



# Development of dashboards for the workflow management system in the ATLAS experiment

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

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

# Outline

- Introduction
- ATLAS production dashboards for workflow management system
- Jobs accounting dashboard
- Jobs monitoring dashboard
- HS06 Reports
- Site-oriented dashboard
- Site Status Board
- Harvester monitoring dashboard
- PQ/CE harvester monitoring dashboard
- iDDS Monitoring dashboard
- Summary

# Introduction

- The UMA (unified monitoring infrastructure) software stack developed by the CERN-IT Monit group provides the main repository of monitoring dashboards
- The adaptation of this stack to the ATLAS experiment began in 2018 to replace the old monitoring system
- Many dashboards were created and updated in Grafana for various user groups and use cases to monitor the workflow management system (PanDA) and computing infrastructure
- The ATLAS Production and Distributed Analysis System (PanDA) is a key component of the ATLAS distributed computing infrastructure
- This presentation is dedicated to the overview of these dashboards in the ATLAS experiment

# ATLAS production dashboards for workflow management system

## *Monit based dashboards*

- Jobs monitoring and accounting
  - Jobs Accounting
  - Jobs Monitoring (Monit ElasticSearch)
  - HS06 Reports
- Sites Monitoring
  - Site-oriented dashboard
  - Site Status Board Overview
  - Site Status Board - Panda Jobs
  - Site Status Board - SAM3

## *Custom Grafana dashboards*

- Harvester monitoring
- PQ/CE harvester monitoring
- iDDS Monitoring

**Monit based dashboards**

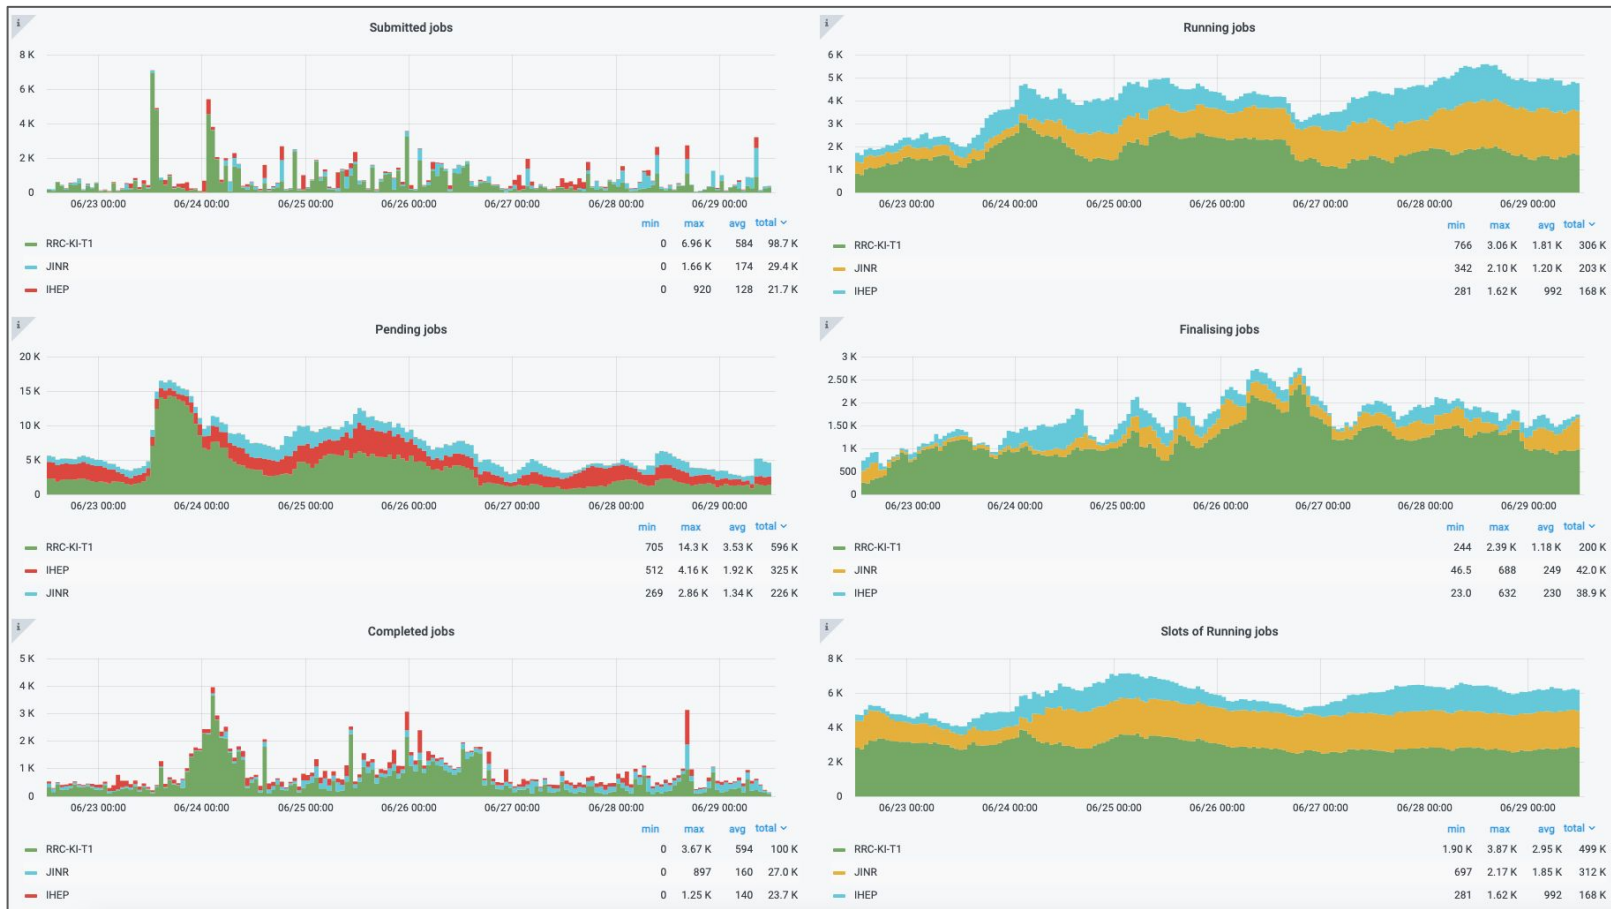
# Jobs accounting dashboard. Overview

- The jobs accounting dashboard is used by shifters, experts and management to spot problems with the workflow management system
- Information about jobs is available since 2010
- Dashboard uses Monit ElasticSearch storage for aggregated data as backend
- Kafka keeps aggregated data into five separate indices for each jobs statuses (completed, finilising, pending, running and submitted jobs) in the ElasticSearch
- CRIC as a topology source
- 87 plots, 25 filters, 23 options to group data
- Binning: 1h, 1d, 1w, 1M

The screenshot shows a complex filter interface for the Jobs accounting dashboard. It consists of multiple rows of dropdown menus and buttons. The first row includes 'Binning' (set to 1h), 'Group by' (set to adactivity), 'Country' (set to All), 'Fed' (set to All), and 'Resource Type' (set to All). The second row includes 'Tier' (set to All), 'Cloud' (set to All), 'Site' (set to All), 'Panda Queue' (set to All), 'Nucleus' (set to All), and 'Cores' (set to All). The third row includes 'Event Serv' (set to All), 'Groups' (set to All), 'Input Data' (set to All), 'Input Proj' (set to All), and 'Output Proj' (set to All). The fourth row includes 'GlobalShare' (set to All), 'Resources Reporting' (set to All), 'ProcessingType' (set to All), 'JobType' (set to All), and a 'Pledges' button. The fifth row includes 'Prodsourcelabel' (set to All), 'JobStatus' (set to All), 'Error category' (set to All), 'Container name' (set to All), and a 'CPU' button. At the bottom, there are three theme selection buttons: 'ATLAS Theme', 'ATLAS Theme (compact legend)', and 'Default'.


- > **General** (6 panels)
- > **Overall** (8 panels)
- > **Overall Pie Charts** (2 panels)
- > **Overall Cumulative Plots** (2 panels)
- > **HS06** (11 panels)
- > **CPU Consumption** (14 panels)
- > **CPU Efficiency** (9 panels)
- > **Processed Data** (13 panels)
- > **Successes/Failures** (15 panels)
- > **Resource Utilization** (4 panels)
- > **For presentations** (3 panels)

# Jobs Accounting dashboard. Jobs plots



# Jobs Accounting dashboard. Efficiency plots

- CPU consumption  
time successful/all  
jobs
- Wallclocktime  
successful/all jobs
- Efficiency based on  
success/all  
accomplished jobs






# Jobs Accounting dashboard. Custom legend styles and plots

## “For presentations”

- ATLAS Theme and ATLAS Theme (compact legend)
- Slots of Running jobs by ADC activity, Resource type
- Slots of Running jobs (HS06) by ADC activity




# Jobs Monitoring dashboard. Overview

- The dashboard provides extended information about completed jobs for 2 months
- It is useful to monitor separate job(s), task(s) or request(s)
- Data is processed in the same as for Jobs accounting dashboard
- Kafka keep aggregated data into dedicated index for completed jobs in 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.jeditaskid	data.reqid	data.statechangetime	data.jobstatus	data.computingsite	data.cpuconsumptiontime
<a href="#">5100926126</a>	<a href="#">25917255</a>	<a href="#">37930</a>	2021-06-28 14:08:30	finished	OU_OSCER_ATLAS	214331
<a href="#">5100926121</a>	<a href="#">25917255</a>	<a href="#">37930</a>	2021-06-28 14:08:29	finished	OU_OSCER_ATLAS	219882
<a href="#">5100926111</a>	<a href="#">25917255</a>	<a href="#">37930</a>	2021-06-28 14:08:28	finished	OU_OSCER_ATLAS	204202
<a href="#">5100926105</a>	<a href="#">25917255</a>	<a href="#">37930</a>	2021-06-28 14:08:28	finished	OU_OSCER_ATLAS	219196


## Panda Failures by ExitCode



# Jobs Monitoring dashboard. Job duration plots

- Status map plot

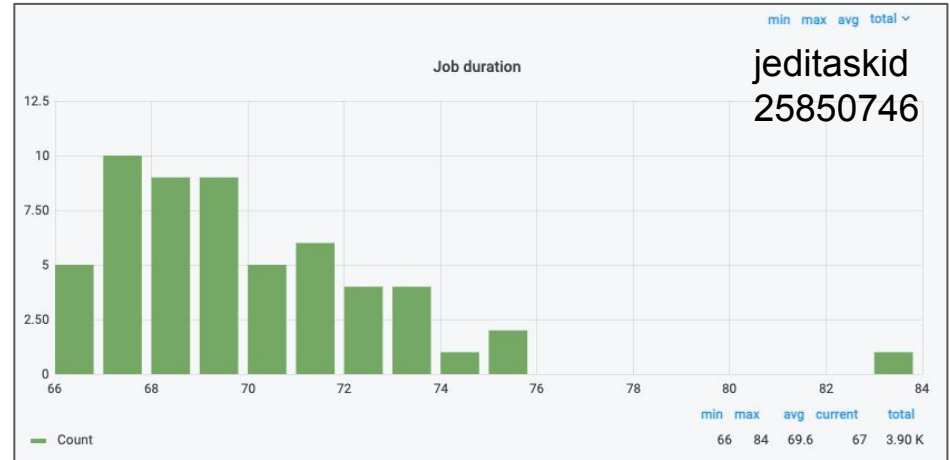
- Green and red buckets are number of the buckets. There is no possibility to display number of jobs in the buckets



- Bar plot

X-axis job duration

Y-axis number of 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 Computingsites, 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 ^
NERSO_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

# Site-oriented dashboard. Overview

- This dashboard combines information from Jobs Accounting and DDM transfer dashboards
- Allows to monitor and analyse computingsites 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

<div>Tier1 + 2CountryAllCloudAllFederationAllSiteAll</div> <div>Choose PanelFast</div>																
	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					
AGLT2	active	online	97.5%	96.6%	98.6%	online	82.0%	83.0%	ok	1						
Australia-ATLAS	active	online	91.7%	100.0%	98.6%	online	72.0%	98.0%	ok	0						
BEIJING-LCG2	active	online	94.3%	78.1%	98.6%	online	0%	99.0%	ok	0						
BNL-ATLAS	active	online	93.1%	98.5%	98.6%	online	82.0%	80.0%	ok	0						
BNLLAKE	active	online	100.0%	100.0%	98.6%	online	94.0%	95.0%	ok	n/a						
BU_ATLAS_Tier2	active	online	98.2%	99.4%	86.3%	online	83.0%	97.0%	ok	0						
CA-SFU-T2	active	online	97.8%	97.9%	98.6%	online	75.0%	97.0%	ok	1						
CA-VICTORIA-WESTGRID-T2	active	online	89.7%	95.2%	98.6%	online	63.0%	98.0%	ok	0						
CA-WATERLOO-T2	active	online	99.2%	98.3%	98.6%	online	88.0%	91.0%	degraded	1						
CSCS-LCG2	active	online	94.1%	84.9%	98.6%	online	87.0%	88.0%	ok	0						
CYFRONET-LCG2	active	online	97.1%	96.3%	98.6%	online	99.0%	93.0%	ok	0						
DESY-HH	active	online	99.3%	95.0%	98.6%	online	82.0%	98.0%	ok	1						
DESY-ZN	active	online	98.3%	87.0%	98.6%	online	95.0%	90.0%	ok	0						
EELA-UTFSM	active	online	95.0%	99.2%	98.6%	online	97.0%	100.0%	ok	2						
FMPH-UNIBA	active	online	94.3%	91.1%	98.6%	online	86.0%	92.0%	ok	0						
FZK-LCG2	active	online	95.2%	97.8%	98.6%	online	84.0%	98.0%	ok	1						
GRIF-IRFU	active	online	97.2%	97.0%	98.6%	online	79.0%	98.0%	ok	n/a						
GRIF-LAL	active	online	56.1%	13.0%	76.1%	test	57.0%	80.0%	n/a	n/a						
GRIF-LPNHE	active	online	98.5%	97.8%	98.6%	online	66.0%	100.0%	ok	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-KI-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

