

Table of Contents

WLCG Operations Coordination Minutes, July 30th 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.....	3
Experiments Reports.....	3
ALICE.....	3
ATLAS.....	3
CMS.....	3
LHCb.....	4
Ongoing Task Forces and Working Groups.....	4
gLExec Deployment TF.....	4
RFC proxies.....	4
Machine/Job Features.....	4
Middleware Readiness WG.....	5
Multicore Deployment.....	5
IPv6 Validation and Deployment TF.....	5
Squid Monitoring and HTTP Proxy Discovery TFs.....	5
Network and Transfer Metrics WG.....	5
HTTP Deployment TF.....	6
Information System Evolution.....	6
Action list.....	6
Specific actions for experiments.....	6
Specific actions for sites.....	7
AOB.....	8
GGUS: How do users (e.g. VO shifters) receive GGUS downtime notifications?.....	8

WLCG Operations Coordination Minutes, July 30th 2015

Highlights

- The new WLCG Operations Portal [is](#) now online.
- Critical vulnerability [affecting](#) RedHat 5, 6 and 7 allows local root exploit. A fix for the affected package is available for Red Hat 7. Workarounds for Red Hat 6 and RedHat 5 described at Redhat documentation [When the patch for Red Hat 6 will be available sites should upgrade their installations, in particular CSIRT is going to check the upgrade testing the Worker Nodes.](#)
- perfSONAR results can be now published via message bus from the OSG collector. Discussing how to run this as a production service.
- Updated FTS performance study [Updated](#)
- Collection of use cases for the WLCG Information System on going.
- Experiments are providing a detailed description of their user support internal organisation to the T0 to better assign GGUS tickets to the correct support units
- Experiments and GGUS team are discussing how to better notify unscheduled downtimes of the GGUS tool

Agenda

- <https://indico.cern.ch/event/393612/>

Attendance

- local: Maria Alandes (Minutes), Maria Dimou, Maarten Litmaath, Andrea Sciabà, Andrea Manzi, Andrea Valassi, Marian Babik, Xavier Espinal, Alessandro Di Girolamo, Maite Barroso
- remote: Alessandra Forti (Chair), Antonio Maria Perez Calero Yzquierdo, Catherine Biscarat, Christoph Wissing, Gareth Smith, Guenter Grein, Felix Lee, Kyle Gross, Massimo Sgaravatto, Renaud Vernet, Thomas Hartmann, Alessandro Cavalli

Operations News

- Following yesterday's GGUS Release, the Did you know? [article](#) this month reminds of existing GGUS features requested at the WLCG Site Survey last autumn, as decided in this meeting a month ago.
- The new WLCG Operations Portal is now online <http://wlcg-ops.web.cern.ch/>. The portal aims at sharing information related to WLCG Operations with Sys admins, Experiments and TF/WG people. For example, Sys admins will be able to find a summary of action items (like the ones presented in the Ops meeting) and useful documentation for the services they have to maintain. Please, check the portal and do not hesitate to send us your feedback! We plan to add some dynamic content with articles covering topics of interest in our community. Stay tuned and help us making the portal useful to everyone and up to date!

Middleware News

- Useful Links:
 - ◆ Baseline Versions

- ◆ MW Issues
- ◆ Storage Deployment
- Baselines:
 - ◆ The end of support for dCache 2.6.x was May 2015. The deadline for decommissioning is 21/09/2015 and starting from 31/08/2015 sites still running dCache 2.6.x will be ticketed. (more details at https://wiki.egi.eu/wiki/Software_Calendars#dCache_v._2.6.x). ~20 instances still running 2.6.x (no T1s).
- Issues
 - ◆ Critical vulnerability affecting RedHat 5, 6, 7 broadcasted by EGI CSIRT (https://wiki.egi.eu/wiki/EGI_CSIRT:Alerts/libuser-2015-07-24) which allows local root exploit. This vulnerability is present in the case of access via a local user account, so only UIs where access is given via local passwd file could be affected. A fix for the affected package (libuser) is available for Red Hat 7, but not for Red Hat 6 yet and it will not be provided for RedHat 5 so sites should apply the workarounds described at <https://access.redhat.com/articles/1537873>.

