

# Applying Big Data solutions for log analytics in the PanDA infrastructure

**A. Alekseev**, F. Barreiro Megino, A. Klimentov, T. Korchuganova, T. Maeno, S. Padolski

26<sup>th</sup> Symposium on Nuclear Electronics and Computing,  
25-29 September 2017, Budva, Montenegro

# Outline

- ATLAS, PanDA, JEDI, BigPanDA
- Reasons for choosing filebeat and logstash and results of implementation for PanDA, JEDI and BigPanDA logs
- Results of integration with BigPanDA monitoring
- BigPanDA logs processing. Notification system for BigPanDA monitoring

# Introduction

- **ATLAS** is one of two general-purpose detectors at the Large Hadron Collider (LHC). 1PB/sec from all sub-detectors reduces to ~1GB/sec raw data which is saved to Data Center.
- **PanDA** (**P**roduction **and** **D**istributed **A**nalysis) is the workflow management system of the ATLAS experiment at the LHC and is responsible for generating, brokering and monitoring up to two million jobs per day across 150 computing centers in the Worldwide LHC Computing Grid.
- **JEDI** (**J**ob **E**xecution and **D**efinition **I**nterface) is an intelligent component in the PanDA system to have capability for task-level workload management.
- **BigPanDA** is a monitoring system which provides a comprehensive and coherent view of the tasks and jobs executed by the system, from high level summaries to detailed drill-down job diagnostics.

# Motivation

- Centralized storage for PanDA server, JEDI and BigPanDA logs
- 400GB daily logs
- Average size of the one log file can be up to 2 GB
- Quick search of errors and potential problems
- Notification features
- **ELK (Elasticsearch+Logstash+Kibana) have more performance than Flume and parses/indexes messages**

