



# Monitoring CERN Data Centres and WLCG Infrastructure

Borja Garrido (CERN IT/CM)  
for the MONIT team

# News and more news

- Last GDB link
- Lots of progress done in the meanwhile
- Some planification updates
- Update on October OpsCoord meeting



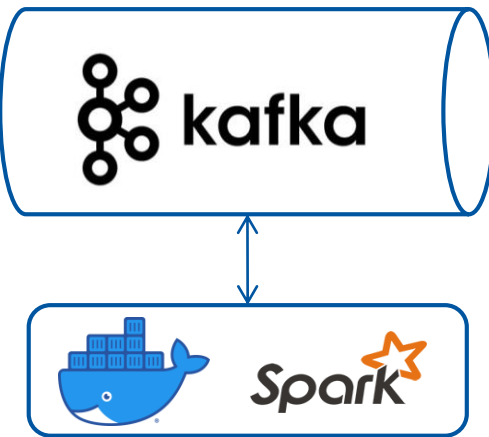
# MONIT Architecture

**Data Sources**  
100+ types  
5M doc/min

FTS Xrootd  
Rucio  
Panda Condor  
AGIS VOfeed  
Data Centre  
Batch OpenStack  
Network Puppet  
Logs Metrics  
Alarms

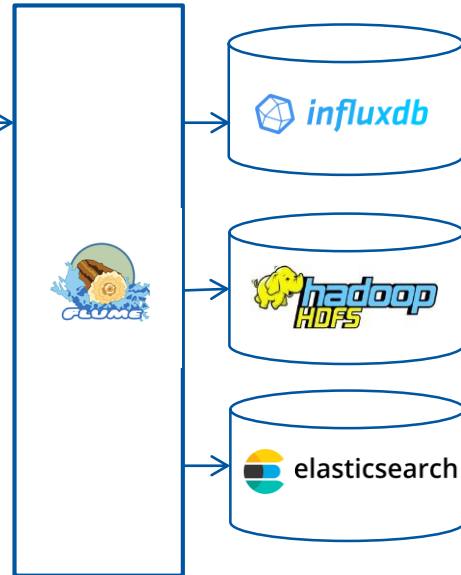


**Transport**  
72h buffer  
separate topics



**Processing**  
20 jobs

**Storage**  
~3TB/day  
compressed



# Overview

DC

Lemon retirement

WLCG

Transfers

Site Monitoring

SSB

ATLAS

DDM Transfers

DDM Accounting

Job Accounting

CMS

Job/Task monitoring

# Data Centre Monitoring



## Context:

- Datacentre monitoring
  - Host, Hardware and Services
- Replacement of Lemon by Collectd
- Work started on 2016
- Base monitoring and alarms available since 2017
- Fully replaced by September 2019



# Lemon retirement



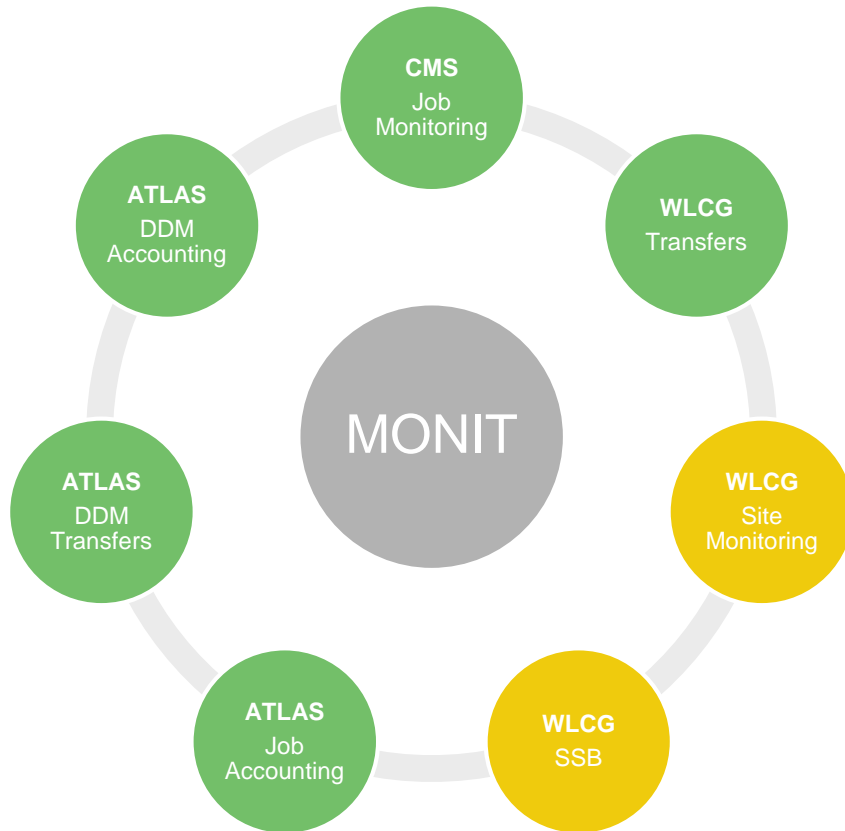
## Done:

