

Table of Contents

WLCG Operations Coordination Minutes, Jan 18th 2018.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Input regarding SAM tests and Site availability reports.....	2
Introduction.....	2
VOs.....	2
ALICE.....	2
Discussion.....	2
ATLAS.....	3
Discussion.....	3
CMS.....	4
Discussion.....	4
LHCb.....	5
Discussion.....	6
Sites.....	6
List of questions.....	6
Site replies.....	6
Short summary table.....	6
Discussion.....	7
KIT.....	7
GridPP survey.....	8
Site summary.....	8
Middleware News.....	8
Discussion.....	9
Tier 0 News.....	9
Tier 1 Feedback.....	9
Tier 2 Feedback.....	9
Experiments Reports.....	9
ALICE.....	9
ATLAS.....	9
CMS.....	10
LHCb.....	10
Ongoing Task Forces and Working Groups.....	10
Accounting TF.....	10
Archival Storage WG.....	10
Information System Evolution TF.....	11
IPv6 Validation and Deployment TF.....	11
Machine/Job Features TF.....	11
Monitoring.....	11
MW Readiness WG.....	11
Network Throughput WG.....	11
Squid Monitoring and HTTP Proxy Discovery TFs.....	11
Traceability WG.....	11
Container WG.....	12
Special topics.....	12
Site availability reports and SAM tests.....	12
Action list.....	12
Specific actions for experiments.....	12
Specific actions for sites.....	12
AOB.....	12

WLCG Operations Coordination Minutes, Jan 18th 2018

Highlights

Outcome of the discussion regarding SAM tests and A/R reports:

The majority of T1 sites and VOs confirmed that they do use SAM tests from the critical profiles and regularly check A/R reports.

Suggestions for improvements:

- Propose policy for accepting A/R recalculation requests. The draft should be reviewed at the next meeting.
- VO should have flexibility in the definition of the critical profile. Though the impact of the changes in the critical profile should be carefully tested by the VO and should be announced in advance to the sites, no approval of the MB is required.
- Make tests as close as possible to the real production flow. ATLAS is planning some work in this direction.
- The proposal to include real production flows in the critical profile was not supported.
- Sites which do have local fabric monitoring like Nagios for example, are recommended to use an API to import test results into the local fabric monitoring. This would help to avoid test failures staying unnoticed for months.
- Transparent navigation from the SAM UI to the log files is required to facilitate test failure debugging. This feature has to be preserved in the new SAM UI being developed by the monitoring team.

The outcome of this discussion will be presented at the MB.

Agenda

- <https://indico.cern.ch/event/696409/>

Attendance

- local: Alberto (monitoring), Alessandro (ATLAS), Gavin (T0), Julia (WLCG), Maarten (WLCG + ALICE), Marian (monitoring + networks), Panos (WLCG), Ryu (ATLAS)
- remote: Alessandra (Manchester + ATLAS), Andrea (WLCG + CMS), Catherine (LPSC + IN2P3), Di (TRIUMF), Eric (IN2P3-CC), Felix (ASGC), Frédérique (LAPP), Gareth (RAL), Giuseppe (CMS), Jeff (NLT1), Jeremy (GridPP), Johannes (ATLAS), Mayank (WLCG), Pepe (PIC + CMS), Peter (Oxford), Renato (LHCb), Stephan (CMS), Ulf (NDGF), Vincent (security), Xavier (KIT), Xin (BNL)
- apologies:

Operations News

- The next meeting will be on *March 1st*

Input regarding SAM tests and Site availability reports

Introduction

presentation [↗](#)

VOs

ALICE

- Do you use SAM tests for operations?
 - ◆ *A: rarely*
 - ◆ If yes
 - ◇ do you use tests of the critical profile?
 - *A: yes*
 - ◇ tests of non-critical profile
 - *A: no*
 - ◇ how often operations team checks the results of SAM tests (daily, weekly, monthly, few times per year, only if updates had been performed)
 - *A: a few times per year, mostly for A/R corrections*
 - ◇ how tests are followed up (ETF Nagios UI, WLCG Dashboard, VO-specific tools, other)
 - *A: mostly MonALISA, sometimes SAM-3 GUI*
- Do you think site Availability/Reliability based on SAM tests is a good metric for site quality estimation?
 - ◆ *A: it is a reasonable metric*
- Any suggestions for changes in a critical profile to make it more realistic for site quality estimation?
 - ◆ *A: making the metrics yet more realistic may have a questionable cost-benefit ratio*
- How often operations team and VO computing management check the WLCG availability reports (each month, as soon as they come out, once each few months, once per year, only if some updates performed, never)
 - ◆ *A: as soon as they come out, to see if A/R corrections might be needed*
- Do you think we (WLCG) should continue to rely on site Availability/Reliability reports:
 - ◆ to estimate the quality of the site?
 - ◇ *A: it is a reasonable metric, among others*
 - ◆ for reporting to funding agencies?
 - ◇ *A: ditto; sites should focus on the **big** messages from such reports and not insist on getting A/R numbers corrected by small amounts; i.e. the machinery should be allowed to have some imperfections that sites can live with*

Discussion

- Maarten: the main message is that SAM is a useful tool that should be allowed to have occasional glitches without an automatically implied consequence that A/R values have to be corrected; in 2017 there were 72 recomputations for ALICE, spread over 14 observed glitches, and differences typically were a few % at most; significant effort was spent just to get reports to look better
- Pepe: SAM A/R reports are complementary to accounting reports that may have their own issues

- Alessandra:
 - ◆ SAM tests ideally ought to be independent of the experiment
 - ◆ A/R correction requests could be accepted only when the expected difference exceeds some threshold, as is done in ATLAS
- Jeff: I have seen percent-level differences highlighted in presentations
- Renato: A/R reports are important for sites to show how they are doing
- Xin: should the thresholds in the A/R reports be changed?
- Alessandra: sites are green already at 90%
- Julia: we should look into a policy on when recomputations are done, viz. when there will be a significant change in the A/R values

ATLAS

presentation [↗](#)

- Do you use SAM tests for operations? *Not now, but soon will be by ADCoS shifters.*
 - ◆ If yes
 - ◇ do you use tests of the critical profile? *Yes*
 - ◇ tests of non-critical profile *No (but investigating to add tests by rucio)*
 - ◇ how often operations team checks the results of SAM tests (daily, weekly, monthly, few times per year, only if updates had been performed) *Currently not, but It will check every day*
 - ◇ how tests are followed up (ETF Nagios UI, WLCG Dashboard, VO-specific tools, other) *WLCG Dashboard*
- Do you think site Availability/Reliability based on SAM tests is a good metric for site quality estimation? *Yes*
- Any suggestions for changes in a critical profile to make it more realistic for site quality estimation? *Adding tests by other file-transfer protocols which match real production*
- How often operations team and VO computing management check the WLCG availability reports (each month, as soon as they come out, once each few months, once per year, only if some updates performed, never) *only if some updates performed*
- Do you think we (WLCG) should continue to rely on site Availability/Reliability reports:
 - ◆ to estimate the quality of the site? *Yes*
 - ◆ for reporting to funding agencies? *Yes*

Discussion

- Jeff: if no one is looking at the tests, how can they be considered a good metric?
- Ryu: only for operations no one was looking yet (that will change); the tests are good metrics for A/R reports
- Julia:
 - ◆ If the tests were made more realistic, they would be considered more useful
 - ◆ Also their inclusion in shift duties will help
 - ◆ The A/R could be calculated from SAM *or* ASAP
- Alessandro:
 - ◆ On the one hand we have atomic SAM tests, on the other the notion of the site being in a working state
 - ◆ Sites need to know the details, which they only can get from simple tests

- ◆ We need both profiles for different uses

- Jeff: as a generic observation, unless there is a really good reason for a test, it may be best just to drop it

CMS

- Do you use SAM tests for operations? *Yes*
 - ◆ If yes
 - ◇ do you use tests of the critical profile? *Yes*
 - ◇ tests of non-critical profile *There are two critical profiles setup for CMS. The CMS_CRITICAL_FULL profile contains all relevant probes for us while CMS_CRITICAL is a subset used for WLCG reporting. (Probes under development/testing are outside the critical profiles.)*
 - ◇ how often operations team checks the results of SAM tests (daily, weekly, monthly, few times per year, only if updates had been performed) *SAM results are looked at several times an hour!*
 - ◇ how tests are followed up (ETF Nagios UI, WLCG Dashboard, VO-specific tools, other) *All of the above: We use the SAM ETF interface for the most recent results. The SAM3 dashboard provides alarms and also historic information and is the most often used interface. With the migration to MonIT and desire/need to combine/correlate SAM with results from other tests, like Hammer Cloud and data transfer results, we developed additional tools. GGUS is used to communicate SAM failures to sites.*
- Do you think site Availability/Reliability based on SAM tests is a good metric for site quality estimation? *CMS adds experiment specific test results to the SAM results to derive an equivalent, more-inclusive site availability. We believe SAM availability to be a fair VO-independent metric for grid performance.*
- Any suggestions for changes in a critical profile to make it more realistic for site quality estimation? *The CMS_CRITICAL_FULL profile is a dynamic combination of SAM tests. We regularly adjust it to changing needs of the experiment. We would not object to the profile used by WLCG be managed more actively.*
- How often operations team and VO computing management check the WLCG availability reports (each month, as soon as they come out,once each few months,once per year, only if some updates performed, never) *Many CMS sites check the WLCG availability reports each month as soon as they come out.*
- Do you think we (WLCG) should continue to rely on site Availability/Reliability reports:
 - ◆ to estimate the quality of the site? *Yes*
 - ◆ for reporting to funding agencies? *Yes*

Discussion

- Stephan, Pepe: do critical profile changes have to be approved by the MB?
- Julia:
 - ◆ In practice the criteria have been relaxed:
 - ◇ New tests should first be introduced via a non-critical profile
 - ◇ Sites should be made aware of new tests to avoid bad surprises
 - ◇ When all looks fine, the experiment can update its critical profile
- Maarten:
 - ◆ Please keep announcing significant changes in this meeting
 - ◇ As e.g. was done for the Singularity tests
 - ◆ We then can include them in MB Service Reports for the record
- Stephan: we would like to see new rules for changing the profile
- Andrea: for example, we want to make an Xrootd test critical

- Julia:
 - ◆ You can go ahead as described
 - ◆ We will discuss this matter in the MB for clarification
- Alessandro:
 - ◆ The tests are about an SLA between the sites and the VO
 - ◆ Could we look into a subset that is common to the experiments?
- Maarten:
 - ◆ The experiments have many differences between them
 - ◇ For example, for the time being Singularity is only critical for CMS
- Renato:
 - ◆ CVMFS is common to all and ought to be tested by all
 - ◆ At my site I see it tested by LHCb, but not by ALICE
- Alessandro: could we envisage a common effort?
- Maarten:
 - ◆ That was the situation years ago, when the ops VO tests were critical
 - ◆ We went to experiment-specific critical tests because of an ever increasing mismatch between what was tested by ops and what was used by experiments
- Alessandra:
 - ◆ There also were issues with the experiment proxies vs. permissions
 - ◇ E.g. for the lcgadmin role
- Stephan:
 - ◆ Regarding A/R vs. accounting reports, the latter also have had inconsistencies
 - ◆ Furthermore, sometimes a VO does not use the full capacity at a site
 - ◆ The SAM reports are therefore complementary
 - ◆ We need to do corrections only a few times per year

LHCb

- Do you use SAM tests for operations? *Yes*
 - ◆ If yes
 - ◇ do you use tests of the critical profile? *Yes they are used to check availability and reliability of a site. It is seen as the metric that everyone agrees on, if a site has bad availability/reliability numbers these are numbers one can refer to when discussing with the site.*
 - ◇ tests of non-critical profile *We have some tests running in the non critical profile (cvmfs, mjj) which we use for checking if sites are compliant with our requirements. E.g we plan to use the cvmfs probe to check if WLCG sites have deployed the newly requested cvmfs mount point. There was a tool which automatically queried the old nagios to extract metrics information and put it into a WLCG dashboard page. After the move to etf this was not ported.*
 - ◇ how often operations team checks the results of SAM tests (daily, weekly, monthly, few times per year, only if updates had been performed) *They are checked to discover discrepancies between the Dirac monitoring portal and WLCG probes. In few cases we are also notified by sites that a certain probe is failing, the last few times these were false negatives, ie. the probe was failing but lhcb operations was ok.*
 - ◇ how tests are followed up (ETF Nagios UI, WLCG Dashboard, VO-specific tools, other) *ETF Nagios UI, WLCG Dashboard*
- Do you think site Availability/Reliability based on SAM tests is a good metric for site quality estimation?
- Any suggestions for changes in a critical profile to make it more realistic for site quality estimation?
- How often operations team and VO computing management check the WLCG availability reports (each month, as soon as they come out, once each few months, once per year, only if some updates

performed, never) *yes results are x-checked*

- Do you think we (WLCG) should continue to rely on site Availability/Reliability reports:
 - ◆ to estimate the quality of the site? *if the probes shall guide operations they give a lower bounds of quality of the site, they are not checking the full operational spectrum of lhcb workflows. If we can agree on some LHCb metrics (e.g. pilots probing the site services) this would be more realistic, the problem I see is that disentangling for these kind of probes the site service from the vo specific failures can be tricky. SAM probes are to some extent more VO agnostic.*
 - ◆ for reporting to funding agencies? *one needs to distinguish between operations and other purposes. The same set of metrics should not serve the two purposes.*

