

# Table of Contents

|                                   |          |
|-----------------------------------|----------|
| <b>Week of 170410.....</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 170410

## 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@cernSPAMNOTNOSPAMPLEASE.ch](mailto:wlcg-operations@cernSPAMNOTNOSPAMPLEASE.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: Julia A (WLCG), Kate (DB, wlcg, chair), Michal (ATLAS), Maarten (WLCG), Vincent (security), Andrea M (fts, MW), Gavin (computing), Jesus (storage)
- remote: Ulf (NDGF), Dimitri (KIT), Andrew P (NLT1), Di Qing (TRIUMF), Gareth (RAL), Kyle (OSG), Luca L (CNAF), Sang Un Ahn (KISTI), Tommaso (CMS), Xin Zhao (BNL), Zoltan (LHCb), David B (IN2P3), Vincenzo (EGI),

Experiments round table:

- ATLAS reports ( raw view) -
  - ◆ Activities:
    - ◇ reprocessing - jobs sometime require a little more memory than 2GB/core
    - ◇ frontier problems - overloading CERN frontier a little - occasional timeouts on machines
  - ◆ Problems
    - ◇ Tier-0 LSF instance was down (OTG:0036888 [↗](#))
    - ◇ no transfer monitoring during Thursday/Friday - oracle issue (OTG:0036910 [↗](#))
- CMS reports ( raw view) -
  - ◆ low production and analysis activities last week
    - ◇ recovering now
    - ◇ expected to ramp up with new 2017 MC + 2016 data rereco
  - ◆ tape tests ongoing (recall + stageout)
  - ◆ Kibana problem with DB monitoring (all pink)
    - ◇ does not seem a Kibana issue, investigating - NOW fixed - issue with feed on DB side
  - ◆ tickets / issues
    - ◇ GGUS:127597 [↗](#) : cpu efficiency @ RAL very low when accessing files via Xrootd. A more widespread test needed, consider the ticket a placeholder where to exchange info.
- ALICE -
  - ◆ very high activity on average, despite major issues at CERN Fri-Sun
  - ◆ CERN: CREAM CEs not working for ALICE since Fri (GGUS:127600 [↗](#))
  - ◆ CERN: HTCondor service in bad shape Fri-Sat (GGUS:127610 [↗](#))
  - ◆ RAL: VOBOX network problem on Sat (GGUS:127611 [↗](#))

Gavin commented both CREAM and HTCondor issues are understood.

- LHCb reports ( raw view) -
  - ◆ Activity
    - ◇ MC Simulation, Data Stripping and user analysis
    - ◇ Staging campaigns are ongoing for Data Stripping.
  - ◆ Site Issues
    - ◇ T0:
      - Transfer errors from the job: could not open connection to srm-eoslhcb.cern.ch (GGUS:127638 [↗](#))
    - ◇ T1:
      - RAL: two alarm tickets are opened during the weekend:
        - CEs were not accepting new jobs (GGUS:127612 [↗](#))
        - SRMs were down (GGUS:127617 [↗](#))

- CNAF: failed contact to the SRM: could not open connection to storm-fe-lhcb.cr.cnaf.infn.it:8444 (GGUS:127608 [↗](#))

Sites / Services round table:

- ASGC:nc
- BNL: ongoing investigation on why dCache space (ATLASDATADISK) fills up quickly recently.
- CNAF:
  - ◆ Deployed new version machine job features on worker nodes.
  - ◆ An unavailable tape made lhcb StoRM go offline, we are working to solve the issue.
- EGI: ntr
- FNAL: ntr
- IN2P3: ntr
- JINR:All OK for the last week. SE Disk and MSS end-points upgraded: dCache 2.13->2.16 and Postgres 9.4->9.6. One of them leads to lowering in SAM at 04-04-2017
- KISTI: ntr
- KIT: a downtime will take place on Thursday
- NDGF:
  - ◆ Short unannounced downtime for storage today, passed quickly with very little impact.
  - ◆ Moved one rack of alice data last week to a new data center, on thursday we got the news that it has to move back.
  - ◆ Preparing a longer downtime in Bergen, Norway for alice data within 2 weeks. The outage will be a 1-2 days for tape, the data on disk is being cached at other locations before the outage.
- NL-T1:
  - ◆ SURFsara scheduled a downtime to implement IPv6 on dCache cluster on May 1 and 2. [https://goc.egi.eu/portal/index.php?Page\\_Type=Downtime&id=22992](https://goc.egi.eu/portal/index.php?Page_Type=Downtime&id=22992) [↗](#)
- NRC-KI: nc
- OSG: ntr
- PIC: nc
- RAL: Ongoing problems with the LHCb Castor instance; We also had some problems with the CMS Castor instance over the weekend (now fixed). A separate problem with a hypervisor crash led to some problems for other services (including one of the CEs and argus).
- TRIUMF: NTR
- CERN computing services:
  - ◆ LSF affected by AFS server issue (now replicated). CREAMs subsequently had problems and got jammed up. Info provider problem affected Alice for longer.
  - ◆ HTCondor was hit by re-delegation bug exhausting memory on CEs - awaiting fix from HTCondor developers.
- CERN storage services:
  - ◆ FTS: upgrade to 3.6.7 completed last week. We had a problem with the DB schema upgrade process (OTG:0036830 [↗](#)), so we imported only the active/submitted transfers back to the tables. LHCb requested then to reimport all transfers present on the DB. We are preparing a new release ( 3.6.8 ), which will speed up the DB schema upgrade before asking BNL/RAL and FNAL to upgrade.
- CERN databases:
  - ◆ LCGR database storage performance was degraded several times from Apr 04, 12:00 to Apr 07, 16:00. It was related to a failed disk in an aggregate that contains LCGR volumes. The reconstruction of the RAID group affected the latency of the storage.
  - ◆ CMSR firewall deployment intervention proposed for April 26th (rolling but at risk)
- GGUS:
- Monitoring:

- ◆ Monthly report distributed waiting for feedback from sites.
- MW Officer: NTR
- Networks: NTR
- Security: NTR

AOB:

- **ATTENTION: next meeting on Tuesday April 18 !**
  - **Have a good Easter [☞](#) break !**
- 

This topic: LCG > WLCGOpsMeetingWeek170410

Topic revision: r28 - 2017-04-10 - MaartenLitmaath



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback