

Table of Contents

WLCG Operations Coordination Minutes, Oct 5th 2017.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Middleware News.....	1
Discussion.....	2
Tier 0 News.....	2
Tier 1 Feedback.....	2
Tier 2 Feedback.....	2
Experiments Reports.....	2
ALICE.....	3
ATLAS.....	3
CMS.....	3
LHCb.....	3
Ongoing Task Forces and Working Groups.....	4
Accounting TF.....	4
Information System Evolution TF.....	4
IPv6 Validation and Deployment TF.....	4
Machine/Job Features TF.....	4
Monitoring.....	4
MW Readiness WG.....	4
Network and Transfer Metrics WG.....	4
Squid Monitoring and HTTP Proxy Discovery TFs.....	5
Traceability WG.....	5
Container WG.....	5
Special topics.....	5
WLCG forum for tape matters.....	5
Theme: Providing reliable storage - KIT.....	5
Action list.....	6
Specific actions for experiments.....	7
Specific actions for sites.....	7
AOB.....	7

WLCG Operations Coordination Minutes, Oct 5th 2017

Highlights

- WLCG forum for tape matters
- Theme: Providing reliable storage - KIT

Agenda

- <https://indico.cern.ch/event/670207/>

Attendance

- local: Alessandro (ATLAS), Andrea M (MW Officer + data management), Gavin (T0), Maarten (WLCG + ALICE), Marian (networks), Mayank (WLCG), Oliver (data management), Vlado (CERN tape storage), Zoltan (LHCb)
- remote: Alessandra D (Napoli), Andrea V (LHCb), Andreas (KIT), Bo (FNAL), Carles (PIC), David B (IN2P3-CC), David C (Glasgow), David Yu (BNL), Di (TRIUMF), Eric (IN2P3-CC), Esther (PIC), Frederic (IRFU), Giuseppe (CMS), Jan Erik (KIT), Julien (CERN tape storage), KIT storage team, Kate (WLCG + databases + CMS), Leo (Sussex), Marc (PIC), Mario (ATLAS), Max (KIT), Natalia (FNAL), Pepe (PIC), Peter (Oxford), Ron (NLT1), Sang-Un (KISTI), Sean (CHPC), Simon (TRIUMF), Thomas (DESY), Tim (RAL), Ulf (NDGF), Vanessa (PIC), Vincent (ATLAS), Vladimir (CNAF), William (CHPC), Xavier (KIT), Xin (BNL)
- apologies: Catherine (Lpsc + IN2P3)

Operations News

- GGUS: some GUI changes will be implemented for the Nov release to improve ticket categorization and prioritization
 - ◆ The ticket category will no longer have a default value (was: Incident)
 - ◆ The ticket priority will no longer have a default value (was: less urgent)
 - ◆ The ticket category Change Request will be renamed Service Request ([doc](#))
- The HEPiX autumn meeting 2017 will be held Oct 16-20 at KEK
 - ◆ [agenda](#)
- The next meeting is planned for Nov 2
 - ◆ Please let us know if that date would present a major issue

Middleware News

- Useful Links:
 - ◆ [Baseline Versions](#)
 - ◆ [MW Issues](#)
 - ◆ [Storage Deployment](#)
- Baselines/News:
 - ◆ dCache 2.13 EOL. 4/12 instances still have to upgrade.

