



# WLCG Service Report

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

~ ~ ~

**WLCG Management Board, 25<sup>th</sup> November 2008**

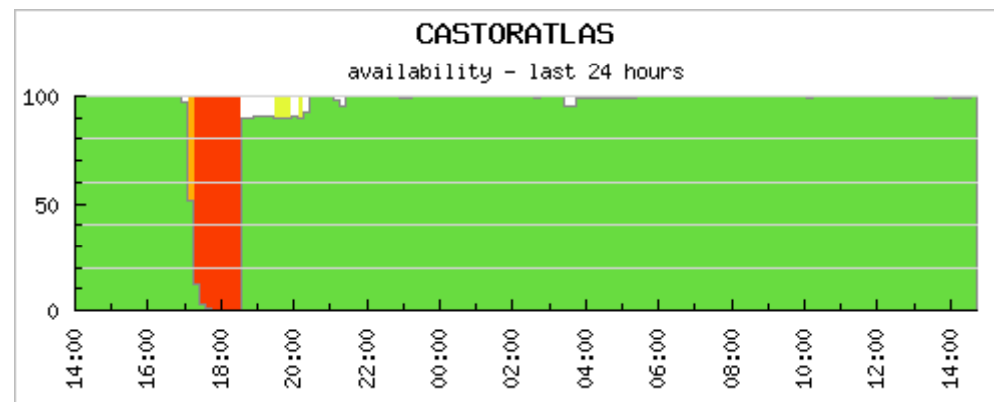
# Summary of the week

- **ASGC – continuing problems with CASTOR service. Current issues are clearly with Oracle service / configuration.**

□ **Some progress since yesterday's call...**

- FZK – clarification of LHCb LFC streaming issues
- CASTOR ATLAS degradation at CERN:
  - Castor Database Jobs were not running on the right cluster node
  - Full report:

<https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortem20081120>



# CERN Cooling Problems – Today!

- Due to a cooling problem the batch system was shutdown. Submission of new jobs is still accepted but the jobs won't be started. A fix is expected by the end of the day.
  - The cooling of the Computer Centre failed this morning around 08:00 due to a major leak in some cooling pipe. In order to limit the temperature increase in the machine room, the batch system was quickly stopped. The ventilation units were also restarted as soon as possible by the cooling expert in order to cool down the room with fresh air only. The batch system will likely remain down for a few hours, until the cooling system is properly fixed.
  - The leak in the cooling installation is fixed (11:40). The cooling system is being restarted but will only be fully operational at the end of the afternoon. Batch nodes are going to be progressively restarted in accordance with the available cooling capacity and with the importance of the nodes (nodes serving important queues will be restarted first).

GGUS: ID#43938 - All transfers to SARA-MATRIX fail - Mozilla Firefox

FileEditViewHistoryBookmarksToolsHelp

BackForwardReloadStopHome

https://gus.fzk.de/ws/ticket\_info.php?ticket=43938&from=search

☆

lhcb detector

Most Visited

EDH - Home (Jamie S...

AudioPresenter

Geneva Weather

LCG Indico

SIMBA2

WHO OMS

WHO/UN Per-diem

PostMortem20081120 ...

WLCGOperationsWeb ...

CERN IT Department

GGUS - /ticket\_search...

GGUS: ID#43938 - ...

subscribe

History Ticket-ID: 43938

Last modifier	Date	Time (UTC)	Action taken/comments
Stephane Jezequel	2008-11-22	07:30	assigned (ROC_North) Problem affects the whole site: NL-T1 Problem affects the whole VO: atlas Sent ALARM mail to mail address nlt1-alarms@biggrid.nl
Ron Trompert	2008-11-22	16:26	Public Diary: working on it
Ron Trompert	2008-11-22	16:26	in progress (ROC_North)
Ron Trompert	2008-11-22	16:45	solved (ROC_North) srm had a full disk due to excessive logging of the srm. Emptied the disk and turned the logging down and restarted everything. Transfers seem to run right now.  Cheers,  Ron
Stephane Jezequel	2008-11-22	17:09	verified (ROC_North)

[Top]

[History]

Done

gus.fzk.de

# Baseline Versions Update

WLCGBaselineVersions - LCG - TWiki - Mozilla Firefox

View History Bookmarks Tools Help

Reload Stop Home <https://twiki.cern.ch/twiki/bin/rdiff/CG/WLCGBaselineVersions?type=history> RSS Star G gorgeous

sited EDH - Home (Jamie S... AudioPresenter Geneva Weather LCG Indico SIMBA2 WHO OMS WHO/UN Per-diem

ortem20081... WLCGOperations... Itinerary 1254442... (1 vs. 2) W... the GGUS shoppin... Surfing Labrador ... EDH - Home (Jami... CERN IT Departm...

