

# Table of Contents

<b>Week of 161017.....</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 161017

## 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: Maria Alandes (chair and minutes), Ivan Glushkov (ATLAS), Maarten Litmaath (ALICE), Fernando Meireles (Computing), Vincent Brillault (Security), Jesus Lopez (Storage), Alexei (LHCb), Alberto Aimar (Monitoring)
- remote: Stefano Belforte (CMS), FaHui Lin (ASGC), Xin Zhao (BNL), Luca Lama and Antonio Falabella (CNAF), Vincenzo Spinoso (EGI), Dave Mason (FNAL), Rolf Rumler (IN2P3), Victor Zhiltsov (JINR), Xavier Mol (KIT), Andrew Pickford and Alexander Verkooijen (NL-T1), Chris Pipes (OSG), Gareth Smith (RAL), Di Qing (TRIUMF)

Experiments round table:

- ATLAS reports (raw view) -
  - ◆ Quiet week - CHEP2016
  - ◆ Ongoing MC12 reprocessing (single core)
  - ◆ Frontiers loaded / brought down on Friday due to nasty reprocessing tasks - under investigation.

Maria asks whether the issues reported with Frontier are related to IN2P3 report. Ivan explains that the task causing these problems was aborted and that an internal investigation is still ongoing, but it looks like it could be related.

- CMS reports (raw view) -
  - ◆ Running at full capacity
  - ◆ Managed to get up to 20% more resources when moved top HTCondor negotiator from CERN to FNAL. Reason for improvement still unclear (2.8 vs 2.6GHz ? PM vs VM ?). Want to be able to get same performance on a CERN based machine, discussion opened with CERN Cloud Team INC: RQF:0654902 [↗](#)
  - ◆ The long delays for data to be presented in Kibana dashboards (aka meter.cern.ch) have been addressed (successfully so far): INC:1156813 [↗](#)

Maria asks whether the Kibana issues reported by CMS were also affecting other experiments. It seems this particular issue was affecting the CMS critical dashboard and it was related to FLUME configuration. Ivan adds that ATLAS has also experienced issues with stale data in Kibana, but they were fixed very quickly. Alberto explains that an upgrade from Kibana 3 to Kibana 4 is ongoing and all these instabilities should disappear in the latest version. Ivan asks whether Kibana 4 will require changes from the experiments and Alberto confirms so. Alberto adds that experiments will be contacted and documentation will be provided soon.

- ALICE -
  - ◆ EOS crashes at CERN and other sites on Thu
    - ◇ Clients unexpectedly used signed URLs instead of encrypted XML tokens
    - ◇ The switch was due to one AliEn DB table being temporarily unavailable
      - and a wrong default (now fixed)
    - ◇ The EOS devs have been asked to support the new scheme and avoid that unexpected requests crash the service
  - ◆ CERN: alarm GGUS:124447 [↗](#) opened Fri evening
    - ◇ none of the CREAM CEs were usable
    - ◇ fixed very quickly, thanks!
- LHCb reports (raw view) -
  - ◆ Activity

- ◇ Monte Carlo simulation, data reconstruction/stripping and user jobs on the Grid
- ◆ Sites
  - ◇ CNAF - failing transfers (GGUS:124481 [↗](#))
  - ◇ RAL - one user disk is down

CNAF confirms the ticket opened 1h before the meeting is in progress although they don't have details from the experts at this moment.

Sites / Services round table:

- ASGC: NTR
- BNL: NTR
- CNAF: As a consequence of the DDN failure of 21st of September, we schedule an 8h intervention for the 1st of November that will impact Atlas, Alice and AMS.
- EGI: NTR
- FNAL: NTR
- GridPP: NA
- IN2P3: FRONTIER problems on last Friday due to some ATLAS jobs brought down our four SQUID servers which are common to ATLAS, CMS, and CVMFS. This apparently revealed a bug in the SQUID middleware, under investigation. We reserved two SQUID servers to ATLAS and the other two for CMS and CVMFS. This and a restart made the service again available for CMS. ATLAS stopped the critical jobs.
- JINR: AAA local redirector (Federation host) upgraded to 4.4.0 from OSG repository. The problem with LHCONE since the mid of the week. Networking division is working on the issue.
- KISTI: NA
- KIT:
  - ◆ Downtime for updating GPFS, dCache and the xrootd redirector for CMS on Wednesday, 19th of October, from 10 to 13 o'clock CEST.
  - ◆ Network intervention on Thursday, 20th of October, from 9:45 to 10:30 o'clock CEST will force a halt for all tape activity.
- NDGF: NA - There will be a short (1h) warning downtime Tuesday for a reboot of some Oslo poolnodes. It should be over in 5 minutes. Some data will be unavailable for a few minutes. Hopefully transfers will be automatically retried and no one will notice.
- NL-T1:
  - ◆ SARA datacenter move: all hardware has arrived in the new datacenter in good condition. No data has been lost. Of the ~3000 disks, only one broke, which was in a RAID6 array, containing non-unique data. We're still faced with a few issues:
    - ◇ Our compute cluster is not yet at full capacity, because of a remote access issue. Around 3800 cores are available out of ~6300; these are the nodes with SSD scratch space.
    - ◇ A non-grid department commissioned new hardware, which forced our network people to commission new Qfabric network hardware, which then forced a software upgrade of the existing Qfabric nodes (version v13 to v14d15). This introduced a bug which broke the 40Gbps ports of our Qnodes of type 3600. The vendor was so kind to lend us Qnodes of type 5100, which don't have this bug. This workaround will enable us to start production without (considerable) delay. But with the type 5100 Qnodes, we've observed a bandwidth limit for some storage nodes, for an unknown reason. Of the 65 pool nodes we have, there are 36 pool nodes with an iperf bandwidth ~10 Gbps per pool node instead of ~23 Gbps that we have measured before. We think however that this bandwidth, aggregated over all pool nodes, will be sufficient for normal usage. Meanwhile we're investigating this issue.
    - ◇ A top-of-rack uplink was unstable. This affected 9 pool nodes. The fibre has been cleaned and reseated. This solved the problem. Our storage is now back in

production.

- NRC-KI: NA
- OSG: NTR
- PIC: NA
- RAL: RAL is aware of the disk problem reported by LHCb. They are working on it.
- TRIUMF: NTR
  
- CERN computing services:
  - ◆ ALARM ticket from Alice re availability of CREAM CEs. CREAMs were rebooted and brought back online; not clear the underlying cause. An info provider issue for one of the CEs was addressed. HTCondor CEs were unaffected.
- CERN storage services:
  - ◆ EOSALICE crash on 13.10 (description in the ALICE report)
  - ◆ EOSATLAS problem early this morning (17.10), the system was set to read-only while recovering, was back at 8:50
  - ◆ CASTOR ALICE We are running part of the capacity in default with Ceph, it's been deployed with a couple of issues that are being taken care of and it will be improved.
- CERN databases: NA
- GGUS: NA
- Monitoring: A new Working Group will be created to communicate with the experiments any progress on the Monitoring Consolidation project.
- MW Officer: NA
- Networks:
  - ◆ ntr
- Security: NTR

AOB:

---

This topic: LCG > WLCGDailyMeetingsWeek161017

Topic revision: r18 - 2016-10-20 - MariaDimou



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback