

Table of Contents

WLCG Operations Coordination Minutes, October 22nd 2015.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Middleware News.....	1
Tier 0 News.....	2
Tier 1 Feedback.....	2
Tier 2 Feedback.....	2
Experiments Reports.....	2
ALICE.....	2
ATLAS.....	3
CMS.....	3
LHCb.....	3
Ongoing Task Forces and Working Groups.....	4
gLExec Deployment TF.....	4
HTTP Deployment TF.....	4
Information System Evolution.....	4
IPv6 Validation and Deployment TF.....	4
Middleware Readiness WG.....	4
Multicore Deployment.....	5
Network and Transfer Metrics WG.....	5
RFC proxies.....	5
Squid Monitoring and HTTP Proxy Discovery TFs.....	5
Action list.....	5
Specific actions for experiments.....	6
Specific actions for sites.....	6
AOB.....	6

WLCG Operations Coordination Minutes, October 22nd 2015

Highlights

- Sites using aliases for their Grid services are, once again, **strongly** encouraged to follow the advice given here [↗](#) to be ready for the (re)introduction of the STRICT_RFC2818 mechanism in Globus. Please contact your Certification Authorities as soon as possible. **Concerned Grid service instances may stop working if no action is taken.**
- The CERN HTCondor pilot service is expected to be declared production-ready on November 2nd.
- dpm-xrootd 3.5.5 is now in EPEL and DPM sites should upgrade to it as soon as possible.
- Sites should still refrain from updating openldap on their BDII's to any version newer than 2.4.39: a fix to the bug affecting newer versions will be available very soon.

Agenda

- <https://indico.cern.ch/event/393618/↗>

Attendance

- local: Andrea Sciabà (minutes), Stefano Perazzini (LHCb), Maarten Litmaath (ALICE), Andrea Manzi (MW Officer), Ulrich Schwickerath (T0), Marian Babik, Eva Dafonte Perez (DB), Alessandro Di Girolamo (ATLAS)
- remote: Pepe Flix (chair), Michael Ernst, Catherine Biscarat, Christoph Wissing (CMS), David Cameron (ATLAS), Di Qing, Dave Mason, Frederique Chollet, Renaud Vernet, Thomas Hartmann, Gareth Smith, Andrea Valassi (LHCb), Antonio Perez Calero Yzquierdo, Hung-Te Lee, Massimo Sgaravatto, Jeremy Coles

Operations News

Middleware News

- Useful Links:
 - ◆ Baseline Versions
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines:
 - ◆ dpm-xrootd 3.5.5 has been pushed to EPEL, fixes an issues affecting CMS AAA. In case the installation has been done via metapackages, the new ones for DPM 1.8.10 are under release via EMI
- Issues:
 - ◆ issue affecting BDII with the latest version of openldap (2.4.40-5/6) on SL6/CentOS6, which lead to slapd crash on TopBDII and ARC-CE resource BDII. Sites have been asked via broadcast message to not upgrade openldap (even if it contains a security fix, also SVG has suggested the same). We are waiting for a fix from RedHat to test, so until the fix is released we ask sites to stay with the version of slapd which is not crashing (2.4.39)
 - ◆ Various issues with FTS affecting RAL and CERN instances related to memory get exhausted have been reported by ATLAS and LHCb. Fixes have been developed and are under testing by the FTS team

- ◆ Regarding the problem reported during the last WLCG ops coord with FTS3 not able to transfer to IPV6 only storages, it has been analysed and the problem has been introduced by a new release of globus-ftp-client. The Issue has been reported to globus <https://globus.atlassian.net/browse/GT-604>
- T0 and T1 services
 - ◆ IN2P3
 - ◇ dCache upgraded to v 2.10.42

Maarten points out that the importance of IPv6 must not underestimated, because already next year there may be sites that have to use IPv6 for their services. It is clarified that the issue with IPv6 is the same initially reported in JIRA:DMC-681 [back](#) in June.

Maarten mentions that the table containing the version of each FTS instance is not accurate. Andrea M. and FTS3 support will ask sites to better maintain the information and find more reliable ways to make it available (e.g. via the FTS3 monitoring).

Tier 0 News

- We need to upgrade LSF because the current version (7.0.6) will no longer be supported by the end of the year. Our first attempt to upgrade the public LSF master nodes failed. The issues have been investigated in collaboration with the vendor (IBM), and have been understood. We will plan a new attempt in the coming weeks, preferably outside data taking. The exact date still needs to be decided on and will be communicated in due time before the intervention will start.
- Status of HTCondor: please see slides attached to the meeting agenda. Main points:
 - ◆ the HTCondor cluster contains 2500 cores, is fully used and accepts only grid jobs (no local submission before 2016)
 - ◆ it is accessible via both ARC-CEs and HTCondor-CEs; the former will be decommissioned at the end of this year
 - ◆ it will be put in production from November the 2nd, when also accounting will be enabled; after that, resources will be gradually moved from LSF to HTCondor and the CREAM CEs eventually retired

Ulrich confirms that at the moment the HTCondor cluster can run only single core jobs. In the near future they will start testing the various tools (DEFrag_DAEMON, DrainBoss, Farlow, ...) which are needed to manage node draining for multicore-enabled clusters.

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- generally high activity
 - ◆ 90k+ for a few hours on Oct 9
 - ◆ Reco, analysis, MC
- CASTOR setup for the heavy ion run
 - ◆ a double pool configuration has been set up
 - ◆ one "fast" pool for the DAQ plus writing to tape
 - ◆ one "slow" pool for T1 transfers, reco and recalls from tape
 - ◆ with CASTOR handling the internal replication.

- ◆ Many disk servers that were prepared for retirement have been added for the heavy ion run.
 - ◇ essentially allowing the whole run to stay on disk until early next year
- ◆ We thank the CASTOR team for their good support in this matter!

Maarten clarifies that the additional disk servers are of an older type (out of warranty and with 1 Gb/s connectivity), but they are good enough for the intended usage. If all goes as planned, we should find that they were not strictly necessary; their availability is certainly convenient and should give protection against unexpected issues in the workflows.

ATLAS

- Slightly less than usual slots filled due to job mixture/brokering/DDM issues
- Problems with FTS: submission hanging, duplicate transfers
- PRODDISK decommissioning going ahead full steam (merging I/O buffer with persistent data token).
In the next week or two we will start telling sites to move the space.
- T0 spillover in final stages of testing

CMS

- Good utilization over the last three weeks
 - ◆ Reached again ~120k sustained over days
- Ticketing Campaign for sites to prepare for update of secrets for Phedex
 - ◆ We do it every few years
 - ◆ Beginning of the week 67 out of 108 contacted sites answered
 - ◆ **Deadline is Nov 3rd**
- Operational issues
 - ◆ GGUS outage on Monday
 - ◆ SAM test submission at CERN
 - ◇ Some Argus problems affected SAM tests in particular GGUS:117001 [↗](#)
 - ◇ CMS reduced pilot life time to 3 days (from 7) to ease scheduling

Maarten explains that the GGUS outage was due to a power cable being cut during works at KIT. It was fixed in two hours and had no impact on LHCOPN or LHCONE.

He also explains that the recurrent failures of the CMS pilot role SAM tests were mainly due to the number of pilot pool accounts being too low, causing constant recycling of account mappings (while normally it should be very rare). That situation led to an increased exposure to a race condition in Argus when a new mapping is added; ticket GGUS:117125 [↗](#) has been opened for the Argus devs to look into that matter. There are still other issues surrounding Argus to be understood, though.

LHCb

- Operations
 - ◆ Run 2 data processing is going on with no major issues. During this week data processing activities have been extended also to some T2 (mesh processing) in order to speed up the processing and reduce the number of jobs to be executed at Tier1s and CERN.
 - ◆ Very high activity, LHCb using all available distributed compute resources available to the VO (130,000 jobs waiting in the central task queue, 40,000 running jobs steadily of which 18,000 at Tier1s)
- Issues
 - ◆ File export from pit was affected by the transfer of large number of small files. Now a merging procedure before transfer has been put in place and is currently used without obvious problems.

- ◆ Last week observed several problems with data transfer to CERN EOS. An alarm ticket has needed to be opened (GGUS:116973). It was quickly fixed by a restart of SRM (bestman2). Now issue seems to be solved but a couple of bugs have been found. These bugs affect BESTMAN and FTS, and long term fix will need to wait for the next release.
- ◆ Issue with slow CERN/LSF worker nodes has been looked into. Fix will be provided in November. Need both Kernel and OpenStack changes.

Alessandro asks if LHCb foresees to directly use the EOS GridFTP service instead of SRM for transfers; Stefano will find out. Also CMS is using GridFTP, but SRM is still used by everybody at least for metadata operations or file deletions.