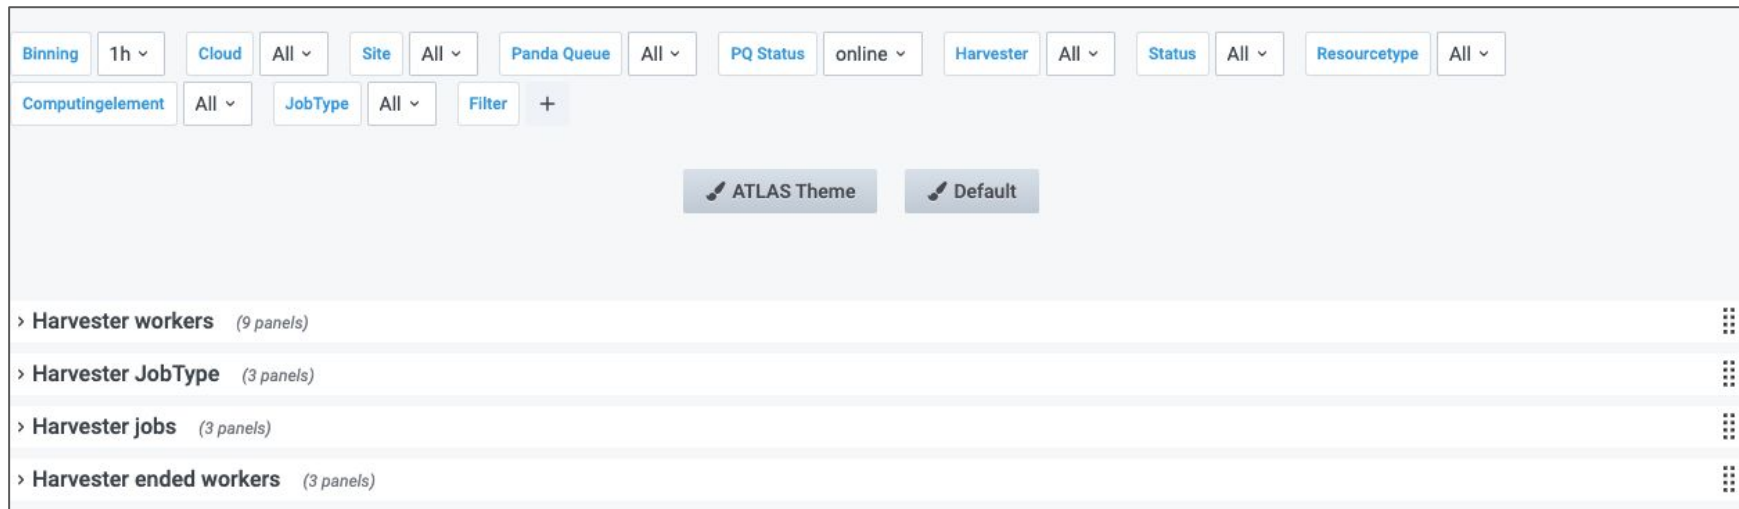
Tier	1 + 2 ▾	Country	All ▾	Cloud	All ▾	Federation	All ▾	Site	All ▾
i SAM3 Site Availability									
i SAM3 SRM test counts									
OKWarningCritical									
AGLT2100.0%28600									
Australia-ATLAS100.0%28600									
BEIJING-LCG2100.0%000									
BNL-ATLAS100.0%28600									
BNLLAKE100.0%000									
BU_ATLAS_Tier287.7%000									
CA-SFU-T2100.0%28800									
CA-VICTORIA-WESTGRID-T2100.0%33600									
CA-WATERLOO-T2100.0%28700									
CSCS-LCG2100.0%28800									
CYFRONET-LCG2100.0%000									
DESY-HH100.0%28600									
DESY-ZN100.0%28700									
EELA-UTFSM100.0%28800									
FMPH-UNIBA100.0%000									
FZK-LCG2100.0%28800									
GRIF-IRFU100.0%000									
i SAM3 GRIFTP test counts									
OKWarningCritical									
28800									
000									
28800									
28600									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									
000									



# Custom Grafana dashboards

# Harvester monitoring dashboard. Overview

- Harvester is a resource-facing service between the PanDA server and collection of pilots
- This dashboard uses Elasticsearch storage as backend
- Information in the storage is copied and updated every 5 minutes using the Logstash
- Two tables in PanDA as a datasource: harvesterworkers and schedconfig (CRIC)
- 18 plots, 10 filters
- Binning: 10m, 30m, 1h, 6h, 12h, 1d, 7d, 14d, 30d



# Harvester monitoring dashboard. Worker evolution

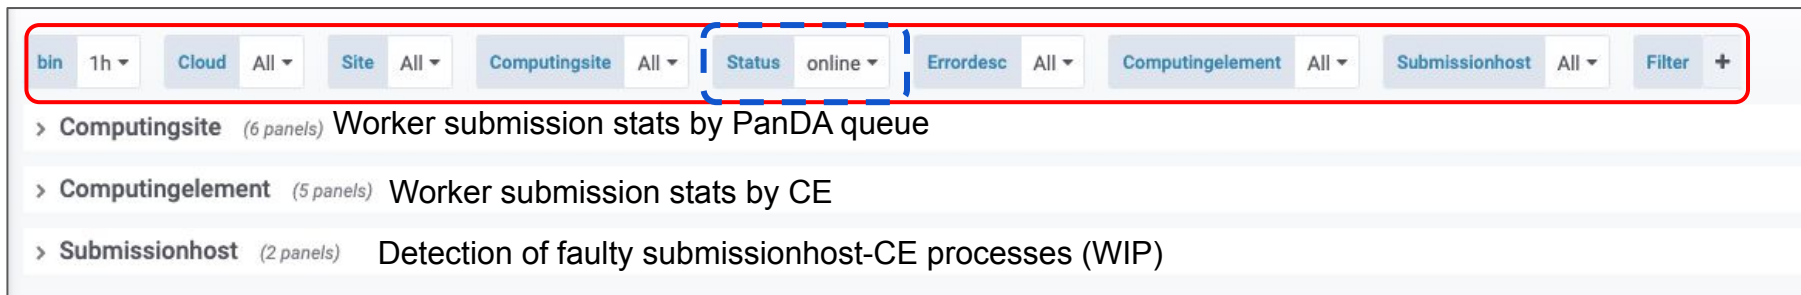


# PQ/CE harvester monitoring dashboard. Overview

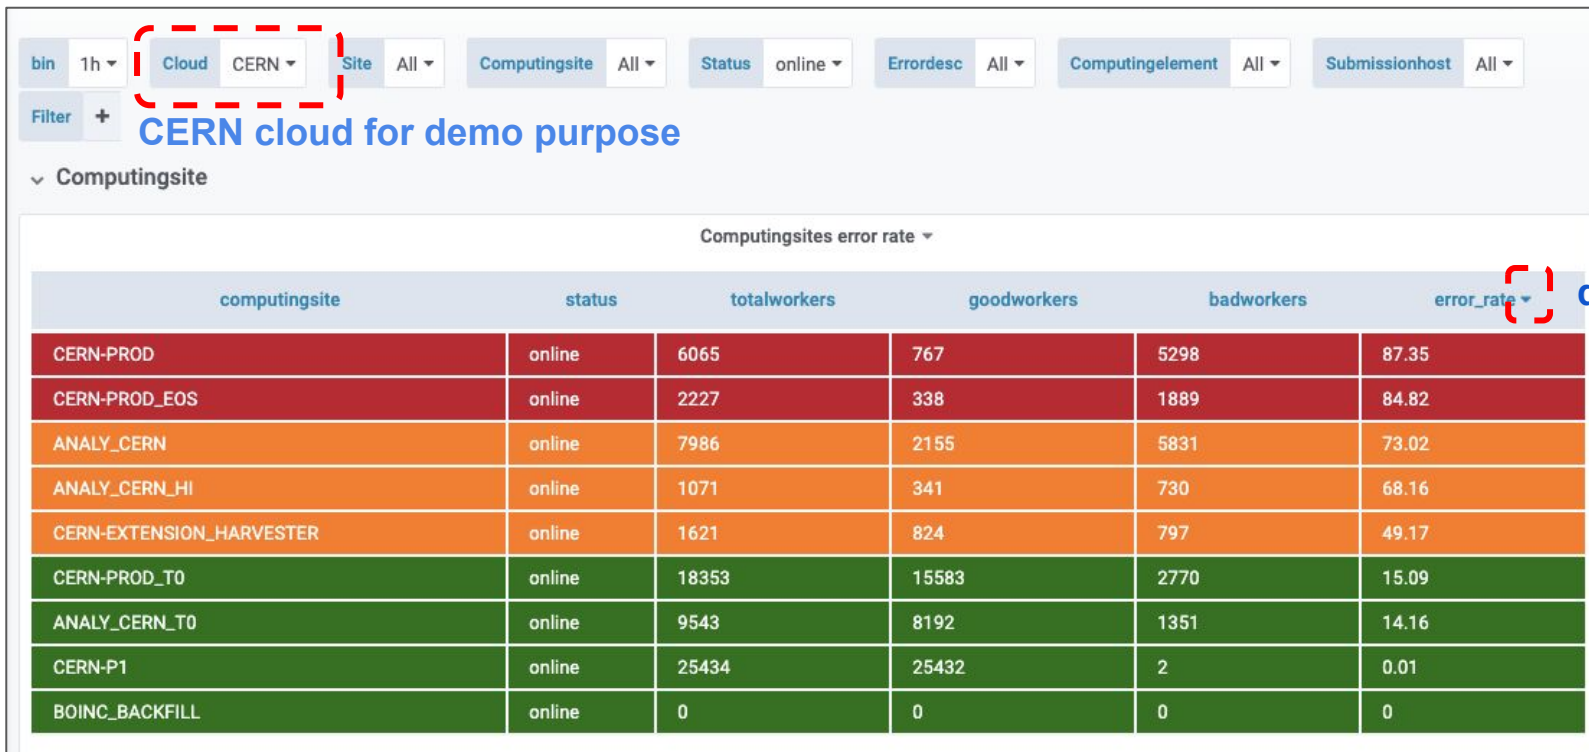
- Monitors Harvester worker submission issues at PanDA queues and CEs. Four components:
  - Elasticsearch: repository with worker information
  - Python script for data extraction and analysis
  - InfluxDB/MySQL: storage of analyzed data
  - Grafana: visualization
- Analyzes workers in final states in the last 1h. “Good workers”: finished status “Bad workers”: failed, cancelled, missed statuses
- 13 plots, 8 filters

