



# WLCG Status Report

CERN-RRB-2008-038

15<sup>th</sup> April, 2008  
Computing Resource  
Review Board



Ian Bird  
LCG Project Leader

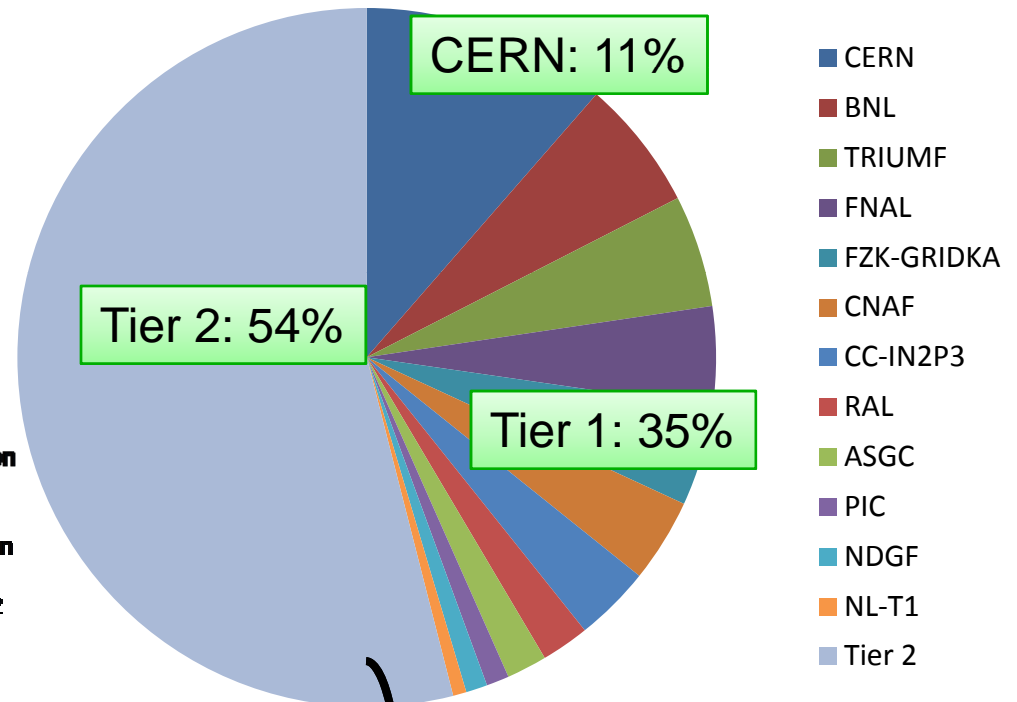





# Recent grid use

- 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




Ian.Bird@cern.ch

**Federations not yet reporting:**

- Finland
- India (IN-INDIACMS-TIFR)
- Norway
- Sweden
- Ukraine



# Recent grid activity



- 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 ...



# SRM v2.2 Deployment


- Deployment plan was defined and agreed last September, but schedule was very tight
- Deployment of dCache 1.8.x and Castor with srm v2.2 was achieved at all Tier0/Tier 1 by December
  - Today 174 srm v2 endpoints are in production
- During February phase of CCRC'08 relatively few problems were found:
  - Short list of SRM v2 issues highlighted, 2 are high priority
  - Will be addressed with fixes or workarounds for May
  - Effort in testing was vital
- Still effort needed in site configurations of MSS – iterative process with experience in Feb & May





# Castor performance - Tier 0

Migration policy mod.




t0export: in from WN, out to tape

Recovered

Ran out of tapes


- CMS:
  - Sustained rate to tape 1.3 GB/s with peaks > 2 GB/s
  - Aggregate rates in/out of castor of 3-4 GB/s
- May:
  - Need to see this with all experiments





# Data transfer

- Each experiment sustained in excess of the target rates (1.3 GB/s) for extended periods.
  - Peak aggregate rates over 2.1 GB/s – no bottlenecks
- All Tier 1 sites were included





27 Mar 08

# 24x7 Support

Done	Late <1 mon	Late >1 mon
colour coding		

27-Mar-08		WLCG High Level Milestones - 2007													
ID	Date	Milestone	Done (green)					Late < 1 month (orange)				Late > 1 month (red)			
			ASG C	CC IN2P3	CER N	FZK GridKa	INFN CNAF	NDGF	PIC	RAL	SARA NIKHE F	TRIUM F	BNL	FNAL	
<b>24x7 Support</b>															
WLCG-07-01	Feb 2007	<b>24x7 Support Definition</b> Definition of the levels of support and rules to follow, depending on the issue/alarm													
WLCG-07-02	Apr 2007	<b>24x7 Support Tested</b> Support and operation scenarios tested via realistic alarms and situations				Apr 2008	Apr 2008				Jan 2008				
WLCG-07-03	Jan 2007	<b>24x7 Support in Operations</b> The sites provides 24x7 support to users as standard operations				Apr 2008	Apr 2008			Mar 2008	Mar 2008	Apr 2008			

Only 6→9 sites have tested their 24 X 7 support, and only 5→7 have put the support into operation

Must be in place for May; understood by Tier 1s now after February experience






# CCRC'08 (Feb) - results

- Preparation:
  - Focus on understanding missing and / or weak aspects of the service and in identifying pragmatic solutions
  - Main outstanding problems in the middleware were fixed (just) in time and many sites upgraded to these versions
  - The deployment, configuration and usage of SRM v2.2 went better than had predicted, with a noticeable improvement during the month
- Despite the high workload, we also demonstrated (most importantly) that we can support this work with the available manpower, although essentially no remaining effort for longer-term work
- If we can do the same in May – when the bar is placed much higher – we will be in a good position for this year's data taking
- However, there are certainly significant concerns around the available manpower at all sites – not only today, but also in the longer term, when funding is unclear



# 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%

- For Tier 2 sites not reporting, 12 are in top 20 for CPU delivered



# Reliability reporting

- Currently (Feb 08) All Tier 1 and 100 Tier 2 sites report reliabilities
- Recent progress: MB set up group to
  - Agreement on equivalence of NDGF tests with those used at EGEE and all other Tier 1 sites – now in production at NDGF
    - Should also be used for Nordic Tier 2 sites
  - Similar process with OSG (for US Tier 2 sites): tests only for CE so far, agreement on equivalence, tests are in production, publication to SAM in progress
    - Missing – SE/SRM testing
    - Expect full production May 2008 (new milestone introduced)
- Important that we have all Tier 2s regularly tested and reporting
- Important that we have correct Tier 2 federation contact to follow up these issues



# Applications Area

- Recent focus has been on major releases to be used for 2008 data taking:
  - QA process and nightly build system to improve release process
- Geant4 9.1 released in December
- ROOT 5.18 release in January
  
- Two data analysis simulation and computing projects in the PH R&D proposal (July 2007) (Whitepaper)
  - WP8-1 - Parallelization of software frameworks to exploit multi-core processors
    - Adaptation of experiment software to new generations of multi-core processors – essential for efficient utilisation of resources
  - WP9-1 - Portable analysis environment using virtualization technology
    - Study how to simplify the deployment of the complex software environments to distributed (grid) resources



# Progress in EGEE-III

- EGEE-III now approved
  - Starts 1<sup>st</sup> 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 priorities on operations, management, scalability, reliability**





# Comments on EGI design study

- Goal is to have a fairly complete blueprint in June
  - Main functions presented to NGIs in Rome workshop in March
- Essential for WLCG that EGI/NGI continue to provide support for the production infrastructure after EGEE-III
  - We need to see a clear transition and assurance of appropriate levels of support; Transition will be 2009-2010
    - Exactly the time that LHC services should not be disrupted
- Concerns:
  - NGIs agreed that a large European production-quality infrastructure is a goal
    - Not clear that there is agreement on the scope
    - Reluctance to accept level of functionality required
  - Tier 1 sites (and existing EGEE expertise) not well represented by many NGIs
- WLCG representatives must approach their NGI reps and ensure that EGI/NGIs provide the support we need



# Power and infrastructure

- Expect power requirements to grow with capacity of CPU
  - This is not a smooth process: depends on new approaches and market-driven strategies (hard to predict) e.g. improvement in cores/chip is slowing; power supplies etc. already >90% efficient
  - No expectation to get back to earlier capacity/power growth rate
- e.g. Existing CERN Computer Centre will run out of power in 2010
  - Current usable capacity is 2.5MW
  - Given the present situation Tier 0 capacity will stagnate in 2010
- Major investments are needed for new Computer Centre infrastructure at CERN and major Tier 1 centres
  - IN2P3, RAL, FNAL, BNL, SLAC already have plans
  - IHEPCCC report to ICFA at DESY in Feb '08







# Summary

- CCRC'08 first phase was seen as a success
  - SRM and MSS deployment was achieved
  - Project and experiments targets were achieved
- Preparations for CCRC'08 phase 2 under way
  - Will be a full test of the entire system – all experiments together
  - Tuning of tape access with real use patterns – may require experiments to reconsider analysis patterns
- Resource ramp-up: based on experiences and problems with 2008 procurements
  - Must ensure in future years that allowance is made for delays and problems
  - Important that the yearly April schedules are met – to be ready for accelerator start ups



# Summary

- Remaining Tier 2 federations must now ensure that they regularly report (and verify) accounting and reliability data
  - Important that we have the correct contact people for the Tier 2 federations
- WLCG – especially Tier 1s – should influence the directions of the EGI Design study
  - Must ensure that we see a clear and appropriate strategy emerging that is fully supported by the NGIs
  - Must engage the NGI representatives in this
- CERN and the Tier1s must ensure that their CC infrastructure can accommodate the pledged capacity beyond 2009