



Accounting and monitoring infrastructure for Distributed Computing in the ATLAS experiment

Aleksandr Alekseev (Institute for System Programming of the RAS), Dario Barberis
(University and INFN Genova), Thomas Beermann (Wuppertal University)

9th International Conference "Distributed Computing and
Grid-technologies in Science and Education" (GRID 2021),
5-9 July 2021, Dubna, Russia

Outline

- Introduction
- CERN MonIT Unified Monitoring Infrastructure (UMA)
- ATLAS Monitoring production dashboards in UMA
- Jobs Dashboards
- DDM Dashboards
- Sites Dashboards
- Summary

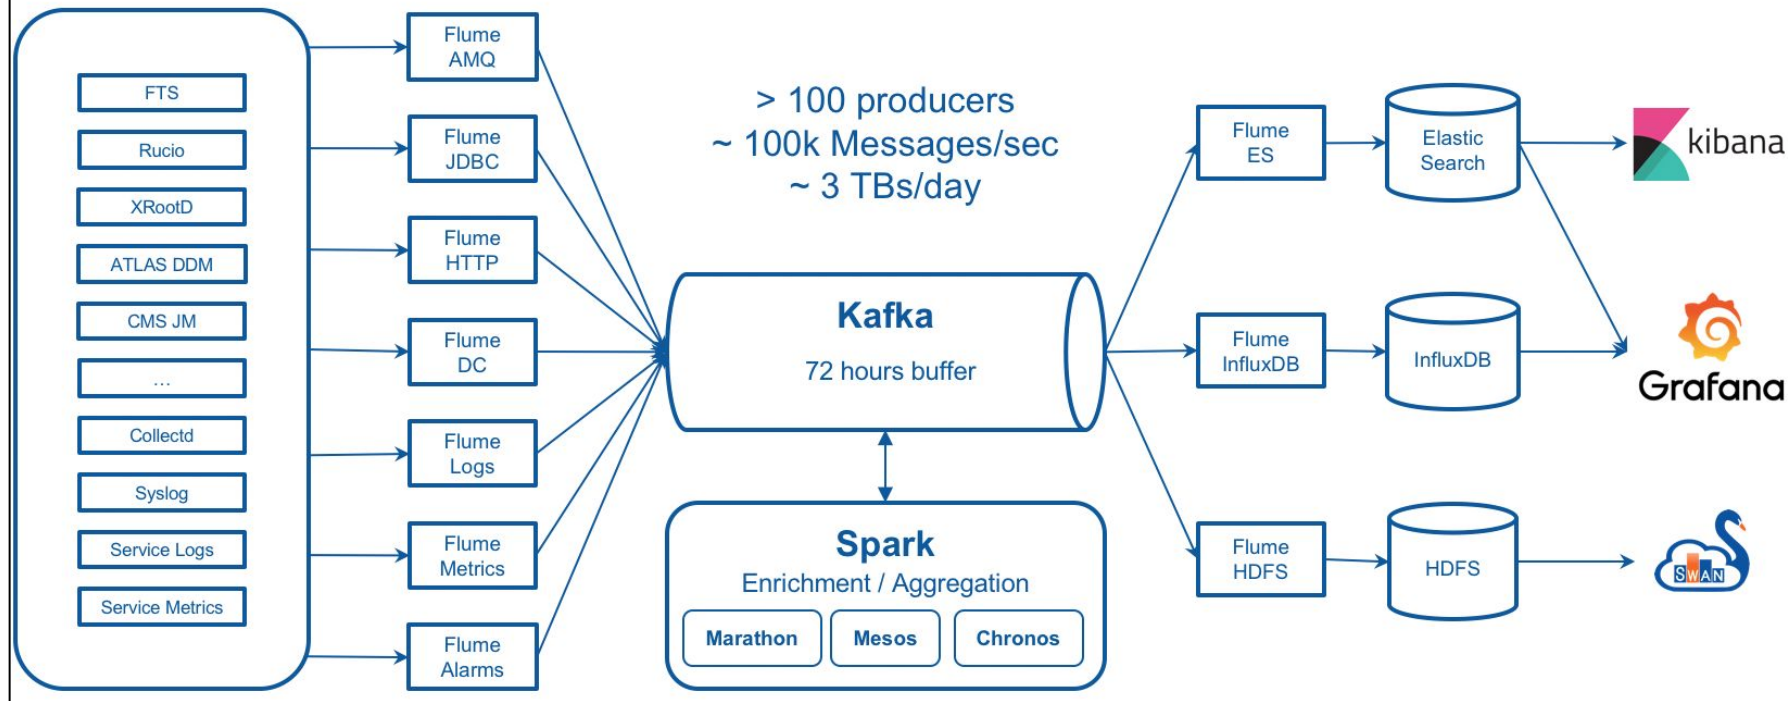
Introduction

- The old monitoring and accounting infrastructure based on custom frameworks was developed by CERN IT and ADC more than 10 years ago to monitor workload management (PanDA) and distributed data management (Rucio) systems
- The old infrastructure had significant limitations in performance, scalability and data visualization
- Since 2016 CERN-IT MonIT group started developing a new unified infrastructure and environment for monitoring and accounting applications based on modern and efficient open-source software stack
- The stack is adapted for the ATLAS experiment and allows the development of dedicated monitoring and accounting applications in Grafana and Kibana visualization environments
- The current state of the monitoring and accounting infrastructure are presented in this talk

CERN MonIT Unified Monitoring Infrastructure (UMA)

Sources > Transport > (Processing) > Storage > Access

[Link](#)



UMA overview

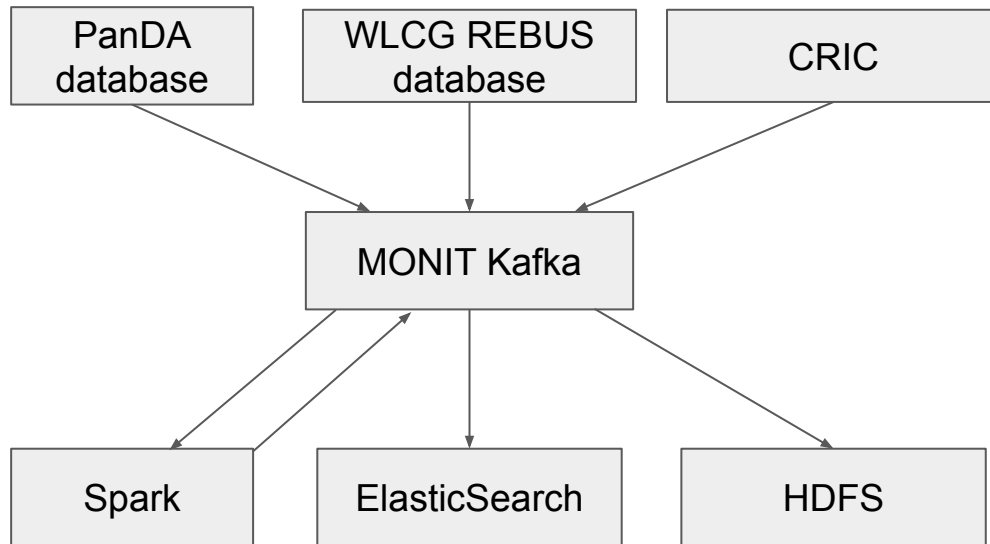
- Data Sources
 - PanDA database
 - Rucio database
 - Rucio trace server
 - Pilot
 - CRIC API
- Collectors
 - ActiveMQ (AMQ)
 - HTTP
 - Java Database Connectivity (JDBC)
- Storages
 - Hadoop File System (HDFS) for long-term data archival and offline analytics
 - Elasticsearch for short-term and long-term data
 - InfluxDB for the medium-term data
- Data access
 - Grafana: access to all data in the MONIT infrastructure
 - Kibana
 - Short-term: access to data for operations
 - Long-term: access to data for accounting
 - Service for Web based Analysis (SWAN)

ATLAS Monitoring production dashboards in UMA

- Jobs monitoring and accounting
 - Job Accounting
 - Jobs Monitoring (Monit Elasticsearch)
 - HS06 Reports
- DDM transfers and accounting
 - DDM Global Accounting (Historical)
 - DDM Global Accounting (Snapshot)
 - DDM Site Accounting
 - DDM Transfers
 - DDM Transfers (Historical data)
- Sites Monitoring
 - Site-oriented dashboard
 - Site Status Board Overview
 - Site Status Board - Panda Jobs
 - Site Status Board - SAM3

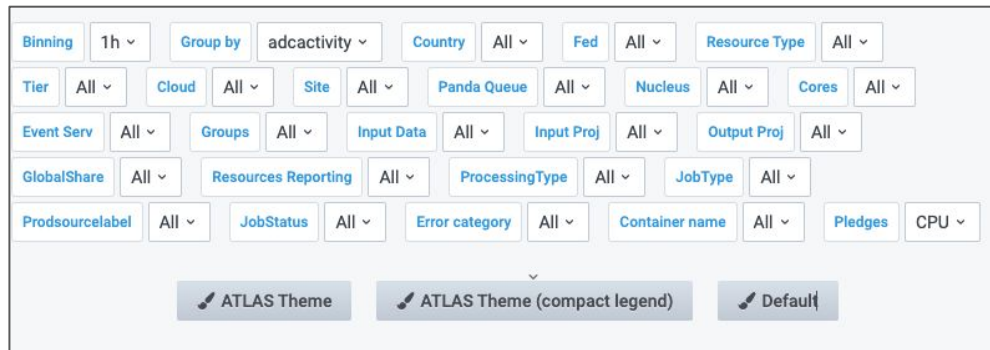
Jobs Dashboard Overview

- Flume JDBC extracts information from jobs tables in PanDA database every ten minutes
- Pledge information from the WLCG REBUS database
- Topology Information is taken from CRIC
- SPARK job for processing data consist of the following steps:
 - In the first step, job computes additional information for each job (ADC activities, error messages, walltime)
 - In the next step, data is enriched with topology information from CRIC
 - In the last step, processed data is written back to Kafka
- Kafka keep aggregated data into five separate indices for each jobs statuses (completed, finilising, pending, running and submitted jobs) in the ElasticSearch
- Aggregated information about jobs also keep into HDFS storage




Jobs Accounting dashboard


- The job accounting dashboard is used by shifters, experts and management to spot problems with the workflow management system
- Special ATLAS styles
- Information about jobs is available since 2010
- 87 plots, 25 filters, 23 options to group data
- Binning: 1h, 1d, 1w, 1M



Efficiency based on success/all accomplished jobs
(by computingsite)



Slots of Running jobs HS06
(by adactivity)




Jobs Monitoring dashboard

- The dashboard provides extended information about completed jobs for 2 months
- It is useful to monitor separate job(s), task(s) or request(s)
- The data is processed in the same as for Jobs accounting dashboard
- Kafka keeps aggregated data into dedicated index for completed jobs in the ElasticSearch storage
- 21 plots, 24 options to group data, 27 filters
- Binning: 10m, 30m, 1h, 6h, 12h, 1d, 7d, 14d, 30d


Jobs list

data.pandauid	data.jedidtaskid	data.reqid	data.statechangetime	data.jobstatus	data.computingsite	data.cpuconsumptiontime
5100926126	25917255	37930	2021-06-28 14:08:30	finished	OU_OSCER_ATLAS	214331
5100926121	25917255	37930	2021-06-28 14:08:29	finished	OU_OSCER_ATLAS	219882
5100926111	25917255	37930	2021-06-28 14:08:28	finished	OU_OSCER_ATLAS	204202
5100926105	25917255	37930	2021-06-28 14:08:28	finished	OU_OSCER_ATLAS	219196

Panda Failures by ExitCode



Number of Successful and Failed jobs



HS06 Reports

- The dashboard is used by management to generate HS06 Reports
- The dashboard uses aggregated data from ElasticSearch index for completed jobs (job accounting dashboard)
- Information is split into separate tables: for Computing sites, Sites, Federations, Tiers
- Data from the tables can be exported to CSV reports using Grafana API and python script

Computing sites

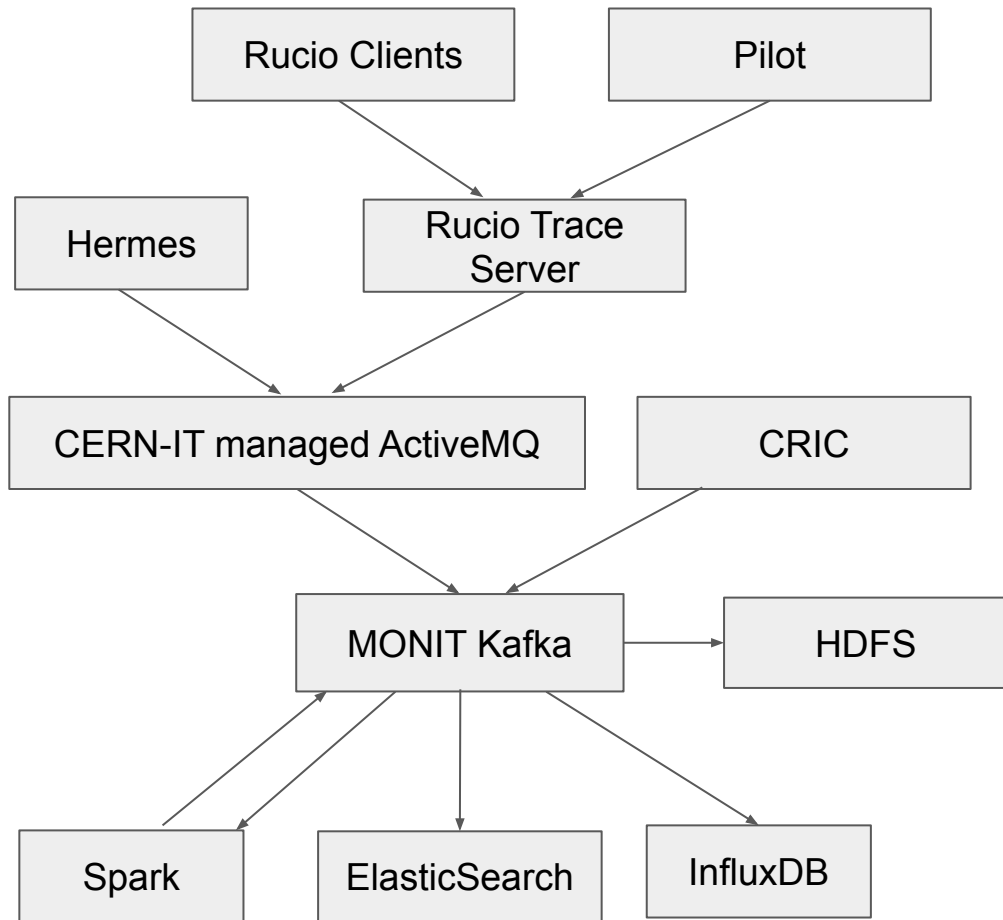
Computingsite	Site	Federation	Tier	Resource	HS06_sec	cpu_in_sec	wc_in_sec	ratio ^
NERSC_Cori_p2_ES_Test	LBNL_DSD_ITB	NON-MOU-Federation	3	hpc_special	2487922770	0	1037807674	0
CERN-EXTENSION_KUBERNETES_TEST	CERN-EXTENSION	CH-CERN	3	cloud	156640	17	15664	0.109
LUNARC	SE-SNIC-T2	SE-SNIC-T2	2	GRID	255859670	430365	39053375	1.10
UIO_CLOUD_LOPRI	NDGF-T1	NDGF	1	cloud	5124628225	28314702	823865456	3.44

Tiers

Tier	HS06_sec	cpu_in_sec	wc_in_sec	ratio ^
3	3227395911963	223121109837	373193368651	59.8
2	11195514192460	650800087912	867061161428	75.1
0	1163183490625	100947515960	129452327052	78.0
1	3571252247759	232805786536	271874012296	85.6
-1	0	103474	103657	99.8

DDM Dashboard Overview

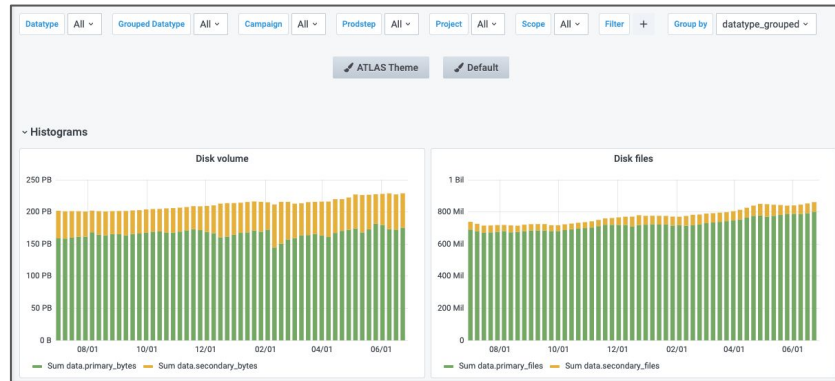
- MONIT infrastructure collects events and traces and injects them into Kafka pipeline
- Furthermore, the infrastructure also regularly reads topology information from CRIC and add them to Kafka
- SPARK job for processing data consist of the following steps:
 - In the first step of data processing. SPARK job transforms the events and traces so they fit in a similar structure
 - In the next step they are enriched with topology information from CRIC
 - In last step everything is aggregated in 1 minute bin by source/destination, activity, cloud, federation, country, tier, etc. and the number of failed and successful transfers are counted
- This data is written back to Kafka and then stored in Elasticsearch/InfluxDB/HDFS



DDM Global accounting dashboards

- Historical dashboard
 - Dataset based. Counts for each dataset how many file replicas are available split by primary/secondary/tier/disk/tape/dataset metadata
 - Shows the evolution of replicas over time
 - 13 plots, 6 filters, 5 options to group data
- Snapshot dashboard
 - Same as the Historical dashboard but shows the numbers for the last week
 - 13 plots + tables from overview part, 6 filters, 5 options to group data
- The data for these dashboards is aggregated once a week

DDM Global accounting (Historical)



DDM Global accounting (Snapshot)


~ Overview

Logical Size	Disk Size	Primary Size	Disk Files	Primary Files
323 PB	229 PB	176 PB	861141756	801006681
Number of Events	Tape Size	Secondary Size	Tape Files	Secondary Files
2422057168088	269 PB	53.4 PB	209308738	60135075
Total Size		Total Files		
498 PB		1070450494		


DDM Site accounting dashboard

- Interesting for operations and site admins
- RSE (Rucio Storage Element) based
- Counts the number of replicas for each endpoint
- Can be aggregated/split by topology/dataset metadata
- For example show volume/number of files for all primary DAOD datasets on US Tier-2 DATADISK
- 233 plots, 17 filters, 18 options to group data


Volume per tombstone



Volume by site



Files by site




DDM transfers dashboards

- The DDM transfer dashboards are mainly used by shifters and experts to monitor the worldwide transfer on the grid and spot problems
- The main dashboard has a granularity of 1 minute for 30 days
- It is using both InfluxDB and Elasticsearch at the same time. The time-series information in the plots is coming from IDB. The transfer details are coming from Elasticsearch
- 35 plots, 18 filters, 17 options to group data
- Binning: 10m, 1h, 6h, 12h, 1d
- Historical transfer dashboard has a minimum granularity of 1h for up to 5 years
- The historical dashboard contains: 39 plots, 16 filters, 23 options to group data. Binning the same

Efficiency (by cloud)

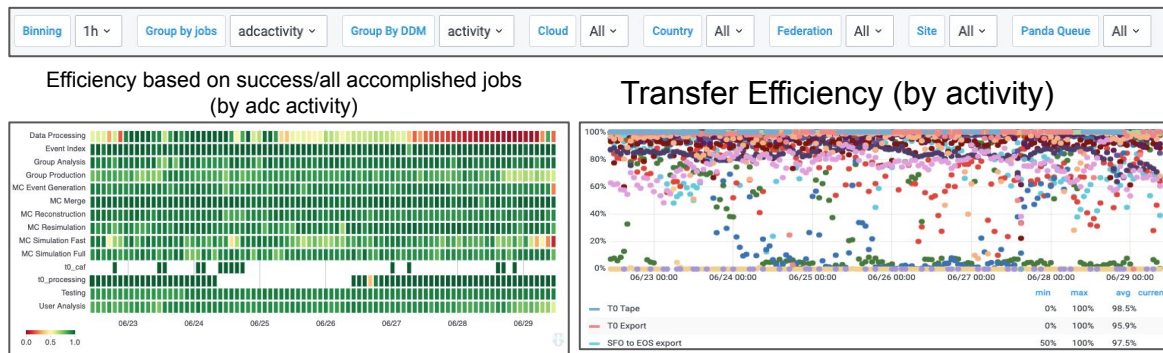
	CA	CERN	DE	ES	FR	IT	ND	NL	RU	TW	UK	US
CA	98%	92%	90%	83%	95%	100%	58%	100%	97%	99%	95%	63%
CERN	94%	99%	93%	100%	99%	100%	100%	100%	100%	100%	98%	98%
DE	87%	94%	63%	75%	96%	91%	97%	96%	96%	93%	95%	53%
ES	76%	85%	70%	90%	93%	99%	89%	99%	99%	96%	92%	76%
FR	91%	96%	84%	78%	92%	97%	94%	94%	96%	94%	91%	78%
IT	83%	99%	91%	98%	98%	99%	100%	100%	100%	100%	96%	88%
ND	86%	99%	74%	63%	97%	94%	98%	100%	100%	98%	99%	95%
NL	99%	100%	80%	99%	98%	100%	99%	100%	100%	100%	96%	75%
RU	100%	100%	92%	100%	100%	100%	100%	100%	100%	100%	94%	70%
TW	98%	92%	85%	44%	97%	65%	75%	99%	100%	-	81%	75%
UK	87%	80%	69%	72%	95%	94%	93%	98%	94%	98%	90%	63%
US	78%	85%	88%	68%	95%	91%	97%	88%	97%	84%	96%	89%

Transfer Efficiency (by dst site)



Site-oriented dashboard

- This dashboard combines information from Jobs Accounting and DDM transfer dashboards
- Allows to monitor and analyse computing sites efficiency
- 8 plots (6 Jobs + 2 DDM), 5 filters, 55 options to group data



CPU Efficiency of good jobs

Group By	CPU consumption	Walltime	CPU Efficiency
FMPH-UNIBA	308.86 Mil	310.63 Mil	0.99
GRIF-LAL	1.24 Bil	1.27 Bil	0.98
RU-Protvino-IHEP	512.60 Mil	526.56 Mil	0.97
UKI-SCOTGRID-GLASGOW	5.15 Bil	5.30 Bil	0.97
WEIZMANN-LCG2	793.15 Mil	815.84 Mil	0.97
IL-TAU-HEP	292.77 Mil	303.54 Mil	0.96
UTA_SWT2	1.23 Bil	1.27 Bil	0.96
UKI-LT2-IC-HEP	1.63 Bil	1.70 Bil	0.96

Site Status Board Overview

Tier		1 + 2 ▾		Country	All ▾	Cloud	All ▾	Federation	All ▾	Site	All ▾			Choose Panel	Fast						
		CRIC Downtime Status		DDM Downtime Status		DDM Transfer Efficiency		SAM3 Site Availability		Panda Queues Status		Jobs Efficiency		Frontier Squid Status		GGUS tickets					
						source	destination					analysis	production								
CP	AGLT2	CP	active	CP	online	CP	97.5%	CP	96.6%	CP	98.6%	CP	online	CP	82.0%	CP	83.0%	CP	ok	CP	1
CP	Australia-ATLAS	CP	active	CP	online	CP	91.7%	CP	100.0%	CP	98.6%	CP	online	CP	72.0%	CP	98.0%	CP	ok	CP	0
CP	BEIJING-LCG2	CP	active	CP	online	CP	94.3%	CP	78.1%	CP	98.6%	CP	online	CP	0%	CP	99.0%	CP	ok	CP	0
CP	BNL-ATLAS	CP	active	CP	online	CP	93.1%	CP	98.5%	CP	98.6%	CP	online	CP	82.0%	CP	80.0%	CP	ok	CP	0
CP	BNLLAKE	CP	active	CP	online	CP	100.0%	CP	100.0%	CP	98.6%	CP	online	CP	94.0%	CP	95.0%	CP	ok	CP	n/a
CP	BU_ATLAS_Tier2	CP	active	CP	online	CP	98.2%	CP	99.4%	CP	86.3%	CP	online	CP	83.0%	CP	97.0%	CP	ok	CP	0
CP	CA-SFU-T2	CP	active	CP	online	CP	97.8%	CP	97.9%	CP	98.6%	CP	online	CP	75.0%	CP	97.0%	CP	ok	CP	1
CP	CA-VICTORIA-WESTGRID-T2	CP	active	CP	online	CP	89.7%	CP	95.2%	CP	98.6%	CP	online	CP	63.0%	CP	98.0%	CP	ok	CP	0
CP	CA-WATERLOO-T2	CP	active	CP	online	CP	99.2%	CP	98.3%	CP	98.6%	CP	online	CP	88.0%	CP	91.0%	CP	degraded	CP	1
CP	CSCS-LCG2	CP	active	CP	online	CP	94.1%	CP	84.9%	CP	98.6%	CP	online	CP	87.0%	CP	88.0%	CP	ok	CP	0
CP	CYFRONET-LCG2	CP	active	CP	online	CP	97.1%	CP	96.3%	CP	98.6%	CP	online	CP	99.0%	CP	93.0%	CP	ok	CP	0
CP	DESY-HH	CP	active	CP	online	CP	99.3%	CP	95.0%	CP	98.6%	CP	online	CP	82.0%	CP	98.0%	CP	ok	CP	1
CP	DESY-ZN	CP	active	CP	online	CP	98.3%	CP	87.0%	CP	98.6%	CP	online	CP	95.0%	CP	90.0%	CP	ok	CP	0
CP	EELA-UTFSM	CP	active	CP	online	CP	95.0%	CP	99.2%	CP	98.6%	CP	online	CP	97.0%	CP	100.0%	CP	ok	CP	2
CP	FMPH-UNIBA	CP	active	CP	online	CP	94.3%	CP	91.1%	CP	98.6%	CP	online	CP	86.0%	CP	92.0%	CP	ok	CP	0
CP	FZK-LCG2	CP	active	CP	online	CP	95.2%	CP	97.8%	CP	98.6%	CP	online	CP	84.0%	CP	98.0%	CP	ok	CP	1
CP	GRIF-IRFU	CP	active	CP	online	CP	97.2%	CP	97.0%	CP	98.6%	CP	online	CP	79.0%	CP	98.0%	CP	ok	CP	n/a
CP	GRIF-LAL	CP	active	CP	online	CP	56.1%	CP	13.0%	CP	76.1%	CP	test	CP	57.0%	CP	80.0%	CP	n/a	CP	n/a
CP	GRIF-LPNHE	CP	active	CP	online	CP	98.5%	CP	97.8%	CP	98.6%	CP	online	CP	66.0%	CP	100.0%	CP	ok	CP	n/a

Site Status Board - Panda Jobs

Tier

0 + 1

Country

All

Cloud

All

Federation

All

Site

All

Choose Panel

Queue Status

Queues

Hammercloud

BNL-ATLAS

online

100%

BNLLAKE

online

100%

CERN-PROD

online

100%

CERN-T0

online

nodefq

FZK-LCG2

online

100%

IN2P3-CC

online

100%

INFN-T1

online

100%

NDGF-T1

online

50%

NIKHEF

test

0%

RAL-LCG2

online

100%

RRC-KJ-T1

online

100%

SARA-MATRIX

online

100%

TRIUMF-LCG2

online

100%

Taiwan-LCG2

online

multdefq

pic

online

no-test

Analysis Jobs

Efficiency

Activated

Running

82.0%

444

3318

94.0%

529

221

79.0%

335

3342

89.0%

66

1051

84.0%

37

1862

79.0%

228

5113

85.0%

91

2135

68.0%

3

1128

39.0%

21

182

73.0%

122

3044

88.0%

56

961

59.0%

2352

404

92.0%

56

670

0%

0

0

75.0%

183

629

Production Jobs

Efficiency

Activated

Running

80.0%

5940

12808

95.0%

168

69

95.0%

6790

5497

93.0%

195

338

98.0%

4094

3232

91.0%

2526

1362

99.0%

1474

883

89.0%

1086

1032

92.0%

328

100

95.0%

1148

1070

92.0%

1027

720

96.0%

930

556

95.0%

2260

3069

68.0%

2552

1653

93.0%

1923

1367

Site Status Board - SAM3

Tier	1 + 2	Country	All	Cloud	All	Federation	All	Site	All
SAM3 Site Availability									
AGLT2									
Australia-ATLAS									
BEIJING-LCG2									
BNL-ATLAS									
BNLLAKE									
BU_ATLAS_Tier2									
CA-SFU-T2									
CA-VICTORIA-WESTGRID-T2									
CA-WATERLOO-T2									
CSCS-LCG2									
CYFRONET-LCG2									
DESY-HH									
DESY-ZN									
EELA-UTFSM									
FMPHi-UNIBA									
FZK-LCG2									
GRIF-IRFU									
SAM3 SRM test counts									
OK	Warning	Critical							
286	0	0							
286	0	0							
0	0	0							
286	0	0							
0	0	0							
0	0	0							
288	0	0							
336	0	0							
287	0	0							
288	0	0							
0	0	0							
286	0	0							
287	0	0							
288	0	0							
0	0	0							
288	0	0							
0	0	0							
SAM3 GRIFTP test counts									
OK	Warning	Critical							
288	0	0							
0	0	0							
288	0	0							
286	0	0							
0	0	0							
0	0	0							
0	0	0							
0	0	0							
0	0	0							
287	0	0							
0	0	0							
0	0	0							
0	0	0							
312	0	0							
0	0	0							
0	0	0							

Summary

- Since 2016 CERN-IT MonIT group started developing a new unified infrastructure (UMA)
- The UMA infrastructure uses modern and efficient open-source solutions such as Kafka, InfluxDB, ElasticSearch, Kibana and Grafana to collect, store and visualize metadata produced by data (Rucio) and workflow management (PanDA) systems
- This software stack is adapted for the ATLAS experiment and allows the development of dedicated monitoring and accounting dashboards in Grafana visualization environment.
- ATLAS Monitoring has 12 production dashboards for Sites, Jobs and DDM monitoring
- The monitoring based on UMA is constantly being improved

Acknowledgements:

This work was partially funded by the Russian Science Foundation under contract No.19-71-30008 (research is conducted in Plekhanov Russian University of Economics)

Thanks!