

Table of Contents

WLCG Operations Coordination Minutes - April 2nd, 2015.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Data Preservation Training.....	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.....	3
gLExec Deployment TF.....	3
SHA-2.....	3
Machine/Job Features.....	4
Middleware Readiness WG.....	4
Multicore Deployment.....	4
IPv6 Validation and Deployment TF.....	4
Squid Monitoring and HTTP Proxy Discovery TFs.....	5
Network and Transfer Metrics WG.....	5
HTTP Deployment TF.....	5
Action list.....	5
AOB.....	6

WLCG Operations Coordination Minutes - April 2nd, 2015

Agenda

- <https://indico.cern.ch/event/380917/>

Attendance

- local: Maria Dimou (minutes), Andrea Sciabà (chairperson), Maarten Litmaath (ALICE), Marian Babik (Network TF), Maite Barroso (Tier0), Alessandro Di Girolamo (ATLAS), Prasanth Kothuri (IT-DB).
- remote: Alexandra Berezhnaya (NRI-KI), Frederique Chollet & Catherine Biscarat (France), Thomas Hartmann (KIT), Hung-Te Lee (ASGC), Renaud Vernet (IN2P3), Ulf Tigerstedt (NDGF), Christoph Wissing (CMS), Jeremy Coles (GridPP), Di Qing (Triumpf), Michael Ernst (BNL), Rob Quick (OSG), Ron Trompert (NL_T1), John Kelly (RAL), Alessandro Cavalli (CNAF).

Operations News

- Analysed suggestions received from the sites during the WLCG survey, which proved to be extremely useful. They have been summarised (<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGSiteSurvey>).

Data Preservation Training

Andrea S. walked through these slides [prepared by J.Shiers](#), including encouragement by I.Bird for T1 managers to join a free-of-charge training course. It will have to be cancelled if not more people join. Anyone can register but we will need to take a decision soon whether the course actually takes place. The idea was to fund it centrally but this might not fly if a sufficient number of WLCG "curation centres" aren't there. People interested to participate should register by April 17 in the agenda page (<https://indico.cern.ch/event/376809/>).

Middleware News

- Useful Links:
 - ◆ Baseline Versions
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines:
 - ◆ Some sites complained about the missing ARGUS PAP 1.6.2 in UMD (it fixes a blocking issue triggered by the latest JAVA upgrade)
 - ◇ a new UMD release is under preparation to include this version, but in the meantime it can be taken from the EMI repo
- MW Issues:
 - ◆ NTR
- T0 and T1 services
 - ◆ CERN
 - ◇ CASTORCMS planned to be upgraded to 2.1.15-3 on 07-Apr-2015
 - ◇ CASTORATLAS planned to be upgraded to 2.1.15-3 on 08-Apr-2015

Alessandro said that the 08-Apr-2015 may be not good for ATLAS but they will take this offline with the CASTOR team.

- ◆ CNAF
 - ◇ All STORM Instances Updated to the last released version (1.11.8) (both frontend and backend).
 - ◇ LHCB instances moved to a single machine (be+fe) deployment on a virtualized environment.
 - ◇ planning to Move to the new virtualized environment also the CMS and ATLAS instances
- ◆ NL-T1
 - ◇ Minor dCache upgrade from 2.10.20 to 2.10.23
- ◆ JINR-T1
 - ◇ Minor dCache upgrade from 2.10.22 to 2.10.23
- ◆ RAL
 - ◇ CASTOR planned to be upgraded to 2.1.14-15 on 08-Apr-2015

Tier 0 News

NTR

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- high activity
- most WLCG VOBOXes at the sites have been switched to RFC proxies
 - ◆ some VOBOXes still need to be upgraded to EMI-3 / UMD-3 first
- CASTOR at CERN: instabilities in the last 2 weeks, affecting both online and offline transfers
 - ◆ Oracle row lock contention due to concurrent activities
 - ◇ bunch size of staging requests has been reduced as a mitigation
 - ◆ disk pool was not being rebalanced, which led to unnecessary garbage collection (fixed)
 - ◆ garbage collection removed newly staged files because they were not (yet) accessed recently!
 - ◇ will be improved
 - ◆ one disk server with a lot of staged files went offline unnoticed (rebooted OK)
 - ◆ others suffered network issues
 - ◆ not enough slots per disk server (fixed)
 - ◆ thanks for the prompt support in these matters!
- team ticket GGUS:112716 [↗](#) for VOMS service managers:
 - ◆ Sao Paulo admin certificate being rejected as expired, while the proxy worked OK in other usage
 - ◆ the cause turned out to be an expired **CRL** of the Brazilian CA!
 - ◆ that in turn was caused by the CRL URL being firewalled for **IPv6**
 - ◇ this is at least the 4th such incident with a CA...
 - ◆ GGUS:112774 [↗](#) opened for the VOMS devs to improve the error message...

