

# CCRC'08 Status Report

Jamie Shiers/Harry Renshall

Grid Deployment Board

6<sup>th</sup> February 2008

# Agenda

- Where are we with preparations for February run?
- What are the remaining outstanding issues?
- What can we expect?
- How we monitor and report progress
- Preparations for May and beyond...
- Summary

# Where are we?

- With respect to previous challenges, it is fair to say that we are much better prepared
- Information flow has been considerably improved, a lot of focused work on ensuring that middleware and storage-ware is ready, ...

## 💣 **But things could still be done better:**

- “Week of stability” didn’t happen – some fixes will only be deployed during February
  - And those are the ones we know about...
- Information flow, particularly to / from Tier2s, needs to be improved
  - **Accurate and timely reporting of problems and their resolution – be they scheduled or unscheduled – still needs to be improved**
- We are still finding quite a few problems – particularly (it seems) in ‘new’ services or those just beginning to be hit hard
- We need to profit as much as possible from the February run
- 👉 **There is no “rocket science” – mainly common sense!**

# 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 MonteCarlo 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.
- The 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.

# 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 have been encouraged to upgrade to these versions
- 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...**
- Almost certainly, additional ones will be needed

# SRM v2.2 Issues

- A meeting was held yesterday to try to find acceptable solutions to a short-list of problems:
  1. Tokens for selecting tape sets
  2. Write access on spaces (follow the DPM strategy)
  3. Tokens and PrepareToGet / BringOnline (pool selection)
  4. Support for true bulk operations (list/remove)
- **Each of these items is quite complex!**
- Target is for “solutions” – which may well be “work-arounds” that can be deployed in production in time for May 2008
- Some issues may need to be revisited in the longer term!
- **Still plagued by technical problems with remote meetings!**

# CCRC'08 Communications

- “Daily” CCRC'08 meetings at 15:00 Geneva time
  - Except Mondays – joint operations meeting at 16:00!
- Fixed agenda, good attendance from CERN teams, less from sites (peaked during ramp-up)
- Some issues around communication with Tier2s
  - Speculation that CCRC'08 didn't involve Tier2s!
- Propose WLCG Tier2 coordinators with specific role of bi-directional communication
  - These people have to agree to perform this role actively!
- Communication with [Asia-Pacific](#) Tier2s being addressed by regular con-calls
  - ~Every two weeks so far, including not only CCRC'08 but also preparation for pre-ISGC [A-P Tier2 workshop](#)
- Similar, but lighter-weight, coordination also with DB community

# 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.



# What can we expect?

- During the ramp-up over the past couple months have uncovered quite a few bugs & configuration issues
- The teams involved have been very responsive and professional in responding to the issues in a timely manner
- It *seems* that the rate of new problems has decreased...
- But we cannot dream of finding no new issues...
- **Use the same follow-up mechanism to fix asap**
- There have also been quite a few problems around the DB services (e.g. streams, ASM, FP precision in instant client...)
- ¿ **Is this related to increased usage, visibility or both?**
- Most likely a case of the “**Fisk phenomenon**” –  
“If it ain't broken it probably hasn't been tested hard enough”

# Critical Service Follow-up

- An on-call service is being setup (in principle from Feb 4<sup>th</sup>) for CASTOR, LFC & FTS
  - “Critical” DB services were also recommended but not included
- Possible targets are given below.
- N.B. follow-up at April WLCG Collaboration Workshop!

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%

# Tier1 / Tier2 Targets

- The MoU says:

*"All storage and computational services shall be "grid enabled" according to **standards** agreed between the LHC Experiments and the regional centres."*


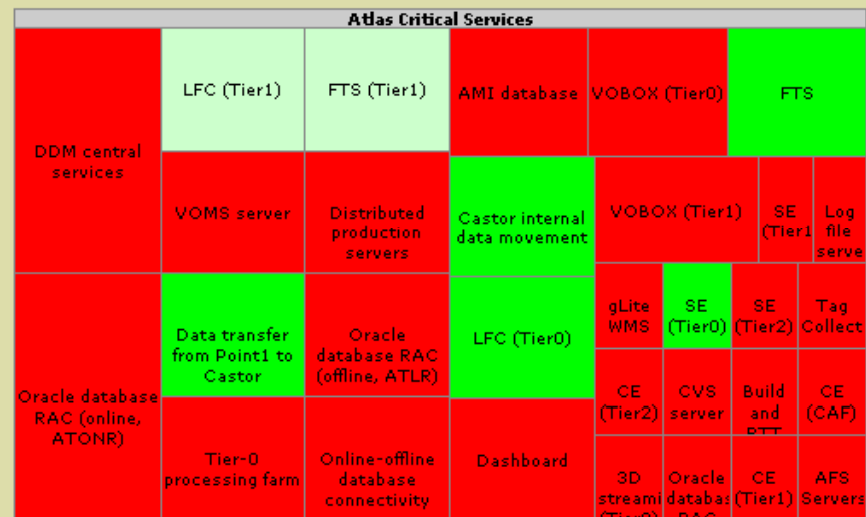
- ¿ **Does this mean that all sites should (by when?) implement the agreed standards for service design / deployment and operation?**
- ↳ **This would certainly have a big impact on service reliability, as well as the ability to perform transparent interventions**
- Some better understanding of how 24x7 is implemented – in particular, common baseline across Tier1s – is needed

# How we monitor & report progress

- For February, we have the following three sets of metrics:
  1. The **scaling factors** published by the experiments for the various **functional blocks** that will be tested. These are monitored continuously by the experiments and reported on at least weekly;
  2. The lists of **Critical Services**, also defined by the experiments. These are complementary to the above and provide additional detail as well as service targets. It is a goal that all such services are handled in a standard fashion – i.e. as for other IT-supported services – with appropriate monitoring, procedures, alarms and so forth. Whilst there is no commitment to the problem-resolution targets – as short as 30 minutes in some cases – the follow-up on these services will be through the daily and weekly operations meetings;
  3. The services that a site must offer and the corresponding availability targets based on the **WLCG MoU**. These will also be tracked by the operations meetings.
- We agreed yesterday to measure whether sites upgraded to the baseline services
  - Explicitly and/or implicitly – some aspects of challenge require these versions to work!
- Some of this reporting is still too “manual” – need to improve!


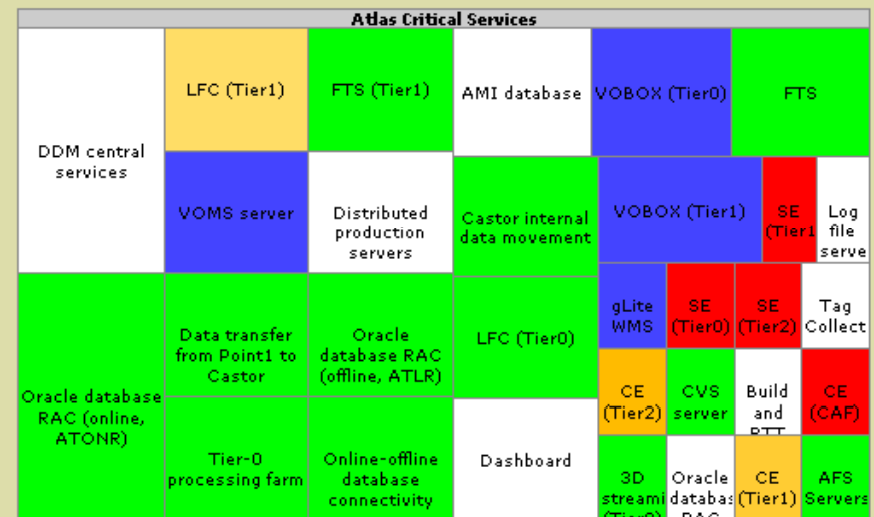
# WLCG CCRC'08 Critical Services "GridMap"

Ticklist Status (updated manually)



Alice **Atlas** CMS LHCb

Test Status (live data)



## 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!**
- ↳ **CCRC'08 "post-mortem" workshop: June 12-13**

# Summary

- As a focusing activity, CCRC'08 has already been very useful
- We will no doubt learn a lot about our overall readiness for 2008 data taking
- **“No surprises” would be a surprise!**  
**Expect the unexpected...**
- There is very little time for any fixes / work-arounds...  
WYSIWYG
- We should make significant progress on automatic measurement of service delivery (e.g. servicemaps)
- **Watch this space!!!**
- **Thanks to all for their efforts so far, as well as those to come!**