and
 - The sites – hopefully sites should have well developed fabric monitors that monitor local and grid services in a common way and trigger alarms directly
 - Receive triggers also from ROC, Operators, VOs, etc.; and tools like SAM
- The central organisation of EGEE ops (the “OCC”) should become:
 - Coordination body to ensure common tools, common understanding, etc.
 - Coordination of operational issues that cross regional boundaries
 - The ROCs should manage inter-site issues within a region
 - Maintain common tools (SAM, accounting, GOCDB, etc.)
 - Of course, the effort in these things may come from ROCs
 - Integration/testing/certification of middleware (SA3)
 - Monitor of SLA's, etc. (and provide mechanisms for WLCG to monitor MoU adherence)
 - ...



... Operations in 2010




- Ideally we should have an operational model with daily operational responsibility at the regional or national level (i.e. the ROCs)
- This will make the organisational transition to NGIs simpler – if the NGIs see their role as taking on this responsibility ...



Sustainability



- Need to prepare for permanent **Grid infrastructure**
- Ensure a high quality of service for all user communities
- Independent of short project funding cycles
- Infrastructure managed in collaboration with National Grid Initiatives (NGIs)
- European Grid Initiative (EGI)






Summary




- We have an operating production quality grid infrastructure that:
 - Is in continuous use by all 4 experiments (and many other applications);
 - Is still growing in size – sites, resources (and still to finish ramp up for LHC start-up);
 - Demonstrates interoperability (and interoperation!) between 3 different grid infrastructures (EGEE, OSG, Nordugrid);
 - Is becoming more and more reliable;
 - **Is ready for LHC start up**
- For the future we must:
 - Learn how to reduce the effort required for operation;
 - Tackle upcoming issues of infrastructure (e.g. Power, cooling);
 - Manage migration of underlying infrastructures to longer term models;
 - Be ready to adapt the WLCG service to new ways of doing distributed computing.

Scheduled = 21539
Running = 25374



21:13:50 UTC



GridPP
UK Computing for Particle Physics