- Issues:
 - ◆ High risk Linux kernel high local root vulnerability broadcasted by EGI SVG (<https://wiki.egi.eu/wiki/SVG:Advisory-SVG-CVE-2017-1000253>)
 - ◆ High risk dnsmasq multiple vulnerabilities broadcasted by EGI SVG (<https://wiki.egi.eu/wiki/SVG:Advisory-SVG--CVE-2017-14491>)
 - ◆ Xrootd 4.7.0 broke xrd-cp command towards dCache sites. (<https://github.com/xrootd/xrootd/issues/593>). Xrootd 4.7.1 is under preparation to fix this issue, please use Xrootd 4.6.1 in the meantime.
 - ◆ Vulnerabilities reported in Xen by EGI SVG <https://wiki.egi.eu/wiki/SVG:Advisory-SVG-CVE-2017-8903>
 - ◆ 2 DPMs still running SL5 have started to fail FTS Atlas transfers via SRM. SL5 is EOL since March 2017 Sites declared a downtime while moving to SL6/C7 (already done in one case).
- T0 and T1 service
 - ◆ BNL
 - ◇ FTS upgraded to v3.7.4
 - ◇ Xrootd 4.7.x under testing
 - ◆ CERN
 - ◇ check T0 report
 - ◇ EOS ALICE major upgrade to v4.1 (Xrootd4.7 and IPv6 support)
 - ◆ IN2P3
 - ◇ Upgraded Xrootd to 4.6.1-1, upgraded dCache to 2.16.46
 - ◆ KIT:
 - ◇ Upgraded dCache for ATLAS to 2.16 on September 20th. Dropped IPv6 support after several issues
 - *Andreas: a special setup of 1 IPv6 subnet had to be rolled back; more testing will need to be done first*
 - ◇ Downtime for updating dCache for CMS fixed on October 11th. LHCb pending
 - ◆ NDGF:
 - ◇ Upgraded dCache to 3.0.28 this week with xrootd fixes for the pools
 - ◆ RRC-KI-T1:
 - ◇ Upgraded dCache to 2.16.30

Discussion

Tier 0 News

- Will move asap remaining CERN Grid capacity fully to HTCondor. CREAM CEs will be retired.
- Alessandro: OK for the grid access, but what about local access to LSF?
- Gavin: that will remain available, but at a level of a few k cores
- Maarten: OK for ALICE
- Zoltan: OK for LHCb
- Giuseppe: OK for CMS

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- The activity levels typically have been very high
- CERN: EOS-ALICE was upgraded to Citrine release with XRootD 4.7
- CERN: Oracle cloud project resources have been used with good results
 - ◆ Up to ~9k cores, located in Phoenix AZ
 - ◆ Mostly a mix of MC and reco jobs
 - ◆ Analysis jobs were briefly enabled: their efficiency was ~3% !
 - ◇ Expected result from the network latency (not the throughput)
 - ◆ Reco was disabled Oct 4 after CERN's commercial 10-Gbit uplink almost got saturated !

ATLAS

- Smooth data processing at Tier0. Mu between 60 and 40 (start -end), processing time in the ballpark of 30-15secs/event. Compute resources completely full.
- Smooth grid production operations with ~300k concurrent running job slots, good contribution from HPCs in the past month with peaks up to 600k.
- Small grid operational hiccups in the authentication service of DDM/Rucio due to service node upgrades to CC7.4 and package dependency problems (RQF:0850843)
- In the next week(s) we will run a full Derivations campaign which will reprocess all the data AOD, inputs about 3-4PB.
- changes in ATLAS Computing : now Torre is CC, Davide Costanzo his deputy. Thanks Simone!
ATLAS Distributed Computing: now ADG is the coord, and Johannes Elmsheuser the deputy ADC coord. Thanks Andrej and Nurcan!

CMS

- High CPU utilization ~200k cores surpassed in the last week (with also T0 involved)
- High transfer activity - no major problems
- CPU efficiency Task Force: work ongoing, some changes already implemented in our production system
- all CMS sites have updated the WLCG IPv6 site survey
- PhEDEx: observed tape staging of files that have disk copies (should be strongly disfavoured)
 - ◆ Correction to PhEDEx tape stage policy is ongoing to fix the problem
- T1 usage sharing (presently 95/5 for prod vs analysis)
 - ◆ Possible increase of analysis sharing under discussion
- Posix/cp stage-out commands need to be changed to gfal2/xrdcp/command to handle certificates for Singularity: few sites affected have been contacted
- Missing files observed recently in the /store/unmerged area at few sites. We remind that our policy is that only files not belonging to active workflows and older than two weeks should be removed
- Migration to HTCondor for GRID submission (as requested by CERN): no problem expected

