



MONIT/ATLAS Meeting

24.06.2020



Outline

- Dashboards migration
- Current status
- Major achievements
- Future plans (MONIT internal)
- Future plans (impact on ATLAS)

Dashboards migration

Objectives of dashboards migration

- Consolidate in common MONIT infrastructure
- Preserve old functionality / Import historical data
- Build an internal community / Contribute upstream

Successful migration of old dashboards

- Activity started in 2017 with MONIT infrastructure deployment
- All dashboard migrated (except WLCG Site Monitoring)
- In parallel, other new ATLAS producers integrated

Old dashboards migration timeline

- **09.2017**- WLCG Transfers
- **05.2019**- ATLAS DDM Transfers
- **05.2019**- ATLAS DDM Accounting
- **07.2019**- ATLAS Job Accounting
- **05.2020**- WLCG Rebus (*in collaboration with CRIC*)
- **06.2020**- ATLAS Site Status Board
- **07.2020**- WLCG Site Monitoring

Matrix (Efficiency)

Efficiency												
	CA	CERN	DE	ES	FR	IT	ND	NL	RU	TW	UK	US
CA	61%	65%	77%	61%	73%	36%	71%	86%	93%	91%	89%	57%
CERN	100%	88%	100%	100%	99%	66%	100%	99%	100%	100%	100%	99%
DE	96%	99%	93%	100%	97%	58%	99%	98%	99%	99%	98%	98%
ES	99%	99%	85%	100%	98%	43%	95%	97%	100%	100%	99%	98%
FR	98%	98%	99%	100%	96%	44%	99%	99%	97%	100%	97%	98%
IT	43%	90%	57%	60%	50%	47%	58%	79%	62%	62%	66%	79%
ND	99%	100%	99%	100%	99%	63%	52%	91%	100%	100%	93%	99%
NL	97%	98%	99%	99%	98%	55%	99%	99%	100%	100%	97%	96%
RU	100%	96%	94%	100%	96%	52%	100%	100%	100%	100%	100%	100%
TW	100%	100%	100%	100%	81%	66%	98%	98%	100%	-	98%	99%
UK	97%	91%	93%	98%	96%	41%	98%	96%	99%	97%	96%	91%
US	100%	96%	89%	99%	98%	73%	97%	97%	98%	100%	95%	99%


- Matrix (Failures) (1 panel)
⋮
- Matrix (Successes) (1 panel)
⋮
- Transfer Table (1 panel)
⋮
- Transfer Plots (6 panels)
⋮
- Transfer Errors (1 panel)
⋮



> General (4 panels)


Overall

Submitted jobs




	min	max	avg	total
User Analysis	1.3 K	32.3 K	11.3 K	1.3618 Mil
Group Production	1.9 K	10.9 K	4.6 K	551.4 K
MC Simulation Full	178	17.4 K	4.1 K	494.2 K

Running jobs




	min	max	avg	total
User Analysis	20.3 K	39.8 K	29.7 K	3.5584 Mil
MC Simulation Full	9.5 K	27.2 K	20.1 K	2.4164 Mil
Group Production	12.4 K	21.8 K	16.5 K	1.9848 Mil

Pending jobs



	min	max	avg	total
User Analysis	32.8 K	113.3 K	70.4 K	8.4492 Mil
Group Production	41.4 K	61.2 K	53.3 K	6.3903 Mil
MC Simulation Full	35.1 K	62.5 K	48.7 K	5.8472 Mil


Finalising jobs



	min	max	avg	total
User Analysis	65.1 K	100.6 K	80.3 K	9.6366 Mil
Group Analysis	6.3 K	15.9 K	10.1 K	1.2124 Mil
Group Production	3.0 K	16.2 K	5.9 K	707.0 K



