How we can improve accounting data quality?

Julia Andreeva, CERN

WLCG Accounting Task Force meeting

The problem to be addressed

- After many years of working on the WLCG accounting system , we still face inconsistencies of data measured by different parties (sites, experiments, central LCG accounting systems like APEL and WSSA)
- These inconsistencies are indicators of the fact that accounting data quality can be improved
- Several ways to tackle the problem:
  - More active involvement of site administrators to check monthly accounting data
  - Enable straight forward way to compare data coming from various sources
  - Defining better workflow and responsibilities for chasing faulty data, inconsistencies, etc…

More active involvement of site administrators in data validation

Current workflow:

1). Monthly accounting reports are generated by the EGI portal based on APEL data for CPU, manually injected T1 only data via REBUS UI for disk and tape usage

2). Reports are sent to sites by WLCG project office

3). Sites are supposed to check reports and complain in case of problem. However, data can not been changed in APEL quickly. Investigation of the issue and fixing data can take months

4). Often the problem is not noticed for a long time and is being discovered while preparing RRB report

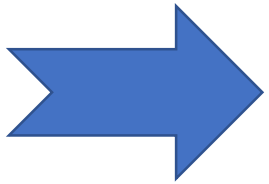

New workflow:

1). Monthly accounting reports will be generated by CRIC **based on data validated by sites.**

2)**.** Sites will get notification with the request to validate auto-generated data. Auto-generated data is coming from WAU (see presentation of Boris) Primary sources :APEL for CPU, WSSA for storage. The validation interface is currently being validated by T1s, next month (November data) , T2 will be included as well

3). Validated data will be pushed from CRIC to WAU. There will be a possibility to see both auto-generated and validated data

4). Inconsistencies (generated vs validated) should be followed up. Need to discuss today, how we go about it.

Accounting data validation workflow

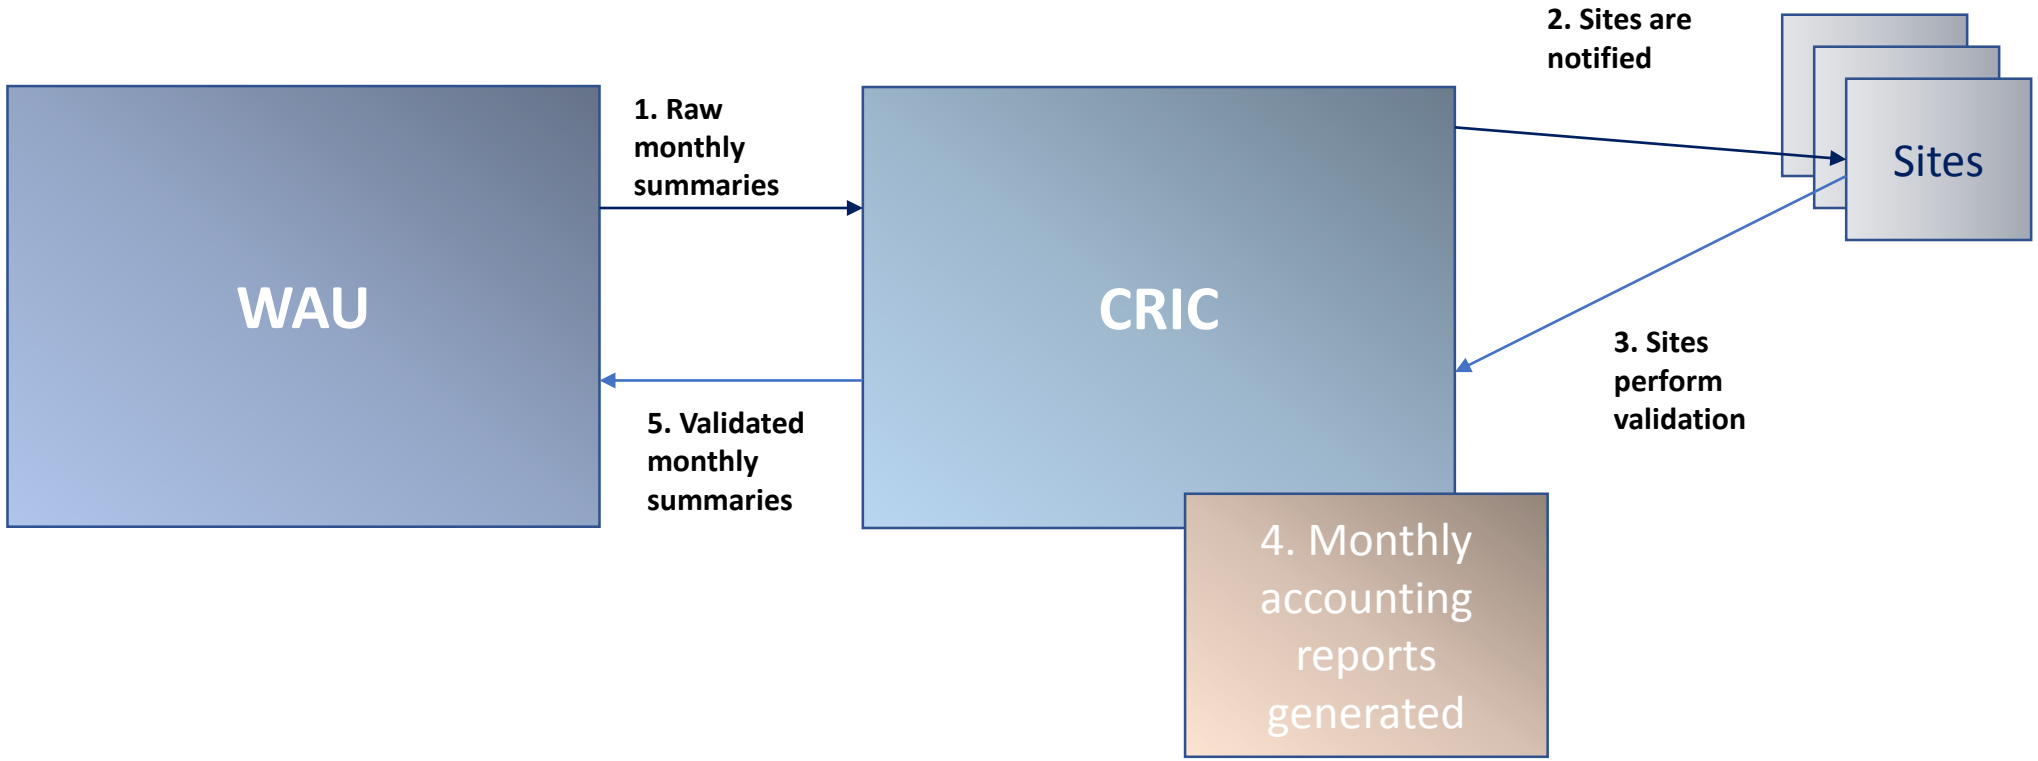

Accounting data validation UI (1)

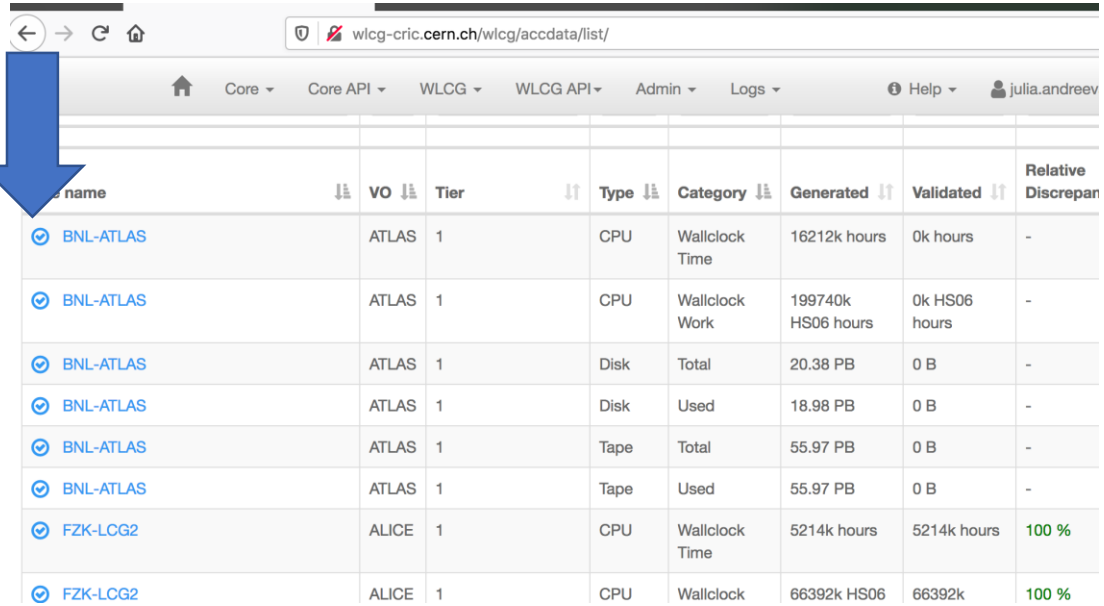

name	VO	Tier	Type	Category	Generated	Validated	Relative Discrepan
BNL-ATLAS	ATLAS	1	CPU	Wallclock Time	16212k hours	0k hours	-
BNL-ATLAS	ATLAS	1	CPU	Wallclock Work	199740k HS06 hours	0k HS06 hours	-
BNL-ATLAS	ATLAS	1	Disk	Total	20.38 PB	0 B	-
BNL-ATLAS	ATLAS	1	Disk	Used	18.98 PB	0 B	-
BNL-ATLAS	ATLAS	1	Tape	Total	55.97 PB	0 B	-
BNL-ATLAS	ATLAS	1	Tape	Used	55.97 PB	0 B	-
FZK-LCG2	ALICE	1	CPU	Wallclock Time	5214k hours	5214k hours	100 %
FZK-LCG2	ALICE	1	CPU	Wallclock	66392k HS06	66392k	100 %

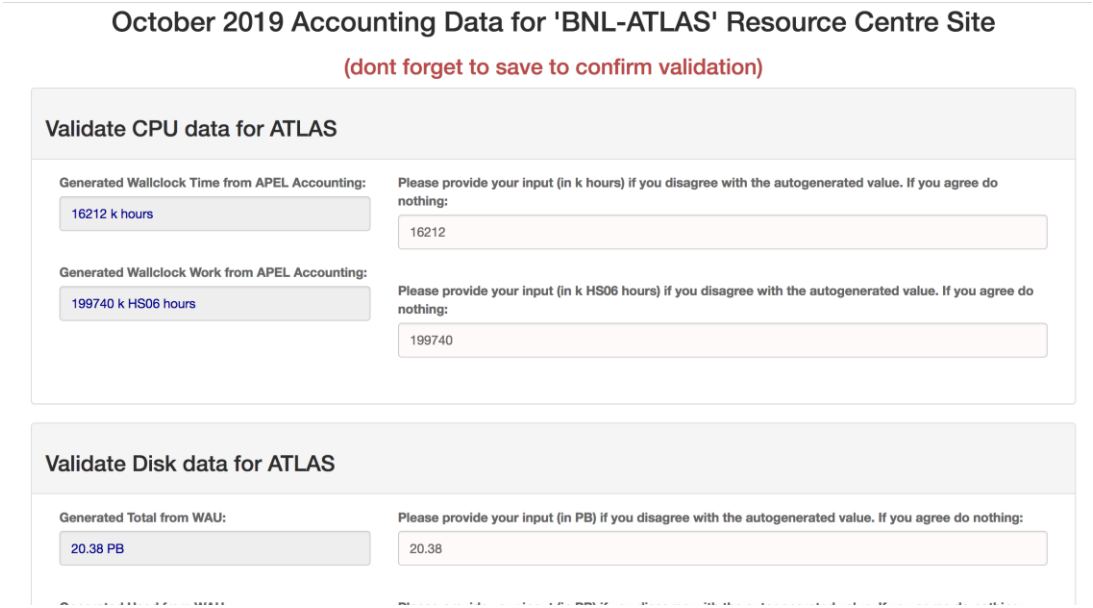

Accounting data validation UI (2)

FZK-LCG2	ALICE	1	CPU	Wallclock Time	5214k hours	5214k hours	100 %	0 hours	2019-10-01
FZK-LCG2	ALICE	1	CPU	Wallclock Work	66392k HS06 hours	66392k HS06 hours	100 %	0 HS06 hours	2019-10-01
FZK-LCG2	ALICE	1	Disk	Total	8.34 PB	8.34 PB	100 %	0 B	2019-10-01
FZK-LCG2	ALICE	1	Disk	Used	7.61 PB	7.61 PB	100 %	0 B	2019-10-01
FZK-LCG2	ALICE	1	Tape	Total	9.94 PB	9.94 PB	100 %	0 B	2019-10-01
FZK-LCG2	ALICE	1	Tape	Used	9.94 PB	9.94 PB	100 %	0 B	2019-10-01
FZK-LCG2	ATLAS	1	CPU	Wallclock Time	7745k hours	7745k hours	100 %	0 hours	2019-10-01
FZK-LCG2	ATLAS	1	CPU	Wallclock Work	98621k HS06 hours	98621k HS06 hours	100 %	0 HS06 hours	2019-10-01
FZK-LCG2	ATLAS	1	Disk	Total	9.83 PB	12.38 PB	125.96 %	-2.55 PB	2019-10-01
FZK-LCG2	ATLAS	1	Disk	Used	9.32 PB	11.83 PB	126.99 %	-2.51 PB	2019-10-01
FZK-LCG2	ATLAS	1	Tape	Total	27.63 PB	27.63 PB	100 %	0 B	2019-10-01
FZK-LCG2	ATLAS	1	Tape	Used	20.9 PB	20.9 PB	100 %	0 B	2019-10-01
FZK-LCG2	CMS	1	CPU	Wallclock Time	4402k hours	4402k hours	100 %	0 hours	2019-10-01
FZK-LCG2	CMS	1	CPU	Wallclock Work	56046k HS06 hours	56046k HS06 hours	100 %	0 HS06 hours	2019-10-01

- User is required to be authorized to edit site level data
- Request of privileges is enabled on the UI and has been tested
- Currently any difference is colored in red. We need to agree on the threshold when we need to follow up

Enable straight forward way to compare data coming from various sources

What we have now

1). Accounting validation application contains:

- APEL monthly summaries (per site /per experiment) retrieved from the EGI portal are
- Experiment monthly summaries, retrieved where possible from the experiment specific systems
- Ratio between the two
- Possibility to see history of comparison

2). Implemented in the SSB framework. Will retire soon since SSB framework is not ported to MONIT

3). Though being useful, was not actively used by the sites.

4). No central effort to check and chase and fix inconsistencies

New scenario

1). For all types of accounting data (CPU, disk and tape usage) WAU will contain data from 3 sources:

- Auto-generated (APEL & WSSA)
- Validated (after validation from CRIC)
- Experiment specific accounting

2). Enable Dashboard to easily spot inconsistencies between all available data sources.

3). Still might not be effective, if there is no agreed workflow of how we follow up on those inconsistencies

Defining better workflow and responsibilities for chasing faulty data, inconsistencies, etc…

- We have no central effort to follow up accounting issues.
- Since accounting metrics are not critical for computing operations , they can stay unnoticed for a long while and are being addressed with low priority
- Can we do better?
- Agreed on the contribution from Olga Kodolova to check monthly comparison reports and create summary of the most problematic sites
- However, we need engagement of the experiment experts

Follow up on inconsistencies (APEL vs experiments)

- Set up twiki page for monthly reports: https://twiki.cern.ch/twiki/bin/view/LCG/DataQualityChecks
- Currently perform exercise only for ATLAS
- Monthly generated excel table attached to the page
- Table with conversion factors comparison (used by APEL vs used by experiments) attached to the page
- Should we rather create a googledoc with excel page data, comments, GGUS tickets, etc…?

	ATLAS					
	HEP-SPEC06 Wall Clock Work			Raw Wall Clock Time		
Site Name	Dashboard	EGI	Ratio	Dashboard	EGI	Ratio
AGLT2	91082286	91339147	100	8310890	8333864	100
ANY						
AUVERGRID						
Australia-ATLAS	23953346	16462268	69	2100922	1447232	69
BEIJING-LCG2	6395065	N/A	N/A	316419	N/A	N/A
BEgrid-ULB-VUB						
BNL-ATLAS	202834779	199740702	98	16462918	16212719	98
BUDAPEST						
BU_ATLAS_Tier2	73338170	67177035	92	6841114	6266514	92
BelGrid-UCL						
CA-SFU-T2	43367256	20379843	47	2980744	975112	33
CA-VICTORIA-WESTGRID-T2	15996361	15899520	99	1510155	1525248	101
CA-WATERLOO-T2	37195668	14646973	39	2558726	700812	27
CBPF						
CERN-PROD	422831976	477484464	113	31042846	39299712	127
CIEMAT-LCG2						
CIT_CMS_T2						
CSCS-LCG2	N/A	33500515	N/A	N/A	3151507	N/A
CYFRONET-LCG2	6652814	6766505	102	468513	476514	102
DESY-HH	83354746	68685030	82	7154587	5958405	83
DESY-ZN	25468067	23535518	92	1433201	1324601	92
EELA-UTFSM	2653553	2713675	102	229439	233709	102
FI_HIP_T2						
FMPhI-UNIBA	5579841	5209708	93	530743	461036	87
FZK-LCG2	132627777	98621516	74	10555936	7745956	73
GLOW						
GR-07-UOI-HEPLAB						
GRIF	80444190	75939504	94	7292756	6882820	94
GoeGrid	11407469	12232014	107	1111178	1191740	107
HK-LCG2	N/A	7289244	N/A	N/A	616520	N/A
Hephy-Vienna						
ICM						
IEPSAS-Kosice	5148105	5057182	98	322577	316866	98
IFCA-LCG2						
IFIC-LCG2	36598669	37042653	101	2822189	3704265	131
IL-TAU-HEP	12323415	8966141	73	1048652	830198	79

Discussion

- How experiment experts are involved and can help in debugging and fixing accounting issues?
- For sites: Alessandro suggested that if we decrease number of checks we require from the sites, there are better chances to get them involved. We might need to think that in the future we create per site monthly report with accounting and availability and require one time action from the site if something is wrong with site data. This is doable, however, will require development effort, since availability and accounting are handled independently
- Other suggestions what we can do better?