- Migration of base metrics/plugins and alarms
- Migration service specific plugins and alarms
  - Including experiments hostgroups as voatlas/vocms
- Stop old Lemon infrastructure

# WLCG and Experiments

**Successfully migrated all  
job and data dashboards**  
*(presented in the next slides)*


Current work on Site  
dashboards migration





# Site Monitoring

## Context:



- SAM3 metrics (ETF + Alice tests)
- Work started on 2018
  - Reduced priority in favor of other migrations
  - Top priority since summer 2019
- Expected to be finished by Q1 of 2020
  - Migration of SAM3 dashboards to MONIT
  - Monthly availability/reliability report generation from MONIT data


# Site Monitoring

## Done:

- SAM3 Data enriched and computed for A/R statistics
  - Critical profiles
- Report generation adapted to MONIT
- New workflow for recomputation requests

## TO-DO:

- Data validation
- Prototype dashboard in Grafana
- Migration of non critical profiles by the MONIT team
- Stop old infrastructure (once everything else is finished)





## Availability of WLCG Tier-0 + Tier-1 Sites ATLAS

Target Availability for each site is 97.0%. Target for 8 best sites is 98.0%

Availability Algorithm: (CREAM-CE + ARC-CE + HTCONDOR-CE + GLOBUS) \* (all SRMv2 + all SRM + all GRIDFTP)

# WLCG MONIT A/R report

August 2019


(same look!)




Mode availability VO atlas Country All Federation All Tier 1 Site BNL-ATLAS Flavour All Profile ATLAS\_CRITICAL Hostname All

> Help (1 panel)



Sites



Endpoints




Tests



- S **site-monitoring-recomputations**
- Project
- Repository**
- Files
- Commits
- Branches
- Tags
- Contributors
- Graph
- Compare
- Charts

Monitoring > site-monitoring-recomputations > **Repository**

master site-monitoring-recomputations / atlas / 2019 / 09 / **IN2P3-LAPP-3.json**




 **[MONIT-2172] Test CI**  
Nikolay Tsvetkov authored 1 month ago

 **IN2P3-LAPP-3.json** 135 Bytes 

```
1 {
2   "dst_site": "IN2P3-LAPP-3",
3   "start_time": 1569275112000,
4   "end_time": 1569275312000,
5   "status": "SCHEDULED",
6   "vo": "atlas"
7 }
```


**WLCG**  
Recomputation  
workflow 

22051812 

   Edit Web IDE Replace **Delete**


# Site Status Board

## Context:



- Complementary metrics to SAM3
  - WLCG
  - ATLAS
  - CMS
- Resumed migration process during summer 2019
  - Meeting with the experiments

# Site Status Board



## Done:

- Discuss with WLCG/ATLAS/CMS migration strategy
- Many SSB metrics already integrated
  - Data transfers and Job monitoring use cases
  - Experiment specific metrics via MONIT provided HTTP endpoint

