

Table of Contents

Week of 160613.....	1
WLCG Operations Call details.....	1
General Information.....	1
Tier-1 downtimes.....	1
Links to Tier-1 downtimes.....	1
Monday.....	2

Week of 160613

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: Ivan (Atlas), Maarten (Alice), Julia A, Kate(SCOD, DB), Maria A. (MW), Vincent (Security), Marian (Monitoring, Network), Fernando (Computing)
- remote: Eric (BNL), Dave (FNAL), Dmytro (NDGF), Onno (NL-T1), Stefan Roiser (LHCb), Tiju (RAL) , Jose (PIC) , Rolf (IN2P3), Andrew (NIKHEF), Di Qing (TRIUMF), Kyle (OSG)

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ Activities:
 - ◇ Derivation production started last week for ICHEP, done on data15/data16. MC derivations are running.
 - ◇ Database and Metadata TIM (13.06 - 15.06), ADC TIM (15.06 - 17.06).
 - ◆ Problems:
 - ◇ NTR.
- CMS reports (raw view) -
 - ◆ Apologies from Stefano, who couldn't attend
 - ◆ Activities:
 - ◇ data taking ongoing. Grid sites utilizes fully.
- ALICE -
 - ◆ NTR
- LHCb reports (raw view) -
 - ◆ Site Issues
 - ◇ T0:
 - VOMS admin interface not reachable on Friday, Alarm ticket GGUS:122068 [↗](#), fixed
 - EOS, SRM not reachable on Friday, GGUS:122080 [↗](#), too many connections, restarted bestman, fixed
 - EOS gridftp problem today, going through single node GGUS:122100 [↗](#), enabled alternative hosts, fixed
 - ◇ T1: GRIDKA: Network (firewall) problems today, GGUS:122106 [↗](#), fixed
 - ◇ All of the issues reported were picked up

Sites / Services round table:

- ASGC: nc
- BNL:
 - ◆ Issues related to VOMS server at CERN were reported, the ticket was not updated but the issue was closed
- CNAF: nc
- FNAL:
 - ◆ EOS problem possibly related with GridFTP issues
- GridPP: nc
- IN2P3: Final reminder of site's maintenance outage tomorrow, see downtime declarations for details
- JINR: ntr
- KISTI:nc
- KIT: nc

- NDGF: one of the sites will reboot storage servers due to new kernels/libc/etc. Some data may be shortly unavailable on Wed., 15 June, 11-12 UTC.
- NL-T1:
 - ◆ SURFsara
 - ◇ Atlas has suffered from tape congestion last week. After the last tape system maintenance, there were some instabilities in the SSH/SCP interface between dCache and the tape system, causing our mass storage system to sometimes fail on some operations. Many of the Atlas files were written to the tape file system, but our HSM script was often unable to check the checksum of these files, reporting back to dCache that the store operation had failed. The interface was fixed on Monday, but problems were not over yet. Subsequent efforts to store these files failed because each effort to overwrite an existing file on the tape FS hung. We think this was caused by a bug in DMF, our HSM system. When scp operations got stuck, they occupied scp slots, so less slots were available for tape operations. This caused a backlog. We fixed it on Friday by using an 'rm -f' on an existing file before overwriting it.
 - ◇ We have found that dCache sometimes starts more than one simultaneous restore per file. We reported it to the dCache developers. We're going to use the tape system cache to increase performance in such events. Before, we bypassed the cache because purging this cache when it was full was an expensive process. See also next:
 - ◇ On 21 to 22 June, we've scheduled a tape maintenance. SSDs will be deployed for the inodes of the tape FS, so that performance is increased and cache purging will no longer be an expensive process.
 - ◇ A friendly reminder: the datacenter move. First 3 weeks of August the tape system will be moved. First 2 weeks of October the other grid systems will be moved. This requires enormous effort from our network experts so they may be slow to respond. Thank you for your understanding.
- NRC-KI: ntr
- OSG: One issue, related to other T1, being worked on
- PIC: ntr
- RAL: Problems with the stability of the control software for the tape library are continuing. Restarting the software has enabled us to deliver a tape service. Some operations (mainly reads) may be delayed.
- TRIUMF: dCache upgrade completed smoothly last Wednesday, it only spent about two hours,
- CERN computing services:
 - ◆ Series of voms-admin bugs exacerbated by the original root full:
 - ◇ High-load from connections on voms2. Double the connections compared to lcg-voms2;
 - ◇ The file-system partitioning was bad on both instances but only affected voms2 as all the extra errors in the voms-admin logs filled up the root partition faster;
 - We did an online-resize of the root disk partition on the voms2 and lcg-voms2 nodes.
 - ◇ On 2016-06-09) voms24 (master) encountered an AUP reminder event storm, the alias switched over to voms25 (slave) but unfortunately it encountered a root-full and load errors and was knocked out of the available members of the alias. The alias switched back to the voms24 master node which with the on-going AUP load, in conjunction with the increased load from clients and users trying to access voms started deadlocking in java and resulted in time-outs;
 - We'd restart the service and it would be okay for a short-while and then the events would build up again.
 - ◇ On 2016-06-10) voms24 exhausted the heap-space with a memory leak, the service failed over to voms25 which proceeded to encounter issues with garbage collection pauses in voms-admin, deadlocks and would terminate database connections before they were complete.
- CERN storage services: ntr

- CERN databases:
 - ◆ All databases had their OS updates applied last week.
 - ◆ 1 instance of LHCb offline database was restarted on Sunday. Some sessions might have failed.
- GGUS: ntr
- MW Officer: ntr
- Security:
 - ◆ iperf3 vulnerability in Perfsonar:
<https://wiki.egi.eu/wiki/SVG:Advisory-SVG-CVE-2016-4303> Marian suggested enabling auto updates, Maarten suggested that it has to be checked with Sonar team as the update had to be forced at least at one site
- Monitoring:
 - ◆ ETF update tomorrow
(<https://cern.service-now.com/service-portal/view-outage.do?from=CSP-Service-Status-Board&n=OT>)
- Network:
 - ◆ GGUS:121687 RAL consistent loss - waiting for an upgrade of the router at RAL
 - ◆ GGUS:121905 BNL to SARA - SARA perfSONARs were fixed Marian reported issue didn't disappear and he will present another report on it
 - ◆ Grid output retrieval failing: Victoria - Prague - asymmetric paths and MTU step down issues
 - ◆ Possible network issue between McGill and BU - gridftp transfers timing out

AOB: Nobody from KIT to report on their network issue

This topic: LCG > WLCGDailyMeetingsWeek160613

Topic revision: r16 - 2016-06-13 - 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