

# Table of Contents

<b>Week of 180806.....</b>	<b>1</b>
WLCG Operations Call details.....	1
General Information.....	1
Best practices for scheduled downtimes.....	1
Monday.....	1

# Week of 180806

## 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 Geneva time 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 Portal [↗](#)
- Whenever a particular topic needs to be discussed at the operations meeting requiring information from sites or experiments, it is highly recommended to announce it by email to [wlcg-scod@cernSPAMNOTNOSPAMPLEASE.ch](mailto:wlcg-scod@cernSPAMNOTNOSPAMPLEASE.ch) to allow the SCOD to make sure that the relevant parties have the time to collect the required information, or invite the right people at the meeting.

## Best practices for scheduled downtimes

- Best practices for scheduled downtimes

## Monday

Attendance:

- local: Borja (Chair/Monitoring), Gavin (Computing), Jiri (ATLAS), Marian (Network), Roberto (Storage), Vincent (Security)
- remote: Dave M (FNAL), David B (IN2P3), Jeff (OSG), John (RAL), Ken (CMS), Luca T (LHCB), Marcelo (CNAF), Pepe (PIC), Sang Un (KISTI), Ville (NDGF), Xavier (KIT), Xin (BNL)

Experiments round table:

- ATLAS reports ( raw view) -
  - ◆ 250 - 350 k grid jobslots used plus some HPC (Titan) resources are also used
  - ◆ a new procedure to restart rucio web servers helped to avoid midnight problems
  - ◆ another 2 sites hit by CREAM CE update issue described in [GGUS:136074](#) [↗](#)
- CMS reports ( raw view) -
  - ◆ Overall another quiet week for the distributed system, as far as this CRC can tell.
  - ◆ Averaged 172k cores for production this week, with a decline in usage over the weekend. Averaged 62k cores for analysis, with a corresponding increase over the weekend.
  - ◆ Interesting T1 interaction of the week: [GGUS:136518](#) [↗](#), CNAF identifies which files were lost due to waterlogged tapes from the flood, files are globally invalidated by CMS, no raw

data loss.

- ◆ Dynamic data management exhibited some unexpected behavior due to deployment of new features; these have been rolled back.

- ALICE -

- ◆ Apologies: ALICE operations experts will not attend today
- ◆ Normal activity levels on average last week
- ◆ CERN: EOS-ALICE incident Sun afternoon caused ~50k jobs to fail (OTG:0045256[↗](#))
  - ◇ Resolved early evening, thanks!
- ◆ IN2P3-CC: 110 TB lost due to RAID problem

There was a question from IN2P3 to know if ALICE will resend lost data, experts will be involved.

- LHCb reports ( raw view) -

- ◆ Activity
  - ◇ Data reconstruction for 2018 data
  - ◇ User and MC jobs
- ◆ Site Issues
  - ◇ CERN: CVMFS problem at RHEA (GGUS:136330[↗](#)); solved
  - ◇ RAL: FTS3 issue (GGUS:136199[↗](#)); in progress
  - ◇ CNAF: data transfer problem (GGUS:136123[↗](#)); in progress

Sites / Services round table:

- ASGC: NC
  - BNL: NTR
  - CNAF: GGUS:136123[↗](#), Issue solved and is stable for 3 days, ticket still open to ensure everything is fine.
  - EGI: NC
  - FNAL: NTR
  - IN2P3: NTR
  - JINR: NC
  - KISTI: NTR
  - KIT: NTR
  - NDGF: Bluegrass site (the Triolith's replacement) downtime continues for some time. About 1/4 of the computational power of NDGF-T1 site is not available.
  - NL-T1: NC
  - NRC-KI: NC
  - OSG: NTR, Jeff will be the the representative for OSG from now on. Marian explained a bit what is the purpose of the meeting.
  - PIC: There was a big network incident originated 3 days ago on midnight, solved now. It created 8 hours of blackout and the experts are still working on understanding the problem.
  - RAL: [UPDATED 15/8/2018] - Tier-1 at RAL have had an unscheduled outage for Echo since Friday evening (10/8/18). The initial cause of this issue is thought to be a combination of client I/O and cluster I/O (moving data onto new hardware), which is causing an excessive amount of memory usage. This in turn is causing machines to start swapping heavily, which is placing even more load on the cluster. This has necessitated pausing all I/O operations on the cluster to try and reduce memory usage while we are restoring the cluster to a healthy state, and putting measures in place to reduce the memory usage on the storage nodes. While this investigation and work is carried out the ongoing downtime has been extended until 12:00 16/08/2018
  - TRIUMF: NC
- 
- CERN computing services: NTR
  - CERN storage services:

- ◆ EOSALICE was stuck Sunday afternoon, it had to be manually restarted and went back online at ~19h40
- CERN databases: NC
- GGUS: NTR
- Monitoring:
  - ◆ Draft reports for the July availability sent around
- MW Officer: NC
- Networks: Since Fri Aug 3rd, outage on the 100g link between SURFsara and NIKHEF (NIKHEF is using SURFnet IP Uplink, KIAE/Russian T2s using CERN or LHCONE). Should have no impact on production, please report any issues to SURFsara or WLCG Network Throughput.
- Security: NTR

AOB:

---

This topic: LCG > WLCGOpsMeetingWeek180806

Topic revision: r14 - 2018-08-15 - DarrenMoore



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