

# Status of CCRC'08

Status of Sites & Services  
Jamie Shiers

~ ~ ~

LCG-LHCC Referees Meeting, 19<sup>th</sup> February 2008

# Agenda

- Brief reminder of timelines, goals & requirements
- “GBU” review of preparations and initial running
- Observations & Recommendations
- Looking ahead to May and beyond...
- Summary



## CCRC'08 – Motivation and Goals

What if:

- LHC is operating and experiments take data?
- All experiments want to use the computing infrastructure simultaneously?
- The data rates and volumes to be handled at the Tier0, the Tier1 and Tier2 centers are the sum of ALICE, ATLAS, CMS and LHCb as specified in the experiments computing model
- Each experiment has done data challenges, computing challenges, tests, dress rehearsals, .... at a schedule defined by the experiment
- This will stop: we will no longer be the master of our schedule...  
.... Once LHC starts to operate.
- We need to prepare for this ... together....

A combined challenge by all Experiments should be used to demonstrate the readiness of the WLCG Computing infrastructure before start of data taking at a scale comparable to the data taking in 2008.

This should be done well in advance of the start of data taking in order to identify flaws, bottlenecks and allow to fix those.

We must do this challenge as WLCG collaboration: Centers and Experiments



## CCRC'08 – Proposed Scope (CMS)

- Test data transfers at 2008 scale:
  - Experiment site to CERN mass storage
  - CERN to Tier1 centers
  - Tier1 to Tier1 centers
  - Tier1 to Tier2 centers
  - Tier2 to Tier2 centers
- Test Storage to Storage transfers at 2008 scale:
  - Required functionality
  - Required performance
- Test data access at Tier0, Tier1 at 2008 scale:
  - CPU loads should be simulated in case this impacts data distribution and access
- Tests should be run concurrently
- CMS proposes to use artificial data
  - Can be deleted after the Challenge

- Two challenge periods agreed:

1. February 4<sup>th</sup> – 29<sup>th</sup> : limited by available resources

- Try to go as far as we can, but expect problems, particularly in areas of new/revised services
- Agreed metrics!

2. May 5<sup>th</sup> – 30<sup>th</sup> : full 2008 production resources at all sites

- Remaining issues “ironed out” – although this may well mean work-arounds / stop-gap solutions

☛ Very compressed timescales & limited manpower!

# CCRC'08 Preparations...

- Monthly Face-to-Face meetings held since time of “kick-off” during WLCG Collaboration workshop in BC
- Fortnightly con-calls with A-P sites started in January 2008
- Weekly planning con-calls ← suspended during February: restart?
- Daily “operations” meetings @ 15:00 started mid-January
- **Quite successful in defining scope of challenge, required services, setup & configuration at sites...**
- **Communication – including the tools we have – remains a difficult problem... but...**
- **Feedback from sites regarding the information they require, plus “adoption” of common way of presenting information (modelled on LHCb) all help**
- We are arguably (much) better prepared than for any previous challenge
- **There are clearly some lessons for the future – both the May CCRC'08 challenge as well as longer term**

# SRM v2.2 Deployment at Sites

- Targets of deploying production SRM v2.2 at Tier0 and “key Tier1” (self-defining) sites by end 2007 was met
- Deployment at Tier2s continues, as per schedule
- ☹ **Obtaining required information from experiments as well as making production instances at sites usable took until “the last minute” (and more...)**
- **Bottom line: this is both new and extremely complex**
- February run will shake-down production versions and has already revealed a few bugs as well as larger concerns
- The “management follow-up” of both production deployment and resolution of outstanding issues has been effective
- ↳ **For the future, such follow-up should be part of the process from the beginning**

EMT gLite 3 coordination (13 February 2008) - Mozilla Firefox


File Edit View History Bookmarks Tools Help

Back Forward Reload Stop Home <http://indico.cern.ch/conferenceDisplay.py?confId=29147> chilli gif

AudioPresenter Geneva Weather LCG Indico SIMBA2 WHO OMS WHO per diem

Mozilla Fir... Mozilla Fir... Mozilla Fir... Mozilla Fir... Mozilla Fir... 92nd LHC... GDB (06 F... ATLAS Sof... CERN HR ... Benchmar... GridView: ... EMT g...

category | view: Indico style | manage | SHIERS, Jamie - lo



## EMT gLite 3 coordination

Wednesday 13 February 2008  
 from 16:30 to 17:30  
 at 600-R-002  
 chaired by: *Oliver Keeble*

**Description:**

**EMT gLite coordination**

**WMS/LB/CE**

- ◆ SL4/VDT1.6 WMS/LB tests - progress and issues
- ◆ [bug #33026](#) - "no compatible resources" problem on SL4 WMS
- ◆ [bug #32345](#) - WMPProxy forward request to WM when dirmanager sigseg
- ◆ [patch #1593](#) - Patch for Bugs 31894, 32200, 29600 (security Hole), 32573 (WMS alias)

**Data Management**

- ◆ [patch #1679](#) - R3.1/SLC4/i386 GFAL 1.10.8

**VOMS/VOMSAdmin**

- ◆ [patch #1657](#) - Upgrade to VOMS 1.8.0 -- System OpenSSL migration, phase 1, plus fixes from patch #1572
- ◆ [patch #1629](#) - VOMS-Admin server 2.0.10-1 & VOMS-Admin client 2.0.6-1
- ◆ [patch #1591](#) - VOMS server update (multiple bug fixes)

**Certification Priorities**

**CCRC**

[Tracked Issues](#)

**New issues**

**AOB**

To join the conference:

To join the conference, click here: <https://audioconf.cern.ch/call/0148752>

To dial in to the conference:

- a. Dial +41227676000 (Main)
- b. Enter access code 0148752

The conference is called "gLite 3.0"

## Where Are We? Resources...

- For the February phase of CCRC'08 with a 55% 'accelerator on' model the Tier 1 cpu needed for full reprocessing is  $39590 * 0.55 = 21775$  KSi2K needed and currently 19041 KSi2K should be available
  - ATLAS will now only perform Tier 1 reconstruction at a few selected sites but will be running intensive Monte Carlo work

For 1 April most sites will now have their full 2008 cpu pledges, a total of 36725 KSi2K.

- **Largest missing is +2500 at NL-T1 due Nov.**
- For disk and tape many sites will catch up later in the year as need expands:
  - 2008 disk requirements are 23 PB and 15.5 PB expected for 1 April while 2008 tape requirements are 24 PB and 15 PB expected on 1 April
- May run of CCRC'08 at 55% 'accelerator on' only requires +1PB of disk and +1.5PB of tape (mostly reusable) so should have no resource problems
  - Delivery issues also putting DB h/w at CERN at risk – running May exercise on 2007 h/w is not a good option – see comments later re conditions!
  - A backup plan is in place (alternate supplier)

# What are the remaining problems?

- A table of "[baseline versions](#)" for clients / servers is linked from the [CCRC'08 wiki](#)
  - This is updated regularly and sites are encouraged to upgrade to these versions – any changes announced through standard channels
- Main outstanding issues:
  - "Space" related features in SRM v2.2 still need more production experience → on-going discussions...
  - Relatively minor outstanding bugs with DM tools...
- **A number of interventions are already foreseen for February...**
  - Additional ones will be needed – to fix problems encountered
  - But these are part of **normal operation** and we have shown that such interventions can be performed in a non-disruptive manner
- **Still some problems that 'arrive' via private e-mail: this cannot work in anything other than 'hero' mode**
- **Follow-up of critical problems: how does this work overnight & (in particular) at weekends? (See later...)**



# Storage Issues

- Attempt to identify and resolve key issues around storage in a pragmatic and timely manner
  - There is extremely little time prior to the May run
  - There is also limited effort – FNAL budget cuts
- Draft “Addendum” to SRM v2.2 MoU:
  - <http://indico.cern.ch/conferenceDisplay.py?confId=28791>
    1. This document is an addendum to the WLCG SRM v2.2 MoU (also referred to as “Usage Agreement”).
    2. It details changes in methods, their signatures and / or behaviour required for proper operation of the SRM v2.2 implementations under production conditions.
    3. (Once approved by the WLCG Management Board), this document has precedence over the MoU for the described methods and associated behaviour.
    4. [ Support issues should also be covered by this addendum ]

# Storage – the way ahead

- The list of key issues was discussed at a dedicated meeting on Monday 11<sup>th</sup> February
  - <http://indico.cern.ch/conferenceDisplay.py?confId=28791>
  - None of the issues are new, however some progress has been made since the previous discussion
- ↳ **Action: experiments to provide their feedback on the list of issues & the relative priority by 22<sup>nd</sup> February**
- **Have scheduled a further meeting during the following week to discuss how – and on what timescale – these issues can be addressed**
- Re-iterate on this list during March's F2F meetings...
  - As well as during WLCG Collaboration workshop...
- 💣 **In the immediate future, work-arounds (sometimes manpower intensive) are to be expected**

# Draft List of SRM v2.2 Issues

Priorities to be discussed & agreed:

- Protecting spaces from (mis-)usage by generic users
  - Concerns dCache, CASTOR
- Tokens for *PrepareToGet/BringOnline/srmCopy* (input)
  - Concerns dCache, DPM, StoRM
- Implementations fully VOMS-aware
  - Concerns dCache, CASTOR
- Correct implementation of *GetSpaceMetaData*
  - Concerns dCache, CASTOR
    - Correct size to be returned at least for T1D1
- Selecting tape sets
  - Concerns dCache, CASTOR, StoRM
    - ¿ by means of tokens, directory paths, ??

# Databases



- We have recently seen quite a few interventions around database services, both at CERN & other sites
  - This may well be due to the combined affect of increasing usage of these services for CCRC'08-related activities, together with additional visibility that this brings
  - ↳ **"Stand-by" databases potentially offer protection at low cost – consequences of prolonged downtimes are significant!**
- **As last year's "Critical Services" review tells us, the key services – as ranked by the experiments – concentrate on those related to data(-base) access**
- **Extension of on-call service to cover databases required for CASTOR, LFC, FTS etc being studied, as well as >8x5 streams coverage**
- 💣 **It is a concern that production conditions DB usage is only part of the May challenge: late testing = late surprises!**

# Databases – Some Details


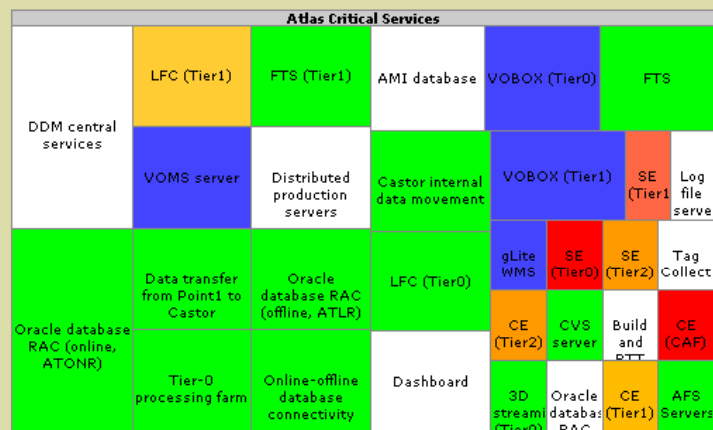
- Streaming from ATLAS online to offline RAC was down for 1 week shortly before challenge officially started (Oracle bug)
  - The issue was escalated with maximum priority – worked on 24x7 until a solution is found (production halted)
    - See later discussion on targets for problem resolution
  - Oracle patches ‘typically’ come on a shorter timescale
    - Which I guess means that they have them already!
    - Above patch existed for 64bit versions of Oracle, not 32
  - Cannot assume that such events will “never” happen again!
  - We must (IMHO) include an analysis of these events as part of post-mortem
  - I would suggest to cover not only Feb 4-29 but all 2008!
- **Goal is to learn as much as possible!**

## WLCG CCRC'08 Critical Services "GridMap"

### Ticklist Status (updated manually)



### Test Status (live data)



Alice Atlas CMS LHCb

CERN-PRDD TRIUMF-LCG2 IN2P3-CC FZK-LCG2 INFN-T1 SARA-MATRIX NDGF-T1 pic Taiwan-LCG2 RAL-LCG2  
BNL-LCG2 USCMS-FNAL-WC1


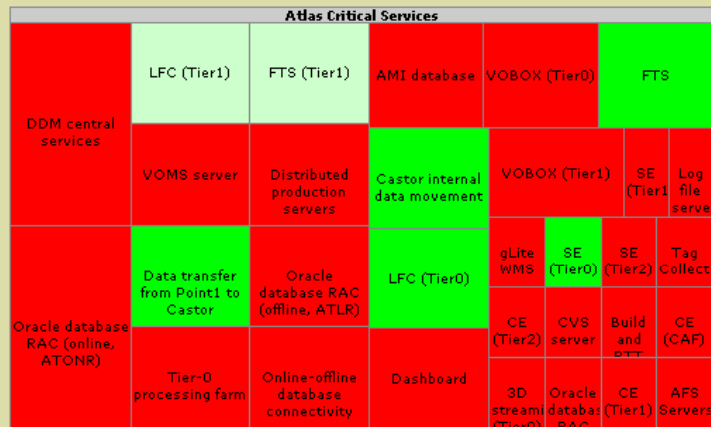
[bookmark settings](#)

|                       |                              |                    |                      |                                |
|-----------------------|------------------------------|--------------------|----------------------|--------------------------------|
| Transfer to Tier-0    | Recording Raw data to CASTOR |                    | Processing at Tier-0 | Transfer from Tier-0 to Tier-1 |
| Processing at Tier-1  | Monte Carlo at T1            | Analysis at Tier-1 | Transfer to Tier-2   | Data recording at Tier-1       |
| Monte Carlo at Tier-2 | Analysis at Tier-2           | Transfer to Tier-1 |                      | Transfer to T2                 |


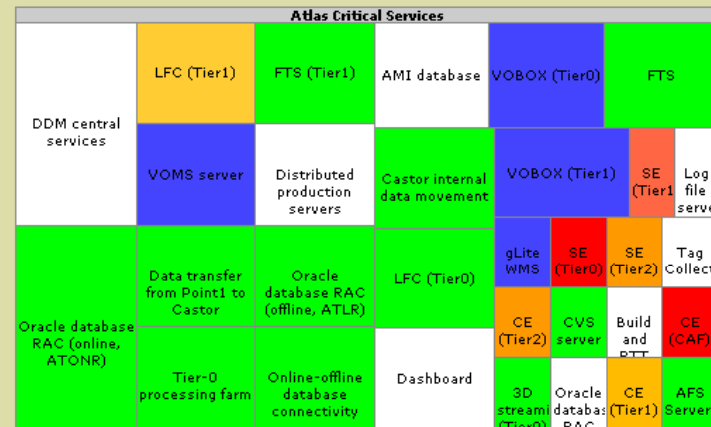
- Show the VO view of the infrastructure
- Two extra ‘maps’
  - Reliability (e.g successful data transfer, jobs, ...)
  - Metrics (MB/s, events/s, ...)
- **Need interaction with experiments to create these two views**
- 👉 **N.B. this is very similar structure to MoU view**
  - merge the two and report to sites on this structure?
- ☺ **“SMART” objectives!**

## WLCG CCRC'08 Critical Services "GridMap"

### Ticklist Status (updated manually)



### Test Status (live data)



Alice Atlas CMS LHCb

CERN-PRDD TRIUMF-LCG2 IN2P3-CC FZK-LCG2 INFN-T1 SARA-MATRIX NDGF-T1 pic Taiwan-LCG2 RAL-LCG2  
BNL-LCG2 USCMS-FNAL-WC1

[bookmark settings](#)






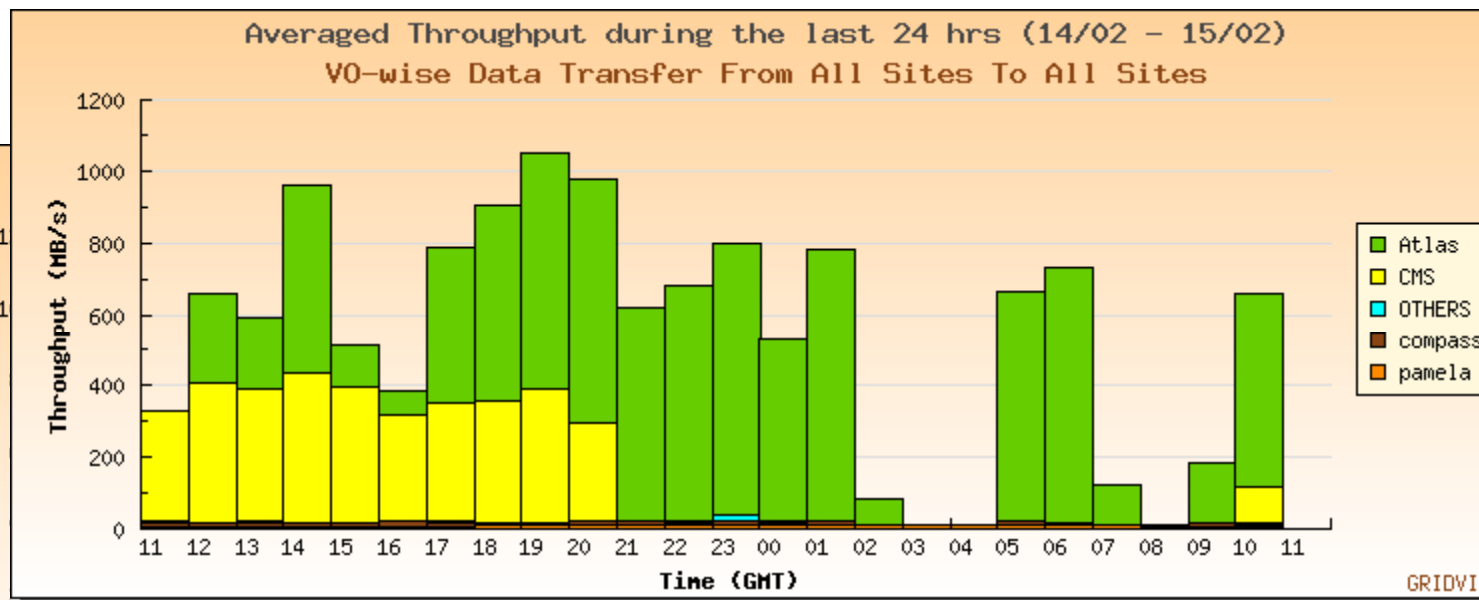
Need to follow-up with experiments on "check-lists" for "critical services" - as well as additional tests



# Critical Service Follow-up

- Targets (not commitments) proposed for Tier0 services
  - Similar targets requested for Tier1s/Tier2s
  - Experience from first week of CCRC'08 suggests targets for **problem resolution** should not be too high (if ~achievable)
    - The MoU lists targets for responding to problems (12 hours for T1s)
  - ¿ Tier1s: 95% of problems resolved < 1 working day ?
  - ¿ Tier2s: 90% of problems resolved < 1 working day ?
- **Post-mortem triggered when targets not met!**

| Time Interval | Issue (Tier0 Services)                       | Target |
|---------------|--|--------|
| End 2008      | Consistent use of all WLCG Service Standards | 100%   |
| 30'           | Operator response to alarm / call to x5011   | 99%    |
| 1 hour        | Operator response to alarm / call to x5011   | 100%   |
| 4 hours       | Expert intervention in response to above     | 95%    |
| 8 hours       | Problem resolved                             | 90%    |
| 24 hours      | Problem resolved                             | 99%    |



# Data Export

- We need to sustain 2008-scale exports for at least ATLAS & CMS for at least two weeks
  - The short experience that we have is not enough to conclude that this is a solved problem
  - The overall system still appears to be too fragile – sensitive to ‘rogue users’ (what does this mean?) and / or DB de-tuning
  - ↳ **(Further) improvements in reporting, problem tracking & post-mortems needed to streamline this area**
- We need to ensure that this is done to all Tier1 sites at the required rates and that the right fraction of data is written to tape
- Once we are confident that this can be repeatedly done, we need to mix-in further production activities
- 💣 **If we have not achieved this by end-February, what next?**
- **Continue running in March & April – need to demonstrate exports at required rates for weeks at a time – reproducibly!**
- Re-adjust targets to something achievable?
  - e.g. reduce from assumed 55% LHC efficiency to 35%?

# Recommendations

- ✓ **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” for **D**esign, **I**mplementation, **D**eployment and **O**peration are followed
- **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**

# And Record Openly Any Problems...

- The intervention is now complete and tier1 and tier2 services are operational again except for enabling of internal scripts.
- Two problems encountered.
  1. A typo crept in somewhere, **dteam** became **deam** in the configuration. Must have happened a while ago and was a reconfiguration problem waiting to happen.
  2. fts103 when rebooted for the kernel upgrade (as were the rest) decided it wanted to reinstall itself instead and failed since not a planned install. Again an accident waiting to happen.
- Something to check for next time.
- Consequently the tiertwo service is running in degraded with only one webservice box. If you had to choose a box for this error to occur on it would be this one.
- Should be running non-degraded mode sometime later this afternoon.
- ☺ **People are actively using the elog-books – even though we will have to review overlap with other tools, cross-posting etc.**






## CCRC'08 F2F Meeting

**Tuesday 04**  
from **09:**  
at CERN ([IT A](#))

[Tuesday 04](#)

### Tuesday 04 March 2008

- 09:30    Introduction & House-keeping (10)
- 09:40    Middleware Review - Problems Encountered & Roadmap / Schedule (45)
- 10:25    Storage-ware Review: Problems Encountered & Roadmap (1h00)
- 11:25    Site review (45)
- 12:10 lunch break
- 13:10    ATLAS Review (30)
- 13:40    CMS Review (30)
- 14:10    ALICE review (30)
- 14:40    LHCb Review (30)
- 15:10    CCRC'08 - Monthly Review of Calendar (20)
- 15:30    Service / Operations Review & Summary (30)

# Preparations for May and beyond...

- Aim to agree on baseline versions for May during April's F2F meetings
- Based on versions as close to production as possible at that time (and not (pre-)pre-certification!)
- ↳ **Aim for stability from April 21<sup>st</sup> at least!**
  - The start of the collaboration workshop...
- This gives very little time for fixes!
- **Beyond May we need to be working in continuous full production mode!**
- ↳ **March & April will also be active preparation & continued testing – preferably at full-scale!**
- **CCRC'08 "post-mortem" workshop: June 12-13**

# Summary

- The preparations for this challenge have proceeded (largely) smoothly – we have both learnt and advanced a lot simply through these combined efforts
  - As a focusing activity, CCRC'08 has already been very useful
  - We will learn a lot about our overall readiness for 2008 data taking
  - We are also learning a lot about how to run smooth production services in a more sustainable manner than previous challenges
- 💣 **It is still very manpower intensive and schedules remain extremely tight: full 2008 readiness still to be shown!**
- 👉 **More reliable – as well as automated – reporting needed**
- Maximize the usage of up-coming F2Fs (March, April) as well as WLCG Collaboration workshop to fully profit from these exercises
- **June on: continuous production mode (all experiments, all sites), including tracking / fixing problems as they occur**