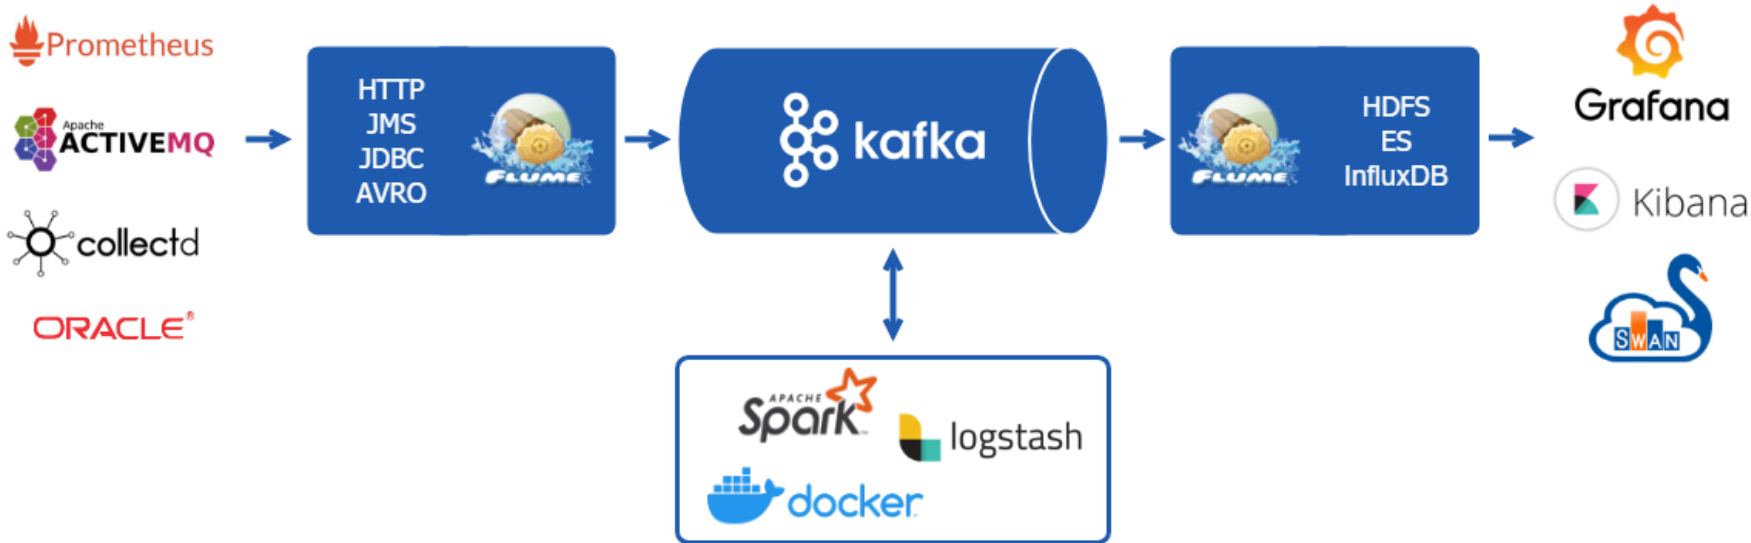
▼ Sites



▼ Endpoints



Current status




ATLAS is using all MONIT services

- *Ingestion*: data integrated via JDBC, AMQ, and HTTP
- *Transport/Processing*: data enrichment and aggregation in Spark
- *Storage/Visualization*: ES, InfluxDB, and HDFS / Grafana and Kibana

Some statistics of ATLAS @ MONIT today

- Number of ATLAS producers: 20 (panda, rucio, etc.)
- Number Grafana dashboards: 196 (using 80 data sources)
- Total data volume in ES: ~5 TB (kibana) + ~2.2 TB (timber)
- Total data volume in HDFS: ~19 TB

ATLAS Producers in MONIT



Major achievements

Last 12 months

Migration of Job Monitoring backend to ES

- Better performance, especially on longer time periods
 - e.g. 6 months plot (1 day bins): from ~3 min to ~3 sec
 - e.g. 10 years plot (1 month bins): now loads in ~20 sec
- Easier to change schema (add field, replace field)
- Out-of-the-box correct time literals (1w, 1M)

New dashboard for HEPSPEC info

- Dedicated dashboard to show historical data

Integrated Rucio Kubernetes logs

Several data classification improvements

- *(few more representative examples)*
- Job Accounting: new ***prodsourcelabel*** tag
- Job Accounting: new ***MC Merge*** value for ***adcactivity*** tag
- Job Accounting: new ***COVID*** value for ***adcactivity*** tag
- DDM Transfers: new ***throughput*** field

Improved FTS efficiency

- From CMS proposal and accepted by all VOs and WLCG
- More accurate computation algorithm

Improved infrastructure operations

- Improved procedure for interventions (e.g. Grafana)
- Better alarming for infrastructure problems

Very good collaboration between teams

- Meetings every other week
- 90 SNOW tickets closed in the last 12 months

Future plans (MONIT internal)

Next 12 months

Keep the infrastructure updated

- Upgrade Kafka and ElasticSearch to latest versions

Move services to Kubernetes

- Good candidates are Mesos, Marathon, and Chronos

Evaluate new interesting technologies

- Mainly in the areas of timeseries storage and data transport
- Alternatives for some use cases where InfluxDB and Flume are used

Redesign Monitoring of Monitoring

- Current duplicate workflow is too complex

Future plans (impact on ATLAS)

Next 12 months

Redesign old GLED collectors

- Old collectors for XRootD data must be retired/replaced

Migrate XSLs producers

- Move producers from old XML to new JSON format

New data retention in HDFS

- For logs default will be 13 months from 1st Sep

Move to CRIC topology

- Replace current AGIS endpoint
- Will be done dashboard by dashboard (FTS first)