Discussion

- Renato:
 - ◆ SAM A/R reports are important for sites to get funding and demonstrate improvements
 - ◆ They can be used by a ROC/NGI as a reason to open tickets against its sites

Sites

List of questions

- Do you use SAM tests for operations?
 - ◆ If yes
 - ◇ do you use tests of the critical profiles of the LHC VOs?
 - ◇ tests of other profiles
 - ◇ do you rely on SAM/ARGO/... results for any non-LHC VOs?
 - ◇ how often do you check the results of SAM tests (daily, weekly, monthly, few times per year, only if some updates performed)
 - ◇ how do you access them (API for importing to site local monitoring, ETF Nagios UI, WLCG Dashboard, VO-specific tools, other)
- Do you run tests rather than the ones of SAM to check availability of your services?
- If yes
 - ◆ what framework are you using?
 - ◆ how different are those tests compared to SAM?
- Do you think site Availability/Reliability based on SAM tests is a good metric for site quality estimation?
- Any suggestions for changes in SAM tests to make them more realistic for site quality estimation?
- How often do you check the WLCG availability reports (monthly as soon as they come out, once each few months, once per year, only if some updates performed, never)
- Do you think we (WLCG) should continue to rely on site Availability/Reliability reports:
 - ◆ to estimate the quality of the sites?
 - ◆ for reporting to funding agencies?

Site replies

See SAMquestionnaireSiteReplies for the details per site

Short summary table

Site	Use tests	How often	Access	Use reports	How often	Continue to rely on reports
------	-----------	-----------	--------	-------------	-----------	-----------------------------

Site	Use tests	How often	Access	Use reports	How often	Continue to rely on reports
ASGC	Yes	Daily	WLCG Dashboard	Yes	Monthly	Yes
BNL	Yes	Daily	WLCG Dashboard	Yes	Monthly	Yes
KIT	Rather not, only if something goes wrong	-	WLCG Dashboard	Not regularly	Just casually	Rather not
NL-T1	Rather not	-	-	Rather not	-	Yes, if there is no false 'down'
NRC-KI	Yes	Daily	ETF Nagios UI, WLCG Dashboard	Yes	Monthly	Yes
RAL	Yes	Daily	Everything what available	Yes	Monthly	Yes
TRIUMF	Yes	Daily	API,ETF Nagios UI,the WLCG Dashboard	Yes	Monthly	Yes
CERN	Yes	Daily	ETF UI, Dashboard, soon API	Yes	When they come out	Yes
IN2P3	Yes	Twice per day	ETF UI, Dashboard, API	Yes	When they come out	Yes
JINR	Yes	Several times per day	ETF UI, Dashboard, Experiment-specific tools	No	Never	No
KISTI	Yes	Daily	ETF Nagios UI and WLCG Dashboard + ALICE specific monitoring	Yes	Monthly, as they come out	Yes
NDGF	Yes	Constantly as it is integrated into local Nagios	API	Yes	Monthly, as they come out	Yes
PIC	Yes	Constantly as it is integrated into local Nagios	API,CMS SSB, WLCG Dashboard	Yes	Monthly, as they come out	Yes

Discussion

KIT

- Xavier:
 - ◆ The ops tests used to be similar for all VOs, but there was a mismatch between them and actual VO operations
 - ◆ We therefore use our own tests
- Julia: would you consider importing the VO tests into your local monitoring?
- Xavier:
 - ◆ Their availability does not seem to be guaranteed?
 - ◆ We would need that in order to avoid getting spurious alarms
- Julia: many T1 import the Nagios results without problems
- Xavier: we will look into it

NLT1

- Jeff:
 - ◆ In the past it was important to follow the SAM tests closely, because of low A/R values at sites in those days
 - ◆ Nowadays the A/R is generally OK and issues are resolved through tickets