## TO-DO:

- WLCG: missing metrics being integrated
- ATLAS: most metrics already integrated, working on dashboard
- CMS: all metrics integrated, working on dashboard, migration plan to be defined soon

# DDM Transfers



## Context:

- Distributed Data Management transfers monitoring
- Work started on 2016
- Production ready since end of 2018
- Official replacement of the old dashboard since May 2019

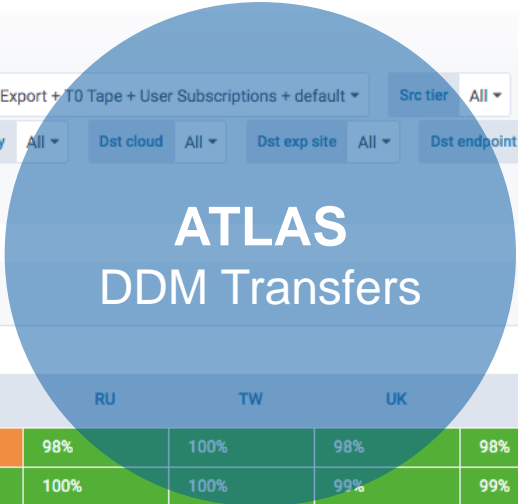


# DDM Transfers



## Done:

- Close collaboration with ATLAS
  - Improve look&feel
  - Increase performance
- Added new functionality (Rucio queues)
- Historical data imported and available in MONIT
- Old infrastructure stopped on September 2019



Binning auto ▾ Group by dst\_cloud ▾

Activity Analysis Input + Data Brokering + Data Consolidation + Data rebalancing + Express + Functional Test + Production Input + Production Output + Recovery + Staging + T0 Export + T0 Tape + User Subscriptions + default ▾ Src tier All ▾

Src country All ▾ Src cloud All ▾ Src exp site All ▾ Src endpoint All ▾ Src token All ▾ Src columns src\_cloud ▾ Dst tier All ▾ Dst country All ▾ Dst cloud All ▾ Dst exp site All ▾ Dst endpoint All ▾

Dst token All ▾ Dst rows dst\_cloud ▾ Include endpoints Exclude endpoints none +

> Overview & Help (1 panel)

Matrix (Efficiency)

Efficiency												
	CA	CERN	DE	ES	FR	IT	ND	NL	RU	TW	UK	US
CA	100%	98%	99%	98%	94%	99%	97%	74%	98%	100%	98%	98%
CERN	100%	97%	100%	98%	99%	99%	100%	84%	100%	100%	99%	99%
DE	90%	99%	98%	90%	89%	99%	96%	66%	98%	99%	95%	87%
ES	99%	99%	100%	92%	98%	100%	98%	87%	100%	99%	98%	83%
FR	94%	87%	92%	85%	87%	95%	98%	55%	97%	98%	95%	89%
IT	99%	99%	96%	94%	95%	98%	99%	68%	100%	100%	85%	87%
ND	100%	99%	100%	62%	94%	98%	100%	60%	100%	100%	98%	67%
NL	100%	99%	99%	93%	98%	100%	99%	90%	91%	99%	99%	93%
RU	100%	100%	100%	100%	97%	100%	100%	63%	100%	100%	94%	96%
TW	100%	98%	100%	98%	97%	100%	89%	42%	100%	-	100%	87%
UK	99%	99%	98%	97%	97%	100%	100%	76%	97%	99%	98%	82%
US	99%	100%	100%	97%	95%	99%	99%	80%	100%	98%	99%	94%

# DDM Accounting



## Context:

- Distributed Data Management accounting
- Work started end of 2018
- Completely new flow
  - ATLAS provides the final metrics
  - MONIT enriches them with topology data
- Official replacement of the old dashboard since May 2019

# DDM Accounting



## Done:

- Integrate global summaries and per site accounting
- Data enrichment with topology metadata
- Historical data imported and available in MONIT
- Old infrastructure stopped on September 2019

