



WLCG Service Report

5 weeks, Nov 15th – Dec 19th

WLCG Management Board, Dec 19th 2017

Katarzyna Dzedziniewicz-Wojcik
CERN

v1.4

Operations Coordination meetings

- These meetings normally are held **once** per month
 - Usually on the first Thursday
- Each meeting has a **standard agenda** plus at least one **dedicated topic**, announced in advance
- Experiments and sites are kindly asked to have the **relevant experts** attend, depending on the topic(s)
- Next meeting:
 - Jan **18th** – Theme: Site availability reports and SAM tests

Operations Coordination highlights (1)

- CMS sites are expected to *deploy Singularity* by March 2018
 - A broadcast from WLCG is foreseen after documentation is finalized
- Theme: Handling consequences of CNAF incident
 - No loss of non-reproducible data
 - CPU resources compensated by other sites
 - Sample reproduction will be the most effort-consuming
 - Single copy of reconstructed data due to RRB policies
 - Full recovery cost to be calculated

Operations Coordination highlights (2)

- Container WG
 - Setuid needed for Singularity in SL6
 - Not needed for RH 7.4 when OverlayFS is not used
 - Not using OverlayFS would require additional work from ATLAS
 - Image distribution will be done unpacked in CVMFS
 - Cannot be done for fast iterations in analysis workflows
- Archival Storage WG
 - First meeting was held
 - Questionnaire for sites is being finalized
- Accounting TF
 - Implementation of Storage Resource Reporting is progressing for all SE types
- Information System Evolution TF
 - CRIC description of computing resources implemented
 - Storage resource description started

Operations Coordination highlights (3)

- SAM and SSB migration to the MONIT infrastructure
 - All SAM data available in MONIT
 - Validation planned for Q1 2018
 - Support for VO-injected SSB metrics will be available in Q1 2018
- Squid Monitoring and HTTP Proxy Discovery TFs
 - All ATLAS Tier 1 and Tier 2 sites are now having their squids monitored with MRTG
 - CMS OSG Opportunistic moved to <http://openhtc.io> aliases for CMS Frontier servers
- IPv6 Validation and Deployment TF
 - T2 deployment ticket campaign is progressing

Operations Coordination highlights (4)

- Network and Transfer Metrics WG
 - perfSONAR 4.0.2 was released
 - perfSONAR workshop held by JISC in November
 - HEPiX WG on SDN/NFV was established and will look into networking R&D topics
- Middleware
 - New UK Science CA release caused issues with many Java-based services for UK users
 - Notably Argus, dCache, VOMS-Admin instances
 - A bug that caused DPM SRM crashes at multiple sites was identified and fixed
 - Advisory on various vulnerabilities allowing privilege escalation (SVG-CVE-2017-5712)

Selected items from Operations (1)

- Overall high activity across the experiments
- ATLAS
 - High CPU utilization (peaks at ~600k cores)
 - EOS/CASTOR export limits reached by high-rate runs
 - Derivation and reprocessing campaigns started
 - Tape data hosted at CNAF were re-replicated from CASTOR
- CMS
 - Continued high utilization of the system
 - Transfer system remains under pressure
 - GEN-SIM samples at CNAF are being reproduced already
 - CERN Glidein Factory issue caused global pool to shrink to 100k cores

Selected items from Operations (2)

- ALICE
 - New record: 162k concurrent jobs
 - CERN: several issues with CASTOR and EOS
 - E.g. XRootD bugs revealed by high load
 - CASTOR was reconfigured for reco at high scale
 - Accidental overload of EOS-PUBLIC by broken grid jobs obtaining a special file from the wrong location
- LHCb
 - Stripping campaign ongoing, stressing tape systems
 - Almost no free disk space left
 - FZK-LCG2: **Alarm** ticket because of low disk space in GridKa (GGUS:132090)

Selected items from Operations (4)

- BNL
 - Site wide connectivity outage due to interruptions of both network circuits
- CNAF
 - Tape library and disk SE recovery preparatory works ongoing
 - Partial site availability is expected early in February
- KIT
 - Several damaged tapes were discovered
 - ~ 4200 files lost
- NL-T1
 - dCache hangs caused by external Zookeeper communication cured by using internal version

Selected items from Operations (5)

- TRIUMF
 - All computing resources migrated from CREAM + Torque to ARC + HTCondor

- CERN
 - Multiple databases were down or degraded due to a storage issue (OTG:0040966)
 - CASTOR recall queue length reached ~1 week due to repacking of tapes and re-replication of CNAF data
 - Urgent recalls can be prioritized on request

Service Incident Reports

- CERN: Database Storage Incident

GGUS summary (5 weeks Nov 13 – Dec 17)

VO	User	Team	Alarm	Total
ALICE	3	6	0	9
ATLAS	35	94	0	129
CMS	198	1	0	199
LHCb	3	26	1	30
Totals	239	127	1	367