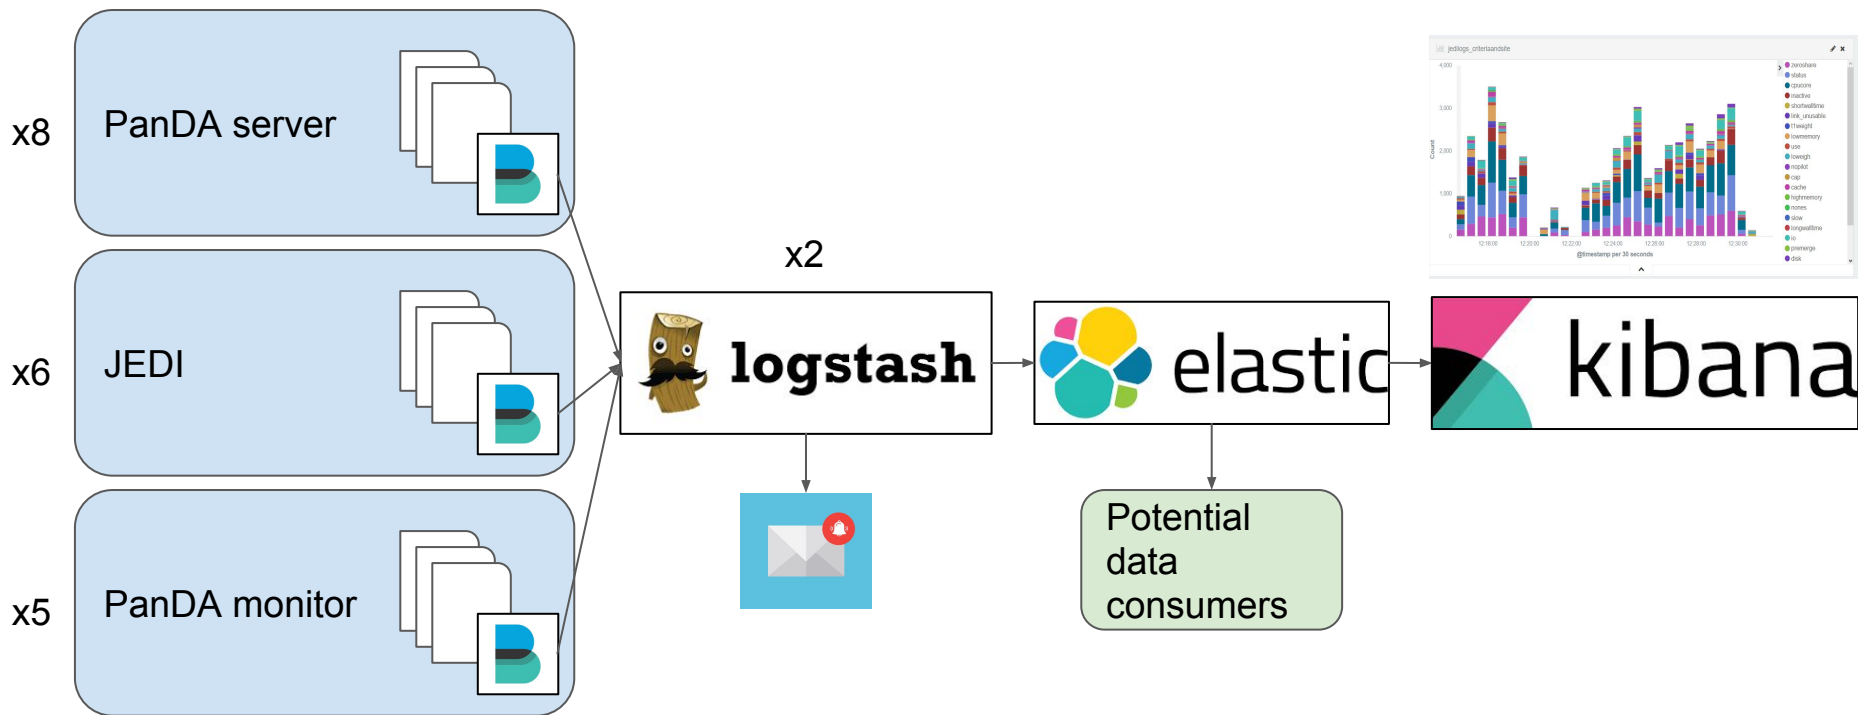
# ELK stack and Flume

- **ELK stack**

- **Logstash** is an special open source software for collecting, filtering and normalizing logs. It is used to collect collect log events from different log types.
- **ElasticSearch** is a distributed open source software for storing and searching information. In our case, Logstash writes all log events into the ElasticSearch repository.
- **Kibana** is an open source data visualization plugin for Elasticsearch. It is used for visualization of data from Elasticsearch cluster.

- **Apache flume (or Flume)** is a distributed service for collecting, aggregating, and moving large amounts of log data.


# PanDA infrastructure



# Scale of the Problem

## PanDA server


~ 210 GB per week in ES (34 types of log filters/files)



Timestamp per 3 hours

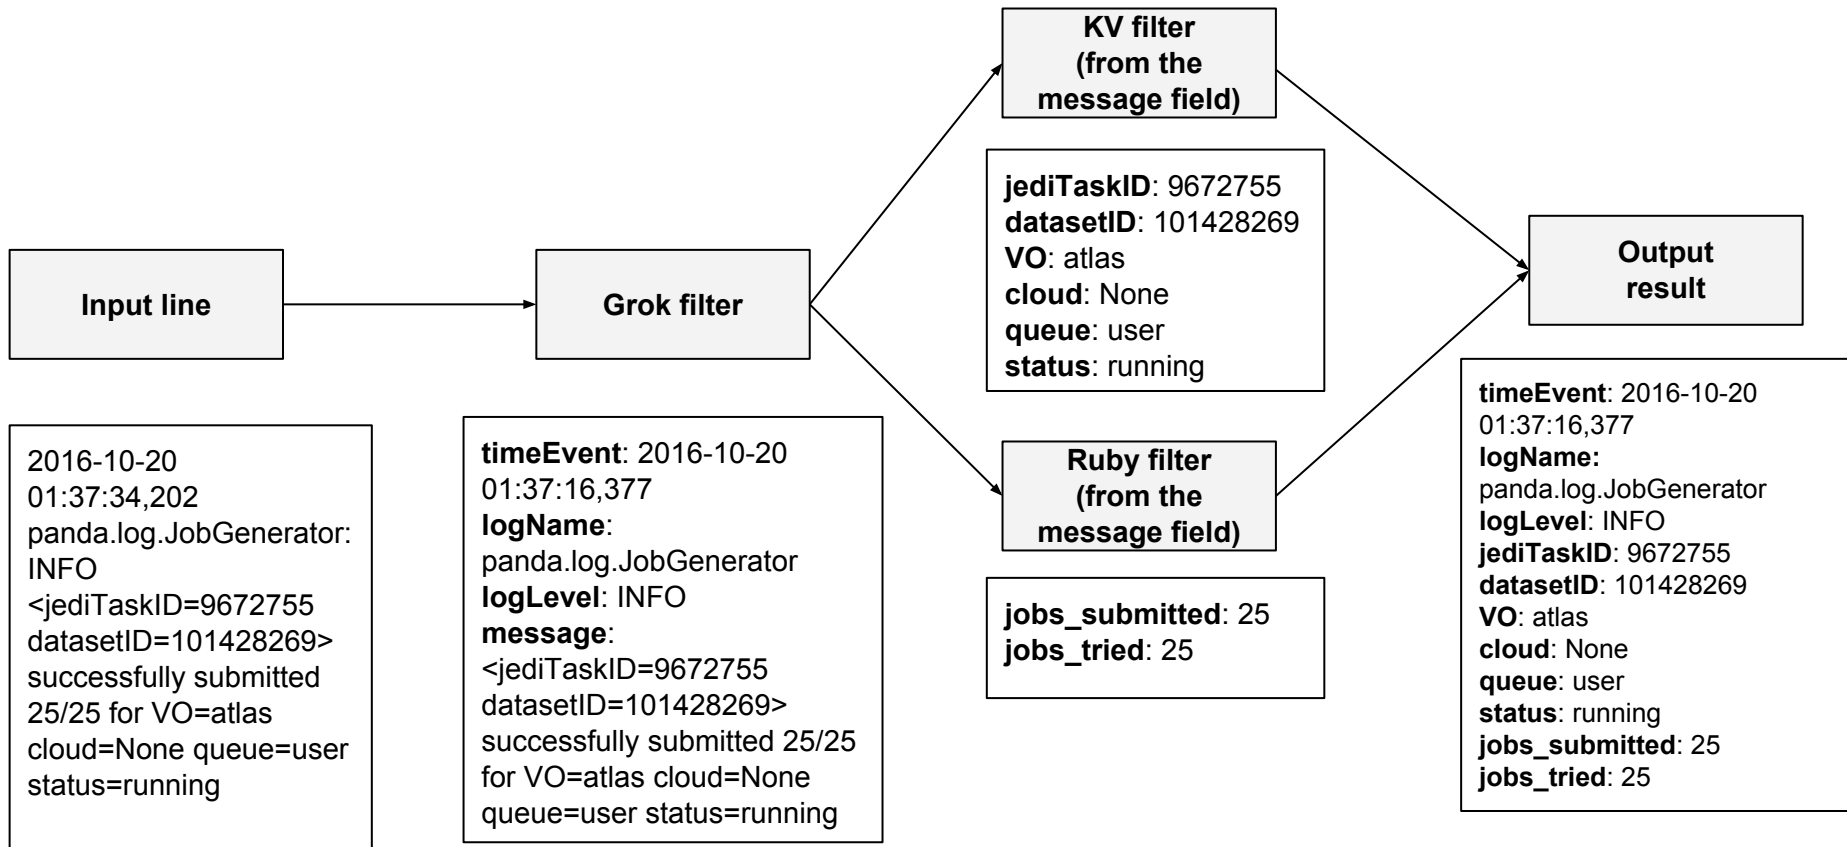
## JEDI

~ 73 GB per week in ES (27 types of log filters/files)