Ongoing Task Forces and Working Groups

gLExec Deployment TF

- NTR

HTTP Deployment TF

Information System Evolution

IPv6 Validation and Deployment TF

Middleware Readiness WG

- This JIRA dashboard shows per experiment and per site the product versions pending for Readiness verification. Details:
 - ◆ JIRA:MWREADY-84 CMS: at Brunel ARC-CE 5.0.3 test stalled - no jobs displayed.
 - ◆ JIRA:MWREADY-85 ATLAS: at NDGF Rucio tests fail since ipv6 has been enabled. so dCache verifications are pending at NDGF
 - ◆ JIRA:MWREADY-90 ATLAS: at TRIUMF dCache 2.10.42 verification completed.
 - ◆ JIRA:MWREADY-91 CMS: at PIC dCache 2.13.9 verification ongoing.
 - ◆ JIRA:MWREADY-87 ATLAS: at Edinburgh for DPM 1.8.10 verification - site manager silent.
 - ◆ JIRA:MWREADY-82 ATLAS: at GLASGOW DPM 1.8.10 CentOS7 verification - missing tests' set-up.
 - ◆ JIRA:MWREADY-81 CMS: at CERN EOS 0.3.129-aquamarine installed on the PPS - verification pending.
- Issue found:
 - ◆ JIRA:MWREADY-89. When moving transfer tests to the FTS CERN Pilot (with a different gfal2 configuration w.r.t FTS@RAL) we discovered an issue with DPM deployed with gridftp redirection. Reported to the devs
- Reminder: Next meeting **October 28th at 4pm CET**. Agenda http://indico.cern.ch/e/MW-Readiness_13

Multicore Deployment

Network and Transfer Metrics WG

- perfSONAR collector, datastore, publisher and dashboard now in production !
- psmad becomes the official dashboard for perfSONAR meshes
- perfSONAR 3.5: 183 sonars were updated, ALL sites are encouraged to enable auto-updates for perfSONAR.
- Detailed report from the WG presented at HEPiX/GDB [↗](#), we will also present status update again at the November's GDB
- ATLAS started processing perfSONAR stream to create a network cost-matrix for use by PANDA with additional use cases in scheduled transfers and dynamic data access
- LHCb also started processing perfSONAR stream and correlates it with the network and transfer metrics in DIRAC
- Next WG meetings will be on 4th of Nov and 2nd of Dec

Alessandro proposes to discuss and agree on a naming convention for the perfSONAR endpoints to be able to produce a consistent picture across VOs. Marian points out that it is possible to define associations on top of the core metrics that use a well defined perfSONAR topology.

RFC proxies

- NTR

Squid Monitoring and HTTP Proxy Discovery TFs

- One minor Squid Monitoring task on the list was completed, number 4. Dates for other unfinished tasks were pushed forward.
- There was keen interest at the CMS Offline & Computing week for the Proxy Discovery deliverables in order to support opportunistic computing. They are blocked by the Squid Monitoring task number 6 but there's some hope that this renewed interest will help progress.

Action list

Creation date	Description	Responsible	Status	Comments
2015-06-04	Status of fix for Globus library (globus-gssapi-gsi-11.16-1) released in EPEL testing	Andrea Manzi	ONGOING	GGUS:114076 ↗ is now closed. However, host certificates need to be fixed for any service in WLCG that does not yet work OK with the new algorithm. Otherwise we will get hit early next year when the change finally comes in Globus 6.1. A broadcast message has been sent by EGI. A first version of the monitoring script has been completed to test SRM instances. 10% of the WLCG SRM instances have been found with incorrect certificates, 10 GGUS tickets opened, CERN CASTOR already fixed.

				JIRA:MWREADY-86 ↗
2015-10-01	Follow up on reporting of number of processors with PBS	John Gordon	ONGOING	
2015-10-01	Define procedure to avoid concurrent OUTAGE downtimes at Tier-1 sites	SCOD team	ONGOING	A Google calendar is not yet available, therefore the only way for the moment is to check GOCDB and OIM for downtimes using the links that will be provided for convenience in the minutes of the 3 o'clock operations meeting

Additional points discussed about the Globus algorithm:

- GFAL2 is not affected by it, as it uses its own logic, but other clients are
- Christoph asked help to test CMS services; to be followed up offline
- It must be stressed that it is essential that Certification Authorities are able to produce service/host certificates with aliases. Site managers are encouraged to put some pressure on their CAs, if needed

Specific actions for experiments

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

Specific actions for sites

Creation date	Description	Affected VO	Affected TF	Comments	Deadline	Completion
2015-09-03	T2s are requested to change analysis share from 50% to 25% since ATLAS runs centralised derivation production for analysis	ATLAS	-	-	a.s.a.p.	ONGOING
2015-06-18	CMS asks all T1 and T2 sites to provide Space Monitoring information to complete the picture on space usage at the sites. Please have a look at the general description and also some instructions for site admins.	CMS	-		None yet	~10 T2 sites missing, Ticket open

AOB

-- AndreaSciaba - 2015-10-20

This topic: LCG > WLCGOpsMinutes151022

Topic revision: r14 - 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