Alessandra comments that the only way the EGI CSIRTS has to check that sites are patched is to send a probe to the WNs and check the RPM version so it would be better for sites to upgrade all the services even if their configuration means they are not vulnerable just for good practice. In particular the WNs to minimize the need of communication with the EGI CSIRTS to explain the site status. It is only 1 rpm and doesn't have dependencies. Checks will not start until the rpm is in the SL6 repository.

- T0 and T1 services
 - ◆ JINR
 - ◇ dCache upgraded to 2.10.36
 - ◆ NL-T1
 - ◇ DPM upgraded to v 1.8.9 at NIKHEF in order to fix a data transfer issue.
 - ◆ PIC
 - ◇ dCache upgraded to 2.10.37

Tier 0 News

- experiment support tickets assigned by mistake to the T0 this week:
 - ◆ CMS, change T3_IT_BolognaCloud image id in the ITB factory configuration: unknown exception occured: https://ggus.eu/ws/ticket_info.php?ticket=115298
 - ◆ Atlas, ATLAS-CERN-Frontier degraded more than 17 hours: https://ggus.eu/ws/ticket_info.php?ticket=115302

Maite informs that the upgrade due to vulnerability is expected to be finished by end of the week.

Maite reminds that a description of the user support structure in the experiments is needed so that CERN understands where to assign tickets that are not relevant for them. See Action Item on this issue at the end of this twiki for more details on the status of this action.

Maria Dimou reminds that when the TPM assigns a ticket to the wrong support unit, anyone with support status can re-assign the ticket to the correct support unit. Maite explains that this is what CERN is trying to do but that the problem is that it is not clear to CERN which support unit is the correct one for each experiment.

Tier 1 Feedback

None

Tier 2 Feedback

None

Experiments Reports

ALICE

- high activity
 - ◆ new record 83k briefly reached on Jul 29
- CERN:
 - ◆ raw data copies from Point 2 to CASTOR were timing out (GGUS:115145 [↗](#))
 - ◇ raw data reconstruction jobs were keeping many disk server slots busy
 - ◇ many more slots appeared to be occupied by stale transfers
 - ◇ to be followed up further with the CASTOR devs
 - ◆ job submissions became really slow multiple times (GGUS:115153 [↗](#) and GGUS:115238 [↗](#))
 - ◇ some issues were cured on the Argus side
 - ◇ the real cause of such problems has not yet been identified
- NDGF reported inefficient data transfers and noise in their logs
 - ◆ due to failed attempts with 2 methods before the 3rd succeeds
 - ◆ Xrootd client only checks if the *source* supports 3rd party copies
 - ◇ also the *destination* should be checked
 - ◇ a bug [↗](#) has been opened for the Xrootd devs
 - ◆ meanwhile a workaround has been applied on the ALICE side

Xavier confirms that the Castor issues mentioned by Maarten are indeed being discussed with developers. Xavier adds that mixed offline and online activities are not that usual and that possible mitigations need to be explored.

ATLAS

- Activity as usual, no major issues.
- working on the ATLAS Tier-0 dedicated cluster to understand "slow" nodes (not nodes with 10% less performances but more going half or one third of the others). Now setup HammerCloud stress test with single core analysis jobs, not able to fully saturate the cluster (now 12k slots out of 14.5k), investigating with experts.
- we are reviewing the ATLAS Central Service monitoring: procedures (and twiki) are being setup
- minor: GGUS unscheduled downtime. It was published in GOCDB, but we didn't know. We suggest an email to atlas-adc-crc at cern.ch is sent out next time.
- minor: kibana meter.cern.ch was down. issue announced in IT Status Board. Discussing with Pedro (itmon team) that ATLAS would like to send a GGUS team ticket for such issues, he agreed.
- in preparation the draft of KB article for the issue of users contacting CERN for ATLAS issues

CMS

- Main production activities
 - ◆ PromptRECO (Tier-0): No major infrastructure problems
 - ◆ DIGI-RECO for Run2 and Upgrade: Using all T1s and round 15 T2s
 - ◆ Continues GEN-SIM production
- Assignment of custodial location of Primary Datasets to T1 sites
 - ◆ One tape copy always at CERN, 2nd tape copy on other T1
 - ◆ First 50ns data went all to CERN and FNAL
 - ◆ Forthcoming data distribution to all T1 sites being iterated

