

Table of Contents

Week of 171113.....	1
WLCG Operations Call details.....	1
General Information.....	1
Tier-1 downtimes.....	1
Links to Tier-1 downtimes.....	1
Monday.....	2

Week of 171113

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 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: Kate (WLCG, DB), Gavin (computing), Ivan (ATLAS), Edoardo (networks), Maarten (WLCG, ALICE), Alberto (monitoring), Julia (WLCG)
- remote: Stefan R (LHCb), Christopher, Dmytro (NDGF), Kyle (OSG), Chi Hsun (ASGC), Xin (BNL), Di Qing (TRIUMF), Gareth (RAL), Dave (FNAL), Marcelo (CNAF), David (IN2P3)

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ CNAF Incident - ATLAS point of view
 - ◇ Crisis Unit formed to handle the case. Includes experts from CNAF, DDM, central production and analysis support.
 - ◇ Working under the assumption of the worst case scenario i.e. Everything in CNAF is lost
 - ◇ DATA:
 - Tape
 - 16.9PB - pledged (9% T1),
 - 9.0 PB - used (4.9 PB - data, 3.6 PB - MC, 0.5 group)
 - We should always have two copies of RAW on tape: Replication from CERN CASTOR to other T1s (only BNL for the moment) of INFN data RAW.
 - Replication goes well. The situation represents a perfect testbed for testing the bandwidth capabilities of replication from CASTOR. Currently the limitation seems to be on CERN network infrastructure side.
 - Disk
 - 5.1 PB - pledged (9% T1),
 - 4.5 PB - used (4.4 PB - data, 0.054 PB - scratch)
 - 37k unique datasets (21.9k NTUP, 4.7k log, 3k DAOD, 2.3k HIST, 2.2k HITS, 1.2k AOD)
 - ◇ Processing:
 - 6k slots pledged (7% T1), 3.8k running in the moment of the incident (2.8k - Derivations, 0.5k - MC Simulations, 0.5k - Analysis). All jobs aborted and left to Panda for further reassignment
 - Central Production
 - All production managers are informed. We are working with the ones affected
 - All tasks assigned to INFN (102 tasks, 66 running) are being paused centrally and aborted and resubmitted by the corresponding production managers.
 - The inputs for the ones which have their inputs only on INFN will be eventually recreated.
 - Analysis
 - All analysis jobs were aborted
 - All users were advised to rerun their jobs if they had any running on INFN. No user complaints till the moment.
- CMS reports (raw view) -
 - ◆ **CMS Computing week ongoing; probably no physical presence at the meeting - our apologies**

- ◆ high level of activity everywhere - eagerly waiting for finalization of RELEASE+CALIBRATIONS in order to start MC17v2 processing + Data17 reprocessing. Yes, we are (very) late.
 - ◇ first step will be production of PreMixed samples. Do not really expect real processing before 1 week.
- ◆ PhaseII TDR samples (HGCAL) finalization. Currently estimating what has become unreachable since on CNAF disks. Might need urgent reinjections
- ◆ Disk situation is increasingly worrying: at the moment we have > 80% of the total disk used by unmovable (by policy) samples. A first O(5PB) deletion should happen in 1-2 days. It will not be the last one.
- ◆ Currently taking PPref data, high impact on T0 and transfer system due to the very high trigger rate. Still in early phase, monitoring what is happening. On the other hand, the High-Lumi pp backlog at T0 has already vanished (CPU), or is stable (transfers).
- ◆ CNAF problem impact for CMS:
 - ◇ **first of all our best wishes for a positive improvement of the situation in Bologna**
 - ◇ CNAF is our 2nd largest Tier-1, and hosting 1/3-1/4 of the remaining free tapes for CMS. At first approximation:
 - We do not think the CPU will really be a problem
 - Tape is a problem mostly for the drastic reduction of free tape. Assuming we will have the tape back online at some point, we can probably survive
 - Disk is a problem since in CMS, due to the design + storage underprovision, we have generally a single copy of everything. So we are effectively missing samples, and we need to plan for their regeneration (PhaseII, mostly)
- ALICE -
 - ◆ Our best wishes to **CNAF** for recovery from the flooding !
 - ◆ Effects on ALICE:
 - ◇ 68 M files unavailable, but the most important ones have replicas elsewhere.
 - ◇ Tape: raw data replication will at least be delayed...
 - ◇ CPU: not a problem.
 - ◆ High to very high activity on average
 - ◆ CERN: severe EOS instabilities and downtime on Fri. (GGUS:131749 [↗](#))
 - ◇ ~1.2 M files were lost from the name space.
 - ◇ The EOS team are recovering them; about one third are done already.
 - ◇ We thank the EOS team for their efforts!
 - ◆ CERN: severe issues with ALICE CI services on OpenStack on Sat. (GGUS:131756 [↗](#))
 - ◇ Cured by ALICE expert after checks and advice from OpenStack team, thanks!
 - ◇ Root cause not yet understood.
- LHCb reports (raw view) -
 - ◆ Activity
 - ◇ New round of stripping validation before launching the campaign.
 - ◆ Site Issues
 - ◇ INFN-T1:
 - Several issues b/c of the site outage in all areas of experiment distributed computing. Currently working on an analysis of the situation also in view of upcoming data processing campaigns.