ATLAS

- activities:
 - ◆ Cosmics data export going smoothly
 - ◆ ready for the LHC beam splashes
 - ◆ MC15 simulation: quite a lot of submissions, the Grid should be completely full
- we are aware of another network incident (degradation) between Triumf and RAL (50KB/s average transfer rate): any more info from the networking team?
- WhiteHat challenge: interesting and useful but agreed with SL that communication to experiment central services should have happened in advance.

Christoph asked what is a WhiteHat challenge. Here is their web page[?]. It contains an MoU and a *Code of Ethics* but it doesn't seem to require an explicit **a priori** information to the experiment about what will happen when. This is what the experiment requires, so that they can explain the sudden heavy load on their services.

CMS

- Cosmic run with magnetic field on (basically) finished
- Main production activities
 - ◆ Tier-1: DIGI-RECO of upgrade MC (~all resources busy)
 - ◆ Tier-2: MC production
- Problems with Squid/Frontier infrastructure
 - ◆ Still not fully understood
 - ◆ Can be just too many (short) jobs starting at time
 - ◆ Can be demanding jobs regarding Squid access
- Scheduling of SAM tests
 - ◆ Tests are competing with other jobs and potentially don't run in time
 - ◆ Investigation to "OR" the metric with e.g. HammerCloud results
- Tape staging exercises
 - ◆ Concluded successfully at T1 sites
 - ◆ A test at CERN is still to come

LHCb

Ongoing Task Forces and Working Groups

gLExec Deployment TF

- gLExec in PanDA:
 - ◆ testing campaign is covering 61 sites out of 94 (one mistakenly included site was removed)
 - ◆ the pilot will be enhanced with yet more debugging to help understand what is going wrong at RAL and RALPP
 - ◇ while it is working OK basically everywhere else!
 - ◆ the work will be presented on a poster at CHEP

SHA-2

- retirement plans for the old VOMS servers
 - ◆ The only reason for keeping the old VOMS server aliases was to allow the old VOMRS registration URLs to forward users to the new VOMS-Admin services:
 - ◇ <https://lcg-voms.cern.ch:8443/vo/alice/vomrs/>[?]
 - ◇ <https://lcg-voms.cern.ch:8443/vo/atlas/vomrs/>[?]
 - ◇ <https://lcg-voms.cern.ch:8443/vo/cms/vomrs/>[?]

◇ <https://lcg-voms.cern.ch:8443/vo/lhcb/vomrs/>

- ◆ While the `lcg-voms.cern.ch` alias still exists (and `voms.cern.ch` too), we need to have special exceptions in the CERN routers to prevent that remote clients might hang if they still try to access the VOMS daemons that used to run on those hosts.
- ◆ Of course, nothing *should* be doing that anymore, but you can be sure there still are *plenty* of places where the old VOMS endpoints are configured along with the new ones. In such cases we want clients to fail over quickly, instead of hanging and timing out.
- ◆ We now have made the transition from VOMRS to VOMS-Admin, and although everything is not perfect yet, we will not be needing to go back.
- ◆ Therefore we finally would like to remove those old aliases and get rid of those routing exceptions.
- ◆ Can we do this "now", or do you see some issue at this time?
- ◆ Meanwhile the special rules have been extended until Tue May 5.
- ◆ Proposal: remove the aliases Mon Apr 27.
 - ◇ nothing changes until after CHEP
 - ◇ allow 1 week to "recover" from CHEP
 - ◇ allow 1 week to deal with the fallout

Machine/Job Features

Middleware Readiness WG

Excellent progress made by the Volunteer sites for selected MW product versions, as planned at our last meeting on March 18th. Thanks to the MW Officer Andrea Manzi for testing with the sites and to the pakiti client developer Lionel Cons for the good technical collaboration with the other EGI-funded experts. Details:

- dCache problem found in v.2.10.18 is solved in v.2.10.23. dCache versions 2.11.14 and 2.12.3 also contain the fix. Triumf is now testing dCache 2.10.23 while NDGF 2.12.3. (both for ATLAS).
- StoRM 1.11.8 has been verified at QMUL for ATLAS.(only one small issue found)
- CREAM-CE 1.16.5 is tested at INFN-Napoli for ATLAS.
- EOS is in the pipeline for testing at CERN for CMS. Please add news in JIRA:MWREADY-40.
- Please do the Actions from our March 18th meeting.
- Remember **Wednesday May 6th at 4pm CEST** is the next MW Readiness WG meeting date.