- Operational Issues
 - ◆ Over subscribed PIC disk space
 - ◇ Production system queried bad source for available disk space
 - ◇ Sorted out with good support from PIC team
 - ◇ Improvements of tools under way
 - ◆ Dataset need by SAM test accidentally removed at several sites
 - ◆ Bad HammerCloud results at many sites under investigation
 - ◇ Appears to be a monitoring issue - not a site problem
- CMS User tickets in SNOW
 - ◆ CMS has a Funtional Element (FE) "CMS Support"
 - ◆ This FE should be used to route CMS related user issues to
 - ◆ Supporters will help directly or forward the user to the appropriate channel

Andrea Sciaba asks how available disk space is calculated. Christoph answers that he thinks this is taken from some of the dashboard columns reporting storage information. Andrea Sciaba says that he is maintaining this information and it's coming from the BDII. Christoph says that in any case this will be changed and will be no longer used.

LHCb

- Operations
 - ◆ Currently finishing a restripping of the Run1 legacy data and of the 50 ns Run2 ramp
 - ◆ Discussion with CERN/LSF team about the queue capabilities, problems found both in LSF and Dirac (GGUS:115027 [↗](#))
 - ◆ Preparations for the 25ns ramp up ongoing.
- Developments
 - ◆ Hammercloud testing for LHCb is currently being re-vitalized. The probe will check the possiblity to run user analysis jobs with protocol access at sites.
 - ◆ perfSonar data extraction from WLCG sources is almost finished, currently working on the publishing of the data into LHCbDIRAC

Maite adds that the LSF problems for LHCb are understood and are related to timing out on the DIRAC side. It's a different problem than the one ATLAS is suffering from.

Ongoing Task Forces and Working Groups

gLExec Deployment TF

- NTR

RFC proxies

- NTR

Machine/Job Features

- A nagios probe checking the availability and sanity of machine / job features has been developed. It's currently running in preprod for the LHCb SAM instance. Results can be seen at <http://cern.ch/go/Gzn8> [↗](#) . LHCb sites providing MJF are
 - ◆ CERN
 - ◆ GRIDKA
 - ◆ LPNHE
 - ◆ Imperial College

Middleware Readiness WG

- Thanks to Edinburgh and Grif for offering to verify selected MW products for Readiness over CentOS7/SL7 this autumn. Progress/comments via our JIRA tracker [?](#).
- A provisional agenda of our next (16/9 at 4pm CEST) meeting is on page http://indico.cern.ch/e/MW-Readiness_12 [?](#). Please send additional items to the e-group [wlcg-ops-coord-wg-middleware](mailto:wlcg-ops-coord-wg-middleware@cern.ch) at cern...

Multicore Deployment

IPv6 Validation and Deployment TF

Squid Monitoring and HTTP Proxy Discovery TFs

- Alastair Dewhurst has implemented the next step in the critical path of development, still testing

Network and Transfer Metrics WG

- Successfully tested publishing of the perfSONAR results to the message bus directly from the OSG collector. Discussing possible SLA to run this as a production service in collaboration with OSG.
- OSG datastore on track to go production at the end of July, this will be a service provided to the WLCG, it will store all the perfSONAR data and provide an API
- Started testing proximity service, which helps to map sonars to storages and thus enables integration of the network and transfer metrics.
- Review of the experiments use cases was presented/discussed at the last meeting, see slides for details (<https://indico.cern.ch/event/393101/> [?](#))
- FTS performance study update - see slides for details (<https://indico.cern.ch/event/393101/> [?](#)), observations from the report so far:
 - ◆ Peak transfer rates between Europe and North America are less asymmetric than they were last month (to be followed up)
 - ◆ Almost all incoming to BNL uses TCP=1 (Alejandro confirmed this is how BNL is configured right now, the other FTS instances use auto-tuning)
 - ◆ CMS T1s have better transfer rates compared to ATLAS and LHCb (to be followed up)
 - ◆ CMS uses TCP=1 more often than ATLAS and LHCb for large files
 - ◆ TCP stream=1 transfer do timeout about 2-3% of the time, however timeouts are concentrated at a few sites.
 - ◆ Throughput dependence on TCP streams possibly understood (see http://egg.bu.edu/lhc/fts/docs/2015-05-26-status/results_so_far.pdf [?](#))
- perfSONAR operations status
 - ◆ Agreed to establish WLCG-wide meshes for top 100 sites (based on the contributed storage and location). This will enable full mesh testing of latencies, traceroutes and throughput (ongoing).
 - ◆ ESNNet interested in the perfSONAR configuration interface developed for WLCG, development design document for an open-source project based it is currently discussed.

HTTP Deployment TF

Information System Evolution

- The first TF meeting took place last week ([agenda](#), [minutes](#))
 - ◆ It was agreed to implement in REBUS a set of easy fixes. For more details, please check REBUS known issues
 - ◆ A set of action items were defined, for more details, please check Task tracking and timeline. A summary below:
 - ◇ Requirements to remove information (Physical CPU) or change how information is collected (HS06) in REBUS will be followed up
 - ◇ Agree on a better definition of Installed Capacities, or even decide to change this name and better use "Available capacities" or something similar
 - ◇ Discuss at the MB the possibility of adding T3s and also publish pledges per sites in REBUS
- A draft document describing use cases from experiments and project activities relying on the information system has been circulated among TF members for their contribution. This will be presented in the future MB (date to be confirmed) although we are aiming to have the document ready by end August

Alessandro explains that ATLAS is collecting the requested feedback for the Use Case document. However, it's very likely that this won't be ready for the end of August. Maria will confirm with Ian when he is interested in having this presented at the MB. In any case, the document won't be presented until all use cases from all experiments are described.

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.

After a question from Alessandra, Maarten explains that secure services behind an alias should check whether their host certificates are compliant with the new algorithm for host cert verification.

Specific actions for experiments

Creation date	Description	Affected VO	Affected TF	Comments	Deadline	Completion
2015-07-02	Provide a description of each experiment computing support structure, so tickets wrongly assigned to the T0 (via SNOW or GGUS) can be properly redirected; evaluate the creation of SNOW Functional Elements	all	n/a	ALICE done. ATLAS draft under discussion. CMS will discuss with Maite next week. LHCb pending	July 30th. Extended to August 20th	~50%

for the experiments, if this is not already the case					
--	--	--	--	--	--

Specific actions for sites

Creation date	Description	Affected VO	Affected TF	Comments	Deadline	Completion
2015-06-18	Some sites have still not enabled multicore accounting	All	Multicore Deployment	Instructions here	a.s.a.p.	Almost DONE. HERE is the list of the remaining still pending sites.
2015-06-04	ALL ATLAS sites implementing a cap to their multicore resources (whether their configuration is dynamic just for a portion of nodes or it is a static partition) should review the cap to give 80% of the ATLAS production resources to multicore. As a reminder the shares for ATLAS jobs are as follows T1: 5% analysis and 95% production; T2: 50% analysis and 50% production. So multicore should get 80% of 95% at T1s and 80% of 50% at T2. More info here	ATLAS	Multicore		None	CLOSED
2015-06-04	LHCb T1s requested to make sure that all the RAW data will be stored on the same tape set in each tape system when it is feasible	LHCb	-	More details in GGUS:114018		CLOSED
2015-06-18	CMS requests an adjustment of the Tier-1 fair share target for the following VOMS roles: /cms/Role=production 90% (was 95%), /cms/Role=pilot 10% (was 5%). Note that for CMS SAM tests the role cms/Role=lcgadmin is used, it basically needs very little fair share but should be scheduled asap to have the test not timing out. Overall at least 50% of the pledged T1 CPU resources should be reachable via multi-core pilots (this is as before - just mentioned for completeness)	CMS			None yet	CLOSED

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

GGUS: How do users (e.g. VO shifters) receive GGUS downtime notifications?

<https://its.cern.ch/jira/browse/GGUS-1454>

This is ongoing, experiments have put comments in the ticket already. They have all acknowledged the need to find a method to inform shifters and supporters to notify unscheduled downtimes in GGUS. No need an action item since the progress is satisfactory.

-- MariaALANDESPRADILLO - 2015-07-27

This topic: LCG > WLCGOpsMinutes150730

Topic revision: r22 - 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