Sites / Services round table:

- ASGC: NTR
- BNL: NTR
- CNAF: Last Thursday CNAF T1 was flooded by the breaking of a water supply line from the street. The Situation is the following:

- ◆ During the Weekend there was an extra shift mainly of Network people to try to restore networking services.
- ◆ The main damage was in the power breakers area.
- ◆ The farm had water up to the second level of machines out of 40 levels (~5%)
- ◆ Some tapes were also under water. Oracle was contacted and it will be seen if it is possible to restore the data in those tapes.
- ◆ The first thing being evaluated is the power supply damage, after that is up and running we can start to see the damage in other systems.
- ◆ It is still too early to estimate a timeline to have at least part of the T1 running at full, but all the necessary actions are being taken.
- EGI: NC
- FNAL: NTR
- IN2P3: NTR
- JINR:
 - ◆ We can not find a reason for HC errors 10&11.11, seems it is not our problem.
 - ◆ Found a few files with bad checksum at Buffer. Under investigation.
 - ◆ For the whole week, CMS pilots were run with a very low number of payload jobs. While 10 cores are allocated to pilot, it runs only 4-6 payload.

- KISTI: nc
- KIT: nc
- NDGF: network downtimes to NDGF are expected on Monday and Thursday, time window 21:00-24:00 for both days, due to routers upgrade and maintenance. In both cases the outage should not last longer than 10 minutes. Downtimes are in GOCDB.
- NL-T1: NTR. Unable to dial in, apologies.
- NRC-KI: nc
- OSG: NTR
- PIC: nc
- RAL: Announced in GOCDB: Castor outage tomorrow (Tuesday) for patching; Outage on Wednesday to upgrade LHCb Castor SRMs.
- TRIUMF: NTR

- CERN computing services: NTR
- CERN storage services:
 - ◆ FTS: FTS upgrade to v 3.7.7, tomorrow 14 Nov from 10:00 to 11:00 CET (OTG:0040777 [↗](#))
- CERN databases: NTR
- GGUS: NTR
- Monitoring:
 - ◆ Final Site Availability monthly report will be generated on the **16 November**.
 - ◆ Please send corrections to the VOs and LCG Office.
 - ◆ Check: <http://wlcg-sam.cern.ch/reports/2017/201710/wlcg/> [↗](#)
- MW Officer: As requested by LHCb, new versions of HEP_OSLibs for EL7 (7.2.3-1) and HEP_OSLibs_SL6 (1.1.2.-1) are available on the WLCG repos. Site supporting LHCb VO should upgrade their WNs.
- Networks: NTR
- Security: NTR

AOB:

This topic: LCG > WLCGOpsMeetingWeek171113

Topic revision: r20 - 2017-11-13 - MaartenLitmaath



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

WLCGOpsMeetingWeek171113 < LCG < TWiki

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback