

Table of Contents

WLCG Operations Coordination Minutes, July 6th 2017.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Middleware News.....	1
Discussion.....	2
Tier 0 News.....	2
Tier 1 Feedback.....	3
Tier 2 Feedback.....	3
Experiments Reports.....	3
ALICE.....	3
ATLAS.....	3
CMS.....	3
LHCb.....	4
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.....	5
MW Readiness WG.....	5
Network and Transfer Metrics WG.....	5
Squid Monitoring and HTTP Proxy Discovery TFs.....	6
Traceability and Isolation WG.....	6
Special topics.....	6
MW deployment forums and feedback.....	6
Theme: Providing reliable storage - IN2P3.....	6
Action list.....	7
Specific actions for experiments.....	8
Specific actions for sites.....	8
AOB.....	8

WLCG Operations Coordination Minutes, July 6th 2017

Highlights

- Operations News
- Presentation and discussion on MW deployment forums and feedback
- Presentation and discussion on Providing reliable storage at IN2P3-CC

Agenda

- <https://indico.cern.ch/event/651433/>

Attendance

- local: Andrea M (MW Officer + data management), Andrea S (IPv6), Gavin (T0), Julia (WLCG), Maarten (WLCG + ALICE)
- remote: Alessandra D (Napoli), Alessandra F (Manchester + ATLAS), Alessandro (CNAF), Brian (RAL), Catherine (LPSC + IN2P3), David B (IN2P3-CC), David C (Glasgow), David M (FNAL), Di (TRIUMF), Eric (IN2P3-CC), Felix (ASGC), Gareth (RAL), Giuseppe (CMS), Javier (IFIC), Jeremy (GridPP), Kyle (OSG), Marcelo (LHCb), Marcin (PSNC), Renaud (IN2P3-CC), Ron (NLT1), Sang-Un (KISTI), Thomas (DESY), Vikas (VECC), Xin (BNL)
- apologies: Marian (networks), ATLAS

Operations News

- WLCG workshop took place from the 19th to the 22nd of June hosted by the University of Manchester. Thank you to our Manchester colleagues, in particular Alessandra for excellent organization. The operations session chaired by Pepe included many important areas like benchmarking, monitoring, IS system evolution and storage space accounting. More details can be found here
- Pre-GDB on *containers* will be held Tue July 11 afternoon
- GDB will be held on Wed July 12
- the next meeting is planned for Sep **14**
 - ◆ please let us know if that date would present a major issue

Middleware News

- Useful Links:
 - ◆ Baseline Versions
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines/News:
 - ◆ Globus EOL in 2018 (<https://www.globus.org/blog/support-open-source-globus-toolkit-ends-january-2018>).

- ◊ So far it looks likely that CERN together with OSG will take over the code maintenance and support in the short term, hopefully with the continued participation of a person from NDGF. In the longer term we will look at how this code should be replaced, in particular gsi and gridftp. Essentially this is a non-issue for now.
- ◆ Perfsonar Baseline moved to v4.0.0 (from last meeting), removed dCache 2.13 from baseline and added dCache 2.16.39
- ◆ dCache 2.13.x EOL on June, only KIT and FNAL among T1s are still running this version.
- ◆ Some new products are expected to be released in UMD4 within this month:
 - ◊ WN metapackage for C7 (GGUS:128753 [↗](#))
 - ◊ new Cream-CE release, only for C7 (GGUS:129068 [↗](#))
 - ◊ dynafed both SL6/C7 (GGUS:128556 [↗](#))
- ◆ As broadcasted by C Aiftimiei, the EMI repos have been shutdown on 15/06.
- Issues:
 - ◆ The issue for EOSCMS and EOSATLAS GSI connection overloading reported during the last meeting has been fixed with a new version of the Xrootd GSI plugin released and deployed.
 - ◆ Critical Vulnerability on VOMS Admin (disclosed this week but already fixed) <https://wiki.egi.eu/wiki/SVG:Advisory-SVG-2017-12543> [↗](#)
 - ◆ Heads-up: A memory allocation vulnerability (Stack Clash) has been found which may allow an authorized user to gain root privileges. <https://wiki.egi.eu/wiki/SVG:Advisory-SVG-CVE-2017-1000364> [↗](#)
 - ◆ Moderate Risk : sudo local root vulnerability <https://wiki.egi.eu/wiki/SVG:Advisory-SVG-CVE-2017-1000367> [↗](#)
- T0 and T1 services
 - ◆ CERN
 - ◊ Castor upgrade to 2.1.16-18 for all VOs, diskserver migration to C7
 - ◊ 2 load balanced HAProxy servers deployed in front of Production FTS
 - ◆ IN2P3
 - ◊ Major dCache upgrade to v2.16.37
 - ◊ Upgrade of Xrootd during the next stop in september
 - ◆ JINR
 - ◊ Minor dCache upgrade 2.16.31 -> 2.16.39 on both instances;
 - ◊ minor xrootd upgrade 4.5.0-2.osg33 -> 4.6.1-1.osg33 for CMS
 - ◆ KISTI
 - ◊ xrootd upgrade from v3 to v4.4.1 for tape
 - ◆ NL-T1:
 - ◊ SURFsara Major dCache upgrade to 2.16.36 on June 6-7
 - ◆ RAL:
 - ◊ Castor stagers updated to 2.1.16-13 and SRMs to 2.1.16-0.
 - ◊ All data now on T10KD drives/media.
 - ◊ Upgrade of FTS "prod" instance delayed due to non-LHC VOs usage of SOAP API. Hope to be able to upgrade during July
 - ◆ TRIUMF:
 - ◊ Major dCache upgrade to v2.16.39

Discussion

Tier 0 News

- Storage: see above
- Batch capacity increases ongoing

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- The activity levels typically have been very high
 - ◆ The average was 112k running jobs, with a new record of 143k on May 29
- CERN
 - ◆ Some fallout from the DNS incident on June 18
- No other major problem

ATLAS

- stable production at 300k cores, about 80k used for derivations.
- derivation production is causing too many transfers, need to further optimize the workflow (eg 70 outputs/ mcore job)
- ongoing ATLAS P1 to EOS to CASTOR data throughput test to fully validate (at approx double of the nominal rate) the data workflow from ATLAS experiment to the tape infrastructure.
- ongoing efforts to understand sites not performing very well (high - wrt to the average of the other sites- wallclock time wasted).

CMS

- CMS Detector
 - ◆ Commissioning progressing
 - ◆ Most effort goes into the new pixel detector
- Processing activities
 - ◆ Overall utilization rather moderate
 - ◆ Finished a RE-RECO of 2016 data
 - ◆ Main MC production campaign for 2017 still in preparation
 - ◆ Small (but urgent) RE-RECOs of recent 2017 data for commissioning
- Sites
 - ◆ Deprecation of stage-out plugins
 - ◆ In contact with sites to test IPv6 readiness of storages
- EOS
 - ◆ Suffered from limitations in GSI authentication capacity - fixed
 - ◆ Identified a source of occasional file corruptions: Improper handling of write recoveries
 - ◇ Can be circumvented by setting environment variable
 - ◇ Details: [GGUS:127993](#)
- EL7 migration
 - ◆ Found some issues with Singularity under some configuration circumstances
 - ◆ Recommendation is to wait with migration, if possible
- Rising interest in CMS to use MPI compute resources for certain generators
 - ◆ Sites, that want to provide such resources should contact Stephan Lammel and Giuseppe Bagliesi

LHCb

- High activity on the grid, keeping an average of 60K jobs
- CERN
 - ◆ The proxy expiration problem in HTCondor CEs still being investigated. (GGUS:129147)

Ongoing Task Forces and Working Groups

Accounting TF

- Progress on the storage space accounting prototype has been reported at the WLCG Workshop
- At the latest Accounting TF meeting in May discussed the plan to add raw wallclock job duration to the accounting portal as a separate metric. Currently wallclock time can contain either raw or scaled wallclock time. APEL colleagues presented EGI work regarding storage space accounting.