Jeremy reported on-going effort within the DPM Collaboration on how to continue DPM Readiness verification, previously done by Wahid/Edinburgh.

Multicore Deployment

IPv6 Validation and Deployment TF

- FTS3 testbed operational, with servers at KIT and Imperial College both working fine
- The following sites activated IPv6:
 - ◆ LHCOPN: CERN, KIT, NDGF, PIC, NL-T1, IN2P3-CC, HIP
 - ◆ LHCONE: CERN, CEA Saclay, IN2P3 -CC, IJS (NDGF site)
- OSG is testing (among other middleware) glideinWMS. The central manager, frontend and schedd machines have to be dual-stack and can talk to IPv4, IPv6 and dual-stack startd's. glideinWMS must specify to wget that it prefers IPv6 (details)
- OSG confirmed that Bestman2 is IPv6-compliant, but srmcp is not (it has not been patched for the extensions needed for IPv6)
- squid 2 is not IPv6-compliant, while squid 3 is. OSG is still using squid 2
- Duncan's dual stack mesh includes several dual-stack perfSONAR instances (~14 sites included) (link)

Squid Monitoring and HTTP Proxy Discovery TFs

- Nothing to report

Network and Transfer Metrics WG

- perfSONAR status
 - ◆ Security: CVE released today for cassandra, which is used by the perfSONAR measurement archive software, esmond. NO action required to protect perfSONAR Toolkit since vulnerable ports are both disabled and firewalled.
 - ◆ perfSONAR 3.4.2 was released and auto-deployed to 163 sonars, there are 42 instances still on 3.4.1. We no longer have any active instances on older versions.
 - ◆ We encourage ALL sites that are still on 3.4.1 to check status of their sonars (mainly disk space) and enable auto updates ASAP.
 - ◆ Significant improvement observed in getting consistently all the needed metrics after this update. The plan is to resume validation in LHCOPN/LHCONE and continue with a ramp up to full mesh latency tests.
 - ◆ Full mesh trace paths now at 80%
- Network performance incidents follow up (proposal):
 - ◆ New mailing list and GGUS SU will be established to follow up, proposed name is wlcg-network-throughput, initial participation will be the same as for the WG mailing list (transfer systems, experiments, perfsonar support, esnet, lhcopn/lhcone).
 - ◆ Experiments can report to the GGUS SU/ mailing list potential network performance incidents/degradations, WLCG perfSONAR support unit will investigate and confirm if this is network related issue. Once confirmed, it will notify relevant sites and will try to assist in narrowing down the problem to particular link(s). Affected sites will be contacted and should open an incident with their network providers. Tracking of the ongoing incidents will be done on the WG page (https://twiki.cern.ch/twiki/bin/view/LCG/NetworkTransferMetrics#Network_Performance_Incidents)
 - ◆ Sites observing a network performance problem should follow their standard procedure, i.e. report to their network team and if necessary escalate to their network provider while informing the wlcg-network-throughput mailing list. If confirmed to be WAN related, WLCG perfSONAR support unit can assist in further debugging of the problem. For the non-technical (policy) issues or if unclear, sites should escalate to the WLCG operations coordination.
- Next WG meeting: 8th of April (<https://indico.cern.ch/event/382622/>)

HTTP Deployment TF

The HTTP TF has been assembled, with e-group and twiki space:

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

The first meeting will be at 16hr on Wed 29th April.

Action list

- The network incident (degradation) between Triumpf and RAL reported by ATLAS will be a case to test the procedure put in place by the network metrics WG.
- CMS instructions to shifters to be changed so that tickets are not opened if just one CE is red.
- Maarten to follow-up with the experiments that the dates for removing the VOMS servers' aliases, as reported in the SHA-2 TF section above, are kept.

AOB

- The next meeting is on April 23rd.

-- AndreaSciaba - 2015-03-31

This topic: LCG > WLCGOpsMinutes150402

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](#)