online queues only by default

Filters



# PQ/CE harvester monitoring dashboard. Submission stats



default sorting

Worker submission stats: error rate = bad workers/total workers


# PQ/CE harvester monitoring dashboard. Error messages

Computingsites errors list ▾					
computingsite	errordesc	ratio_computingsite	count ▾	ratio_error	
ANALY_CERN	Condor HoldReason: None ; Condor RemoveReason: The system macro SYSTEM_PERIODIC_REMOVE expression '((NumJobStarts >= 1 && JobStatus == 1)    (NumJobStarts > 1 && JobStatus == 2))    ((JobRunCount >= 1 && JobStatus == 1)    (JobRunCount > 1 && JobStatus == 2))	73.02	5782	99.16	
CERN-PROD	Condor HoldReason: None ; Condor RemoveReason: The system macro SYSTEM_PERIODIC_REMOVE expression '((NumJobStarts >= 1 && JobStatus == 1)    (NumJobStarts > 1 && JobStatus == 2))    ((JobRunCount >= 1 && JobStatus == 1)    (JobRunCount > 1 && JobStatus == 2))	87.35	5289	99.83	
CERN-PROD_T0	Condor HoldReason: None ; Condor RemoveReason: The system macro SYSTEM_PERIODIC_REMOVE expression '((NumJobStarts >= 1 && JobStatus == 1)    (NumJobStarts > 1 && JobStatus == 2))    ((JobRunCount >= 1 && JobStatus == 1)    (JobRunCount > 1 && JobStatus == 2))	15.09	2478	89.46	
CERN-PROD_EOS	Condor HoldReason: None ; Condor RemoveReason: The system macro SYSTEM_PERIODIC_REMOVE expression '((NumJobStarts >= 1 && JobStatus == 1)    (NumJobStarts > 1 && JobStatus == 2))    ((JobRunCount >= 1 && JobStatus == 1)    (JobRunCount > 1 && JobStatus == 2))	84.82	1887	99.89	
ANALY_CERN_T0	Condor HoldReason: None ; Condor RemoveReason: The system macro SYSTEM_PERIODIC_REMOVE expression '((NumJobStarts >= 1 && JobStatus == 1)    (NumJobStarts > 1 && JobStatus == 2))    ((JobRunCount >= 1 && JobStatus == 1)    (JobRunCount > 1 && JobStatus == 2))	14.16	1307	96.74	
CERN-EXTENSION_HARVESTER	submission failed: Exception OSError: [Errno 28] No space left on device	73.02	789	99.12	
ANALY_CERN_HI	Condor HoldReason: None ; Condor RemoveReason: The system macro SYSTEM_PERIODIC_REMOVE expression '((NumJobStarts >= 1 && JobStatus == 1)    (NumJobStarts > 1 && JobStatus == 2))    ((JobRunCount >= 1 && JobStatus == 1)    (JobRunCount > 1 && JobStatus == 2))	68.16	730	100.00	
CERN-PROD_T0	Condor HoldReason: HTCondor-CE held job due to no matching routes, route job limit, or route failure threshold; see 'HTCondor-CE Troubleshooting Guide'; Worker canceled by harvester due to held too long or not found	21.26	291	33.37	
ANALY_CERN	Condor HoldReason: Network error talking to schedd, probably an authorization failure ; Worker canceled by harvester due to held too long or not found	73.66	47	0.86	


