

Table of Contents

Week of 160620.....	1
WLCG Operations Call details.....	1
General Information.....	1
Tier-1 downtimes.....	1
Links to Tier-1 downtimes.....	1
Monday.....	2

Week of 160620

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 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

ALICE	ATLAS	CMS	LHCB
	BNL	FNAL	

Monday

Attendance:

- local: Maarten (ALICE), Julia A, Kate(SCOD, DB), Maria A. (MW), Marian (Monitoring, Network), Miguel M.(Computing)
- remote: Andrew (Nikhef), David C. (ATLAS), Eric (BNL), Dave (FNAL), Dmytro (NDGF), Tiju (RAL) , Jose (PIC) , Rolf (IN2P3), Di Qing (TRIUMF), Kyle (OSG), Zoltan (LHCb) , Sang Un (KISTI), Francesco (CNAF)

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ Activities
 - ◇ LHC and data processing running full steam
 - ◇ Derivation production on MC finished
 - ◇ Expecting ramp up in analysis as summer conferences approach
 - ◆ Problems
 - ◇ EOS "an end-of-file was reached globus_xio: An end of file occurred" Thurs/Fri last week and after restart/upgrade last night GGUS:122208 [↗](#)
 - SRM space reporting is also not working since the upgrade GGUS:122233 [↗](#)
 - ◇ RRC-KI-T1: Lost storage DB on 8 June. Affected is just ATLAS disk-based dCache instance, not other VOs, not tape. Last backup was 19 April. Recovery of the DB finished on 17th June, but still many files missing. ATLAS will run a consistency check to determine what is lost.
 - Would be nice to know WLCG expectations on DB backups
 - Discussion:
 - David reported that Atlas is worried about DB backups in T1s.
 - Maarten said that we are all surprised and a service incident report is expected. This also could serve as a wake-up call for other T1s. There are no specific backup recommendations, there are availability thresholds. This will be discussed in the next MB tomorrow. Atlas will also be represented and the issue can be highlighted.
 - Kate: DB team provides advice to T1 DBAs, this can include backup advice.
 - M: dCache is a popular solution, there should be strategies in place. Lessons should be learned from this issue.
 - Maria A: Is only Atlas affected? Alice and LHCb might be as well.
 - M: Alice was not aware of this incident
 - Julia A: It was raised multiple times, by different channels that the ops meetings should be attended.
 - Maria A: The meeting is as short and as rare as possible
- CMS reports (raw view) -
 - ◆ **apologies, I will most probably not be connected, boarding on a plane**
 - ◆ uneventful week: LHC taking tons of data, MC production full speed. No major issue.
 - ◆ nothing to report: all the open tickets are for specific issues, already being taken care of.
- ALICE -
 - ◆ NTR
- LHCb reports (raw view) -
 - ◆ Activity
 - ◇ Monte Carlo simulation, data reconstruction and user jobs on the Grid

◆ Site Issues

◇ T0:

- EOS gridftp problem GGUS:122100 [↗](#), did not appear during the weekend.
- wrong configuration of CERN worker nodes: GGUS:122187 [↗](#)

Sites / Services round table:

- ASGC: nc
- BNL: ntr
- CNAF: Tomorrow (21st June 16) at 18.00 we scheduled a down (GOCDB warning) on our geographic network for about 2 hours because of an update of the router. The down doesn't affect the access from/to LHCOPN and LHCONE, but it has an impact on several services, included DNS and site BDII.
- FNAL: ntr
- GridPP: nc
- IN2P3: The maintenance outage last week went well except for a dCache configuration problem on our side which obliged us to extend the dCache and batch downtime by four hours (until 10pm). Alice could start as planned, though.
- JINR: T1 was OK except short network drop. Not long enough to get reflected in CMS Availability/Reliability. Some issues with air conditioning caused unstable RAM and disks functioning in some hosts. The problem is under control and being resolved.
- KISTI: ntr
- KIT: nc
- NDGF: one of the CEs has filesystem problems. No ETA yet. The affected endpoint has been put into the downtime.
- NL-T1: ntr
- NRC-KI: nc
- OSG:ntr
- PIC:ntr
- RAL: Problems with the stability of the control software for the tape library are continuing. Working with vendor to fix problems. The operational impact on LHC VOs is low.
- TRIUMF: ntr

- CERN computing services: ntr
- CERN storage services:
- CERN databases: One instance of CMSONR went down on Friday due to NIC issues. The machine is stopped and being tested. The DB is running on 3 nodes, so it is not under stress.
- GGUS: ntr
- MW Officer: ntr
- Security:
- Network:
 - ◆ GGUS:121687 [↗](#) RAL consistent loss - waiting for an upgrade of the router at RAL
 - ◆ GGUS:121905 [↗](#) BNL to SARA - SARA perfSONARs were fixed. Consistent loss to other T1s (KIT, PIC, CERN), mainly outbound, but also inbound, informed OPN. Suggested to postpone further investigation until SARA moves to the new data centre.
 - ◆ Grid output retrieval failing: Victoria - Prague - asymmetric paths and MTU step down issues, resolved.
 - ◆ Possible network issue between McGill and BU - gridftp transfers timing out - issue with storage, resolved.

AOB:

This topic: LCG > WLCGDailyMeetingsWeek160620
 Topic revision: r14 - 2016-06-21 - EygeneRyabinkin



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