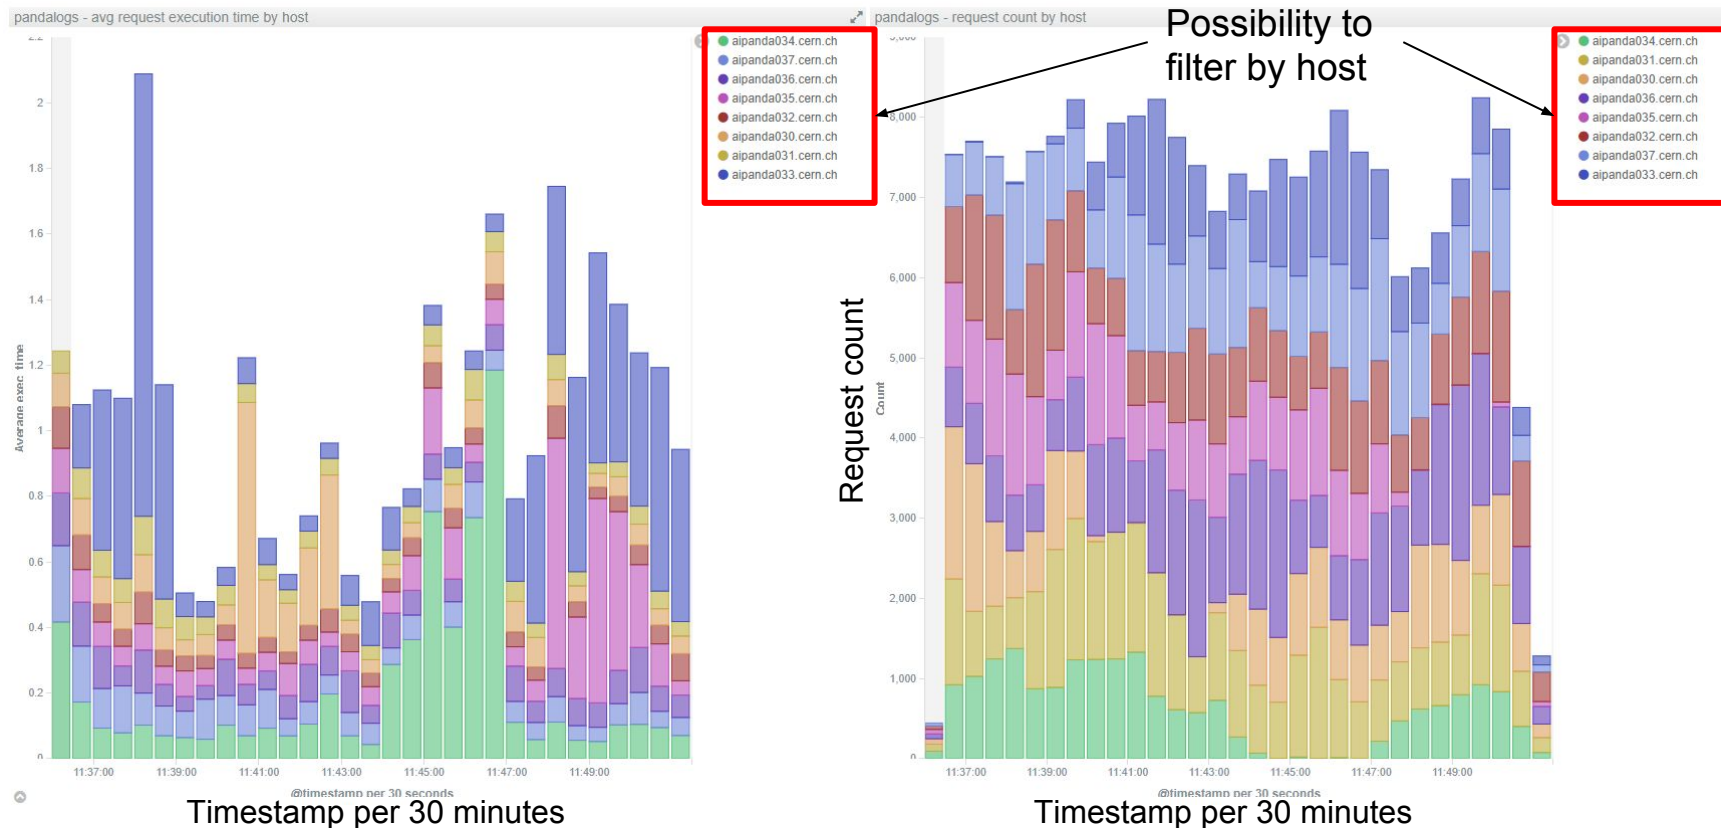
Timestamp per 3 hours

# Example of log processing





# Examples: Panda server health metrics



# Integration with BigPanDA monitoring. Logs page

Logs	Tasks and Jobs	Prod job brokerage	Analy job brokerage	Throttle: queued vs running
------	----------------	--------------------	---------------------	-----------------------------