Information System Evolution TF

- The IS evolution plans and progress in CRIC development have been presented at the WLCG Workshop in Manchester

IPv6 Validation and Deployment TF

- Andrea S:
 - ◆ we will prepare a campaign for T2 sites to start looking into their IPv6 preparations
 - ◆ it will be started at a small scale, to gain experience before all sites are contacted
 - ◆ we probably need a GGUS support unit and a mailing list
 - ◆ the text sent to the sites needs to be very clear
 - ◆ we aim to have dual-stack deployment of storage services at the vast majority of sites by the end of Run 2
- Julia:
 - ◆ there should be a communication channel for sites to share experiences
 - ◆ a Twiki page would be helpful for recipes etc.
- Julia: did the IPv6 session at the workshop go OK?
- Andrea S:
 - ◆ there were ~30 people in the hands-on session
 - ◆ the exercises were easy and went well

Machine/Job Features TF

Current status

MJF hosts (all sites) total: 158

- Hosts OK: 25
- Hosts WARNING: 15
- Hosts CRITICAL: 112

The warnings/errors are of just a few types (configuration mistakes), and it looks like no much effort is required to correct them. Namely: WARNING

- Warning Key hs06 absent (or empty): 11
- Warning Key max_swap_bytes absent (or empty): 4

CRITICAL

- Error Environment variable MACHINEFEATURES not set: 98
- Error Environment variable JOBFEATURES not set: 2
- Error Key total_cpu absent (or empty): 10
- Error Key cpu_limit_secs absent (or empty): 2

Propagation of MJF to other experiments requires some amount of work. In particular, Antonio (aperez@picNOSPAMPLEASE.es) wrote about CMS:

- CMS SI have worked with glideinWMS (our pilot system) developers to incorporate the information published as MJF into our pilots (where available). So potentially we could add new features (such as job-masonry, but also signaling job/node shutdown times) when the rest of dependencies are solved. One of those dependencies of course will be the deployment of MJF to the CMS sites not shared with LHCb.

Monitoring

- the status and plans were presented during the WLCG Workshop in Manchester:
 - ◆ Status, plans of WLCG Unified Monitoring [↗](#)
 - ◆ ElasticSearch Service at CERN [↗](#)
 - ◆ Monitoring in Experiments (Summary) [↗](#)

MW Readiness WG

This is the status of jira ticket updates [↗](#) since the last Ops Coord of 20170518:

- MWREADY-146 [↗](#) - dCache 2.16.34 verification for ATLAS @ TRIUMF with IPV6 as well - completed (there has been a problem when TRIUMF updated the production unfortunately not spotted in the testing instance)
- MWREADY-145 [↗](#) - The latest version of the WN metapackage for C7 has been released (v 4.0.5 - renamed to wn), and tested by Liverpool. The metapackage is under inclusion in UMD4 (GGUS:128753 [↗](#))
- MWREADY-147 [↗](#) - ARC-CE 5.3.1 under testing at Brunel.
- MWREADY-148 [↗](#) - New CREAM-CE for C7: we agreed with M. Sgaravatto to do the testing for CMS at LNL.

Network and Transfer Metrics WG

- Detailed WG update presented as part of the network session at the WLCG workshop in Manchester [↗](#)
- perfSONAR 4.0 was released on 17th of April
 - ◆ 194 nodes updated so far
 - ◆ ES/Kibana dashboard showing perfSONAR infrastructure status in testing
- WLCG/OSG network services

- ◆ New central mesh configuration interface (MCA) in production (<http://meshconfig.grid.iu.edu>)
 - ◇ Accessible to mesh administrators only - please contact wlcg-perfsonar-support@cern.ch to request access
- ◆ New monitoring based on ETF in production (https://psetf.grid.iu.edu/etf/check_mk/)
- ◆ New OSG collector handling multiple backends (Datastore, CERN ActiveMQ and GOC RabbitMQ) in production
- ◆ New LHCOPN grafana dashboards done in collaboration with CERN IT/CS and IT/MONIT in testing
 - ◇ Now with open access at <http://monit-grafana-open.cern.ch/dashboard/db/lhcopn?orgId=16>
- ◆ Additional perfSONAR dashboards to be added soon
- Throughput call was held on Wed May 24th at 4pm CEST (<https://indico.cern.ch/event/640627/>) mainly focusing on review of new production services

Squid Monitoring and HTTP Proxy Discovery TFs

- CMS frontier at CERN is now using <http://grid-wpad/wpad.dat> with IPv6 in production. ATLAS frontier at CERN has been all this time randomly using squids at Geneva and Wigner, regardless of the location of the worker nodes, causing much traffic to go over the long distance links. They are now making plans to start using <http://grid-wpad/wpad.dat> to select local squids.

