

# Table of Contents

<b>Week of 160704.....</b>	<b>1</b>
WLCG Operations Call details.....	1
General Information.....	1
Tier-1 downtimes.....	1
Links to Tier-1 downtimes.....	1
Monday.....	2

# Week of 160704

## WLCG Operations Call details

- At CERN the meeting room is [513 R-068](#).
- For remote participation we use the Vidyo system. Instructions can be found [here](#).

## General Information

- The purpose of the meeting is:
  - ◆ to report significant operational issues (i.e. issues which can or did degrade experiment or site operations) which are ongoing or were resolved after the previous meeting;
  - ◆ to announce or schedule interventions at Tier-1 sites;
  - ◆ to inform about recent or upcoming changes in the experiment activities or systems having a visible impact on sites;
  - ◆ to provide important news about the middleware;
  - ◆ to communicate any other information considered interesting for WLCG operations.
- The meeting should run from 15:00 until 15:20, exceptionally to 15:30.
- The SCOD rota for the next few weeks is at [ScodRota](#)
- General information about the WLCG Service can be accessed from the [Operations Web](#)
- Whenever a particular topic needs to be discussed at the daily meeting requiring information from site or experiments, it is highly recommended to announce it by email to [wlcg-operations@cern.ch](mailto:wlcg-operations@cern.ch) to make sure that the relevant parties have the time to collect the required information or invite the right people at the meeting.

## Tier-1 downtimes

Experiments may experience problems if two or more of their Tier-1 sites are inaccessible at the same time. Therefore Tier-1 sites should do their best to avoid scheduling a downtime classified as "outage" in a time slot overlapping with an "outage" downtime already declared by another Tier-1 site supporting the same VO(s). The following procedure is recommended:

1. A Tier-1 should check the downtimes calendar to see if another Tier-1 has already an "outage" downtime in the desired time slot.
2. If there is a conflict, another time slot should be chosen.
3. In case stronger constraints cannot allow to choose another time slot, the Tier-1 will point out the existence of the conflict to the SCOD mailing list and at the next WLCG operations call, to discuss it with the representatives of the experiments involved and the other Tier-1.

As an additional precaution, the SCOD will check the downtimes calendar for Tier-1 "outage" downtime conflicts at least once during his/her shift, for the current and the following two weeks; in case a conflict is found, it will be discussed at the next operations call, or offline if at least one relevant experiment or site contact is absent.

## Links to Tier-1 downtimes

<a href="#">ALICE</a>	<a href="#">ATLAS</a>	<a href="#">CMS</a>	<a href="#">LHCB</a>
	<a href="#">BNL</a>	<a href="#">FNAL</a>	

# Monday

Attendance:

- local: Kate D-W (Chair and minutes, DB), Maria Alandes (WLCG), Renato Santana (LHCb), Maarten Litmaath (ALICE), Nils Hoimyr (Computing), Maria Dimou (GGUS), Andrea Manzi (MW Officer), Vincent Brillault (Security), Marian Babik (Monitoring and Networking), Xavi Espinal (Storage), Jesus Lopez (Storage), Matthieu Ecosse (Storage)
- remote: Oliver Gutsche (CMS), Ulf Tigerstedt (NDGF), Alexander Verkooijen (NL-T1), Francesco Noferini (CNAF), Rolf Rumler (IN2P3), John Kelly (RAL), Jose Flix Molina (PIC), Di Qing (TRIUMF), Dario (ATLAS)

Experiments round table:

- ATLAS reports ( raw view) -
  - ◆ Activities:
    - ◇ LHC and data processing running full steam. Tier-0 saturated.
    - ◇ (Re)Testing T0 spillover (i.e. processing of T0 data on the Grid). Still some small differences being followed up.
  - ◆ Problems:
    - ◇ SARA had network problems just before the week-end, now seems fixed, but no report from the site (yet).
- CMS reports ( raw view) -
  - ◆ Data taking continues at full scale, putting the infrastructure under a lot of stress because the uptime is larger than expected (spent 59% in stable beam last week)
    - ◇ We still have problems with outgoing transfers. We think that EOS high load conditions produce file open errors which let the FTS optimizer back off on parallel transfers (we saw transfer rates as low as ~1 MB/s to T1 sites). This hurts the outgoing rates significantly. Started to deactivate the optimizer in FTS and set the parallel transfers manually, saw increase in transfer rates (back up to ~100 MB/s for individual T1 sites)
    - ◇ We will continue to work with the experts: [GGUS:122415](#)
    - ◇ Comment from Xavi: probable cause is the 4x16GB/s limit for gridFTP. One door added, next one will be added in an hours, around GB/s. The workflow needs to be checked as a next step.
  - ◆ SAM tests: ETF to SAM-3 synchronization stopped on July 1st but recovered itself over the weekend ([GGUS:122489](#))
  - ◆ GGUS not reachable with CERN certificates since Monday morning, July 4th, experts are working on this, more information in GGUS report
  - ◆ We are aware of the fiber optics replacement campaign in the CC (Tue Jul 05 10:00:00 CEST 2016)
- ALICE -
  - ◆ SARA: network problems started Thu afternoon (team ticket [GGUS:122438](#))
    - ◇ jobs running OK again since Fri afternoon, thanks!
  - ◆ CERN: intermittent CASTOR errors last week and ongoing, particularly for reading
    - ◇ this has hampered reco jobs
    - ◇ the devs are looking into it, thanks!
    - ◇ Comment from Xavi: the the issue with failing 1% is being fixed
  - ◆ GGUS web service was refusing CERN Grid CA certificates as "expired" today
    - ◇ at least as of 00:33, not clear when it started
    - ◇ most likely the **CRL** had expired on the web service side

