




OPS and Facilities session

OPS and Facilities session during WS

- Organized in two slots on Tuesday and Wednesday, 1.5 hours each
- Topics included in the agenda have been discussed with the community (input collected from the LHC VOs, suggestions from the sites, set of proposed topics reviewed at the WLCG OPS coordination meeting)
- Focus on the topics which are important from the WLCG Ops Coordination perspective, that is have impact on the infrastructure evolution and WLCG operations, where currently work is ongoing and progress update is needed
- In preparation for the OPS and Facilities session WLCG OPS launched site questionnaire in order to get site view on some points which were discussed.
- Very useful discussions during the session and coffee breaks

Tuesday session


Overview of the ongoing changes and OPS campaigns o 
Julia Andreeva

Improving LHCONE security and Use of Jumbo frames 
Edoardo Martelli

Transition from VOMS Admin to IAM. Token profile defini 
Maarten Litmaath

APEL status and plans *Thomas Dack* 
Hoersaal, DESY 16:40 - 17:00

AUDITOR *Michael Boehler* 
Hoersaal, DESY 17:00 - 17:15

Follow up on the outcome of the WLCG review on the ac 
Panos Paparrigopoulos

Network

- In order to increase the trust in LHCONE, propose to tag prefixes announced to LHCONE with BGP communities. Network and WLCG OPS Coordination teams are launching tagging campaign to be accomplished by spring 2025.
- CERN decided to go for testing of the JUMBO frames on some production server. CERN network team is pushing for JUMBO frames at the WLCG infrastructure. There are different opinions regarding benefits vs operational overhead. More discussion is needed to decide how we can progress. Probably starting with several pioneer sites.

Transition from VOMS admin to IAM. Token profile definitions

Conclusions from Maarten's presentation

- Various items related to tokens concern many of us in the next months
 - VOMS-Admin EOL – final deadline June 30
 - IAM usability for VO administration by LHC experiments
 - High-priority issues are being worked on for a **next release** in June
 - HA options for LHC experiment IAM instances – not to be rushed
 - Data management: lessons learned from DC24
 - Aiming to reach the next level of token usage in the second half of this year
 - HTCondor CE versions that no longer support GSI – along with moving to EL9
 - APEL adjustments for tokens – short vs. medium term
 - GUT Profile WG progress toward a new VO attribute – for accounting and more
 - Version 2.0 of the WLCG token profile – to signal where we intend to go
 - More deployment and operations know-how – with IAM being in the center
 - More use of auxiliary services – gradually benefiting more use cases
- ... while we keep gaining experience with tokens everywhere!

Accounting. APEL

- APEL has been reorganized. The goals are to improve management of the project, adopt agile methods to improve service delivery, improve user support
- New APEL Python3/EL8 has been released. WLCG Ops coordination will start deployment campaign together with EGI operations
- Ongoing work on APEL version required for tracking HEPSCORE deployment. Deployment of this version to production is planned in two coming months

Accounting. Auditor.

- Interesting development by the site of the accounting system aiming on the accounting of the opportunistic resources. Though not limited to the opportunistic resources
- Plugin-based. Implementation already exists for different type of resources, including popular batch systems as SLURM and HTCondor.
- What is important , it is already integrated with APEL and therefore can be considered by the sites, in particular those which send normallized (aggregated accounting records). CERN is interetsed to have a look

Accounting. Follow up on the accounting and pledge management review

- Development in WAU and CRIC are following up the input received from the sites, VOs, WLCG ops and WLCG project office
- Plots which allows to show pledge evolution, not just snapshot are enabled . Almost final version is ready:

<https://monit-grafana.cern.ch/d/e9c3d51f-f377-4552-bdbf-799c3b15270d/dev-pledged-required?orgId=20>


- By the next pledge definition loop, the new pledge definition UI will be deployed
- Work on new dashboards allowing to compare APEL and experiment-specific accounting details progressing

New plots in WAU

Pledged data for CPU



Required data for CPU



Comparison between pledged and required CPU (all Tier and all VOs)



Comparison between pledged and required CPU for Tier (Tier-0)



Relative difference between Required and Pledged for CPU for all



Relative difference between Required and Pledged for CPU for Tier (Tier-0)



Wednesday session

New WLCG helpdesk. Status and plans.

Pavel Weber



Hoersaal, DESY

09:00 - 09:20

Handling of high memory jobs. Whole-node scheduling. Standard 16-core instead of 8-core slots. VO and site perspective.

Antonio Perez-Calero Yzquierdo et al.



Hoersaal, DESY

09:20 - 10:30

New WLCG help desk

- After 20 years of stable operations GGUS is going to retire in the end of 2024. Read only mode after the end of 2024
- Will be replaced by a new system which meet requirements of WLCG and EGI
- Beased on Zammad open source technology
- Core workflows are almost done. User, Team, Alarm workflows are in place. Multisites to be accomplished
- AAI intergration is done. CERN ServiceNow integration is ongoing, FNAL ServiceNow integartion is planned
- Few training and QA sessions are planned in October and November
- Before 1.10.24 pilot phase, early adopters are welcome
- After migration of SU by the 1.10.24 everyone will be invited to join new system

- **Discussion on handling high memory jobs, whole node scheduling, 16 cores (vs 8 cores) standard**
 - Experiments have slightly different approach. CMS and ALICE are handling complexity of scheduling payloads with different requirements inside pilot. ATLAS prefers to pass specific requirements to the batch system and allow it to handle them
 - More cores or whole node scheduling as well as long job time limit provide more flexibility for the VOs and might allow to maximize resource utilization. Though there are concerns expressed by the sites regarding potential accounting issues, CPU efficiency, compatibility with existing infrastructure and availability of the workflows allowing co-existence of payloads with different number of cores
 - Clearly there is a need for further study to develop proper strategy and recommendations regarding most efficient core count for job allocation. People participating in preparation of this discussion propose to follow up by launching a dedicated task force or working group bringing together sites and experiment experts

Questions?