



CCRC Status & Readiness for Data Taking

Jamie.Shiers@cern.ch

~ ~ ~

June 2008

~ ~ ~

June F2F Meetings

Are We Ready?

- Based on the experience in the February and May runs of the Common Computing Readiness Challenge, an obvious question is
 - ¿ “Are we ready for LHC Data Taking?”
 - any time mid-July 2008 on...
- The honest answer is probably:
- “We are ready to **face** LHC Data Taking”
- There are subtle, but important, differences between the two...

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) & 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!

On WLCG Readiness

- The service runs smoothly – most of the time
- Problems are typically handled rather rapidly, with a decreasing number that require escalation
- Most importantly, we have a well-proved “Service Model” that allows us to handle anything from “Steady State” to “Crisis” situations
- We have repeatedly proven that we can – typically rather rapidly – work through even the most challenging “Crisis Situation”
- Typically, this involves short-term work-arounds with longer term solutions
- It is essential that we all follow the “rules” (rather soft...) of this service model which has proven so effective...



Pros & Cons – Managed Services

☺ **Predictable service level and interventions; fewer interventions, lower stress level and more productivity, good match of expectations with reality, steady and measurable improvements in service quality, more time to work on the physics, more and better science, ...**

☹ **Stress, anger, frustration, burn-out, numerous unpredictable interventions, including additional corrective interventions, unpredictable service level, loss of service, less time to work on physics, less and worse science, loss and / or corruption of data, ...**



We need to be here. Right?

CCRC '08 – How Does it Work?

- Experiment “shiffters” use Dashboards, experiment-specific SAM-tests (+ other monitoring, e.g. PhEDEx) to monitor the various production activities
- Problems spotted are reported through the agreed channels (ticket + elog entry)
- Response is usually rather rapid – many problems are fixed in (<) < 1 hour!
- A small number of problems are raised at the daily (15:00) WLCG operations meeting
- **↳ Basically, this works!**
- We review on a weekly basis if problems were not spotted by the above → fix **[+ MB report]**
- **➤ With time, increase automation, decrease eye-balling**

Baseline Versions for May CCRC'08

Storage-ware – CCRC'08 Versions by Implementation

CASTOR: SRM: v 1.3-21, b/e: 2.1.6-12

dCache: 1.8.0-15, p1, p2, p3 (cumulative)

DPM: (see below)

StoRM 1.3.20

Missed DBs - a similar table is needed!

M/W component	Patch #	Status
LCG CE	Patch #1752	Released gLite 3.1 Update 20

Missed FTM - required at WLCG Tier1s

FTS (11)	Patch #1671	Released gLite 3.0 Update 41
gFAL/lcg_utils	Patch #1738	Released gLite 3.1 Update 20
DPM 1.6.7-4	Patch #1706	Released gLite 3.1 Update 18

Baseline Versions for May CCRC'08

Storage-ware – CCRC'08 Versions by Implementation

CASTOR: SRM: v 1.3-21, b/e: 2.1.6-12

dCache: 1.8.0-15, p1, p2, p3 (cumulative)

DPM: (see below)

StoRM 1.3.20

M/W component	Patch #	Status
LCG CE	Patch #1752	Released gLite 3.1 Update 20
FTS (T0)	Patch #1740	Released gLite 3.0 Update 42
FTS (T1)	Patch #1671	Released gLite 3.0 Update 41
FTM	Patch #1458	Released gLite 3.1. Update 10
gFAL/lcg_utils	Patch #1738	Released gLite 3.1 Update 20
DPM 1.6.7-4	Patch #1706	Released gLite 3.1 Update 18

CCRC'08 – Conclusions (LHCC referees)

- The **WLCG** service is running (reasonably) **smoothly**
- The functionality **matches**: what has been tested so far – and what is (known to be) required

↳ We have a good baseline on which to build

- (Big) **improvements** over the past year are a good indication of what can be expected over the next!
- (Very) detailed **analysis** of results compared to up-front metrics – in particular from experiments!

WLCG Operations

I am sure that most – if not all – of you are familiar with:

- Daily “CERN” operations meetings at 09:00 in 513 R-068
- Daily “CCRC” operations meetings at 15:00 in 28 R-006
(will move to 513 R-068 if the phone is ever connected)
- In addition, the experiments have (twice-)daily operations meetings → continuous operations in control rooms
- These are pretty light-weight and a (long) proven way of ensuring information flow / problem dispatching

The WLCG Service

- Is a combination of services, procedures, documentation, monitoring, alarms, accounting, ... and is based on the “WLCG Service Model”.
- It includes both “experiment-specific” as well as “baseline” components
- No-one (that I’ve managed to find...) has a complete overview of all aspects
- **It is essential that we work together and feel joint ownership – and pride – in the complete service**
- It is on this that we will be judged – and not individual components...
- (I’m sure that I am largely preaching to the choir...)

CCRC '08 – Areas of Opportunity

- Tier2s: MC well run in, distributed analysis still to be scaled up to (much) larger numbers of users
- Tier1s: data transfers (T0-T1, T1-T1, T1-T2, T2-T1) now well debugged and working sufficiently well (most of the time...); reprocessing still needs to be fully demonstrated for ATLAS (includes conditions!!!)
- **Tier0: best reviewed in terms of the experiments' "Critical Services" lists**
 - These **strongly emphasize** data/storage management and database services!
 - ↳ We know how to run stable, reliable services
 - IMHO – these take **less** effort to run than 'unreliable' ones...
 - **But they require some minimum amount of discipline...**

Data / Storage Management

- ‡ **There is no doubt that this is a complex area**
 - And one which is still changing
 - We don't really know what load and / or access patterns 'real' production will give
 - But if there is one area that can be expected to be 'hot' it is this one...
 - **IMHO (again) it is even more important that we remain calm and disciplined – “fire-fighting” mode doesn't work well – except possibly for very short term issues**
- **In the medium term, we can expect (significant?) changes in the way we do data management**
- 💣 **Given our experience e.g. with SRM v2.2, how can this be achieved in a non-disruptive fashion, at the same time as (full-scale?) data-taking and production???**

CCRC '09 - Outlook

- SL(C)5
- CREAM
- Oracle 11g
- SRM v2.2+ +
- **Other DM fixes...**
- SCAS
- [new authorization framework]
- ...
- 2009 resources
- 2009 hardware
- Revisions to Computing Models
- EGEE III transitioning to more distributed operations
- Continued commissioning, 7+7 TeV, transitioning to normal(?) data-taking (albeit low luminosity?)
- New DG, ...

Count-down to LHC Start-up!

- Experiments have been told to be ready for beam(s) at any time from the 2nd half of July
- CCRC '08 [ran] until the end of May...
- Basically, that means we have **ONE MONTH** to make any necessary remaining preparations / configuration changes for first data taking (**which won't have been tested in CCRC...**)
- **There are quite a few outstanding issues – not clear if these are yet fully understood**
- ↳ **There must be some process for addressing these**
- If the process is the 'daily CASTOR+SRM operations meeting', then each sub-service absolutely must be represented!
- 💣 **Almost certainly too many / too complex to address them all – prioritize and address systematically.**
- Some will simply remain as "features" to be addressed later – and tested in CCRC '09?

Post-Mortems Required

- RAL power micro-cut (8.5 h downtime of CASTOR)
 - See next slide
- NIKHEF cooling problems (4 day downtime of WNs)
- CERN CASTOR + SRM problems
 - The postmortem of the CERN-PROD SRM problems on the Saturday 24/5/2008 (morning) can be found at <https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortemMay24> . The problem affected all endpoints.
 - Problems on June 5th (5 hour downtime): <https://prod-grid-logger.cern.ch/eelog/Data+Operations/13>

RAL Power Cut

- We lost power on one phase at about 07:00, but by the time pagers went off on-call staff were already in transit to RAL and were not able to respond until normal start of working day (which is within our 2 hour target out of hours).
- We suffered a very short (I am told milliseconds) loss of power/spike that took out one whole phase. As we have no UPS at RAL (will have in new machine room) this caused crash/reboot of over 1/3 of our disk servers.

- Restart commenced about 09:00
- CASTOR Databases Ready 12:00
- Disk Servers Ready 13:45
- Last CASTOR Instance Restarted 16:44

- So - about 2 hours because we were out of hours and had to respond and assess.
- 3 hours for ORACLE concurrent with 4:45 for clean disk server restart
- 3 hours for CASTOR restart and testing before release

- **Additionally, experience highlights the potential gap in the on-call system when people are in transit**

CCRC'08 – Conclusions

- The **WLCG** service is running (reasonably) **smoothly**
- The functionality **matches**: what has been tested so far – and what is (known to be) required **&** the experiments are **happy!**

↳ We have a good baseline on which to build

- (Big) **improvements** over the past year are a good indication of what can be expected over the next!
- (Very) detailed **analysis** of results compared to up-front metrics – in particular from experiments!