

Table of Contents

Week of 191021.....	1
WLCG Operations Call details.....	1
General Information.....	1
Best practices for scheduled downtimes.....	1
Monday.....	1

Week of 191021

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 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: Alberto (Monitoring), Aris (Storage), Borja (Chair, Monitoring), Ivan (ATLAS), Kate (WLCG), Maarten (ALICE), Marcelo (LHCb, CNAF), Vincent (Security)
- remote: Andrew (Nikhef), Caio (CMS), Dave (FLAB), Darren (RAL), Di (TRIUMF), Sang Un (KISTI), Xavier (KIT)

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ Networking incidents feedback
 - ◆ EOS is not reliable enough for a critical monitoring web services. Latest example - OTG:0052868 [↗](#) . Recommendations?

On the networking question, Maarten acknowledged this is having an impact on everyone. There is not much more information available than reported on the OTGs (See Network service section). Agreed Network team should be more proactive on communicating these issues.

About EOS, pointed by Maarten since this is becoming more a critical service, interaction should be through

GGUS tickets that will ultimately be reported on the monthly Management Board. Aris will raise ATLAS question internally to the team.

- CMS reports (raw view) -
 - ◆ **Due to the CMS Offline and Computing workshop in parallel nobody from CMS can attend the call**
 - ◇ Drop a mail, if something needs follow up
 - ◆ Business as usual - no major issues

- ALICE -
 - ◆ NTR

- LHCb reports (raw view) -
 - ◆ Activity:
 - ◇ MC, user jobs and data restripping.
 - ◇ Continuing staging (tape recall) at all T1s
 - ◆ Issues:
 - ◇ GRIDKA: Issue with Data transfers Saturday morning, fixed after a few hours. Investigating for lost files.

Sites / Services round table:

- ASGC: NTR
- BNL: NC
- CNAF: ATLAS 0 size files being investigate (GGUS:143682 [↗](#))
- EGI: NC
- FNAL: NTR
- IN2P3: NC
- JINR: 1.9 PB added to T1_RU_JINR_Disk
- KISTI: NTR
- KIT:
 - ◆ The dCache database of the SE dedicated to LHCb died on Saturday near 4 a.m (GGUS:143699 [↗](#)). The on-call engineer triggered an expert about 09:30 a.m. in the morning and dCache was switched to a warm stand-by database node. Since about 10:15 a.m. LHCb operations were able to resume back to normal operations. However, we are still investigating whether any files have been forgotten by dCache because of the database switch.
 - ◆ Next Wednesday at 8 a.m. CEST, we will reboot one of KIT's border routers with a NX-OS update. This will cause a network interruption for a segment of our WNs of maybe up to 20 minutes. Those WNs have been taken out of the workload management system for the time being, so no jobs should be lost. A downtime at-risk has been added to GOC-DB for your information (GOCDB:27882).
- NDGF: NC
- NL-T1: The Nikhef cvmfs clients were continuously switching over to the fallback cvmfs squids at cern rather than using our local squid servers (GGUS:143535 [↗](#)). This was behaviour stopped when the IPv6 issue at RAL was solved (GGUS:43402 [↗](#)).

Maarten asked if there was any consequence that drove the investigation of this behaviour. Reply from Andrew state they saw more network traffic going to the worker nodes. Conclusion is it might be worth to report this unexpected behaviour to developers.

- NRC-KI: NC
- OSG: NC
- PIC: Not able to connect today. Sorry. We will update to dCache 5.2.5 on 12/11. We will announce the SD soon.

- RAL: NTR
- TRIUMF: NTR

- CERN computing services: NC
- CERN storage services: NTR
- CERN databases: NC
- GGUS: NTR
- Monitoring:
 - ◆ Final availability reports sent for September
- MW Officer: NC
- Networks: Multiple services at CERN affected by a major network incident (OTG:0052774[↗](#), OTG:0052818[↗](#))
- Security: NTR

AOB:

This topic: LCG > WLCGOpsMeetingWeek191021

Topic revision: r20 - 2019-10-21 - 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