LHCb

- LHCb is running continuously 60k jobs: Data reconstruction, data stripping, Monte Carlo simulation and user analysis
- pre-staging of 2015 data has been started. This will increase the load on the T0/T1 sites. The reprocessing of these data will start in the beginning of November.
- LCG.Oracle.cern has been created for using external cloud resource provided by Oracle
- After the LHC technical stop the data processing applications are depending on the existence of the `git` command in the worker nodes. This requires sites to update or install HEP_OSlibs meta package

Ongoing Task Forces and Working Groups

Accounting TF

Information System Evolution TF

- CRIC development is progressing well.
- An update is foreseen to be presented in the GDB [GDB](#) next week.

IPv6 Validation and Deployment TF

- Maarten:
 - ◆ the IPv6 deployment broadcast was sent
 - ◆ the GGUS ticket campaign will be launched in the next few weeks
 - ◆ the tickets are to keep track of the situation at sites, timelines, potential issues etc.

Machine/Job Features TF

Monitoring

MW Readiness WG

NTR

Network and Transfer Metrics WG

- WG update will be presented at HEPiX and LHCOPN/LHCONE workshop (co-located)
- perfSONAR YouTube channel at <https://www.youtube.com/channel/UCjK-P49pAKK9hUrrNbbe0Sg>
- perfSONAR 4.0.1 auto-deployed to 197 instances (21 are already on centos7)
 - ◆ Port 443/https is now used as a controller port for pscheduler and needs to be open on central firewalls
 - ◆ Some sites suffer from an MA access issue after the upgrade, this is being followed up
- perfSONAR 4.0.2 is planned to be released in November
 - ◆ Brings new SNMP plugin that can be used to retrieve local site router traffic
- WLCG/OSG network services
 - ◆ New documentation is in preparation and will be hosted at <https://opensciencegrid.github.io/networking/>
 - ◆ OSG collector handling multiple backends (Datastore, CERN ActiveMQ and GOC RabbitMQ) now in production
 - ◇ GOC will distribute raw data to 3 different locations, FNAL for tape archive, Nebraska for long-term ES storage, Chicago for short-term ES storage
 - ◆ Preparing new LHCOPN and perfSONAR dashboards in collaboration with CERN IT/CS and IT/MONIT
- WLCG Network Throughput Support Unit: see twiki for summary of recent activities.

- HNSciCloud will create its own perfSONAR mesh to follow up on the network performance btw. providers and sites

Squid Monitoring and HTTP Proxy Discovery TFs

- Nothing to report

Traceability WG

Container WG

Special topics

WLCG forum for tape matters

presentation [↗](#) by Oliver

presentation [↗](#) by Vlado

- Alessandro:
 - ◆ the ATLAS tape exercises revealed that past metrics etc. were forgotten
 - ◆ the experiments should be involved in the forum early on
 - ◆ too many layers between tape experts and experiment data management experts
- Vlado:
 - ◆ the tape infrastructure at a site may be shared with other customers
 - ◆ a site would also need to take their interests into account
- Alessandro:
 - ◆ to some extent tape systems should be like batch systems:
experiments do not want to know too many internal details
 - ◆ OTOH it is not yet clear what we can get out of them exactly
 - ◆ an iterative approach would be good
- Vlado:
 - ◆ one challenge lies in the different SW used across the tape sites
 - ◆ discussion of requirements would be welcome
- Alessandro:
 - ◆ could we use the FTS to abstract tape?
- Pepe:
 - ◆ we need to involve the operations and data management people in the experiments

Theme: Providing reliable storage - KIT

presentation [↗](#)

- Vladimir: are the NSD servers shared between VOs?
- Andreas: yes; the NSD servers have given us no problems

- Vladimir: how many XRootD servers for ALICE?
- Andreas: 6 or so for the disk-only SE

