



WLCG Service Report

4 weeks, Jan 19 – Feb 16

WLCG Management Board, 16th Feb 2016

Maarten Litmaath
CERN

v2.0

Motivation for changes in the WLCG operations meetings

- WLCG Infrastructure more stable and mature
- Well established operations procedures inside the experiments where many of operations tasks are performed
- Very positive experience with task forces and working groups focused on particular operational areas
- The WLCG operations coordination team reviewed current operations meetings and proposes changes to improve the effectiveness of these meetings

Proposed changes in frequency

- Reduce the frequency of the meetings
- Ops (3pm) meetings
 - 2 → 1 time per week
- WLCG operations coordination
 - biweekly → monthly

Proposed changes in content

- Most routine operations tracking should happen in 3pm meetings
- **WLCG operations coordination meetings become topical**
 - *proposals for new operations related task forces or closing existing task forces are reviewed*
 - *coordination of the operational work related to the integration of new technologies is performed*
 - *implications of important management decisions for operations are discussed*
 - *common topics of interest for WLCG, EGI, OSG operations are discussed*
 - *long standing operations issues and progress of resolving these issues are reviewed*
 - *interesting operational experiences from sites are presented, etc...*
- We keep the current report structure on Twiki and will ask for written reports to be provided as they are now
 - *But not necessarily go through all of them during the meeting*

Proposed changes in attendance

- Involvement and attendance
 - Taking into account the fact that now a lot of the WLCG-related work is spread along different groups in the department and lot of operational effort is provided in the experiments ...
 - ... we aim to involve more experts from the department outside the WLCG operations team and from the experiments, and will try to extend the attendance from the site representatives and our EGI and OSG peers

- We propose to introduce the changes as of the 1st of March

Operations Coordination highlights (1)

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

- SAM-Nagios hosts on SL5 are to be decommissioned in a few weeks
 - ETF-Nagios hosts on SL6 will take over
 - Currently being validated by the experiments
 - RFC proxies are used instead of legacy proxies
 - The probes and the SAM-3 A/R engine essentially remain the same
- WLCG now is prepared for the Globus GSI change in the service certificate validation algorithm
 - All identified affected services now have compliant certificates

Operations Coordination highlights (2)

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

- First **UMD-4** release for **CentOS/EL 7** available since Jan 15
 - <http://repository.egi.eu/2016/01/15/release-umd-4-0-0/>
 - BDII site and top, dCache, FTS-3, ARC CE
 - BDII and FTS-3 validated in MW Readiness WG
- Batch system **memory limits handling** for jobs
 - <https://twiki.cern.ch/twiki/bin/view/LCG/BatchSystemsConfig>
 - Current practices are described for CERN and all Tier-1
 - Recipes are provided for a variety of batch systems

Selected items from Operations (1)

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

- High activity for all experiments
- ALICE
 - EOS-ALICE instabilities due to being ~full
 - Ad-hoc cleanups done
- ATLAS
 - Jamboree Jan 27-29
 - <https://indico.cern.ch/event/440821/>
 - CREAM CE instabilities at various sites due to HTCondor bug in job lease code ([GGUS:119258](#))
 - Mitigated while being fixed by the developers
- LHCb
 - MD5 proxies vs. the latest Java updates: see page 12

Selected items from Operations (2)

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

- CMS
 - **Alarm** [GGUS:118938](#) because of **AFS token failures**
 - **Alarm** [GGUS:119030](#) for **voms-proxy-init failures** on Ixplus
 - Due to upgrade from OpenJDK 7 to Oracle JDK 8
 - Fixed by changing to OpenJDK 8 instead
 - GlideinWMS factories at CERN and FNAL could not create any new CMS Tier-0 VM ([GGUS:119312](#))
 - Caused by stuck condor_gridmanager daemons in the GlideinWMS hosts
 - CERN: physical machine emergency request procedure under discussion
 - CMS xrootd SAM test mandatory as of Feb 8

Selected items from Operations (3)

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

- CERN
 - Slow WN remedy has made **good progress**
 - Most WN hypervisors have been upgraded to **OpenStack Kilo**
 - The remainder will be done by the end of Feb
 - **EPT** (Extended Page Tables) with **2 MB pages**
 - **NUMA + pinning**, except on older HW that does not support it
 - The virtualization overhead is **less than 5%**
 - **AFS token outage** Jan 18-19
 - Unexpected side effect of a LAN DB cleanup
 - It mainly affected users
 - Grid jobs were not affected
 - Experiments are advised to identify non-trivial AFS dependencies
 - Try to remove them or discuss such cases with the AFS service managers

Selected items from Operations (4)

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

- NLT1
 - Reminder of SARA equipment move to **new data center** in Oct
 - 2 weeks outage expected
 - dCache space manager timeouts finally understood and **fixed**
 - Reduction of deadlock timeout is needed in PostgreSQL < 9.3
 - The back-off algorithm in dCache will be **enhanced** with some randomization

Java update incident (1)

- Latest Java 6/7/8 updates **disable MD5 by default**
 - Can still be re-enabled in `/usr/lib/jvm/jre/lib/security/java.security`
 - Many services in WLCG are based on Java
- LHCb were the first to run into **SRM errors** at a few sites
 - **Alarm** [GGUS:119233](#) for EOS, [GGUS:119294](#) (IC), SARA
 - Their oldest long-lived proxy delegations were still using MD5
 - **Fixed** by renewing them through the current version of their client
- RWTH-Aachen reported **CMS activities were also affected**
 - SAM tests
 - Pilots from at least one GlideinWMS factory
 - Some fraction of PhEDEx transfers

Java update incident (2)

- CMS SAM tests for WN are submitted through a version of **HTCondor < 8.3.2**
 - It contains a very old CREAM client that uses **MD5** in proxy delegations
 - The old SAM-Nagios hosts on SL5 are to be decommissioned soon
 - The new **ETF-Nagios** hosts on SL6 are OK with **HTCondor >= 8.4.4**
- ATLAS WN tests are submitted similarly, but none are affected
- SAM instances for ALICE and LHCb do not use HTCondor
- WLCG broadcasts urged admins to **handle Java updates carefully**
 - Another will be sent when SAM is safe again

Service Incident Reports

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

- **PIC tape incident**

- Since 2015 Q3, T10KD technology is deployed at PIC and ATLAS were enabled to write on it.
- On Dec 17th PIC detected a problematic file to be read from T10KD tape (internal migration) and decided to stop ATLAS writing to T10KD.
- The 80 used T10KD cartridges then were read completely – 8 tapes presented problematic files to be read.
- Data was migrated from these problematic tapes, by using dedicated firmware from Oracle for data recovery.
- In the end only 20 files could not (yet) be recovered, only half were unique copies, none of them critical (confirmed by ATLAS).
- **A lot of time and effort** have been spent to understand the cause of this problem. Tapes have been sent to Oracle US - it might be that we have a problematic drive, it's under investigation.
- A SIR will be provided later.

GGUS summary (4 weeks 18/1/2016 – 14/2/2016)

by Maria Dimou

VO	User	Team	Alarm	Total
ALICE	2	1	0	3
ATLAS	14	47	0	61
CMS	161	2	2	165
LHCb	1	29	1	31
Totals	178	79	3	260