Difference: WLCGBaselineVersions (1 vs. 2)

**Revision 2** 24 Nov 2008 - Main.Gav

**Line: 1 to 1**

META TOPICPARENT name="WLCGOperationsWeb"

**Baseline versions of services and client tools for WLCG operations**

The table below lists the recommended versions for WLCG production. It does **not** necessarily reflect the latest versions of packages available in repository.

Service	Version	Availability	Notes
---------	---------	--------------	-------

**Changed:**

- < FTS 2.0 gLite 3.0 / SL3
- > FTS 2.1 gLite 3.1 / SL4 patch 2643 - planned for release Wed 27th November

# ASGC

- Current recommendation from CERN experts is to do a complete reinstallation of Oracle
    - Preferred option is a cluster configuration (RAC) with the Oracle Automatic Storage Manager (ASM) to manage storage used for data – not OCFS2 as currently(?) used
    - This requires a certain level of Oracle experience, but which in any case is commensurate with that required to run a service of this type
    - A possible **stop-gap** solution – but which would further increase diversity – would be a “single instance” – i.e. non-RAC installation to re-establish at least *some* service...
  - An appropriately qualified pro-active DBA (~1 FTE) is required asap together with participation in CASTOR external operations and Distributed Operations meetings
- + More detail in slide notes & in attachment to agenda page.

# Network Communications

- There are repeated complaints about difficulties (lack?) of communication between various teams involved in providing / using network services
- Not clear how / if these are addressed by the recent proposal on network operations presented to November GDB
- Problem needs to be owned and pursued...

# Recommendations

- Regular (quarterly?) WLCG<->Oracle technical reviews should be established to prioritize and review open Service Requests
  - We need to make sure that problems that affect (degrade, stop) the service get fixed on an acceptable timescale
    - Days / weeks maybe ok; Months / years is not...
- T0 / T1 storage services: urgent to minimize complexity / diversity and streamline operations / stabilize services
  - The dCache workshop being organized at FZK in January and the next CASTOR F2F w/s at RAL in February could be central to this
  - **Recommended configuration(s), routine operations, procedures and “best practices” need to be documented and shared**



- CERN: 6 instances (4x LHC, Public, cernT3)
  - Each instance: stager + DLF + SRM DB schemas, each on a different cluster (2 nodes)
  - Name server on yet another cluster
  - 38 machines!
- RAL: 4 different stagers + DLF + SRM, one Name server, two clusters
- CNAF and ASGC: one castor instance (Stager + DLF + SRM + Name server), one cluster

# Issues Seen

## 1. SQL executing in wrong schema

1. 14k files lost on LHCb, synchronization suspended
2. SR opened, but bug not seen by Oracle on 10.2.0.3
  1. Redo logs inconclusive
  2. Difficult to reproduce

## 2. ORA-600 sometimes when delete on id2type table

## 3. DML partition lock (ORA-14403)

## 4. Violation of primary key constraint ORA-00001

➤ Not seen by any other CASTOR service!

# Proposal for CASTOR DB

- Need to come up with common CASTOR DB architecture and share deployment/diagnostic procedures
  - Dedicated tutorial for CASTOR DBAs?
- Agree on a Oracle version for CASTOR and establish a clear plan for evolution with realistic validation period
  - Synchronized with version for distributed DBs cluster (3D) to avoid divergence
- Need a dedicated test setup similar to T1 setups
  - Reproduce and fix Oracle problems
  - Establish performance metrics via a common stress test

# Conclusions

- Diversity and complexity are the enemies of stability and reliability
- We need to establish sustainable data management operations before being swamped by the Tsunami of user analysis