- Vladimir: don't your XRootD servers get overloaded by remote clients?
- Andreas:
 - ◆ we see lots of load, but the servers are able to handle it
 - ◆ 4k connections or so per server should still be OK
 - ◆ some of the load may actually come from migrations to new HW
 - ◆ we need to study the load more
- Maarten:
 - ◆ such load may be caused by remote raw data reconstruction or analysis jobs
 - ◆ most jobs run close to where their data is, but remote accesses may happen
 - ◆ please let ALICE know when there is too much load due to remote jobs

- Vladimir: what about data access by CMS? [see the AOB]
- Andreas: their data is in dCache

- Vlado: you did not include a description of your tape system?
- Andreas:
 - ◆ the main purpose of this series of talks concerns disk-only systems
 - ◆ our tape system low-level infrastructure (power etc.) looks similar
 - ◆ we do not make 2 copies of the tape data
 - ◆ we do not have a standby server for TSM
 - ◆ our HPSS service is not yet being used for the T1

- Alessandro: should ATLAS try to use the `file` protocol for data access by jobs?
- Andreas:
 - ◆ the storage currently is not mounted on the WN
 - ◆ for dCache that could in principle be done through NFSv4
 - ◆ however, we are satisfied with how data access is done today

Action list

Creation date	Description	Responsible	Status	Comments
01 Sep 2016	Collect plans from sites to move to EL7	WLCG Operations	Ongoing	<p>The EL7 WN is ready (see MW report of 29.09.2016). ALICE and LHCb can use it. NDGF plan to use EL7 for new HW as of early 2017. Other ATLAS sites e.g. Triumf are working on a container solution that could mask the EL7 env. for the experiments which can't use it. Maria said that GGUS tickets are a clear way to collect the sites' intentions. Alessandra said we shouldn't ask a vague question. Andrea M. said the UI bundle is also making progress.</p> <p>Jan 26 update: this matter is tied to the EL7 validation statuses for ATLAS and CMS, which were reported in that meeting.</p> <p>March 2 update: the EMI WN and UI meta packages are planned for UMD 4.5 to be released in May</p> <p>May 18 update: UMD 4.5 has been delayed to June</p> <p>July 6 update: UMD 4.5 has been delayed to July</p> <p>Sep 14 update: UMD 4.5 was released on Aug 10, containing the WN; CREAM and the UI are available in the UMD Preview repo; CREAM client tested OK by ALICE</p>
			Pending	

03 Nov 2016	Review VO ID Card documentation and make sure it is suitable for multicore	WLCG Operations		Jan 26 update: needs to be done in collaboration with EGI
26 Jan 2017	Create long-downtimes proposal v3 and present it to the MB	WLCG Operations	Pending	May 18 update: EGI collected feedback from sites and propose a compromise - 3 days' notice for any scheduled downtime; Oct 5 update: as both OSG and EGI were not happy with the previous proposals, and as this matter does not look critical, we propose to create best-practice recipes instead and advertise them on a WLCG page
06 Jul 2017	Ensure a forum exists for discussing tape matters	WLCG Operations	DONE	
14 Sep 2017	Followup of CVMFS configuration changes, check effects on sites in Asia	WLCG Operations	Pending	

Specific actions for experiments

Creation date	Description	Affected VO	Affected TF/WG	Comments	Deadline	Completion
---------------	-------------	-------------	----------------	----------	----------	------------

Specific actions for sites

Creation date	Description	Affected VO	Affected TF/WG	Comments	Deadline	Completion
---------------	-------------	-------------	----------------	----------	----------	------------

AOB

- Vladimir:
 - ◆ we saw thousands of connections from CMS jobs at CERN to our XRootD servers
 - ◆ we opened a ticket for CMS but did not see much progress yet
- Giuseppe:
 - ◆ the trouble seems to have been due to jobs from a single user
 - ◆ we have updated the ticket and will follow up further
- Xin: do we need to equip our WN with IPv6?
- Alessandro: it is not needed; if you wish to try it out, ATLAS can help
- Maarten: WLCG has only asked for storage to be made dual-stack

This topic: LCG > WLCGOpsMinutes171005

Topic revision: r16 - 2018-02-28 - MaartenLitmaath



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