

# Table of Contents

<b>Week of 200622.....</b>	<b>1</b>
WLCG Operations Call details.....	1
General Information.....	1
Best practices for scheduled downtimes.....	1
Monday.....	1

# Week of 200622

## 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
- 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@cern.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:
- remote: Andrew (Nikhef), Ben (Computing), Borja (Chair, Monitoring), Concezio (LHCB), David (IN2P3), Dave (FNAL), Eric (Storage), Eva (Databases), Jens (NDGF), Julia (WLCG), Maarten (ALICE), Pavlo (ATLAS), Pepe (PIC), Vincent (Security), Xin (BNL)

Experiments round table:

- ATLAS reports ( raw view) -
  - ◆ Storage issues at: UKI-LT2-QMUL GGUS:147558 [↗](#), UNIBE-LHEP GGUS:147556 [↗](#), UKI-NORTHGRID-LANCS GGUS:147553 [↗](#), INFN-T1 GGUS:147462 [↗](#)
  - ◆ CERN Frontiers were down during part of the week
- CMS reports ( raw view) -
  - ◆ **It's (virtual) CMS Computing Workshop and likely nobody from CMS can call in**
  - ◆ Bad CMS workflow caused storage overload at CC-IN2P3
    - ◇ Clarified within CMS and bad WF got aborted - no GGUS ticket
  - ◆ Otherwise no major items to report

- ALICE -
  - ◆ NTR
- LHCb reports ( raw view) -
  - ◆ completing stripping of heavy ions collision data
  - ◆ MC production as usual
  - ◆ planning for DB outage of Saturday
  - ◆ GGUS down -- some impact on operations

Sites / Services round table:

- ASGC: NC
- BNL: FTS was upgraded last Monday, to version 3.9.4.
- CNAF: NTR
- EGI: NC
- FNAL: Two incidents with tape cartridges in entirely different libraries last week. We had an RFID chip in a ~4 year old T10kC cartridge fail. Believe this is recoverable, but data is unavailable until Oracle re-finds the file boundaries and recovers. On Friday though we had a more catastrophic failure of a few month old LTO8 tape in our IBM library. Tape itself snapped and wound within the drive. Removed entangled cartridge and drive and shipped to IBM for further investigation. In that one expect about 192 CMS files will be lost.
- IN2P3: during maintenance on last Tuesday :
  - ◆ CREAM-CEs have been decommissioned.
  - ◆ upgrade of HTCondorCE to 4.1.0: incident on June 19th with HTCondor job router -> impact on ALICE during the afternoon.
  - ◆ dCache: upgrade to 5.2.22 and rollback to version 5.2.21 (last pools for ATLAS/CMS analysis done today midday)
  - ◆ new endpoints for LHCb on dCache to distinguish disk and tape. Registered in GOC DB.
- JINR: NTR
- KISTI: NC
- KIT: NC
- NDGF: NTR
- NL-T1:
  - ◆ The slow directory listings at the Sara dCache (reported last week) are understood. A user has 3 million files in one dir, that is a bit of a challenge. We fixed this by increasing `pnfsmanager.limits.list-threads` from 2 (default) to 10. This slightly increases the database load but nothing it can't handle.
  - ◆ The Sara tape backend was down from Saturday until this afternoon. It is operational now.
- NRC-KI: NC
- OSG: NC
- PIC: PIC Tier 1 will be in scheduled downtime on Tuesday June 30th, from 08:00 to 14:00 (CERN and local time), in order to perform upgrades on the compute (HTCondor) and storage (dCache and Enstore) services. As usual, access to the CPU farm will be closed right before the start of the SD.
- RAL: NTR
- TRIUMF: NC
- CERN computing services:
  - ◆ DB intervention affecting some Compute services
    - ◇ Hammercloud OTG:0057297 [↗](#)
    - ◇ SLURM OTG:0057335 [↗](#)
    - ◇ VOMS (OTG to be added by Ben)
    - ◇ BOINC OTG:0057334 [↗](#)
- CERN storage services:
  - ◆ On 23rd of June, Final closure of ATLAS area in CASTOR : access to the ATLAS tree of

CASTOR will be definitely blocked, before final migration of data to CTA. OTG:0057293 [↗](#)

- ◆ On 25th of June, the EOSCTA ATLAS instance will be opened for reads and writes  
OTG:0057317 [↗](#)
- ◆ On 27th of June, all CASTOR and CTA instances will be down due to a database intervention  
OTG:0057292 [↗](#)
- ◆ On 27th of June, many FTS instances will be unavailable due to DB storage intervention  
OTG:0057307 [↗](#)
  
- CERN databases:
  - ◆ Major intervention on Saturday 27th affecting Oracle and DBoD databases. More details:
    - ◇ OTG:0057263 [↗](#) (Database storage)
    - ◇ OTG:0057252 [↗](#) (DBoD)
    - ◇ OTG:0057201 [↗](#) (Oracle)
  - ◆ Eva: The intervention is scheduled in order to avoid the incident from May 27th happening again. Safest way to perform changes (affecting DBOD and Oracle) is to stop the full rack.
  - ◆ Julia: It was discussed in the Physics Services meeting and given that job submission could be affected, should CERN be in scheduled downtime?
  - ◆ Ben: VOMS proxy renewal won't work, but regular jobs scheduled should be fine
  - ◆ Maarten: The computing service won't be down and EOS won't be either, so no downtime is needed. Tape would only require an "at-risk" downtime, for which we don't have to bother here.
  - ◆ Julia: Any action to take by LHCb previous to the intervention?
  - ◆ Concezio: Given the LHCb has a critical dependency on the DBs, people are preparing beforehand (gracefully emptying the queues)
  - ◆ Julia: Can you survive without the downtime?
  - ◆ Concezio: Can't say right now, will come back with an answer
  - ◆ Ben: If there is anything from batch that can be done, we are happy to help
  - ◆ Concezio: The downtime won't be needed since the queues are going to be drained
  
- GGUS:
  - ◆ Access via CERN Grid CA certificates was refused from Sun afternoon till Mon morning.
    - ◇ Due to an expired CRL.
    - ◇ The problem was announced on the `wlwg-operations` list Sun afternoon.
    - ◇ The GOCDB entry for GGUS provides the contact e-mail to be used for such cases.
  
- Monitoring:
  - ◆ We will be sharing the draft reports from SiteMon (new infrastructure) with site managers during the week after green light from experiment representatives.
  - ◆ Monitoring will be down during the DBOD intervention on Saturday OTG:0057351 [↗](#)
  
- MW Officer:
  - ◆ Julia: There were questions about the recommended HTCondor CE version to run, Maarten, could you clarify it?
  - ◆ Maarten: As sites have been deploying such services at different times, there are several versions currently running that work fine. As far as features are concerned, there is not a required minimum version at the moment, so the only requirement is that it has to be supported by the HTCondor team. The minimum has been set to 3.2.0 for now.

- Networks: NC
- Security: NTR

AOB:

---

This topic: LCG > WLCGOpsMeetingWeek200622  
Topic revision: r25 - 2020-06-23 - BorjaGarridoBear



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.  
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)