

# Table of Contents

<b>WLCG Operations Coordination Minutes, May 2, 2024.....</b>	<b>1</b>
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Special topics.....	1
HammerCloud status and plans.....	1
Discussion.....	1
htgettoken + HashiCorp Vault as a Service for Managing Grid Tokens.....	2
Discussion.....	2
Middleware News.....	3
Discussion.....	3
Tier 0 News.....	4
Tier 1 Feedback.....	4
Tier 2 Feedback.....	4
Experiments Reports.....	4
ALICE.....	4
ATLAS.....	4
CMS.....	4
LHCb.....	4
Task Forces and Working Groups.....	5
Accounting TF.....	5
Migration from DPM to the alternative solutions.....	5
Information System Evolution TF.....	5
IPv6 Validation and Deployment TF.....	5
Monitoring.....	5
Network Throughput WG.....	5
WG for Transition to Tokens and Globus Retirement.....	5
Action list.....	5
Specific actions for experiments.....	5
Specific actions for sites.....	5
AOB.....	5

# WLCG Operations Coordination Minutes, May 2, 2024

## Highlights

- HammerCloud status and plans
- htgettoken + HashiCorp Vault as a Service for Managing Grid Tokens
- Update on EL9 vs. SHA-1
- Transition from VOMS-Admin to IAM

## Agenda

<https://indico.cern.ch/event/1411424/>

## Attendance

- local: Lorenzo (HammerCloud), Maarten (ALICE + WLCG), Steve (HammerCloud)
- remote: Alexander (ATLAS), Andrea (CMS + WLCG), Benjamin (ATLAS), Borja (monitoring), Christoph (CMS), Dave D (FNAL), David B (IN2P3-CC), Eva (CERN IT-DA), Federica (IAM devs), Frédérique (LAPP), Jan (LHCb), Mario (ATLAS), Panos (CMS + WLCG), Petr (ATLAS + Prague), Stephan (CMS), Thomas (DESY)
- apologies:

## Operations News

- the next meeting is planned for June **13** !

## Special topics

### HammerCloud status and plans

see the presentation [↗](#)

- Steve also introduces Lorenzo, who has just joined the team and will be working on HammerCloud as well as the ETF
- Lorenzo has studied computational physics in Padova and has previously worked on ILC-DIRAC for FCC studies

## Discussion

- Stephan:
  - ◆ what is the plan for the back-ends?
  - ◆ will you work with Andrea on those?
- Steve:
  - ◆ presuming you meant the submit nodes: we will work on those with Andrea
- Stephan:

- ◆ do you have an estimate for the interface with Python-3 job submission tools?
- Steve:
  - ◆ no
- Andrea:
  - ◆ the amount of work is difficult to estimate
  - ◆ worked on it 2 years ago, but had to abandon it
  - ◆ Python-3 packages were uploaded, but untested
  - ◆ Puppet manifests were far from finished
- Stephan:
  - ◆ will this work have more priority as of now?
- Steve:
  - ◆ yes

## htgettoken + HashiCorp Vault as a Service for Managing Grid Tokens

see the presentation [↗](#)

### Discussion

- Petr:
  - ◆ as HashiCorp will be bought by IBM,  
could that lead to a licensing problem for Vault?
- Dave:
  - ◆ we could fork the code if needed
  - ◆ as there will not be many instances, we might even pay
- Petr:
  - ◆ you might be charged per token
- Thomas:
  - ◆ a fork is already available at <https://openbao.org/> ↗
- Petr:
  - ◆ does CMS plan to integrate Vault + htgettoken into their SW?
- Dave:
  - ◆ for some use cases
  - ◆ for CRAB it does not seem to be needed
- Petr:
  - ◆ would users actually use the htgettoken command?
- Dave:
  - ◆ it would be called through wrappers
- Petr:
  - ◆ users would have to know about access tokens for specific roles?
- Dave:
  - ◆ there could be scripts, which can be common for popular use cases
- Maarten:
  - ◆ one of the ideas behind the token transition is that ordinary users  
should not have to know anything about tokens
  - ◆ it is good to see various auxiliary services already in production,  
that we can take advantage of for LHC experiments and related VOs
  - ◆ the token transition timeline allows about 2 years still for  
the user experience to get sorted out in each experiment

## Middleware News

- Useful Links
  - ◆ WLCGBaselineTable
- Baselines/News
  - ◆ While the UMD-5 [for EL9](#) is not ready yet, the BDII auxiliary rpms have been made available from the WLCG repository [\[link\]](#)
    - ◊ For EL9 and EL8
    - ◊ The BDII itself is available from EPEL
- EL9 vs. SHA-1
  - ◆ the situation is summarized here
  - ◆ further discussion in this fetch-crl ticket [\[link\]](#)
    - ◊ the problem is with OpenSSL instead
  - ◆ DigiCert are not going to update their root CA (presumably OK in browsers)
    - ◊ might be removed from IGTF if nothing important depends on it
      - today there are 84 WLCG users with such certificates
        - 74 in ATLAS
      - ◊ also is the issuer of 2 TERENA CAs
        - today there are 418 WLCG users with such certificates
          - 178 in ATLAS
          - 149 in CMS
        - ◊ LHC experiment users could switch to CERN Grid CA certificates
        - ◊ but there could also be many **services** with affected certificates
      - ◆ to be followed up further
        - ◊ hopefully we can at least push our classic CAs faster toward SHA-2

## Discussion

- Petr:
  - ◆ should we still rely on EGI for the UMD?
- Maarten:
  - ◆ there have been unfortunate delays due to various causes
  - ◆ the UMD has several important advantages over EPEL etc.
  - ◆ hopefully we will soon be able to start profiting from UMD-5
  - ◆ meanwhile we can e.g. use the WLCG repository as a stopgap
- Thomas:
  - ◆ could multi-hop transfers be considered to reduce the number of hosts that need to have SHA-1 configured?
- Petr:
  - ◆ we only do multi-hop transfers for special cases
- Maarten:
  - ◆ on a large scale it would imply bottlenecks
- *a discussion then followed about what we can do to make progress*
- Jan:
  - ◆ can we do a campaign to move users and services off SHA-1 CAs?
  - ◆ find out how much each VO is affected?
- Maarten:
  - ◆ will follow up with IGTF to see if classic CAs can be pushed
  - ◆ will try to get an idea of the number of services per VO

- ♦ intend to provide updates regularly

## Tier 0 News

## Tier 1 Feedback

## Tier 2 Feedback

## Experiments Reports

### ALICE

- normal to high activity on average in the last weeks
- an XRootD server I/O performance problem was badly affecting analysis trains at KIT
  - ♦ thanks very much to the KIT experts for resolving that with urgency!

### ATLAS

- Everything going smooth
  - ♦ Good job mix with some larger reconstruction and group production campaigns
  - ♦ 650k slots on average, with several 750 peaks
  - ♦ 2M-3.5M file transfers / 3-6PB volume per day
  - ♦ Tape consolidation campaign slowly easing off, 2-3 weeks tail left
- ~4k slots on SWT2 now extended by Google
- Increase of ARM resources, total ~50k slots now
  - ♦ CERN, GLASGOW, INFN-T1, SWT2\_GOOGLE extension
- ALMA9 migration going well and almost finished

### CMS

- overall smooth data taking and computing operations
  - ♦ core usage between 310k and 580k cores
  - ♦ due to HPC/opportunistic contributions
  - ♦ almost all production activities now Run 3
  - ♦ back to a more usual production/analysis split of about 3:1
  - ♦ various monitoring outages the last weeks.
- work on SRM to REST migration for tape endpoints continues
  - ♦ four done / four remaining
- phasing out SRMv2/GSIftp/gridFTP at sites
  - ♦ 20 done / 20 remaining
- three remaining DPM sites to migrate
- token migration progressing steadily
- waiting on python3 version/port of HammerCloud
- we will ask sites to remove x86-64-v1 microarchitecture worker nodes with the end of SL7, i.e. June 30, for CMS (five sites, less than 100 Worker nodes)
- online Oracle DB support issue ten days ago

### LHCb

- *Jan will have items reported as needed next time*

# Task Forces and Working Groups

## Accounting TF

### Migration from DPM to the alternative solutions

## Information System Evolution TF

## IPv6 Validation and Deployment TF

Detailed status here.

## Monitoring

## Network Throughput WG

## WG for Transition to Tokens and Globus Retirement

- timeline [↗](#) for the transition from VOMS-Admin to IAM services for VO management
  - ◆ the legacy VOMS servers are being removed from the vomses configuration files used to create VOMS proxies
    - ◊ on QA hosts at CERN since today
    - ◊ on production hosts at CERN planned for Tuesday May 7
    - ◊ wlcg-voms rpms will be updated accordingly
  - ◆ only the VOMS endpoints of the IAM services on OpenShift will be used
    - ◊ they have been used in production since 2 years by ATLAS and CMS, since ~2 months by ALICE and LHCb
    - ◊ all corresponding LSC files are in production since 2 years
  - ◆ switches from VOMS-Admin to IAM are planned for this month
    - ◊ per experiment, when it gives the green light
  - ◆ ultimate deadline: end of June!

## Action list

Creation date	Description	Responsible	Status	Comments
---------------	-------------	-------------	--------	----------

## Specific actions for experiments

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

## Specific actions for sites

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

## AOB

---

This topic: LCG > WLCOpsMinutes240502

Topic revision: r8 - 2024-05-03 - 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