JEDI	
Type	Level(count)
<a href="#">atlasprodtaskbroker (?)</a>	INFO(14687) DEBUG(2345)
<a href="#">atlasanaljobbroker (?)</a>	INFO(4216525) ERROR(4063)
<a href="#">atlasprodjobbroker (?)</a>	INFO(919857) ERROR(678)
<a href="#">jobgenerator (?)</a>	DEBUG(1378069) INFO(262862) WARNING(2) ERROR(4771)
<a href="#">atlasprodjobthrottler (?)</a>	INFO(83318) WARNING(15829) ERROR(1)
<a href="#">jobsplitter (?)</a>	DEBUG(495025)
<a href="#">taskrefiner (?)</a>	
<a href="#">postprocessor (?)</a>	
<a href="#">taskcommando (?)</a>	ERROR(2202)
<a href="#">atlastasksetrunner (?)</a>	INFO(606763) ERROR(15)

PANDA Server	
Type	Level(count)
<a href="#">retrialmodule (?)</a>	INFO(36999) ERROR(3)
<a href="#">dbproxyfiltered (?)</a>	INFO(70979) ERROR(13357)
<a href="#">serveraccess (?)</a>	INFO(8202664)
<a href="#">entry (?)</a>	INFO(7990110) ERROR(36)
<a href="#">userif (?)</a>	INFO(72)
<a href="#">pilotrequest (?)</a>	DEBUG(3717655) INFO(1235203)
<a href="#">servererror (?)</a>	INFO(22364)
<a href="#">activatorlog (?)</a>	DEBUG(43772)
<a href="#">copyarchive (?)</a>	DEBUG(7088)
<a href="#">finisher (?)</a>	DEBUG(604927)

