



# WLCG Service Report

5 weeks, Apr 19 – May 24

WLCG Management Board, 24<sup>th</sup> May 2016

Maarten Litmaath  
CERN

v1.0

# Operations Coordination highlights (1)

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsCoordination>

- Decommissioning of remaining **SL5** services
  - Deadline May 31
  - [https://wiki.egi.eu/wiki/SL5\\_retirement](https://wiki.egi.eu/wiki/SL5_retirement)
- Detailed review of Network and Transfer Metrics WG
  - Ongoing activities
    - **perfSONAR** very useful in solving recent network issues
  - Future R&D projects
- Traceability & Isolation WG kick-off pre-GDB
- Accounting review TF in preparation

# Operations Coordination highlights (2)

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsCoordination>

- dCache SRM at NDGF-T1 (and other sites) **stopped working** when the OCSP service of the CERN CA was unavailable due to a firewall misconfiguration
  - dCache now has OCSP support by default disabled
  - CRLs are deemed sufficient in our environments

# Operations Coordination highlights (3)

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsCoordination>

- MW Readiness WG asked T0 and T1 sites for feedback on the proposal to deploy **Pakiti** on WLCG service hosts
  - To allow easy identification of rpm version issues
    - MW rpms for which better versions are available
    - Mismatches, in particular between head nodes and disk servers
    - And *possibly*: system rpms with security issues
  - Strict access controls would be applied
- Conclusions
  - **Pakiti** will continue to be used on the MW Readiness **volunteer site testbeds** to track the rpms of products under verification.
  - Various sites which replied are using **Pakiti** *locally*.
  - Most do **not** want rpm information to be sent anywhere else.
  - These matters will be discussed further and followed up in the Information System Evolution TF.

# Selected items from Operations (1)

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsMeetings>

- **High activity** for all experiments
- **ETF** on **SL6** has replaced SL5 SAM-Nagios
  - ALICE, CMS and LHCb as of April 25
  - ATLAS as of May 10
  - Intermittent CREAM job submission failures at a few ATLAS/CMS sites are being investigated

# Selected items from Operations (2)

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsMeetings>

- ALICE
  - Successful T1-T2 workshop in Bergen, Apr 18-20
  - CASTOR incident starting Apr 28 late evening
    - DAQ could no longer write to CASTOR
    - Cured during the night, code was fixed later
  - Both EOS-ALICE head nodes now have 512 GB RAM
  - SAM critical profile changed to include selected **MonALISA metrics** as of May 1<sup>st</sup>
- LHCb
  - **Alarm** ticket GGUS:121038:
    - CERN CASTOR SRM returns a wrong URL for xroot protocol
    - A code regression was identified and fixed by the developers

# Selected items from Operations (3)

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsMeetings>

- ATLAS
  - BNL-SARA transfers were slow ([GGUS:120957](#))
    - Packet loss on the BNL-CERN link was caused by malfunctioning hardware on an ESnet router at CERN. Fixed Apr 25<sup>th</sup>.
    - Congestion on the link CERN-SARA: fixed with the addition of a second 10G link on Apr 29<sup>th</sup>.
  - **Alarm ticket** [GGUS:121134](#):
    - Submission to CERN CREAM CEs timed out
    - Cured by restarting HTCondor on the pilot factories
- CMS
  - Mostly business as usual

# Selected items from Operations (4)

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsMeetings>

- **KIT:** **downtime** May 9-11 for cooling infrastructure maintenance and tape infrastructure migration
- **SARA:** tape **downtime** Jul 29 – Aug 15 to move tape storage system to new data center
- **PIC:** the **T10KD** drives and tapes are back in use (see SIR)
- **ASGC:** perfSONAR shows **connectivity** not reaching 10 Gbps
  - A local router was found to have limited buffer space
  - Further investigations ongoing ([GGUS:119820](#))
- **RAL:** **Stratum-1** issues May 16-18
  - Disk space issues on the production Stratum-1 HA cluster
  - Standby Stratum-1 took over, but did not have the correct routing
  - Led to job failures at least for ALICE



# Service Incident Reports

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGServiceIncidents>

## ■ PIC tape incident

- In October 2015 PIC started to use T10KD technology for ATLAS.
- On Dec 17<sup>th</sup> PIC detected a problematic file to be read from T10KD tape (internal migration) and decided to stop ATLAS writing to T10KD.
- The 79 used T10KD cartridges then were read completely – 8 tapes presented problematic files to be read.
- Data was migrated from these problematic tapes, by using dedicated firmware from Oracle for data recovery.
- 3 cartridges had to be sent to Oracle and the last 21 missing files were finally **recovered** by them.
- The problematic drive wrote its data **“off-track”** in a sense that other drives could not read the data back, depending on the built-in tolerances.

# GGUS summary (5 weeks 18/4/2016 – 22/5/2016)

by Maria Dimou

VO	User	Team	Alarm	Total
ALICE	1	2	0	3
ATLAS	23	80	1	104
CMS	182	2	0	184
LHCb	2	17	1	20
Totals	208	101	2	311

