

# WLCG Accounting: status, plans and data validation

Julia Andreeva, CERN

GDB 11.12.2019

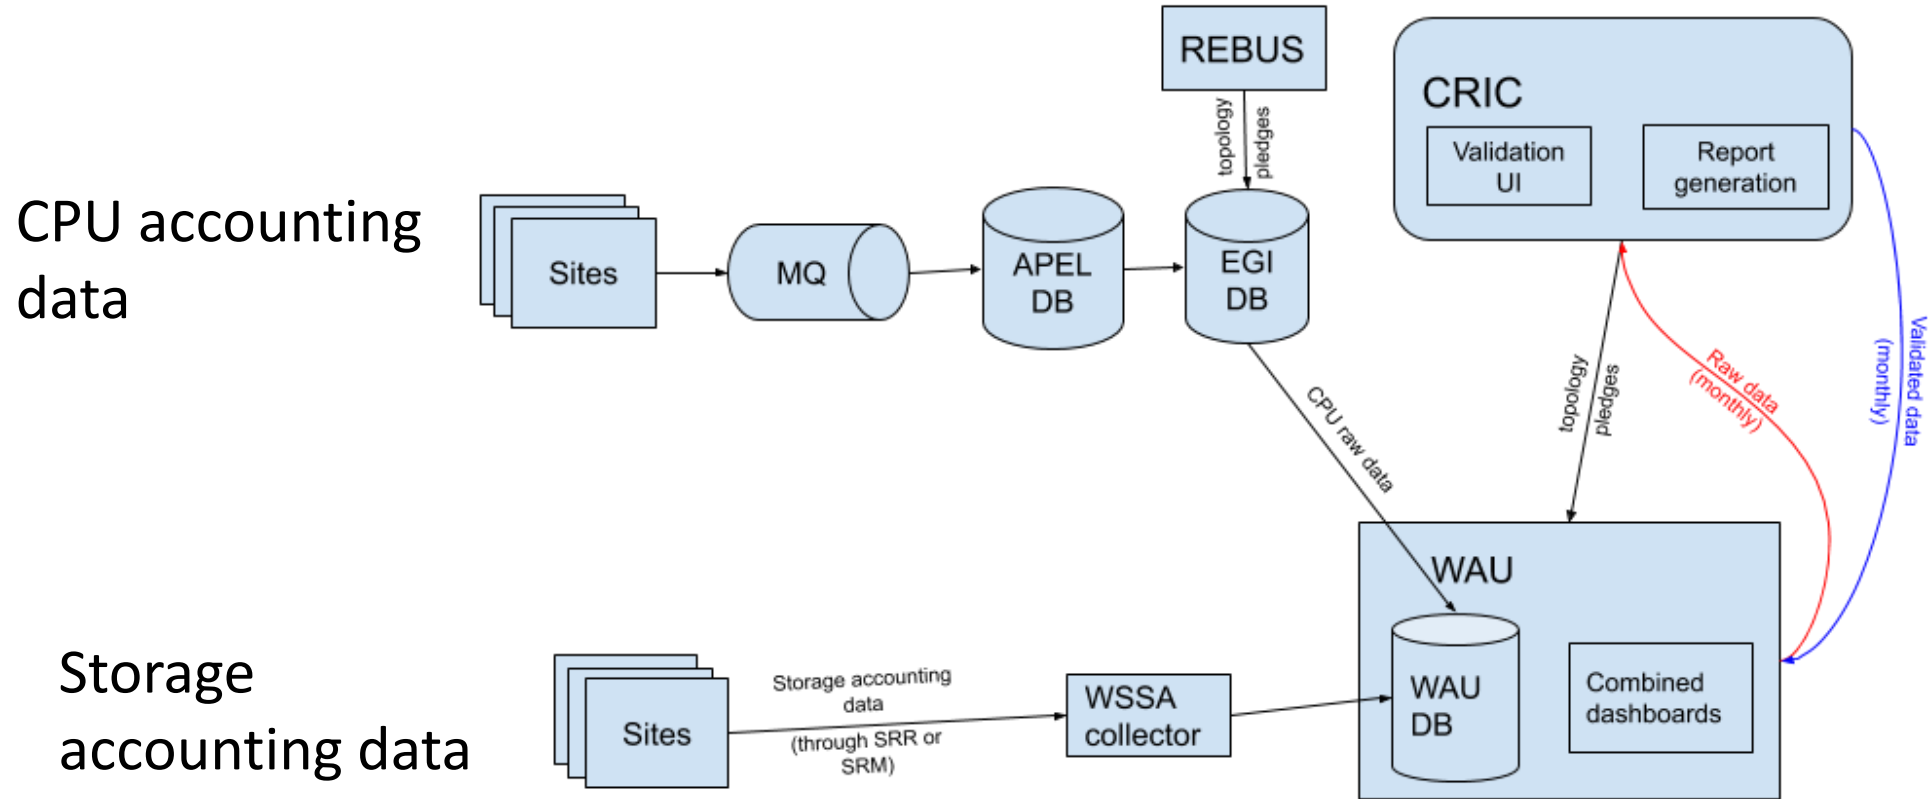
# WLCG accounting improvements

- New WLCG Accounting Utility (WAU) has been deployed in production
- New accounting report generation has been developed in CRIC
- Progressing with SRR deployment at the WLCG sites which would allow to switch storage space accounting from SRM and xrootd queries to SRR
- New accounting data validation workflow has been enabled and tested with T1s
- Got an agreement of how WLCG operations, sites and VOs can work together to improve quality of the accounting data

# WLCG accounting improvements

- **New WLCG Accounting Utility (WAU) has been deployed in production** (Developed by Boris Vasilev and Dimitrios Christidis)
- New accounting report generation has been developed in CRIC
- Progressing with SRR deployment at the WLCG sites which would allow to switch storage space accounting from SRM and xrootd queries to SRR
- New accounting data validation workflow has been enabled and tested with T1s
- Got an agreement of how WLCG operations, sites and VOs can work together to improve quality of the accounting data

# WAW data flow



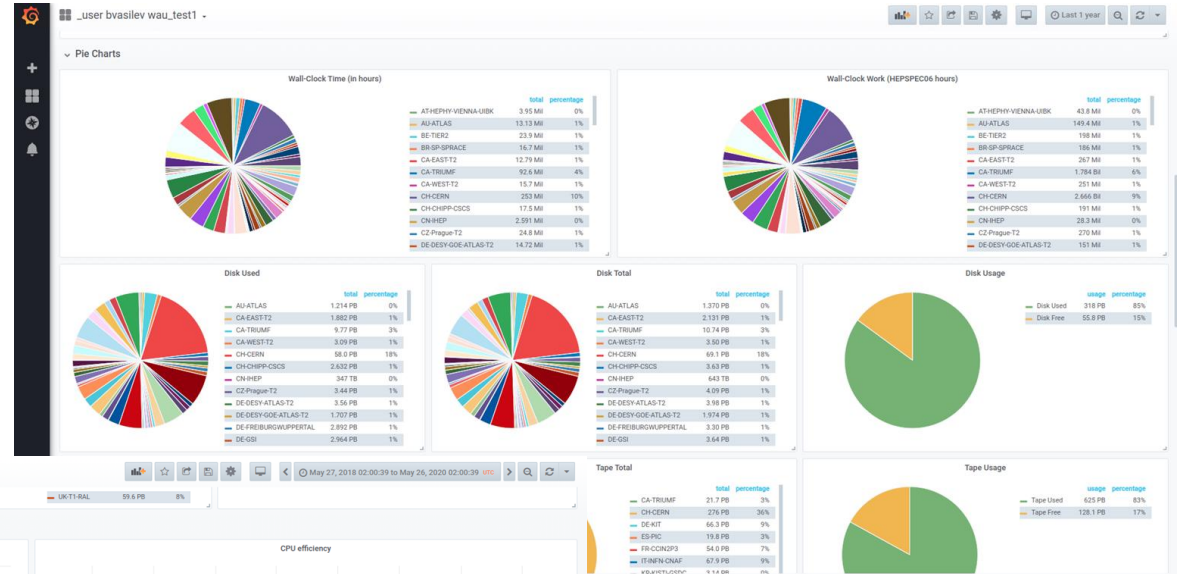
# WAU features

- Foresee 3 main views of WAU Dashboard
  - [Auto-generated data](#) (in production)
  - [Validated data](#) (prototype)
  - Comparison of measurements from different sources – auto-generated, validated, experiment-specific (under development)
- Plan to add accounting data from the experiment-specific systems
- Data can be grouped/filtered by site, federation, country, VO, tier
- Table form, pie charts and historical distributions, comparison with pledges are available

# WAW Dashboard (1)

Summary

federation	WCT (hours)	WCW (HEPSPC66 hours)	Disk Used	Disk Free	Disk Total	Tape Used	Tape Free	Tape Total
AT-HEPHY-VIENNA-UIBK	8 MI	88 MI	-	-	-	-	-	-
AU-ATLAS	13 MI	149 MI	1.21 PB	156.28 TB	1.37 PB	-	-	-
BE-TIER2	48 MI	395 MI	-	-	-	-	-	-
BR-SP-SPRACE	17 MI	186 MI	-	-	-	-	-	-
CA-EAST-T2	13 MI	267 MI	1.88 PB	249.13 TB	2.13 PB	-	-	-
CA-TRUMF	93 MI	2 BI	9.77 PB	970.70 TB	10.74 PB	16.43 PB	5.24 PB	21.67 PB
CA-WEST-T2	31 MI	502 MI	3.09 PB	408.86 TB	3.50 PB	-	-	-
CH-CERN	1 BI	11 BI	58.00 PB	11.15 PB	69.14 PB	247.60 PB	28.35 PB	275.95 PB
CH-CHIPP-CSCS	70 MI	763 MI	2.63 PB	994.37 TB	3.63 PB	-	-	-
CN-HEP	8 MI	85 MI	346.03 TB	296.09 TB	642.73 TB	-	-	-
CZ-Prague-T2	50 MI	540 MI	3.44 PB	651.37 TB	4.09 PB	-	-	-
DE-DESY-ATLAS-T2	-	-	3.56 PB	425.44 TB	3.98 PB	-	-	-
DE-DESY-GOE-ATLAS-T2	15 MI	151 MI	1.71 PB	267.32 TB	1.97 PB	-	-	-
DE-DESY-LHCb	24 MI	425 MI	-	-	-	-	-	-
DE-DESY-RWTH-CMS-T2	178 MI	2 BI	-	-	-	-	-	-
<b>Total</b>	<b>6 BI</b>	<b>66 BI</b>	<b>317.79 PB</b>	<b>55.79 PB</b>				



Many thanks to Adrian Coveney for sharing code to make colour scheme consistent in grafana

# WAU Dashboard (2)

CPU data (from EGI)

Summary

Storage data (from WSSA)


federation ▲	WCT (hours)		WCW (HEPSPEC06 hours)			Disk Used			Disk Free			Disk Total			Tape Used			Tape Free			Tape Total		
CA-TRIUMF	121 Mil	2 Bil				8.76 PB	926.25 TB	9.69 PB				16.15 PB	5.32 PB	21.47 PB									
CH-CERN	2 Bil	18 Bil				56.31 PB	10.45 PB	66.76 PB				242.25 PB	36.78 PB	279.04 PB									
DE-KIT	323 Mil	4 Bil				22.09 PB	2.62 PB	24.70 PB				52.13 PB	13.26 PB	65.39 PB									
ES-PIC	91 Mil	1 Bil				4.01 PB	659.47 TB	4.67 PB				16.26 PB	3.72 PB	19.98 PB									
FR-CCIN2P3	333 Mil	4 Bil				14.06 PB	1.56 PB	15.62 PB				42.44 PB	10.14 PB	52.59 PB									
IT-INFN-CNAF	321 Mil	4 Bil				17.10 PB	4.00 PB	21.10 PB				49.36 PB	18.25 PB	67.61 PB									
KR-KISTI-GSDC	44 Mil	456 Mil				2.54 PB	1.22 PB	3.76 PB				3.14 PB	0 B	3.14 PB									
NDGF	88 Mil	1 Bil				8.91 PB	418.49 TB	9.33 PB				10.03 PB	4.95 PB	14.98 PB									
NL-T1	111 Mil	2 Bil				6.27 PB	831.49 TB	7.11 PB				16.73 PB	3.92 PB	20.65 PB									
NRC-KI-T1	105 Mil	1 Bil				9.82 PB	1.55 PB	11.37 PB				8.55 PB	5.83 PB	14.37 PB									
RU-JINR-T1	90 Mil	1 Bil				-	-	-				7.05 PB	1.93 PB	8.98 PB									
TW-ASGC	40 Mil	502 Mil				4.66 PB	830.17 TB	5.49 PB				-	-	-									
UK-T1-RAL	321 Mil	3 Bil				15.19 PB	2.91 PB	18.10 PB				43.88 PB	15.71 PB	59.59 PB									
US-FNAL-CMS	257 Mil	3 Bil				-	-	-				63.59 PB	16.71 PB	80.30 PB									
US-T1-BNL	229 Mil	3 Bil				17.63 PB	1.92 PB	19.56 PB				47.46 PB	3.00 PB	50.46 PB									

Total

WCT (hours)	WCW (HEPSPEC06 hours)	Disk Used ▲	Disk Free	Disk Total	Tape Used	Tape Free	Tape Total
4 Bil	49 Bil	187.35 PB	29.90 PB	217.25 PB	604.30 PB	134.11 PB	738.41 PB


# WAWU Dashboard (3)

Wall-Clock Time (in hours)




	total	percentage
CA-TRIUMF	120.7 Mil	7%
CH-CERN	432 Mil	26%
DE-KIT	80.7 Mil	5%
ES-PIC	30.2 Mil	2%
FR-CCIN2P3	83.2 Mil	5%
IT-INFN-CNAF	80.3 Mil	5%
KR-KISTI-GSDC	43.7 Mil	3%
NDGF	43.9 Mil	3%
NL-T1	18.5 Mil	1%
NRC-KI-T1	35.0 Mil	2%
RU-JINR-T1	90.3 Mil	5%
TW-ASGC	40.0 Mil	2%

Wall-Clock Work (HEPSPEC06 hours)




	total	percentage
CA-TRIUMF	2.243 Bil	11%
CH-CERN	4.57 Bil	23%
DE-KIT	1.027 Bil	5%
ES-PIC	369 Mil	2%
FR-CCIN2P3	910 Mil	4%
IT-INFN-CNAF	914 Mil	5%
KR-KISTI-GSDC	456 Mil	2%
NDGF	604 Mil	3%
NL-T1	304 Mil	2%
NRC-KI-T1	417 Mil	2%
RU-JINR-T1	1.336 Bil	7%
TW-ASGC	502 Mil	2%

Disk Used




	total	percentage
CA-TRIUMF	8.76 PB	5%
CH-CERN	56.3 PB	30%
DE-KIT	22.1 PB	12%
ES-PIC	4.01 PB	2%
FR-CCIN2P3	14.06 PB	8%
IT-INFN-CNAF	17.1 PB	9%
KR-KISTI-GSDC	2.541 PB	1%
NDGF	8.91 PB	5%
NL-T1	6.27 PB	3%
NRC-KI-T1	9.82 PB	5%
TW-ASGC	4.66 PB	2%
UK-T1-RAL	15.2 PB	8%

Disk Total




	total	percentage
CA-TRIUMF	9.69 PB	4%
CH-CERN	66.8 PB	31%
DE-KIT	24.7 PB	11%
ES-PIC	4.67 PB	2%
FR-CCIN2P3	15.6 PB	7%
IT-INFN-CNAF	21.1 PB	10%
KR-KISTI-GSDC	3.76 PB	2%
NDGF	9.33 PB	4%
NL-T1	7.11 PB	3%
NRC-KI-T1	11.37 PB	5%
TW-ASGC	5.49 PB	3%
UK-T1-RAL	18.1 PB	8%

Disk Usage




	usage	percentage
Disk Used	187 PB	86%
Disk Free	29.9 PB	14%

Tape Used




	used	percentage
CA-TRIUMF	16.1 PB	3%
CH-CERN	242 PB	39%
DE-KIT	52.1 PB	8%
ES-PIC	16.3 PB	3%
FR-CCIN2P3	42.4 PB	7%
IT-INFN-CNAF	49.4 PB	8%
KR-KISTI-GSDC	3.14 PB	1%
NDGF	10.03 PB	2%
NL-T1	16.7 PB	3%
NRC-KI-T1	8.55 PB	1%
RU-JINR-T1	7.05 PB	1%
UK-T1-RAL	43.9 PB	7%

Tape Total



	total	percentage
CA-TRIUMF	21.5 PB	3%
CH-CERN	279 PB	37%
DE-KIT	65.4 PB	9%
ES-PIC	20.0 PB	3%
FR-CCIN2P3	52.6 PB	7%
IT-INFN-CNAF	67.6 PB	9%
KR-KISTI-GSDC	3.14 PB	0%
NDGF	14.98 PB	2%
NL-T1	20.7 PB	3%
NRC-KI-T1	14.37 PB	2%
RU-JINR-T1	8.98 PB	1%
UK-T1-RAL	59.6 PB	8%

Tape Usage




	usage	percentage
Tape Used	604 PB	82%
Tape Free	134.1 PB	18%




# WAU Dashboard (4)


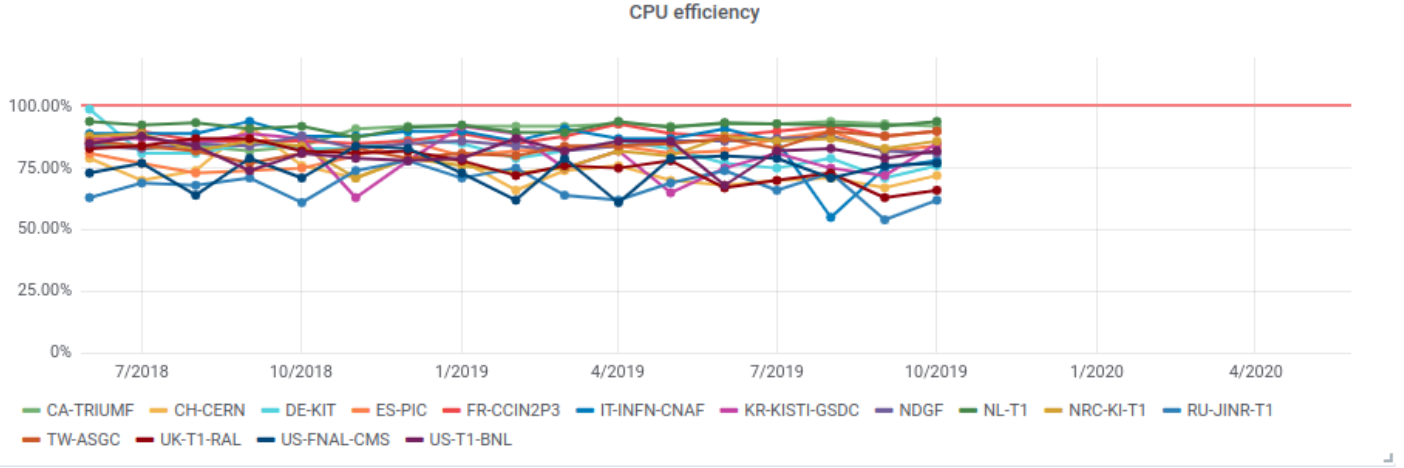
CPU MoU Commitment (HEPSPEC06 hours)



Disk Total



# WAW Dashboard (5)



# WLCG accounting improvements

- New WLCG Accounting Utility (WAU) has been deployed in production
- **New accounting report generation has been developed in CRIC** (Boris Valilev, Panos Paparrigopoulos)
- Progressing with SRR deployment at the WLCG sites which would allow to switch storage space accounting from SRM and xrootd queries to SRR
- New accounting data validation workflow has been enabled and tested with T1s
- Got an agreement of how WLCG operations, sites and VOs can work together to improve quality of the accounting data

# Accounting report generation in CRIC

- CRIC takes over accounting report generation from the EGI portal
- In difference with the current reports include not only CPU consumption , but also disk and tape usage metrics for all sites based on WSSA data
- Use data validated by site admins
- Beginning of the next year will start to send new reports in parallel with the official ones, plan to switch completely to the new ones by the end of spring

# WLCG accounting improvements

- New WLCG Accounting Utility (WAU) has been deployed in production
- New accounting report generation has been developed in CRIC
- **Progressing with SRR deployment at the WLCG sites which would allow to switch storage space accounting from SRM and xrootd queries to SRR**
- New accounting data validation workflow has been enabled and tested with T1s Got an agreement of how WLCG operations, sites and VO's can work together to improve quality of the accounting data

# SRR deployment campaign

- WLCG Operations runs campaign of upgrade WLCG storage to the versions supporting requirements of the DOMA TPC subtask. It also includes enabling Storage Resource Reporting at all WLCG storage
- Started with DPM. 26 sites have already enabled SRR. 20 others are in progress and will do by the end of January next year.
- Next step is dCache migration campaign & SRR enabling. Just started. ~40 sites to go.
- Starting spring next year will proceed with all other storage solutions

# WLCG accounting improvements

- New WLCG Accounting Utility (WAU) has been deployed in production
- New accounting report generation has been developed in CRIC
- Progressing with SRR deployment at the WLCG sites which would allow to switch storage space accounting from SRM and xrootd queries to SRR
- **New accounting data validation workflow has been enabled and tested with T1s** (Panos Paparrigopoulos, Boris Vasilev, Dimitrios Christidis)
- **Got an agreement of how WLCG operations, sites and VOs can work together to improve quality of the accounting data**

# How we can improve accounting data quality?


- 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:

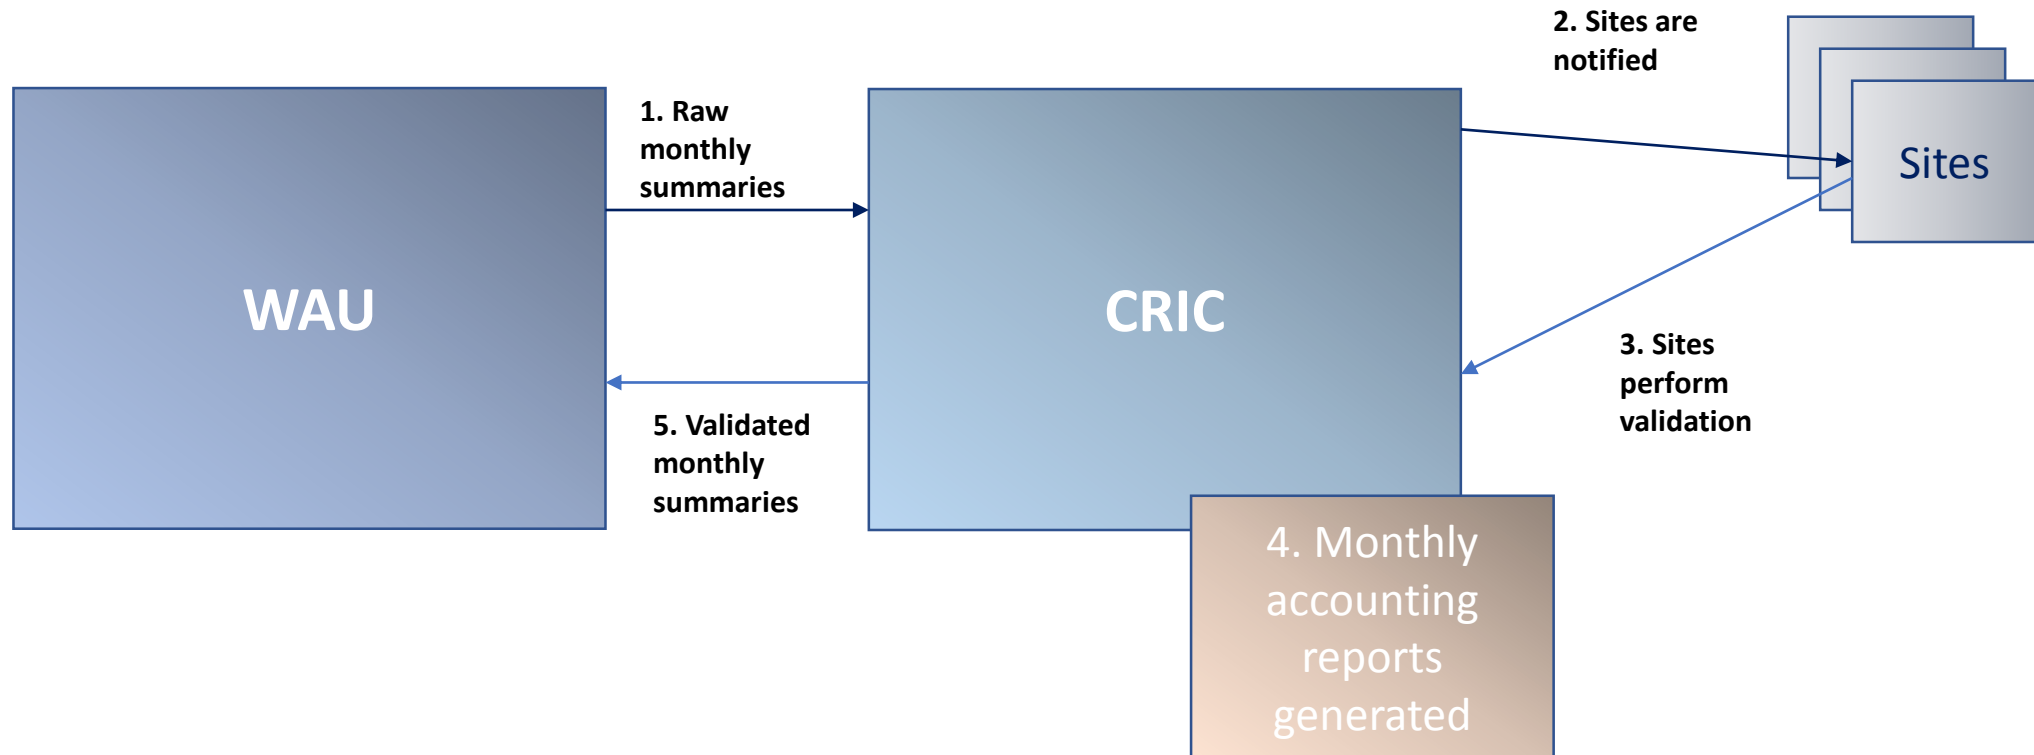
- 1). Monthly accounting reports are generated by the EGI portal based on APEL data for CPU, manually injected T1 only data via REBUS UI for disk and tape usage
- 2). Reports are sent to sites by WLCG project office
- 3). Sites are supposed to check reports and complain in case of problem. However, data can not be 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



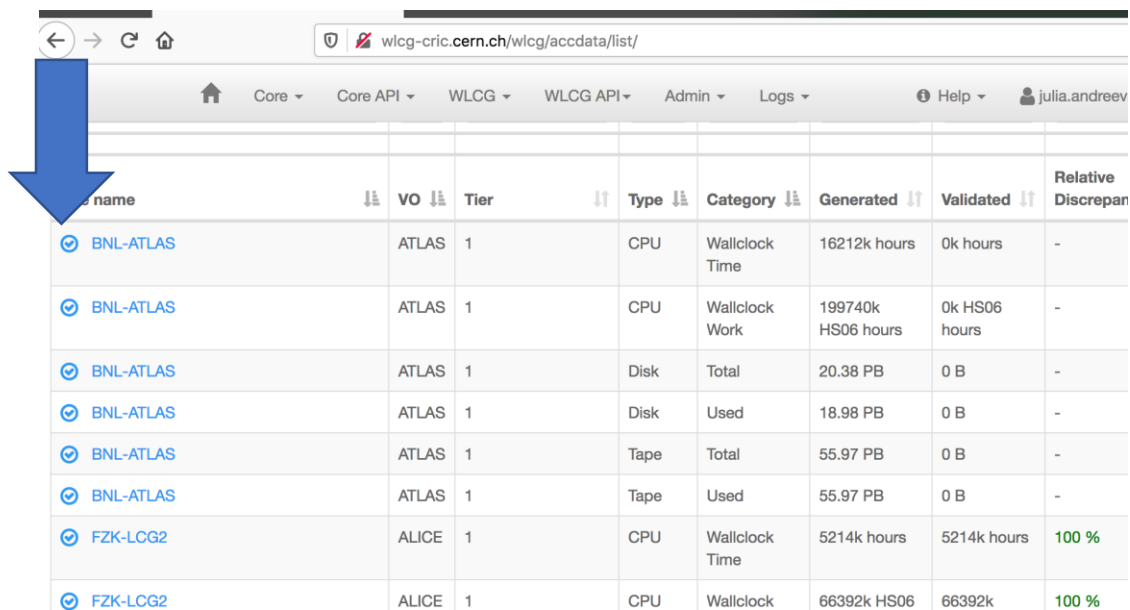
## New workflow:

- 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 %

## October 2019 Accounting Data for 'BNL-ATLAS' Resource Centre Site

(dont forget to save to confirm validation)

### Validate CPU data for ATLAS

Generated Wallclock Time from APEL Accounting:

16212 k hours

Please provide your input (in k hours) if you disagree with the autogenerated value. If you agree do nothing:

16212

Generated Wallclock Work from APEL Accounting:

199740 k HS06 hours

Please provide your input (in k HS06 hours) if you disagree with the autogenerated value. If you agree do nothing:

199740

### Validate Disk data for ATLAS

Generated Total from WAU:

20.38 PB

Please provide your input (in PB) if you disagree with the autogenerated value. If you agree do nothing:

20.38

# 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
- Who/when/what info is logged and can be used for debugging of eventual problems
- Currently any difference is colored in red.
- Relative difference more than 25% should be followed 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...

- So far we did not have central effort to follow up on accounting issues.
- Since accounting metrics are not critical for computing operations , problems can stay unnoticed for a long while and are being addressed with low priority
- Can we do better?
- Recently agreed on the contribution from Olga Kodolova from MSU to check monthly comparison reports and create summary of the most problematic sites
- **At the last WLCG Accounting Task Force meeting got an agreement with all 4 experiments that VO experts would follow up on issues of the monthly summaries**

# Next steps

- Validate November data including also T2s in the loop
- Finalize accounting report generation and start sending new reports along with the official ones starting early next year
- Enable import of the experiment-specific accounting data to WAU and CRIC.
- Enable WAU view to compare accounting metrics from various sources (auto-generated, validated by site admins, experiment-specific)
- Improve notification for site admins inviting them for accounting data validation (direct link for a given site form)
- Setup a new procedure for regular monthly checks of the accounting metrics and following up on eventual problems