- LHCb reports ( raw view) -
  - ◆ Activity
    - ◇ Monte Carlo simulation, data reconstruction/stripping(not too much) and user jobs on the Grid
  - ◆ Site Issues
    - ◇ T0:
      - Wrong configuration of worker nodes: GGUS:122187 [↗](#) (in progress)
    - ◇ T1: SARA:
      - Unscheduled Downtime. Problems with network. "at risk" until tomorrow(5/7).

Sites / Services round table:

- ASGC: nc
- BNL: nc
- CNAF: Last week we declared a down for CMS on the srm-endpoint. The down affected also the CMS queue on the CE but we couldn't declare a down for the CE because of the other VOs. It is not clear for us how to declare a down for a single CE queue without affecting the other VOs and this doesn't allow us to inform the experiment properly.
  - ◆ Francesco asked if it is possible to announce a downtime for a single VO. Maarten replied it is not possible for a service that is not dedicated to a VO. It might be fixed in the future. Maria A asked about the best practice in such case. Maarten: Downtime on SE should be declared and experiments should be able to adapt the workload. More details can be announced to the affected experiment.
- FNAL: nc
- GridPP: nc
- IN2P3: NTR
- JINR: NTR
- KISTI: nc
- KIT: nc
- NDGF: Wednesday 06:00 - 1400 one of our subsites will have a network outage as the main 10G switch must be moved. Another site will reboot their pools into newer kernels and dcache, but that's just a short blip for alice/atlas taking less than 10 minutes. The downtime will probably be shorter.
  - ◆ Atlas activity was so high (around 20GB/s) it was saturating links and computing was struggling
- NL-T1: A failover test for the upcoming datacenter move caused a hardware failure of one the firewalls. Our other firewall should have taken over but it had a bad configuration. Also one of our core routers had some technical issues. The problem started on Thursday 14:30 CEST and lasted until Friday 17:00 CEST. The fail-over tests have been repeated over the weekend with some degree of success. One of our firewalls is still down, repairs are continuing. We expect it to be back on line this week. There will be a few short at risk downtimes this week in order to do more failover tests.
- NRC-KI: nc
- OSG: nc
- PIC: NTR
- RAL: Tape control software is still having problems, but this has almost no operational effect. Otherwise ntr.
- TRIUMF:
  
- CERN computing services: ntr
- CERN storage services: 1 additional door for CMS were already added.
- CERN databases: CMSONR 1 node still down - HW is being fixed.SSB entry will be updated.
- GGUS: All users with CERN certificates were unable to access GGUS between 00h30 and 11:00am today. Maria D. emailed wlcg-operations at 09:57am as users' reports were coming in. Guenter, with help from Maarten, corrected the recently re-written scripts of the GGUS configuration engine which

were incomplete. The issue is related to a similar recent incident and is not likely to re-appear.

- Monitoring:
  - ◆ Final reports for the availability of May were sent around
  - ◆ LHCb VO feed will be stopped because of DIRAC move. It's being followed up with LHCb.
- MW Officer:
  - ◆ dteam VO VOMS server CA change :  
<https://operations-portal.egi.eu/broadcast/archive/id/1395>
- Networks: ntr
- Security: ntr

AOB: (Maria D.) Middleware Readiness WG meeting this Wed and Ops Coord this Thu. Agendas on <http://indico.cern.ch/category/4372/>

---

This topic: LCG > WLCGDailyMeetingsWeek160704

Topic revision: r17 - 2016-07-05 - VictorZhiltsov



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.  
or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback