



# CCRC'08 Phase Two

[Jamie.Shiers@cern.ch](mailto:Jamie.Shiers@cern.ch)

~ ~ ~

LCG GDB, 2<sup>nd</sup> April 2008

# Introduction

- The 2<sup>nd</sup> phase of CCRC'08 is from Monday 5<sup>th</sup> May to Friday 30<sup>th</sup>.
  - This is immediately preceded by a long w/e (Thu-Sun) at CERN
- May is also a month with quite a few public holidays
  - ~1 per week...
- ↳ **Everyone is keen that we improve on February in terms of being ready well in time**
  - And in terms of defining what "ready" actually means...
- There were detailed presentations on middleware, storage-ware, databases, experiments plans & requirements at yesterday's F2F
  - We can summarize this in a couple of tables, on the CCRC'08 wiki, as was done for February
- **I will also mention communication – but at the end!**

↳ **N.B. I have not attempted to summarize the presentations of the experiments – please refer to these for further info!**

# CCRC'08 – Objectives

- Primary objective (**next**) was to demonstrate that we (sites, experiments) could run together at 2008 production scale
  - **This includes testing all “functional blocks”:**
    - Experiment to CERN MSS; CERN to Tier1; Tier1 to Tier2s etc.
- Two challenge phases were foreseen:
  1. **February :** not all 2008 resources in place – still adapting to new versions of some services (e.g. SRM – **later**) & experiment s/w
  2. **May:** all 2008 resources in place – full 2008 workload, all aspects of experiments' production chains

💣 **N.B. failure to meet target(s) would necessarily result in discussions on de-scoping!**

- Fortunately, the results suggest that this is not needed, although **much** still needs to be done before, and during, May!

# What Did We Achieve? (High Level)

- Even **before** the official start date of the February challenge, it had proven an extremely **useful** focusing exercise, in helping understand missing and / or weak aspects of the service and in identifying **pragmatic** solutions
- Although later than desirable, the **main bugs** 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 (**more later...**)
- 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 (**e.g. post EGEE III**)

# Main Lessons Learned

- ☺ **Generally, things worked reasonably well...**
- **Still improvements in communication are needed!**
  - Tools still need to be streamlined (e.g. elog-books / GGUS), and reporting automated
  - ¿ **Service dashboards – should be in place before May...**
  - F2Fs and other meetings working well in this direction!
- **Pre-established metrics extremely valuable!**
  - As well as careful preparation and extensive communication!
- ↳ **Now in continuous production mode – this will continue – as will today's infrastructure & meetings**

# Recommendations – mid-February!

- ✓ To improve communications with Tier2s and the DB community, 2 new mailing lists have been setup, as well as regular con-calls with Asia-Pacific sites (time zones...)
- Follow-up on the lists of “**Critical Services**” must continue, implementing not only the appropriate monitoring, but also ensuring that the WLCG “standards” are followed for **D**esign, **I**mplementation, **D**eployment and **O**peration
- **Clarify reporting and problem escalation lines (e.g. operator call-out triggered by named experts, ...) and introduce (light-weight) post-mortems when MoU targets not met**
- **We must continue to improve on open & transparent reporting, as well as further automations in monitoring, logging & accounting**

↳ We should foresee “data taking readiness” challenges in future years – probably with a similar schedule to this year – to ensure that full chain (new resources, new versions of experiment + AA s/w, middleware, storage-ware) is ready

# Future CCRC'08 Meetings

- (Not discussed yesterday due to time...)
- There is a HEPiX at CERN May 5-9 and a GDB on May 14.
- ¿ ***Do we need a CCRC'08 F2F in May?***
  - *expecting the answer no*
- In June (12 - 13) there is a 2-day "post-mortem" workshop
- Not to mention the WLCG Collaboration workshop 21 - 25 April!
- Pointers to the agendas for these on yesterday's F2F agenda...

- Hardware upgrade – important to get this in place prior to CCRC'08 to be consistent with setup for this summer's LHC data taking
- 64bit Oracle has been on the validation RACs since December 2007 and will be used for the new hardware
- Oracle 10.2.0.4 deployment – on validation RACs this week
- Schedule looks too tight for use in May – need joint decision on whether to postpone until after 2008 data taking or not...



# LFC DB backend deployment: Current Understanding (TBC)

- Site deployment is a **site** decision – subject to meeting requirements of experiments (**negotiation**)
- LHCb requires Oracle backend for LFC in order for current **replication** mechanism from central global catalog at CERN to local R/O replicas to work
- ATLAS has stated requirement of **deleting** O(100K) entries (catalogs + storage) per day – see separate talk on performance goals and results. Primary requirement, however, is **operational**.
- Survey of sites is given in following **table**:

# WLCG Tier1 Sites

Site (ISO standard)	Experiments (ATLAS, LHCb)	Comments
CA-TRIUMF	ATLAS	MySQL, plan to migrate to Oracle in 2009.
DE-KIT	ATLAS, LHCb	
ES-PIC	ATLAS, LHCb	
FR-CCIN2P3	ATLAS, LHCb	
IT-INFN-CNAF	ATLAS, LHCb	Oracle
NDGF	ATLAS	<i>r/s?</i> (LFC on <b>PostgreSQL(?)</b> or MySQL)
NL-T1	ATLAS, LHCb	Oracle (+MySQL for non-HEP VOs)
TW-ASGC	ATLAS	
UK-RAL	ATLAS, LHCb	MySQL → Oracle in 2-3 weeks (ATLAS?)
US-BNL	ATLAS	LRC. Evaluating LFC. b/e decision later.
US-FNAL	N/A	N/A

# Storage Baseline Versions for CCRC08 in May



Implementation	Version	Addresses
CASTOR	SRM v 1.3-17, backend 2.1.7	Possible DB deadlock scenarios; srmLs return structure now conforms; Various minor DB fixes; Fix for leaking sockets when srmCopy attempted; correct user mappings in PutDone; log improvements; better bulk deletion(?)
dCache	1.8.0-15-based CCRC'08 branch	
DPM	1.6.10	Group writable directories(*); Support for MacOSX; IPv6 support
StoRM	1.4.0	SL4/32bit; New thread management; GSOAP 2.7

# SRM v1.1 Services – R.I.P. ??

- **When can sites de-commission SRM v1.1 services?**
  - Pre-May is presumably not realistic... (i.e. end-April)
- Experience from May will presumably (re-)confirm that SRM v2.2 is (extremely fully) ready for business...
- Decide at June post-mortem workshop?

‡ **Set tentative target of end June?**

¿ *Is this too aggressive? Too mild?*

➤ **Discussion...**



# Middleware Patch Status

Component	Patch#	Fixes
LCG CE	<a href="#">Patch #1752</a>	(Significant) Performance enhancement
FTS	<a href="#">Patch #1740</a> <a href="#">Patch #1671</a> (fallback)	Update to transfer-url-copy; gridFTP session handling supports CASTOR2's gridFTP2; (SRM copy space token mixup)
FTM	(new module)	Makes transfer logs and current FTA config available (FTM web server) Improves cleanup of old transfer logs Available via WLCG ( not yet in the gLite release)
gFAL/lcg_utils	<a href="#">Patch #1738</a>	Fixes 10 bugs
DPM	1.6.10	Fixes some long standing issues and introduces dicom interface (1.6.7-4 Fixes the umask problem and some minor issues)

# Communication

- Is an on-going issue...
  - We have made – and continue to make – a lot of progress
  - It is unlikely that everyone will 'agree' about every detail
  - And we must – as has been repeatedly pointed out – be both pragmatic and efficient
- **Let's focus on the things that seem to be broken, whilst continuing to make progress in other areas...**
- For example, it is a perfectly fair comment that the daily WLCG operations meeting (@15:00) is 'like joining a CERN meeting remotely'.
- It *is* useful for service – experiment (- site?) interaction – we should try and make it more so
- ¿ *Make more people aware of the 'minutes' and weekly summaries*
- ¿ *More pro-active use of wlcg-xxx-contacts roles?*

# Week-Daily Operations Calls

- ***Not*** on Mondays (joint operations at 16:00)
- ***Not*** during F2F meetings (1<sup>st</sup>(?) week / month)
- ***Short; focussed; minutes available***
- ***Agenda & call details:***  
<https://twiki.cern.ch/twiki/bin/view/LCG/AgendaandCallDetails>
  1. Dial +41227676000 (Main) and enter access code 0119168, or
  2. To have the system call you, click [here](#)
- Minutes (one wiki page per week) are summarized to the MB of the following week (see following)

# Current activities – Services

Service	Issue
DBs	Two major activities. 1) Preparation of new RAC 5 and 6 hardware (both now on critical power) to migrate experiment services next week. 2) Upgrade of integration RACs to Oracle 10.2.0.4. This version will probably not be used for CCRC'08(?) More details in notes...
3D	CNAF down for the entire week due to maintenance. Streams replication to CNAF has been split and stopped, it will be resumed this week. ATLAS streams replication between offline and online database stopped on Good Friday due to an error from the streams apply process caused by an user application mis-configuration. The replication was resumed after a few hours.
CASTOR	The new CASTOR version (2.1.7) will be ready for pre-production certification on the 1st of April as originally planned. All tests and stress tests are successful.
SRM	A bug has been found in the CERN SRM v2 databases. This is in fact fixed in a rollout due on Monday so the rollout has been brought forwards.
FTS	Some Apache patches are in the pipeline for FTS servers.
CE	The fixes deployed at CERN to the LCG CE that reduce their load by an order of magnitude are being packaged for external distribution and should be ready next week.
DM	Patches to the 1.6.7 version of lcgutils should also be ready for distribution next week.



## Current activities – ATLAS

Day	Issue
Wed	Throughput testing with junk data. Build up T0 to T1 transfers to 150% nominal, adding 50% each day (150% over the w/e) CNAF has a downtime so their share is going to BNL. NDGF have requested a reduced rate due to insufficient resources. SRMv2 is being used at CERN.
Thu	Some Castor fixes at the end of the day. Try to reach 100% of the rate today then 200% on Friday then throttle back to 100% for the weekend. A current problem is they no longer see BNL on the ATLAS dashboard.
Fri	1) Yesterday the export rate was supposed to be increased to 100% of nominal. Problem in the ATLAS T0 machinery and not enough LSF jobs were submitted. 2) Display problem with the dashboard. Some entries end up in the production dashboard but should be in the T0 dashboard and vice versa. 3) SARA is not getting data on disk since disk is full (being cleaned up now). NDGF does not get data on disk since disk is full (can not be cleaned up centrally since NDGF does not use LFC - the only supported catalog for the central deletion tools). NDGF people should do the cleaning (alerted).
LFC	ATLAS are planning a bulk CERN site change to about 60K LFC entries. The consensus was that this is a relative simple operation that should only take a few tens of minutes of real time. <b>To be scheduled!</b>

## Current activities – CMS

Experiment	Day	Issue
CMS	Mon	
	Tue	
	Wed	Will continue their T0 to T1 exports to overlap with ATLAS but there is a downtime coming up at FZK.
	Thu	Asking how to get actions from IT on non-urgent issues during long holiday periods such as Easter. The use case there was to allow some dataops team members to access the DAQ worker nodes to look at LSF issues. These are non-urgent requests but which cannot wait for weeks. It was agreed to discuss this at the next F2F <sup>2</sup> meeting.
	Fri	

# Current activities – ALICE

ALICE	Activity
	Main activities for this week:
	Pass 1 reconstruction of RAW data taken during February/March commissioning exercise (CERN T0)
	Replication of specific RAW datasets to named T2s for fast detector-specific analysis (calibration/software tuning)
	MC production ongoing all centres

## Current activities – LHCb

Experiment	Day	Activity
LHCb	Mon	
	Tue	testing their stripping workflow.
	Wed	
	Thu	are evaluating the requirements for local worker node disk space in case they switch to copying production input data to local WN disk rather than reading via remote I/O. Current estimate is 5 GB. They will be running some production to exercise the manpower of their stripping workflow.
	Fri	

# Summary

- The February run of CCRC'08 was largely **successful** and introduced the important element of up-front and **measurable** metrics for many aspects of the service
- We still have a lot to do to **demonstrate** full 2008 readiness – **May and beyond will be a busy time with no let up in production services prior to then**
- We have developed a way of planning and operating the service that works **well** – re-enforce this and build on it incrementally

➤ **This is it – the WLCG Production Service!**