CERN-EXTENSION\_HARVESTER: No space left on device

# PQ/CE harvester monitoring dashboard. Submission rate history

Error rate for computingsites

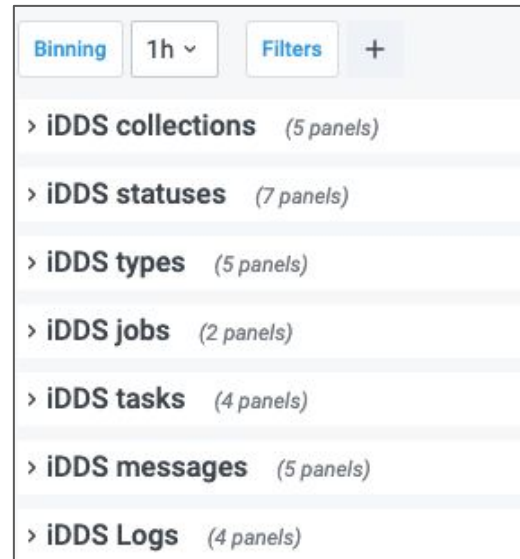


Total bad vs good workers



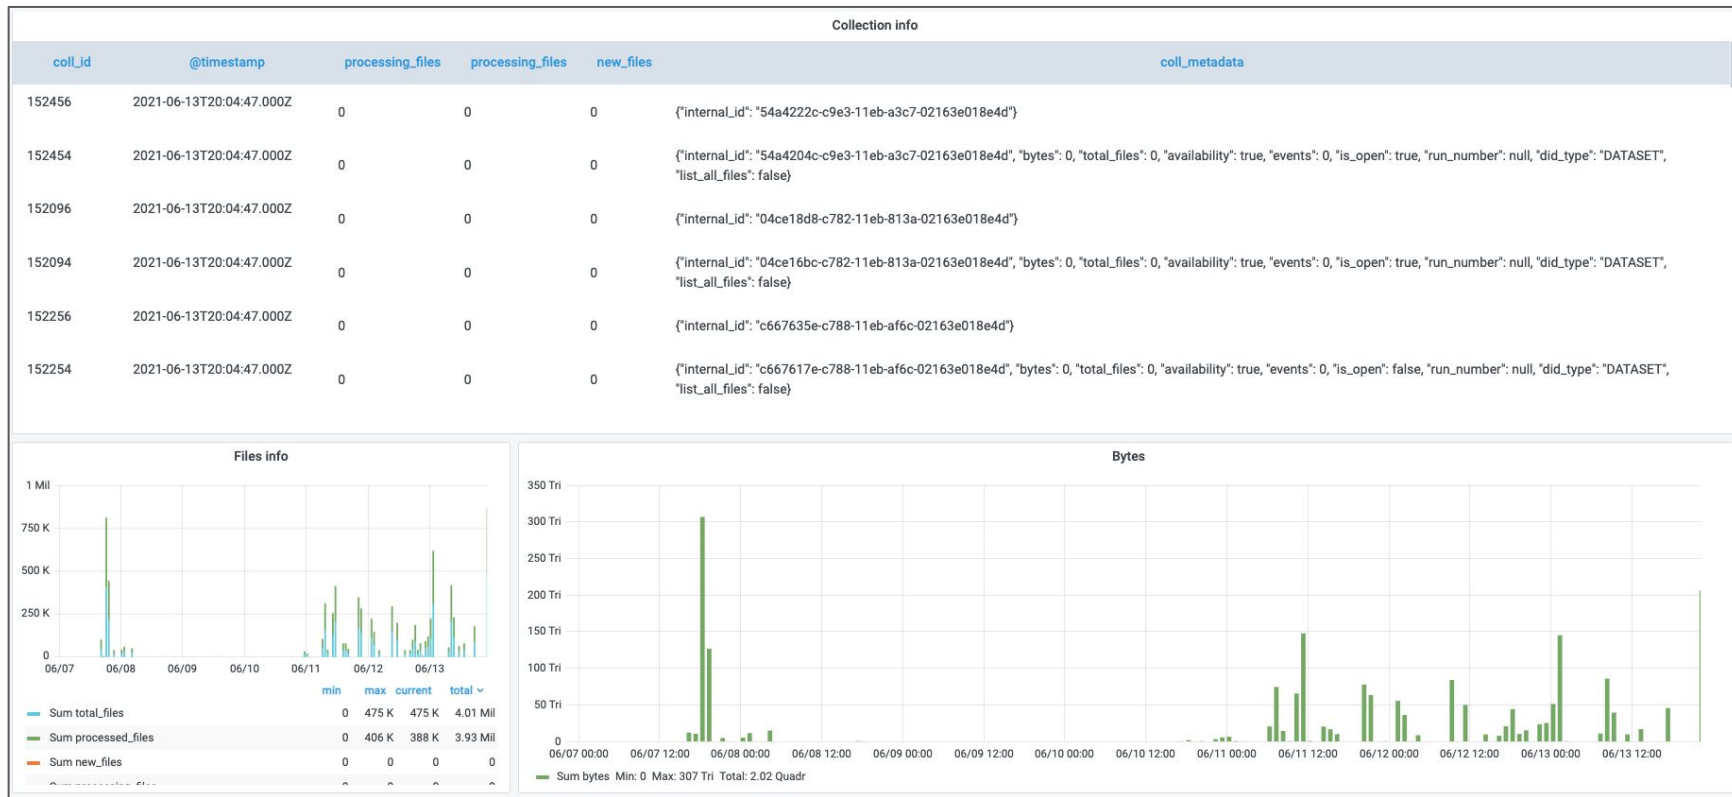
# iDDS Monitoring dashboard. Overview

- intelligent Data Delivery Service (iDDS) is proposed to intelligently transform and deliver the needed data to a processing workflow in a high granularity
- Monitoring has been created in the Monit Grafana to help debug and improve the service
- All iDDS data is available in ElasticSearch storage
- Information in the storage is copied and updated every 5-10 minutes using the Logstash
- Dumps tables from the database:
  - Information from all iDDS tables
  - Information about iDDS jobs from all jobs tables (not only completed jobs) in PanDA database + from jedi\_task and schedconfig tables. Six months data available
  - Information about iDDS tasks from jedi\_task table in PanDA database + from jedi\_taskparams table. Six months data available
- Information about errors in iDDS is aggregated from iDDS logs using Filebeat and Logstash






# iDDS Monitoring. Information about collections



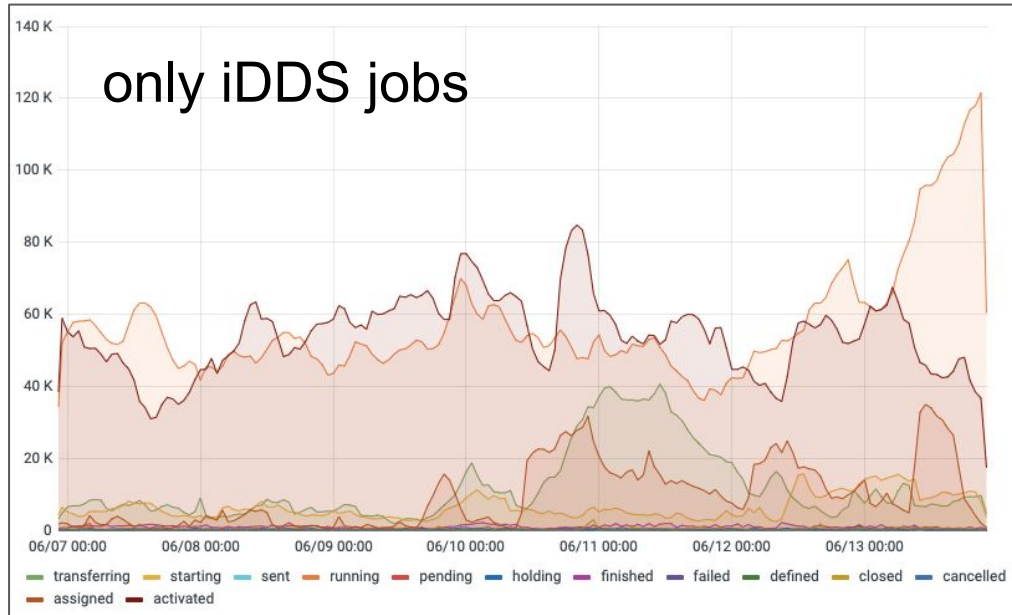
# iDDS Monitoring. iDDS tasks accounting and monitoring

iDDS task statuses




# iDDS Monitoring. iDDS jobs monitoring

iDDS jobs statuses




iDDS jobs statuses. Pie chart



# iDDS Monitoring. iDDS logs

Error messages by iDDS host



# Summary

- ATLAS started using the UMA software stack in 2018 to replace the old monitoring system
- Many dashboards were created in Grafana to monitor PanDA system, computing infrastructure and PanDA's components (iDDS, Harvester)
- Two types of dashboards: Monit based dashboards and Custom Grafana dashboards
- Jobs Accounting, Jobs Monitoring, HS06 Reports, Sites monitoring, Harvester monitoring dashboards are available for the different user groups in ATLAS
- 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!