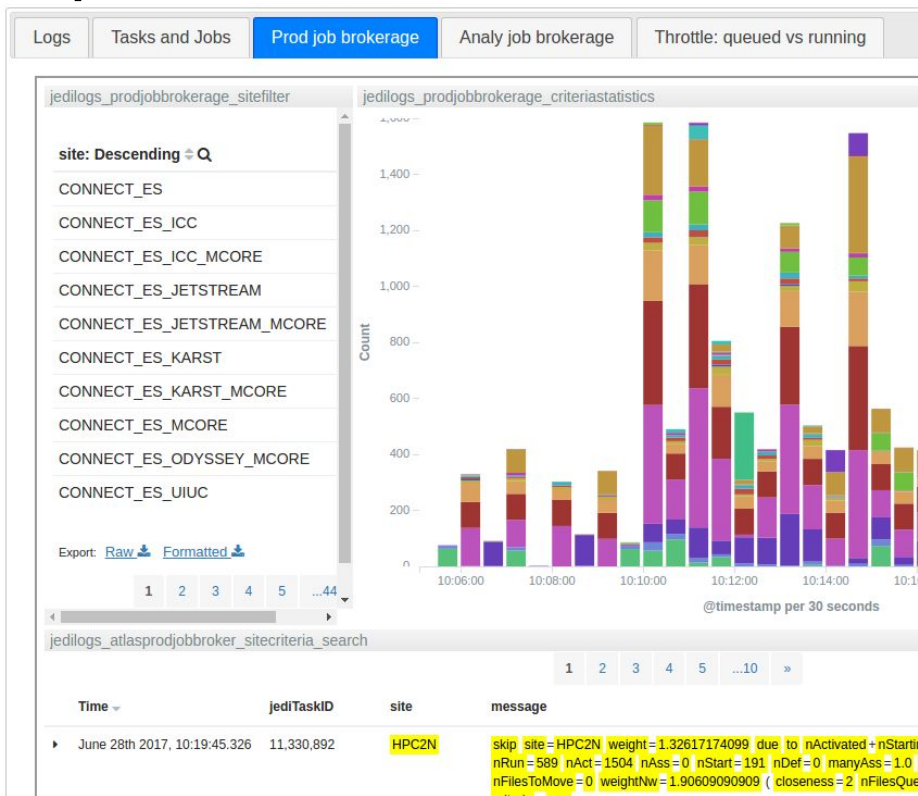
**Tooltip**

Generates tasks in JEDI from definitions found in DEFT

Log types ordered by interest

Log level and number of entries

# Integration with BigPanDA monitoring. Dashboards implementation



# Notification system. Example of notification message

Reply Reply All Forward X Junk Close

**Error from aipanda107.cern.ch**  
noreply@mail.cern.ch [noreply@mail.cern.ch]  
Sent: 04 September 2017 07:33  
To: Aleksandr Alekseev; Tatiana Korchuganova; Siarhei Padolski

**Error description**

```
2017-09-04 05:33:27,025 Internal Server Error: /esatlaslogger/ Traceback (most recent call last): File
"/data/wenaus/virtualenv/twrpm/lib/python2.7/site-packages/django/core/handlers/exception.py", line 39, in inner response = get_response(request) File
"/data/wenaus/virtualenv/twrpm/lib/python2.7/site-packages/django/core/handlers/base.py", line 249, in _legacy_get_response response =
self._get_response(request) File "/data/wenaus/virtualenv/twrpm/lib/python2.7/site-packages/django/core/handlers/base.py", line 187, in _get_response
response = self.process_exception_by_middleware(e, request) File "/data/wenaus/virtualenv/twrpm/lib/python2.7/site-
packages/django/core/handlers/base.py", line 185, in _get_response response = wrapped_callback(request, *callback_args, **callback_kwargs) File
"/data/wenaus/virtualenv/twrpm/pythonpath/core/./core/views.py", line 9438, in esatlasPandaLogger cat = pandaDesc[agg['key']][0] KeyError:
u'panda.log.Watcher'
```

More information here:

```
->https://es-atlas.cern.ch/kibana/app/kibana#/discover?_g=(refreshInterval:(display:Off,pause:!f,value:0),time:(from:'2017-09-
04T05:31:00.000Z',mode:absolute,to:'2017-09-04T05:35:00.000Z'))&_a=(columns:!(request),index:'atlas_bigpanda-*',interval:auto,query:(query_string:
(analyze_wildcard:!t,lowercase_expanded_terms:!f,query:'*')),sort:!(('@timestamp',desc))<-
```

## Detailed information

# Notification system. Information about error in ES



Total entries  
within 4  
minutes

September 4th 2017, 12:31:22.000 /worldjobs/

September 4th 2017, 12:33:27.025 -

September 4th 2017, 12:33:25.000 /esatlaslogger/

Django log  
with error  
description

Table	JSON	View summary
@timestamp	September 4th 2017, 12:33:27.323	
@version	1	
_id	AV5LYSGnBX3BKxyzfki	
_index	atlas_bigpanda-2017.09.04	
_score	-	
_type	port80_access	
agent	"Mozilla/5.0 (Windows NT 10.0; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/60.0.3112.90 Safari/537.36 OPR/47.0.2631.71"	
auth	-	
beat.hostname	aipanda107.cern.ch	
beat.name	aipanda107.cern.ch	
beat.version	5.4.1	
bytes	8796	
clientip	109.123.147.226	
host	aipanda107.cern.ch	

Apache log  
with user  
request

# Summary

- **Logstash infrastructure is set up to process PanDA server, JEDI and BigPanda logs**
  - 67 filters for PanDA, JEDI and BigPanDA logs are ready and in production
- **Improve the lifetime model of the logs**
  - for the moment we are only keeping all of them 20 days
- **The system has a lot of potential to monitor, investigate and understand PanDA and ProdSys**
  - The infrastructure is ready and several dashboards for real-life issues have been implemented
  - We need to continue looking/understanding the data and identify further necessary visualizations
- **Results were integrated with BigPanDA monitoring**
  - The special page with was created in BigPanDA monitoring
  - Further integration points with BigPanDA will be added
- **Notification system was created**
  - Error messages are sent to the BigPanDA developers immediately
  - Error messages are stored in ES

Thank you for your attention!