



# CCRC'08 Weekly Update

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

~ ~ ~

**WLCG Management Board, 13<sup>th</sup> May 2008**

# Service Review – Week 19

- The [minutes](#) of last week's daily calls run to 8 pages!
- In general, ATLAS, CMS and LHCb all active
- ALICE preparing for 3<sup>rd</sup> Commissioning exercise
  - Primary requirement is upgrade of VO Boxes – status at <https://pcalimonitor.cern.ch:8443/admin/linux.jsp>
- More sites participated than in recent weeks
- More questions (and answers...)
- ☺ **In general, problems resolved rather rapidly**
- Mailing lists / elogs / hyper-news etc all very active – including over (long) weekend!
- Progress on various infrastructure issues (for GDB...)
- ☺ **Most services running smoothly – thanks to all concerned! (At all sites...)**

# CERN Services

See IT Status Board: <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/>

- Monday: [Intervention](#) on VOMS services (lcg-voms.cern.ch and voms.cern.ch)
- Tuesday:
  - DB services in SLS: <http://cern.ch/sls/> -> databases -> experiment databases including streams
    - ☹ “Detailed Service Status”, just below IT status board on IT home page contains link to SLS, but Databases are listed under “Services for engineers” and the drill-down is about as clear as mud).
      - CERN DB monitoring pages (next) – only visible within CERN
  - [vo-operator-alarm@cern.ch](mailto:vo-operator-alarm@cern.ch) mailing lists created / configured.
  - RSS feeds for elog & GGUS published + feedback URL
- Wednesday: phone in 513 R-068 to be installed next week!
- Thursday: gLite 3.1 update 22 delayed due to EGEE II-III transition meeting – now next Tuesday (i.e. today!)
- Friday:
  - C2ALICE on 15<sup>th</sup> May(?);
  - AFS problem in AFS22 - ~95% of volumes recovered, ~5% (232) to be restored from backup tape (affected job monitoring application);
  - C2CMS – patch 2.1.7-6 applied at 10:00, on C2PUBLIC Tuesday(?);
  - streams monitoring down – being investigated

# LCG Database Monitoring

Last update: 11:05 Europe/Paris - 13.May - @420 itime (auto-refresh in 60 seconds)

## Database and Streams availability

LCG Production RAC


LCG Validation RAC   
LCG Development RAC

## Useful links


[Physics Databases Wiki](#)  
[Development advices](#)  
[Oracle Documentation](#)

## LCG Production RAC (by hour, last week)


### OS Load



### Host CPU Utilisation



### Physical Reads

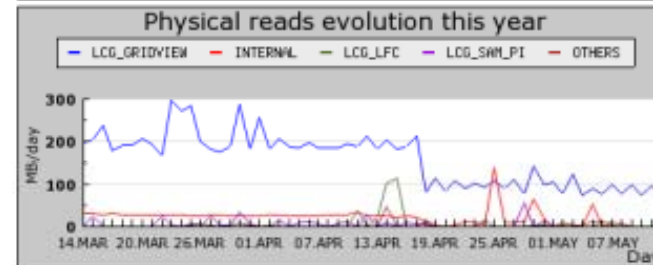

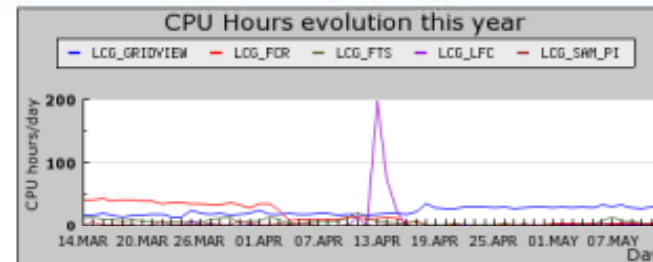
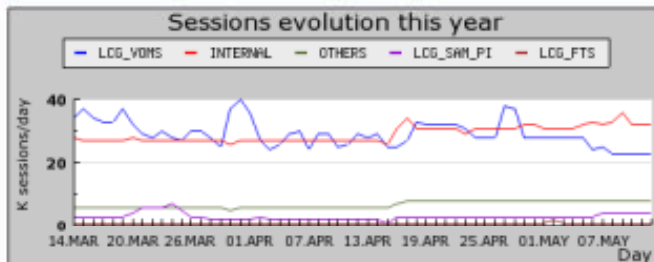


### Physical Writes



## Top 5 applications evolution (by day, last 2 months)

### LCG Production RAC



## Sites – Key(?) Issues

- RAL: power outage in RAL region on Tuesday morning. Quite some time bringing back all services. This also affected GOCDDB – DB corruption (see slide from Andrew).
  - ~8 hours, apart from GOCDDB
- NDGF: ATLAS through-put issues. Faster disks in front of tape (this week); more tape libraries (summer)
- BNL: disk space issues. Had not yet received FY08 resources prior to start of CCRC'08 (1.5PB raw).
  - **Cleanup / additional disk space (110TB) plus deployment of 31 Thumpers (FY08 resources)!**
- NL-T1: dCache library inconsistency – solved Wednesday.
- FZK: transfer problems (LHCb) rapidly solved – SRM load
- PIC: tiny files from ATLAS SAM tests going to tape!
  - **Configure test directory to be not migrated (as other sites...)**

# RAL Area Power Issue

From Andrew Sansum:

- *“We lost one phase at about 07:10 and had the last CASTOR instance up by 15:45 by which time the Tier-1 was fully operational (other than GOCDB).*
- *We lost time to two main items:*
  1. *Some disk servers required full fsck after journal recovery failed (got these by 13:45)*
  2. *Database servers generally started OK but a few had problems and we didn't have everything until 14:45*
- *I doubt we can go faster without more UPS. In the long term (Q12009) this will be mainly solved by the new machine room but in the short term we will look at some UPS for part of the CASTOR core, however in general we have no UPS on the current machine room.”*

# Experiments

- I am not going to attempt to summarize the summaries of Kors / Simone / Daniele et al
- These are the things that struck me, both attending the daily meetings and re-reading the notes today...
  - Some corrections / additions received – slides updated
- LHCb: startup delayed due to issues on online system. Massive cleanup of existing data at 2Hz. Deployment of needed s/w successful. Things running 'happily' by end of week.
- CMS: quite some problems with CASTOR configuration (garbage collection policies etc.) Some actions (done) on CMS side. Discussions with experts. "Reaching stable operations after 4 days"
- ATLAS: next week's focus is T1-T1. Requirements circulated. Post-mortem of week 1 this week. In general OK modulo site problems reported above. LFC intervention to correct SURLs being prepared.

# Summary

- Activity clearly exceeds that of February – and it's still week 1 of the challenge!
- ☺ **Problem resolution (significantly?) faster than February**
- ☹ **But there are still problems to be resolved...**
  - A large amount of information is available from the experiments
  - Can we really expect to condense this into a single “dashboard” any time soon?
  - Are the sites getting the information that they need from the experiments?
  - Are the experiments getting what they need from the sites?
- 👉 **If not, what can we do about it? (Now and later...)**