Traceability and Isolation WG

Special topics

MW deployment forums and feedback

[presentation](#)

- Gavin: we take HTCondor unchanged
- Maarten: but you enhance it e.g. with the BDII info provider; furthermore, the matter is not just about patches, but deployment in general
- Julia:
 - ◆ the `fts3-steering` list is a good example, though only involving VOs and devs
 - ◆ in general the fora would need to allow VOs, sites and devs to participate
 - ◆ feedback from sites should be collected and made easily available for others
 - ◇ deployment documentation, workarounds etc.
- Maarten:
 - ◆ the MW Readiness WG is the right place to have such things organized
 - ◆ in the Sep meeting we will have a checkpoint on the progress

Theme: Providing reliable storage - IN2P3

[presentation](#)

- Maarten: do you have some services permanently available on a UPS?
- IN2P3-CC:
 - ◆ the whole building is on a UPS with a minimal lifetime of about 30 minutes
 - ◆ its main function is to allow switching to the other power line transparently
 - ◆ if needed, to extend the lifetime we can start switching off all the WN etc.

- Julia: how often do you see file losses from tape?
- IN2P3-CC:
 - ◆ typically a few files per month
 - ◆ such incidents tend to get revealed during repack operations
- Xin: couldn't most such files be recovered by the vendor?
- IN2P3-CC:
 - ◆ we usually try other ways to recover the files first (other tapes or copy from another site)
 - ◆ even if the vendor manages to recover part of the data, the files typically are corrupted

- Vikas: what are your RAID group disk size and rebuild times?
- IN2P3-CC:
 - ◆ each disk is 6 to 8 TB, the next ones will be 10 TB
 - ◆ we have ~145 TB per server
 - ◆ the rebuild time is ~24h
 - ◆ we need to rebuild 1 or 2 times per year
- Vikas: 24h is a rather big window for another disk to fail as well...
- Maarten: various parameters need to be taken into account and optimized together; in the end there will always be a calculated risk...

- IN2P3-CC:
 - ◆ for the evolution of our tape system we see 2 options:
 - ◇ move to IBM Jaguar, which would imply replacing the whole library
 - ◇ move to LTO, which so far we have only used for backups in TSM
 - ◆ we would like to discuss such matters e.g. in HEPiX
 - ◆ get an idea on reliability experiences at other sites
- Alessandro:
 - ◆ we have the same matter to deal with at CNAF
 - ◆ we have had meetings with several vendors (IBM, Quantum, Spectra Logic)
 - ◆ we heard some sites are staying with T10KD for the time being
 - ◆ LTO may not be good enough for heavy stage-in and -out operations
 - ◆ we support the revival of the tape forum to discuss these things
- Julia:
 - ◆ we will first follow up with the owner of the existing list
 - ◆ we will ensure there will be a forum and announce it

Action list

Creation date	Description	Responsible	Status	Comments
01 Sep 2016	Collect plans from sites to move to EL7	WLCG Operations	Ongoing	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. Triumpf 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. Jan 26 update: this matter is tied to the EL7 validation statuses for ATLAS and CMS, which were reported in that meeting. March 2 update: the EMI WN and UI meta packages are planned for UMD 4.5 to be released

				in May May 18 update: UMD 4.5 has been delayed to June July 6 update: UMD 4.5 has been delayed to July
03 Nov 2016	Review VO ID Card documentation and make sure it is suitable for multicore	WLCG Operations	Pending	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
18 May 2017	Follow up on the ARC forum for WLCG site admins	WLCG Operations	In progress	
18 May 2017	Prepare discussion on the strategy for handling middleware patches	Andrea Manzi and WLCG operations	In progress	
06 Jul 2017	Ensure a forum exists for discussing tape matters	WLCG Operations	New	

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

This topic: LCG > WLCGOpsMinutes170706

Topic revision: r18 - 2018-02-28 - 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