- ◆ Recently there was a misleading drop in the NIKHEF A/R due to test issues
 - ◆ It led to work at the site and in the VO, a lot of excitement for nothing
 - ◆ For real issues experiment shifters can open tickets, after careful investigation
 - ◇ For example, when many sites are affected, the problem likely is elsewhere
 - ◆ We will look into importing Nagios results again, as was done in the past
 - ◇ We had to flag many experiment-specific tests to be ignored because unreliable
 - ◇ The result was that only parts of the critical profiles were checked locally
 - ◆ Funding agencies probably will continue to look at the reports
-
- Alessandra: do you look at the reports?
 - Jeff:
 - ◆ Since that recent matter
 - ◆ I normally look at our infrastructure, the number of jobs etc.
 - Maarten: the importance of A/R reports to funding agencies depends on the country
 - Julia:
 - ◆ When the funding agency does not care, the site need not care
 - ◆ However, the tests are useful to look at and spot anomalies
 - Alessandro: the reports should be checked, but they are not perfect
 - Jeff: useless work should be avoided
 - Alessandra: should the SAM tests just be for operations and we scrap the reports?
 - Julia:
 - ◆ Some sites can just ignore the A/R reports
 - ◆ But also they are advised to keep an eye on the test results

GridPP survey

presentation [↗](#)

- Jeremy: overall the SAM machinery is considered very useful

Site summary

presentation [↗](#)

- Julia:
 - ◆ One message is that logfiles are very important
 - ◆ That will be taken as input for Alberto's monitoring team

Middleware News

- Useful Links:
 - ◆ Baseline Versions [↗](#)
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines/News
 - ◆ UMD 4.6.0 has been released both for SL6 and C7. Lots of updates, in particular first releases of CreamCE and UI for C7 (more details at <http://repository.egi.eu/2017/12/18/release-umd-4-6-0/> [↗](#))
 - ◆ A new version of DPM (1.9.2) fixing a vulnerability is available on EPEL. EGI is integrating this version in UMD then it will be set as baseline.
- Issues:
 - ◆ https://wiki.egi.eu/wiki/SVG:Meltdown_and_Spectre_Vulnerabilities [↗](#)

- ◊ Kernel updates fixing Meltdown and Spectre variant 1 need to be applied ASAP
- ◆ CMS is contacting sites running DPM as their configuration for AAA federation is incorrect and generates read loops (apparently the issue has been always there..). A new configuration fixing this issue has been designed at GRIF_LLJ. Sites will be then contacted to perform the needed changes.
- ◆ A recent upgrade on EPEL6-testing (<https://bodhi.fedoraproject.org/updates/FEDORA-EPEL-2018-71db8f6f28>) which includes new version of bouncycastle and canl-java, is breaking voms-client installations on UI and WN nodes installed from the UMD repository, CREAM-CE and ARGUS. UMD URT has been informed in order to fix the issue before this upgrade will be pushed to EPEL stable.
- T0 and T1 services
 - ◆ CERN
 - ◊ EOS-ATLAS upgraded to Citrine 4.2
 - ◊ EOS-CMS upgrade to Citrine planned for next week.
 - ◆ FNAL:
 - ◊ FTS upgraded to 3.7.7
 - ◆ JINR
 - ◊ Minor upgrade: dCache 2.16.53 -> 2.16.57; upgrade xrootd 4.7.1 -> 4.8.0
 - ◆ NDGF:
 - ◊ dCache upgraded to 3.0.38
 - ◆ RAL:
 - ◊ Ceph upgrade to Luminous 12.2.2 planned for 24/1/18.
 - ◊ RAL-Echo dual stack by February 2018.

Discussion

Tier 0 News

- Ongoing patching for Spectre/Meltdown

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- The activity levels have been very high on average
 - ◆ Records on Dec 25: 175k on the grid, 100k at CERN
 - ◆ Thanks to the sites!
- Central services
 - ◆ 20h grid activity fallout due to a corrupted file system

ATLAS

- Stable grid production over the last weeks with ~250k concurrent running job slots over the christmas holidays (without HLT farm) and now back up to ~350k concurrent job slots including HLT farm and contributions from HPCs.
- One production and data management incident during the xmas holidays: not well defined event generation task with many very short jobs and outputs to a single site overloaded the file transfer

system (FTS). Expert intervention before the new year mitigated the problem

- Re-derivation campaign of data15/16/17 during November and December finished successfully before the start of the new year
- First half of the data17 reprocessing from RAW finished successfully after only about 2 weeks processing about 2.2b event during the christmas break. The second half of the data17 reprocessing is expected to start within the next days.
- A large campaign of mc16d digitization and reconstruction is on-going since mid of December
- Spectre/Meltdown: There will be performance losses due to the related linux kernel updates and there are first indications that workflows with higher I/O might be more affected. Very preliminary numbers: Simulation -2%, MC Digitization+Reconstruction: -5%

CMS

- Spectre/Meltdown issue
 - ◆ first tests show only small CPU loss for reconstruction and Monte Carlo workflows
 - ◆ detailed measurements in progress, thanks to CERN-IT for providing two identical machines, one with/one without patch!
 - ◆ patching and reboots at sites and CERN an operational disruption nevertheless
- Tier-0 tape backlog worked down to about 2 PB
- Production busy with 2017 data re-processing and Monte Carlo samples
 - ◆ both Tier-0 and HLT resources used by production
 - ◆ successful processing over the holidays
 - ◆ 2017 data re-processing almost complete
 - ◆ enough work to keep CPU resources, ~200k cores, busy for a while
 - ◆ increased number of single-core jobs
 - ◆ issues with tape writing/staging at CERN greatly improved with the fix deployed by the Castor team and switch to xrootd for all CERN internal transfers
- we remind that we plan to process 2018 data on SL/CentOS 7
 - ◆ this requires Singularity at sites by March 2018
- small but steady progress on IPv6 storage access
- CNAF recovery on schedule
 - ◆ thanks to KIT and others for providing extra CPU resources in the interim

LHCb

- Spectre/Meltdown: Reboot was (and will be) followed by operations
- Productions are at full steam
- Waiting for CNAF recovery to finalise productions

Ongoing Task Forces and Working Groups

Accounting TF

- In collaboration with Archival Storage WG tape accounting information is being integrated in the new WLCG Storage Space accounting system
- Next WLCG Accounting Task Force Meeting ins next Thursday 25th of January
<https://indico.cern.ch/event/696513/>

Archival Storage WG

Information System Evolution TF

- Nothing to report

IPv6 Validation and Deployment TF

Progress in IPv6 deployment at Tier-2 sites:

Status	7/12/17	16/1/18
no reply	38%	22%
on hold	20%	29%
in progress	31%	35%
done	11%	14%

Detailed status [here](#).

Machine/Job Features TF

Monitoring

MW Readiness WG

This is the status of jira ticket updates [🔗](#) since the last Ops Coord of 20171207:

- MWREADY-152 [🔗](#) - DPM 1.9.1/1.9.2 tested at GRIF. Found a regression on 1.9.1 fixed and released on v1.9.2.

Network Throughput WG

- perfSONAR 4.0.2 - 190 instances updated out of which 53 are already on CC7
 - ◆ **WLCG broadcast** will be re-sent next week to remind sites of the **upcoming important dates** and new documentation
 - ◆ perfSONAR 4.1 release, planned in Q1 2018 will no longer ship SL6 packages
 - ◆ EOL for SL6 support in Q3 2018
 - ◆ All sites are encouraged to upgrade to CC7 as soon as possible
- WLCG/OSG network services
 - ◆ On-going work on improving OSG collector and broadcasting results via GOC's RabbitMQ
 - ◆ Revised documentation is available at <https://opensciencegrid.github.io/networking/> [🔗](#)
 - ◆ LHCOPN and perfSONAR dashboards done in collaboration with OSG, UC, CERN IT/CS and IT/MONIT are available at <http://monit-grafana-open.cern.ch/dashboard/db/home?orgId=16> [🔗](#)
- WLCG Network Throughput Support Unit: see twiki for summary of recent activities.

Squid Monitoring and HTTP Proxy Discovery TFs

- NTR

Traceability WG

Container WG

Special topics

Site availability reports and SAM tests

See above

Action list

Creation date	Description	Responsible	Status	Comments
01 Sep 2016	Collect plans from sites to move to EL7	WLCG Operations	Ongoing	[older comments suppressed] Dec 7 update: Tier-1 plans are documented in the Nov 2 minutes. Jan 18 update: CREAM and the UI were released in UMD-4 on Dec 18.
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	DONE	[older comments suppressed] Oct 5 update: as both OSG and EGI were not happy with the previous proposals, and as this matter does not look critical, we propose to create best-practice recipes instead and advertise them on a WLCG page. Dec 7 update: a first version is prominently linked from WLCGOpsWeb and the WLCGOpsMeetingTemplate.
14 Sep 2017	Followup of CVMFS configuration changes, check effects on sites in Asia	WLCG Operations	Pending	

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 > WLCGOpsMinutes180118

Topic revision: r29 - 2018-02-28 - MaartenLitmaath



Copyright &© 2008-2024 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