



Survey on WLCG accounting and pledge management tools

Julia Andreeva (CERN)

GDB
10.01.2024

Goal and scope of the survey

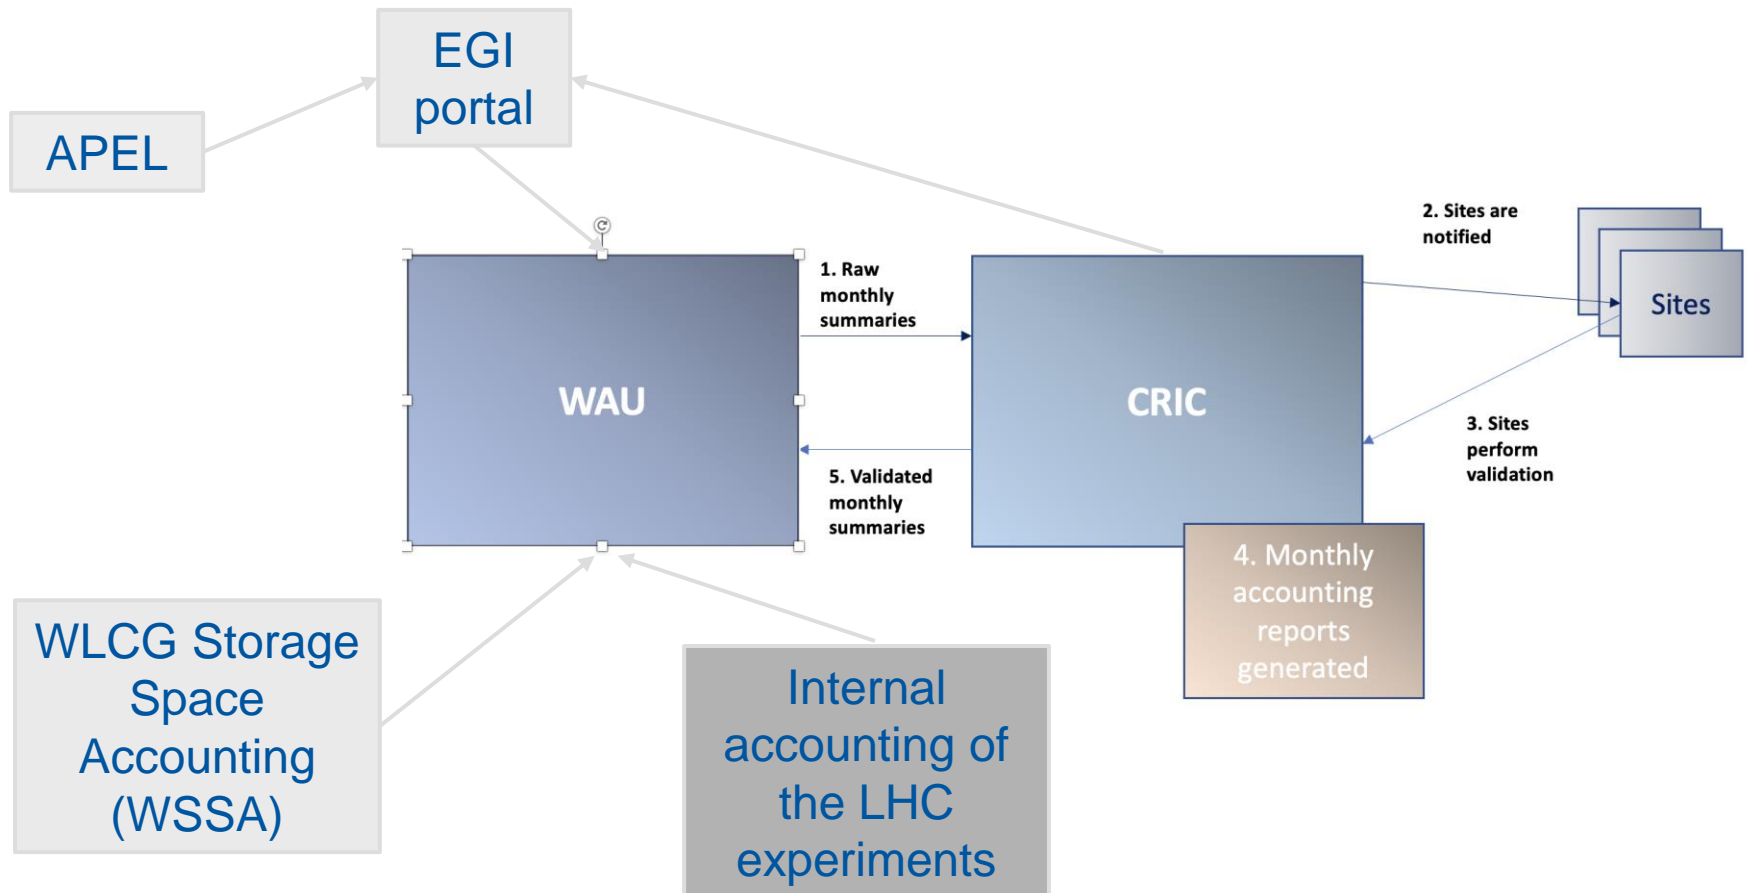
- The goal is to understand whether existing tools and procedures satisfy the needs of the community and to identify possible areas for improvements.
- Input is being collected from the LHC VOs, WLCG Project Office, WLCG Operations Coordination, sites and people who deal with pledge definition procedures.
- The outcome of the analysis of the survey input is used to define the development plans of the involved tools and possible improvements on the procedures

Site survey

Site survey. We asked sites about



- APEL : data quality and support
- EGI Accounting Portal (functionality, usability)
- WLCG Accounting Utility (WAU)
- Accounting validation procedure and related UI
- Monthly accounting reports (data quality and layout)
- Whether WLCG should spend effort on the improvements of the accounting and pledge management tools and procedures. On what exactly
- Any proposals/comments/suggestions related to the topic

WLCG Accounting workflow




Participation

- 23 sites, 5 T1s, 16 T2s, 1 T3

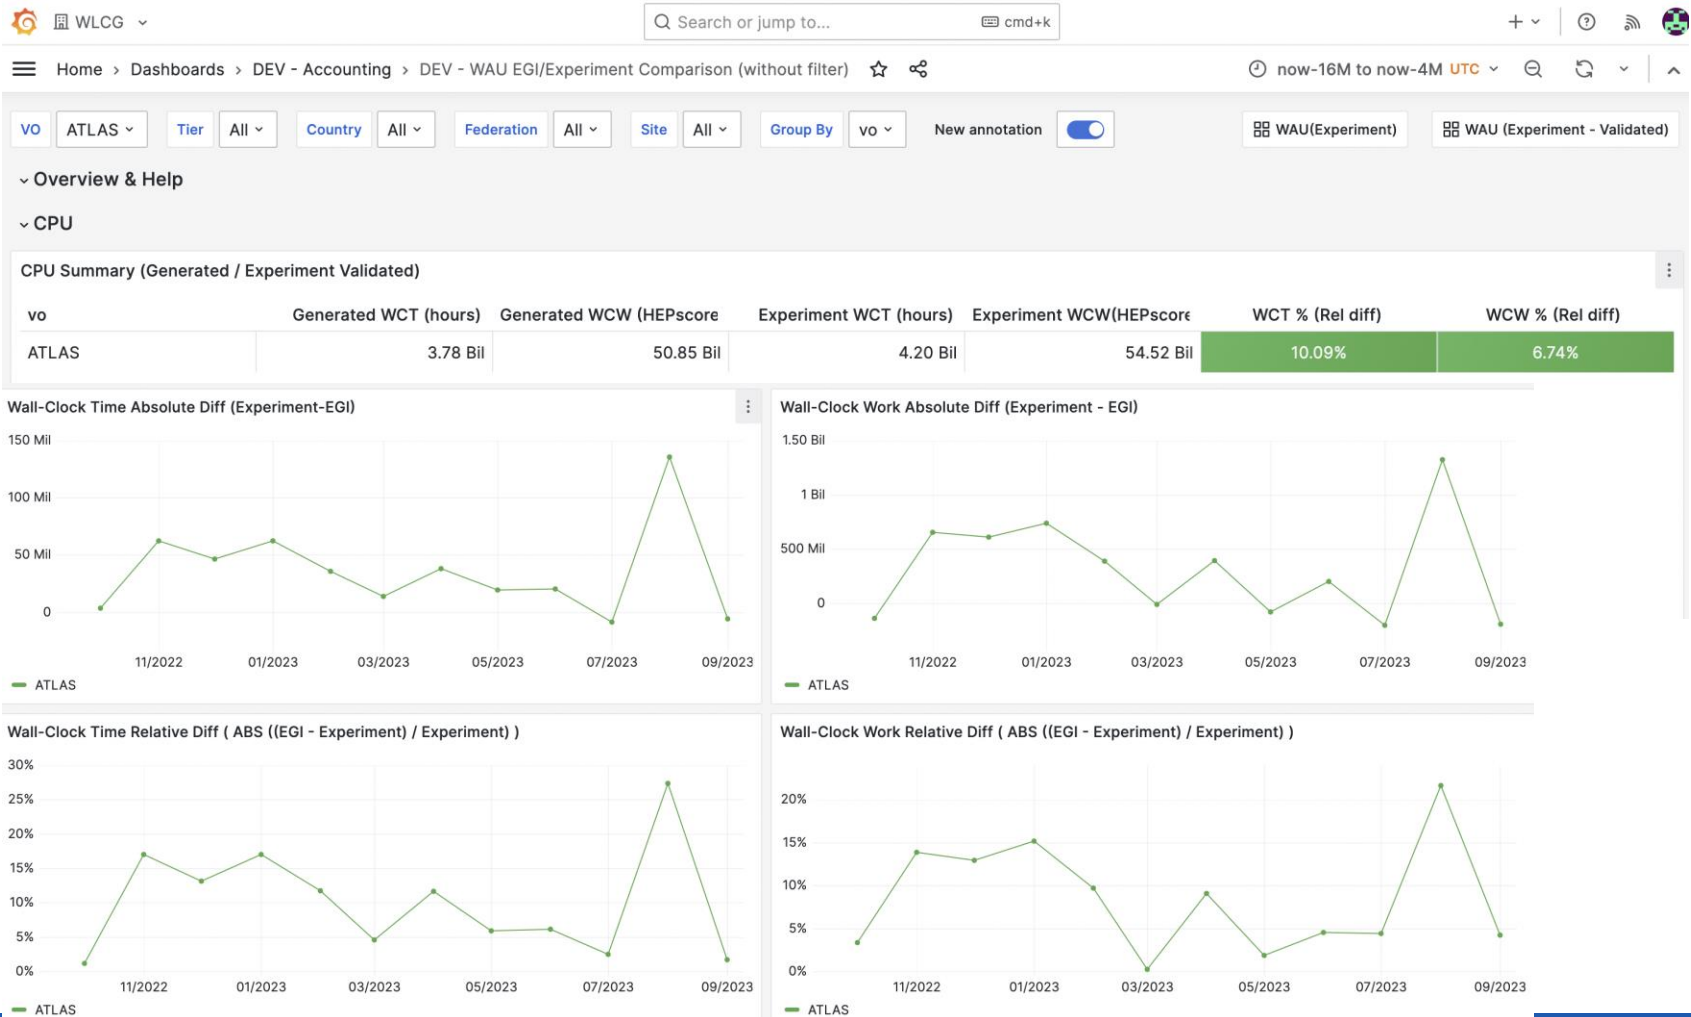


Thanks a lot to everyone who provided their input!

APEL (data quality)




Comparison experiment vs EGI in WAU(1)



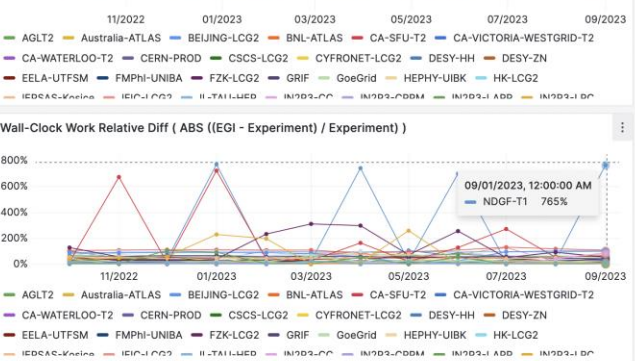
Comparison experiment vs EGI in WAU(2)

CPU Summary (Generated / Experiment Validated)

site	Generated WCT (hours)	Generated WCW (HEPscore)	Experiment WCT (hours)	Experiment WCW(HEPscore)	WCT % (Rel diff)	WCW % (Rel diff)
BNL-ATLAS	226.29 Mil	2.87 Bil	219.81 Mil	2.78 Bil	2.95%	3.18%
CA-SFU-T2	43.67 Mil	912.62 Mil	45.48 Mil	892.02 Mil	3.99%	2.31%
CA-VICTORIA-WESTGRID-...	64.22 Mil	764.17 Mil	56.08 Mil	666.46 Mil	14.51%	14.66%
CA-WATERLOO-T2	20.95 Mil	437.77 Mil	21.03 Mil	412.34 Mil	0.42%	6.17%
CERN-PROD	432.03 Mil	5.40 Bil	515.45 Mil	5.28 Bil	16.18%	2.26%
CSCS-LCG2	70.95 Mil	1.10 Bil	81.47 Mil	1.26 Bil	12.92%	12.57%
CYFRONET-LCG2	62.41 Mil	718.34 Mil	54.89 Mil	654.67 Mil	13.70%	9.73%



II-Clock Time Relative Diff (ABS ((EGI - Experiment) / Experiment))




Wall-Clock Work Relative Diff (ABS ((EGI - Experiment) / Experiment))

!T Experiment Validated (GRID Hours)

WCW Experiment Validated (GRID hours)



APEL (support)




- Good
- Average
- Bad
- No opinion

Some sites expressed concerns about APEL support level. Very long delays for handling tickets, lacking specification and documentation, missing python3 support


EGI accounting portal

Do you use it on regular basis?




Comments and suggestions for site usability. Will follow up with portal developers

Does it satisfy your needs?



WLCG Accounting Utility (WAU)

Do you know about this application?




- Yes
- No
- has been unknown to me so far (link to dashboard without tenant, i.e., have to...
- not accessible without cern account
- "Internal server error" when attempting to connect and "User not authorized /...
- Those links don't work, page not found
- I get "dashboard not found" error for both of these links. For me the correct...

Majority of sites do not know about application. Some of those who tried the link provided in the questionnaire got an authentication problem (requires CERN SSO).


Some of sites questioned its usefulness

Monthly accounting validation


Do you perform validation regularly ?



Do you consider it straightforward?



Do you consider it useful?




Overall positive feedback.
Some suggestions for UI improvements
Several questions why it is needed,
wether it replaces automation?

Why we need validation

- **It does not replace** automatic generation of the accounting data. In case APEL sensor works well at the site and the site has enabled SRR for storage, site does not need to do anything
- It is **extremely useful** in case sites have temporary problems with APEL, or sites which did not enable SRR. It allows to fix data for monthly accounting reports while the problem is being investigated and solved
- Is there a risk that the site provided wrong data from the local accounting? **Yes, such risks exist.** Therefore, WLCG OPS keeps a track of the difference between generated and validated data in WAU and submits tickets against the site if we see constant and high deviation
- If the site does not have local accounting, is validation useless? No, site still has an ability to compare latest data with the previous months to see whether latest metrics do make sense. Though some sites complained that this functionality is missing, it is there and instruction how to do it is included in every monthly validation call
- Since we see that sites do have some questions about validation we clearly need to improve documentation explaining the purpose of validation and provided functionality

Monthly accounting reports


Do you check reports regularly?



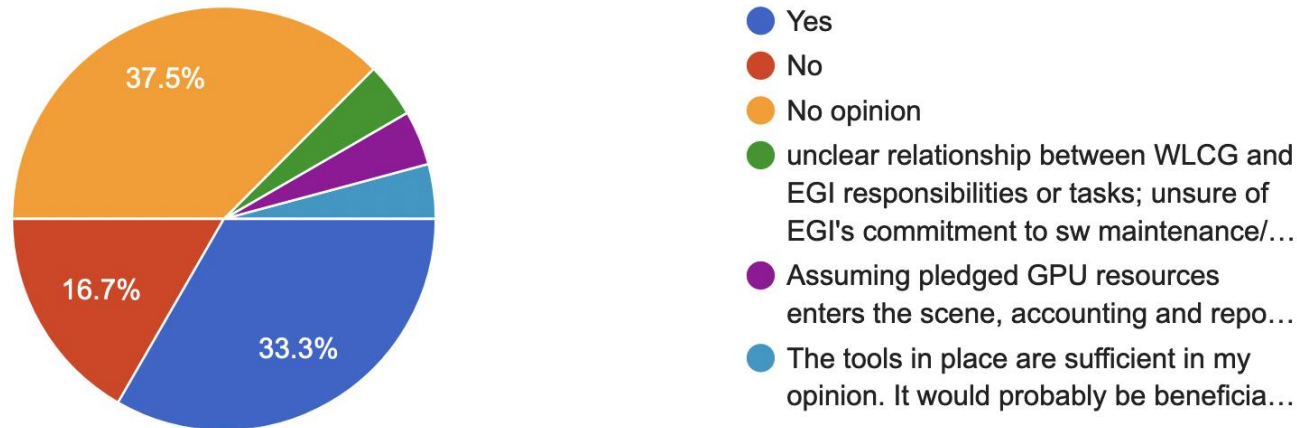
Suggestions for machine parsable retrieval for local integration.

Complains that it is not possible to correct report retroactively (done on purpose)

Are metrics and layout OK?



Does WLCG needs effort to improve accounting infrastructure and what exactly



Main areas mentioned as those required effort:

- Efforts needed to adapt accounting with evolution of authentication procedures (i.e. X509 dismissal, replaced by several token families: wlcg iam tokens, scitokens, egi check-in tokens). Reporting new architectures (GPU, non x86_64 CPUs)
- APEL
- Documentation
- Promotion of the tools developed by various people, to streamline development

VO survey

We asked LHC VOs

- Overall procedure for definition of the VO requirements and yearly pledges
- CRIC UI for pledge and VO requirements definition
- WLCG Accounting tools (APEL, EGI Accounting Portal and WAU) (functionality, usability)
- Monthly accounting reports (data quality and layout)
- Whether WLCG should spend effort on the improvements of the accounting and pledge management tools and procedures. On what exactly
- Any proposals/comments/suggestions related to the topic

Procedure for definition of the VO requirements and yearly pledges

- Overall positive input: "Procedure is well defined"
- Suggestins for improvements:
 - Deadline for pledge submission should not be earlier than the end of September (sometime internal review process of the FA is still ongoing)
 - Communication to the sites of the pledge approval can be improved (now it is FA responsibility)

CRIC UI for pledge and VO requirements definition

- Got a lot of suggestions for additional functionality for CRIC UI for VO requirements and pledge definition:
 - Improvements to the tables for VO requirements and pledge definition,
 - Add plots which would allow to see VO requirements and pledge evolution over years, historical plot showing fraction committed by a given FA, historical plot with VO requirements and pledge evolution by country

WLCG Accounting tools (APEL, EGI Accounting Portal and WAU)

- Though for everyday operations experiments do not rely on WLCG accounting tools, these tools play a role in preparation of the reports submitted to C-RSG and LHCC referees.
- **Concerns:**
 - Due to the fact that central accounting tools are not used in production, according to ATLAS some wrong data in the central accounting tools can stay unnoticed for a while. ATLAS also has concerns regarding data quality. Though application in WAU which allows to compare experiment data with data in WAU/EGI accounting portal shows pretty good consistency (see slides 8 and 9).
 - ATLAS pointed to the CPU efficiency numbers higher than 100%. The reason is that depending on site configuration even a single core job can get more than one core and number of cores can change along job life time. The best approximation for wallclock time which is used in the denominator of the CPU efficiency formula, is runtime of the job multiplies by number of requested cores. If number of cores the job got is higher than the number of requested cores, efficiency can easily go over 100%. There is no good solution for this issue, since the current algorithm is the best we can apply.
 - Experiments (ATLAS, LHCb) would be interested to crosscheck their accounting data with data in the EGI portal. It is not always straightforward (different granularity, not always clear whether all consumption of a given set of resources is reported to the EGI portal, different benchmarking factors (DIRAC) applied). WLCG operations works on WAU application for comparison. Currently it has been enabled for ATLAS (development instance), we plan to enable it for all experiments.
 - Though monthly validation procedure allows to correct faulty or missing generated metrics, validated metrics are not fed back into EGI accounting portal
 - Some complains about EGI portal usability

Monthly accounting reports (data quality and layout)

- Experiments confirmed that they check reports on regular basis
- ALICE suggested to speed up publication of the accounting reports. Current constraints are:
 - Delay in the update of the EGI portal, no point to start validation earlier than 10 days after the end of the month
 - At least 10-15 days are required for site admins to perform validation, and then about one week to follow up by WLCG Ops missing or suspicious data
- ALICE also thinks that readability of the tables for each experiment can be improved. Will clarify concrete suggestions from ALICE
- No comments regarding suggestions for improvements of the accounting reports from other experiments

Whether WLCG should spend effort on the improvements of the accounting and pledge management tools and procedures.

- ALICE: suggestions for improvements of some tables for VO requirements and pledges, and improve readability of experiment tables in the accounting reports
- ATLAS: believes that *'a trustable, user friendly, solid accounting portal would simplify the work done by all the WLCG experiments.'*
- CMS : *no suggestions*
- LHCb: *we would be in favour of a central accounting system, but we understand this requires a lot of development and validation work vs the experiments' accounting systems.*
- All experiments asked for historical plots for VO requirements and pledges in CRIC

Input from WLCG Ops Coordination

Input from WLCG Operations Coordination(1)

- Please, provide your feedback regarding central accounting systems: APEL, [EGI accounting portal](#) and [WAU](#).
 - What about information quality? *According to our crosschecks with the experiment accounting systems , data is rather trustworthy . WAU application which provides functionality to compare ATLAS and CMS data, shows pretty good consistency for ATLAS. We plan to enable comparison for LHCb and ALICE.
 - What about UI / UX?
 - No particular concerns
 - Is there any missing functionality?
 - System is not flexible. Changing/adding any reported metric takes too much effort and too much time
 - More comments
 - Slow processing of tickets and requests
- Do you believe WLCG needs to invest in the improvement of the central accounting systems? Justify your answer, please.
 - - We understand that improving flexibility of APEL and EGI portal would require full redesign and might not be possible given current priorities (enabling accounting for new architectures), however we believe, it is required in a long term
- Please, provide your feedback regarding monthly accounting reports.
 - Does interface for generating of the accounting reports serve your needs and is user friendly?
 - No suggestions for improvements

Input from WLCG Operations Coordination(2)

- Are you checking reports on a regular basis?
 - They are checked by WLCG Ops before generation of the official version which is announced and published in the WLCG repository. We chase all suspicious or missing information contacting the sites. Response from the sites is some time extremely slow which delays report generation
- What about information quality?
 - WLCG Ops tries to follow up on eventual problems, in general we believe it is good
- What about report layout?
 - Fine
- Is anything missing?
- More comments
- Please, provide your input regarding monthly accounting validation
 - Do you think that it did improve the quality of the accounting data?
 - We are convinced that introducing of validation considerably improved data quality. Since monthly validation has been introduced no major long lasting problems. Still we strongly encourage all sites put effort in enabling automatic generation of the accounting metrics, check them during validation and fix automatic generation if it provides wrong results.

Input from WLCG Project Office

Input from WLCG Project Office

- Focused on CRIC functionality for VO requirements and pledge definition
- Main points:
 - Single dedicated interface to enter Y+1 pledges (now there are two slightly different pages which is confusing)
 - Expire CRIC local accounts after 1 year (now they stay active forever)
 - Generate report on the pledge table, now a lot of manual work is required to generate such report
 - Many minor suggestions for UI usability improvements

What is next

- CRIC and WAU are under responsibility of the WLCG Ops Coordination team. Suggestions and requests for these applications from all parties taking part in the review will be prioritised, time required for implementation estimated, working plan agreed. Some of suggestions are already being implemented.
- What concerns APEL and EGI accounting portal the requests will be also documented and discussed in the framework of the WLCG Accounting Task Force with the corresponding development teams.