Overview

Logical Size  
**258 PB**

Disk Size  
**191 PB**

Primary Size  
**159 PB**

Disk Files  
**802893968**

Primary Files  
**773227217**

Number of Events  
**1942594375**

Tape Size  
**242 PB**

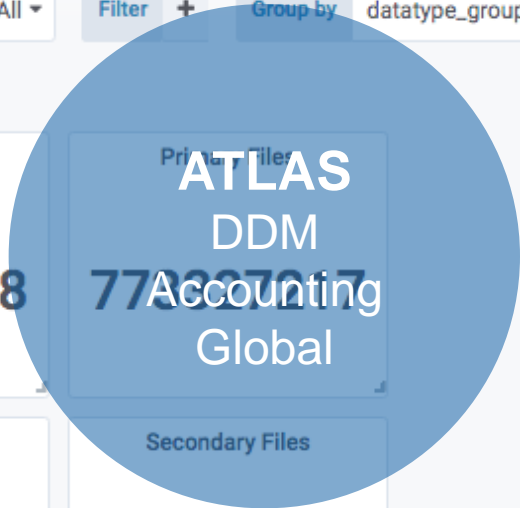
Secondary Size  
**32.7 PB**

Tape Files  
**198547501**

Secondary Files  
**29566751**

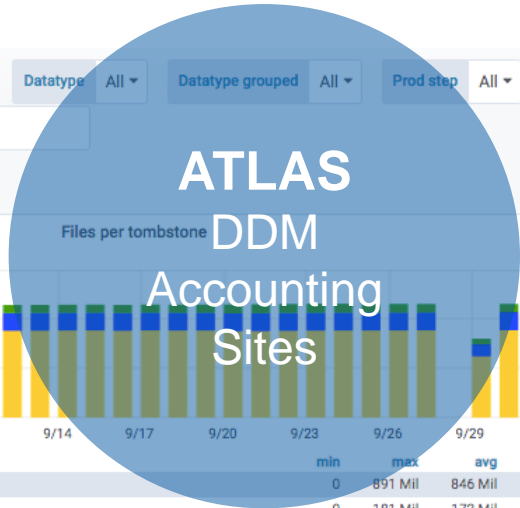
Total Size  
**433 PB**

Total Files  
**1001441469**

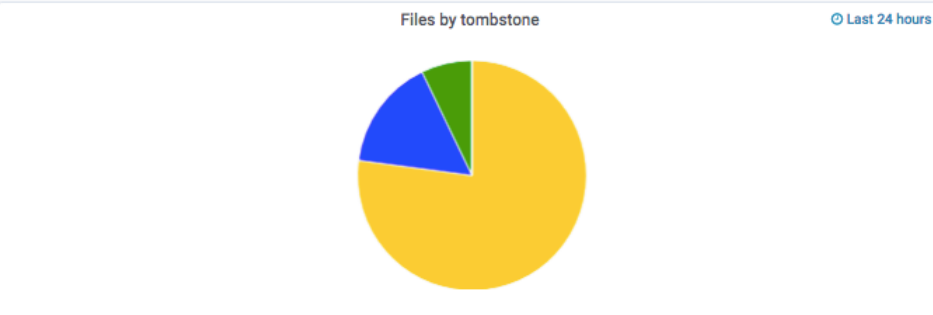
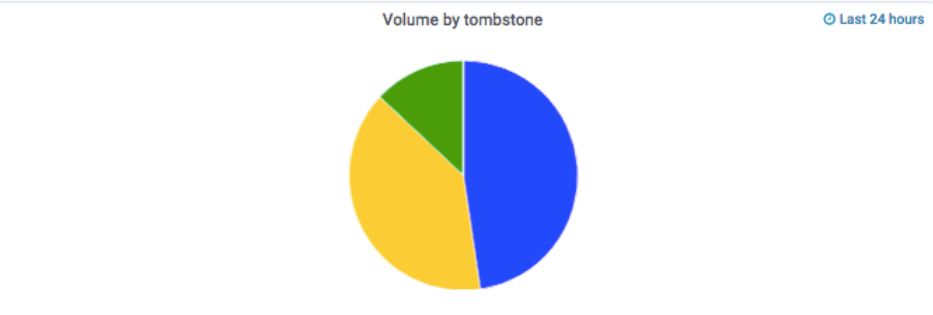
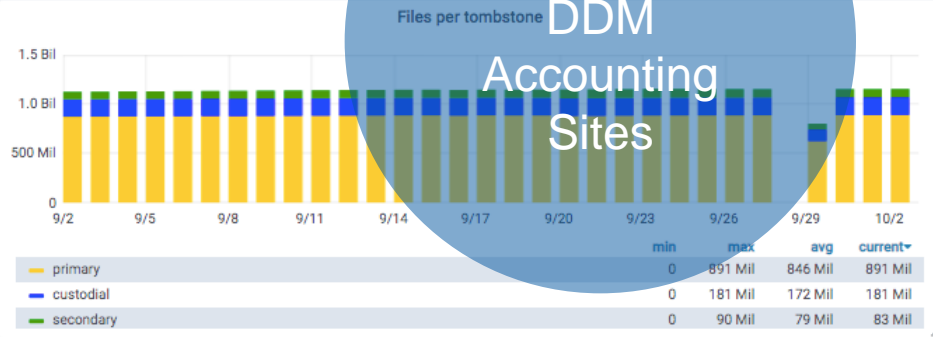
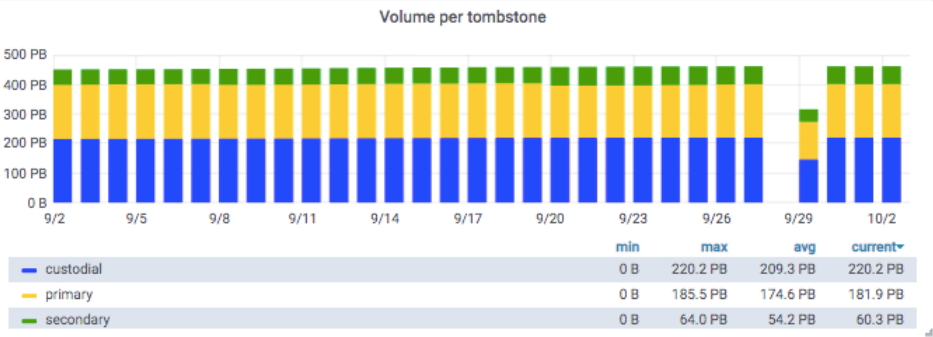


Group by: tombstone | Tier: All | Country: All | Cloud: All | Site: All | RSE: All | Token: All | Account: All | Campaign: All | Datatype: All | Datatype grouped: All | Prod step: All

Provenance: All | Scope: All | Stream name: All | Tombstone: All | Include endpoints: | Exclude endpoints: none



Overview



# Job Accounting



## Context:

- Job accounting information coming from Panda database
- Work started in 2016
- Big improvements during 2018
- Flow stable since end of 2018
- Official replacement of the old dashboard since July 2019

# Job Accounting



## Done:

- New functionality (closed and finalizing states...)
- Added missing data classifications (BOINC sites, Nucleus...)
- Historical data imported and available in MONIT
- Old infrastructure stopped on October 2019

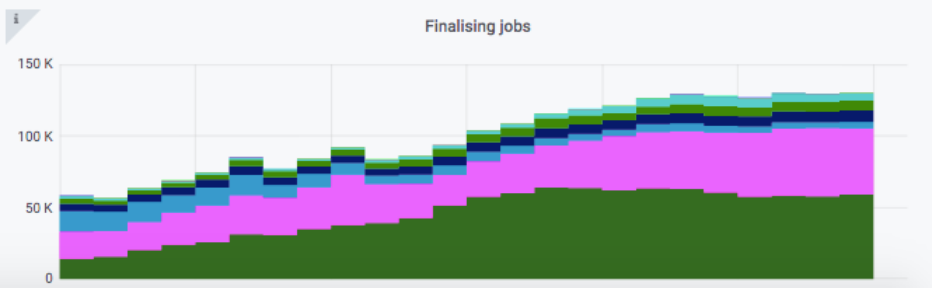
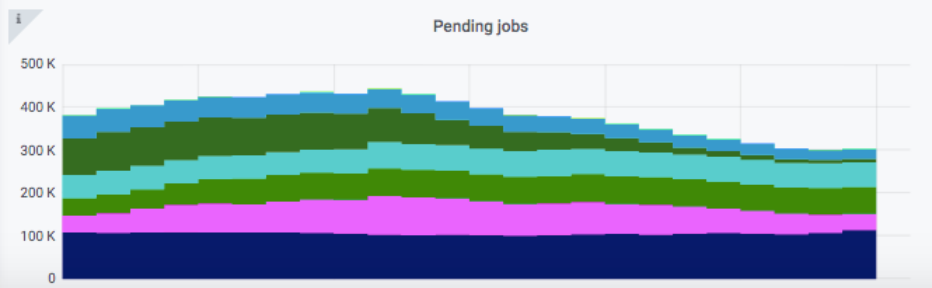
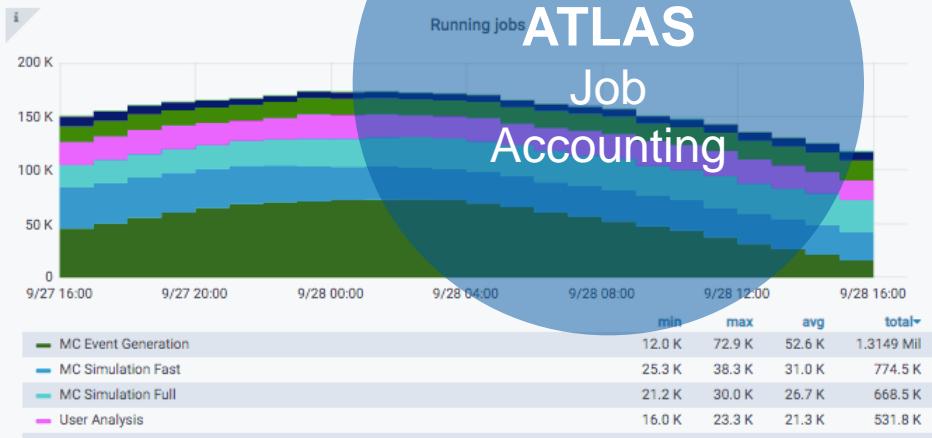
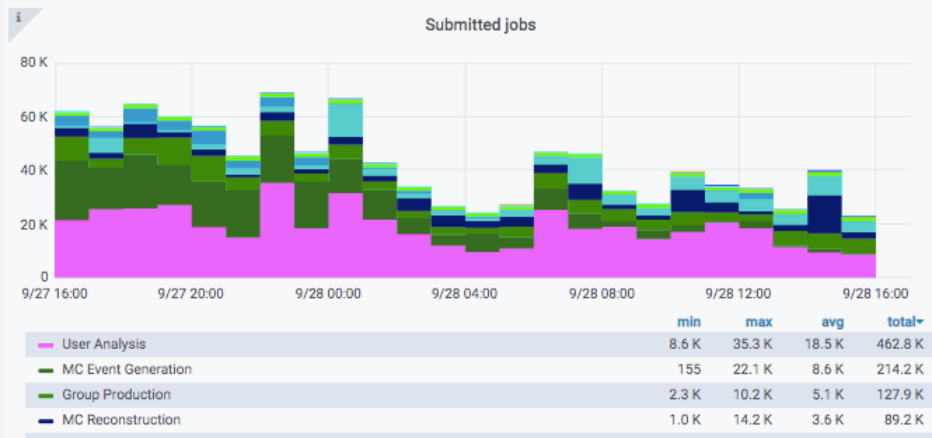
## TO-DO:

- Performance/look&feel improvements



Binning 1h ▾ Group by adactivity ▾ Country All ▾ Fed All ▾ Resource Type All ▾ Tier All ▾ Cloud All ▾ Site All ▾ Panda Queue All ▾ Nucleus All ▾ Cores All ▾ Event Serv All ▾  
 Groups All ▾ Input Data All ▾ Input Proj All ▾ Output Proj All ▾ GlobalShare All ▾ Resources Reporting All ▾ ProcessingType All ▾ JobType All ▾ JobStatus All ▾ Error category All ▾


Overall Status




# Job/Task Monitoring

## Context:

- Condor jobs/task monitoring
- Work started in 2016
- Big improvements during 2018
- Official replacement of the old dashboard since July 2019



# Job/Task Monitoring



## Done:

- Data integrated and validated inside MONIT
  - New flow defined together with CMS
  - Different backends for different use cases
- Historical data imported and available in MONIT
- Old infrastructure stopped on October 2019


## TO-DO:

- Addition of new tags

Group by CMS\_JobType Tier All Site All Type All CMS Job Type All Primary Data Tier All Binning 1h Filters +

**2038442** completed


Completed jobs



Analysis Cleanup Harvesting LogCollect Merge Processing Production Repack Unknown


**134226** running jobs

Running jobs

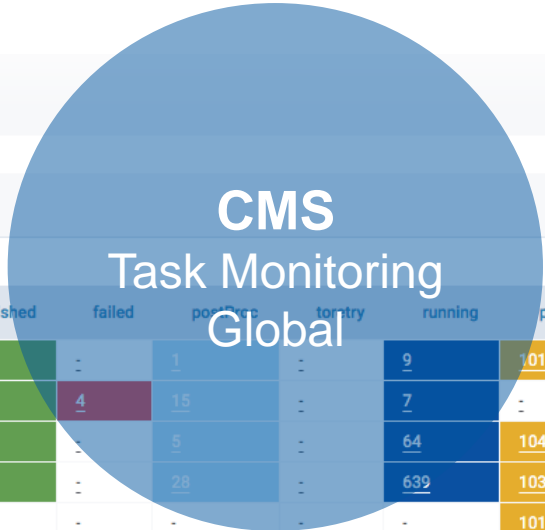


Analysis Cleanup Harvesting LogCollect Merge Processing Production Repack Unknown

**567500** pending jobs  
**CMS**  
 Job Monitoring  
 Pending jobs



Analysis Cleanup Harvesting LogCollect Merge Processing Production Repack Unknown



CMS

Task Monitoring

Global

Select User abenecke ▾ Site All ▾ task All ▾ Filters +

> Overall Job state count for user abenecke (last retry only) (1 panel)

▾ Jobs State by Task for user abenecke (last retry only)

Task for user abenecke

Task ▾	Total	finished	failed	pending	to retry	running	pending
<a href="#">191002_082045:abenecke_crab_DoubleMuonRun2016H-03Feb2017_ver2-v1</a>	1032	9	-	1	-	9	1013
<a href="#">191002_081740:abenecke_crab_DoubleMuonRun2016H-03Feb2017_ver3-v1</a>	129	103	4	15	-	7	-
<a href="#">191002_081446:abenecke_crab_DoubleMuonRun2016G-03Feb2017-v1</a>	1170	56	-	5	-	64	1045
<a href="#">191002_081148:abenecke_crab_DoubleMuonRun2016F-03Feb2017-v1</a>	1900	194	-	28	-	639	1039
<a href="#">191002_081135:abenecke_crab_QCD_Pt-15to7000_TuneCUETP8M1_Flat_13TeV_pythia8</a>	1016	-	-	-	-	-	1016
<a href="#">191002_080821:abenecke_crab_DoubleMuonRun2016E-03Feb2017-v1</a>	2779	458	2	70	5	1242	1002
<a href="#">191002_080517:abenecke_crab_DoubleMuonRun2016D-03Feb2017-v1</a>	1884	229	-	67	6	568	1014
<a href="#">191002_080201:abenecke_crab_DoubleMuonRun2016C-03Feb2017-v1</a>	2375	321	-	68	-	967	1019
<a href="#">191002_075926:abenecke_crab_TprimeBToTH_M-650_RH_TuneCUETP8M1_13TeV-madgraph-pythia8</a>	40	40	-	-	-	-	-
<a href="#">191002_075904:abenecke_crab_DoubleMuonRun2016B-03Feb2017_ver1-v1</a>	434	204	-	38	-	192	-
<a href="#">191002_075802:abenecke_crab_TprimeBToTH_M-600_RH_TuneCUETP8M1_13TeV-madgraph-pythia8</a>	40	39	-	1	-	-	-
<a href="#">191002_075612:abenecke_crab_DoubleMuonRun2016B-03Feb2017_ver2-v2</a>	4106	1883	-	198	12	996	1017
<a href="#">191002_074618:abenecke_crab_TprimeBToTH_M-700_RH_TuneCUETP8M1_13TeV-madgraph-pythia8</a>	14	-	-	13	-	1	-

> Graphs by Site (Completed Jobs) (6 panels)

> Graphs by ExitCode (Completed Jobs) (6 panels)

> Graphs by State (2 panels)

> Histograms (4 panels)

> Histograms (completed without retries) (4 panels)

∨ Jobs - last retry only

Job details for 191002\_082045:abenecke\_crab\_DoubleMuonRun2016H-03Feb2017\_ver2-v1 -last retry only

id	Retry	Status	ExitCode	Submit	Start	Finish	Wall time	Site	Job LOG	PostJob Log
<u>1</u>	0	Pending	-	Oct-02 08:23:24	-	-	-	Unknown	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>2</u>	0	Pending	-	Oct-02 08:23:24	-	-	-	Unknown	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>3</u>	0	Pending	-	Oct-02 08:23:24	-	-	-	Unknown	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>4</u>	0	Pending	-	Oct-02 08:23:24	-	-	-	Unknown	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>5</u>	0	Pending	-	Oct-02 08:23:24	-	-	-	Unknown	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>6</u>	0	finished	0	Oct-02 08:23:25	Oct-02 08:46:37	Oct-02 10:18:24	01:31:47	T2_ES_CIEMAT	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>7</u>	0	finished	0	Oct-02 08:23:25	Oct-02 08:46:37	Oct-02 11:12:03	02:25:26	T2_ES_CIEMAT	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>8</u>	0	Running	-	Oct-02 08:23:25	-	-	-	T2_ES_CIEMAT	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>9</u>	0	Running	-	Oct-02 08:23:30	-	-	-	T2_ES_CIEMAT	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>10</u>	0	Running	-	Oct-02 08:23:30	-	-	-	T2_ES_CIEMAT	<a href="#">JobLog</a>	<a href="#">PostJob</a>
<u>11</u>	0	Running	-	Oct-02 08:23:30	-	-	-	T2_ES_CIEMAT	<a href="#">JobLog</a>	<a href="#">PostJob</a>



# Others

## WLCG:

- CRIC integration in MONIT
  - As a topology information system replacing VOFeed
  - As site capacities/pledges information system replacing REBUS
- Test topology from CRIC with CMS Rucio integration
- REBUS retirement
  - Waiting for official migration to CRIC

# Others

## CMS:

- VOCMS logs integrated into monit-timber
- Rucio data integration and enrichment
- QA workflows for easier tests in flows
- Alarming functionality



# Summary

DC	Lemon retirement		
WLCG	Transfers	Site Monitoring	SSB
ATLAS	DDM Transfers	DDM Accounting	Job Accounting
CMS	Job/Task Monitoring		

# Summary

- MONIT being used as the production tool for the different use cases
  - All dashboards in one place: [monit-grafana.cern.ch](https://monit-grafana.cern.ch)
- Progresses done in many different tasks
  - Many old components decommissioned
  - Current focus on Site Monitoring / SSB (WLCG)
- Completed doesn't mean work stops
  - Discussing/Implementing improvements and new features
  - Done outside the migration process

# Thank you!

[monit.cern.ch](http://monit.cern.ch)

