The BDP Infrastructure for monitoring and analysing the ATLAS experiment processing activities at INFN-CNAF Tier-1

Giacomo Levrini, University of Bologna giacomo.levrini@studio.unibo.it
Aksieniia Shtimmerman, CNAF INFN ashtimmerman@infn.it

CHEP 27th International Conference on Computing in High Energy and Nuclear Physics. 21-25 October 2024. Krakow, Poland

The CNAF Tier1 in the Italian supercomputing distributed infrastructure.

1 Tier1:
- 1.2MHSpec CPU power equivalent to 79'000 cores
- 80 PB disk space
- 65 PB tape library

Italy provides Tier-1, Tier-2 and Tier-3 facilities to the ATLAS collaboration. The Tier-1, located at CNAF, Bologna, is the main center, also referred to as the regional center. The Tier-1, together with the other Italian centers, provides both resources and expertise to the ATLAS computing community and manages the so-called Italian Cloud of Computing.

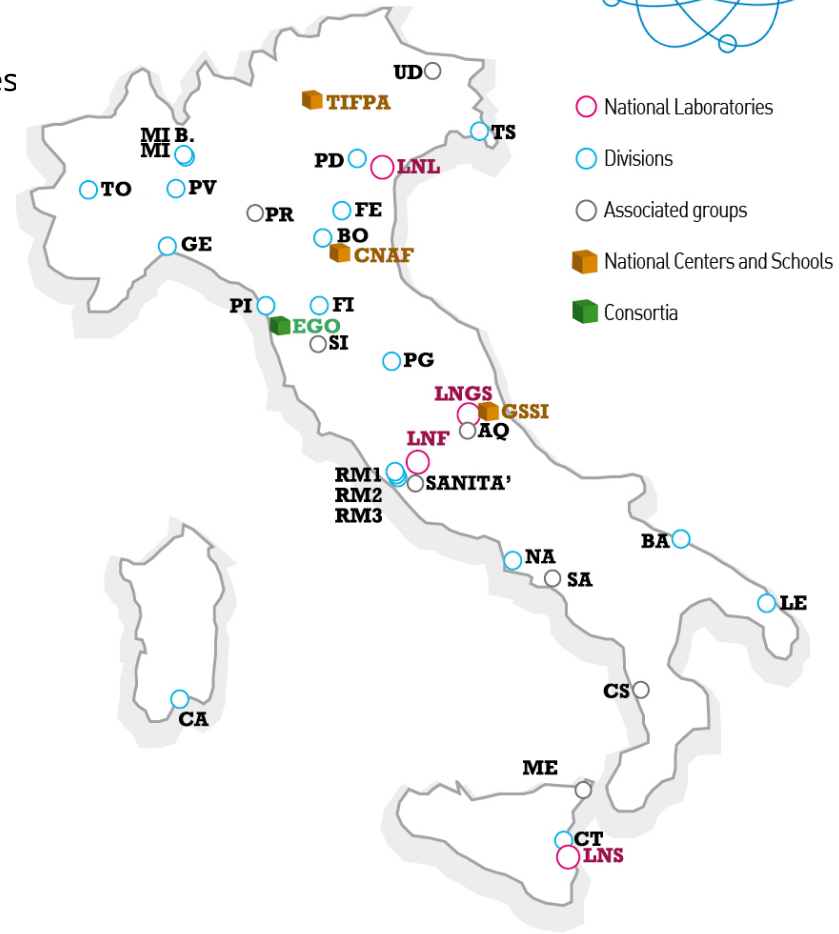

CNAF Big Data Platform

The CNAF group at INFN has implemented a Big Data Platform (BDP) infrastructure, designed for the collection and the indexing of log reports form CNAF facilities. The infrastructure is an ongoing project at CNAF and it is available for the italian groups working in high energy physics experiments.

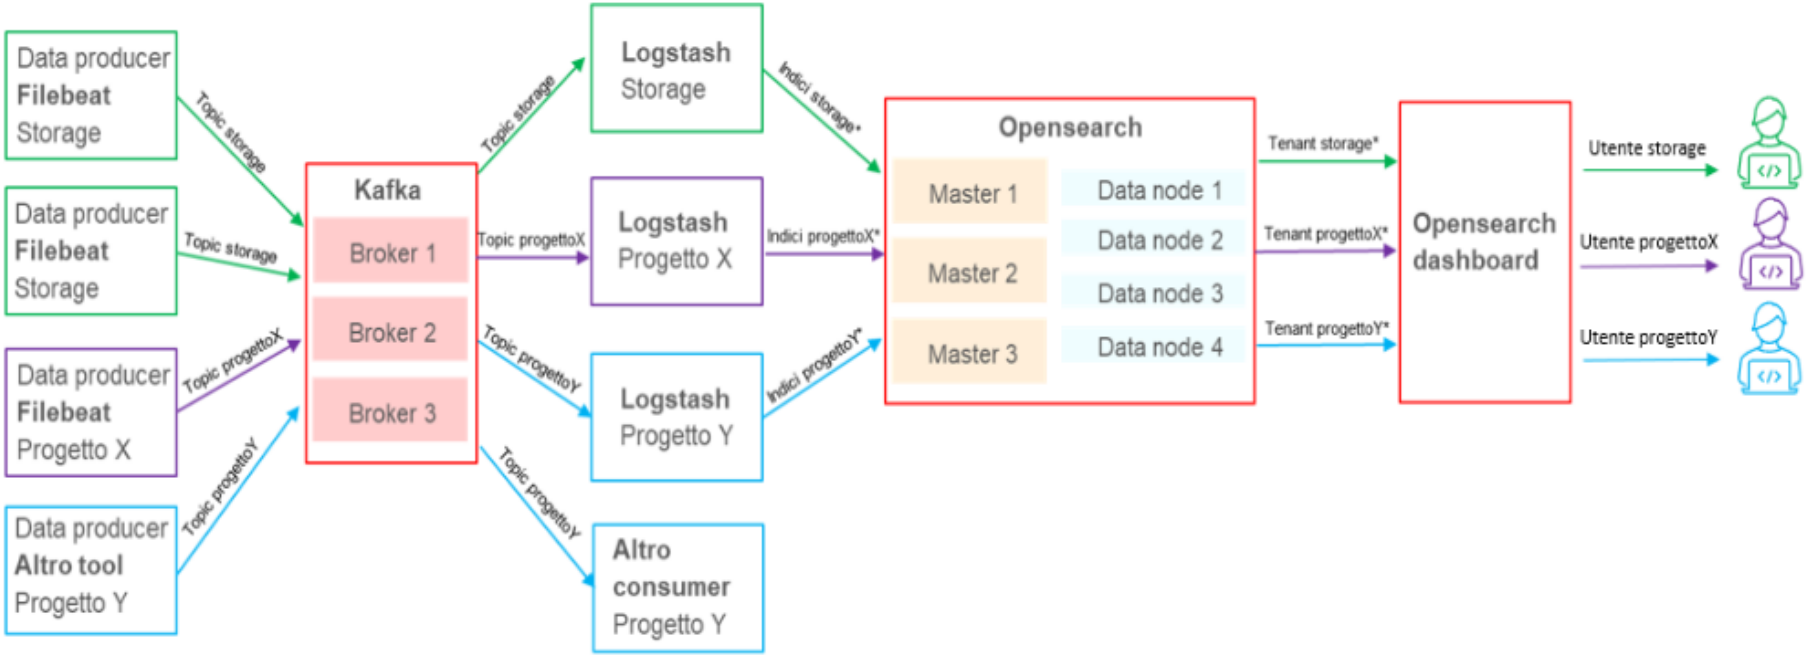

ATLAS jobs log in CNAF BDP

Within this framework, the first data pipeline was established for the ATLAS experiment at CERN, using input from the ATLAS Distributed Computing System PanDa.

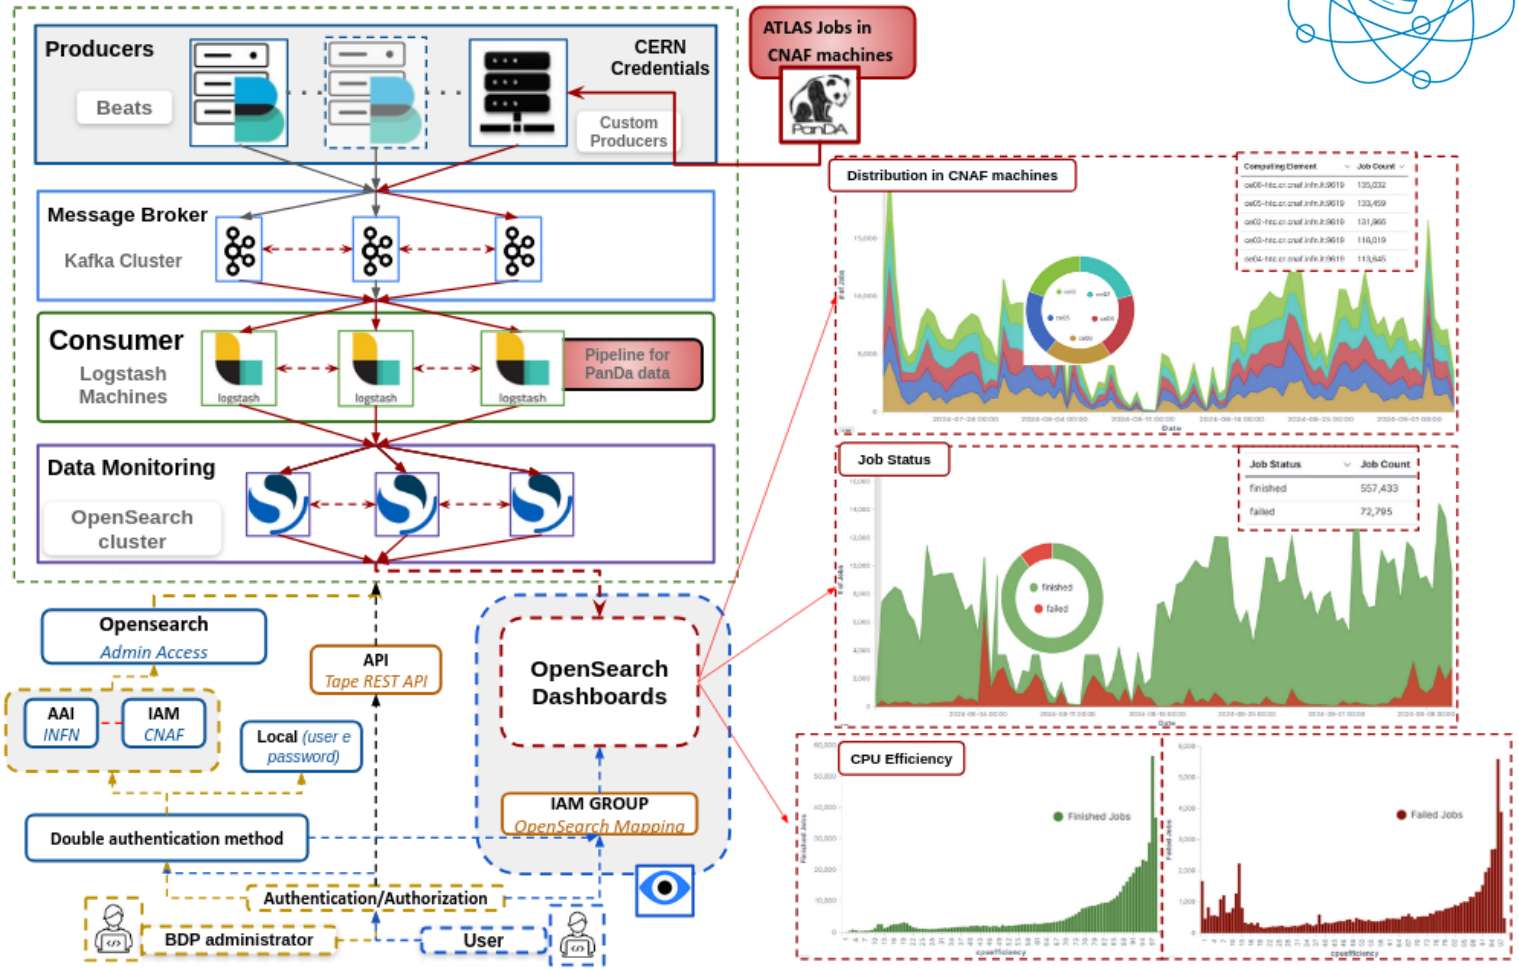

Structure of the BDP

ATLAS Job addressed to Tier-1 machines in Bologna are selected with a PanDA API, a summarizing log report is ingested in the data indexing pipeline in a compact JSON format with a custom producer.

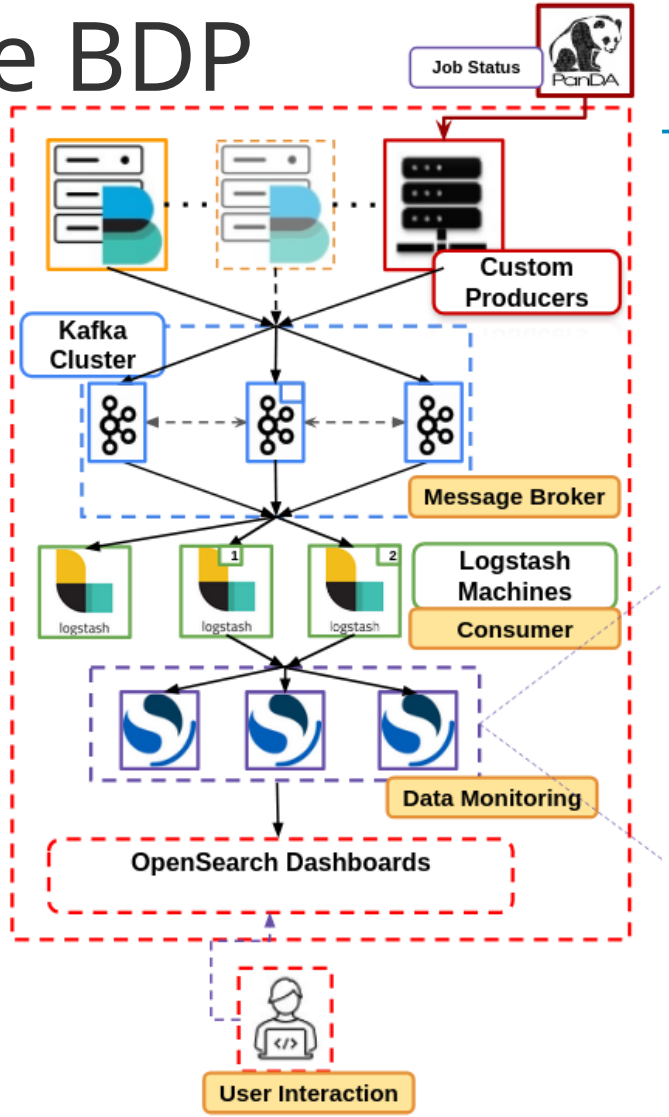

Pipeline BDP

This pipeline focuses on the ATLAS computational job data processed by the Italian INFN Tier-1 computing farm.

```
[aashtimmerman@cnlog-logstash02 ~]$ ll /etc/logstash/config
total 64
-rw-r--r-- 1 root     root   2178 Jul 23 10:55 filters.conf
-rw-r--r-- 1 root     root    654 Jul 15 12:03 input.conf
-rw-r--r-- 1 12011476 users  1833 Jul 19  2023 jvm.options
-rw-r--r-- 1 12011476 users  7437 Jul 19  2023 log4j2.properties
-rw-r--r-- 1 root     root    342 Jul 15 11:15 logstash-sample.conf_orig
-rw-r--r-- 1 root     root     31 Jul 15 11:31 logstash.yml
-rw-r--r-- 1 12011476 users 15104 Jul 11 11:11 logstash.yml_orig
-rw-r--r-- 1 root     root    448 Jul 15 12:58 output.conf
-rw-r--r-- 1 root     root    106 Jul 15 11:32 pipelines.yml
-rw-r--r-- 1 12011476 users  5204 Jul 11 11:10 pipelines.yml_orig
-rw-r--r-- 1 12011476 users  1696 Jul 19  2023 startup.options
```

```
logstash$ cat /etc/logstash/config/pipelines.yml

 pipeline.id: main
 path.config: "/etc/logstash/config/*.conf"
 pipeline.workers: 2
 queue.drain: false
logstash0$ cat /etc/logstash/config/input.conf
input {
  kafka {
    codec => json
      bootstrap_servers => "kafka01.cr.cnaf.infn.it:9192,kafka02.cr.cnaf.in
fn.it:9192,kafka03.cr.cnaf.infn.it:9192"
      sasl_jaas_config => "org.apache.kafka.common.security.plain.PlainLogi
nModule required username='******' password='*****';"
      security_protocol => "SASL_SSL"
      sasl_mechanism => "PLAIN"
      ssl_truststore_location => "/root/truststore.jks"
      ssl_truststore_password => "******"
      group_id => "atlas"
      auto_offset_reset => "earliest"
      topics => ["atlas"]
      id => "logstash"
  }
}
```

Data Logs from PanDa

The system has been operational and effective for several years, marking our initiative as the first to integrate job information directly with the infrastructure.

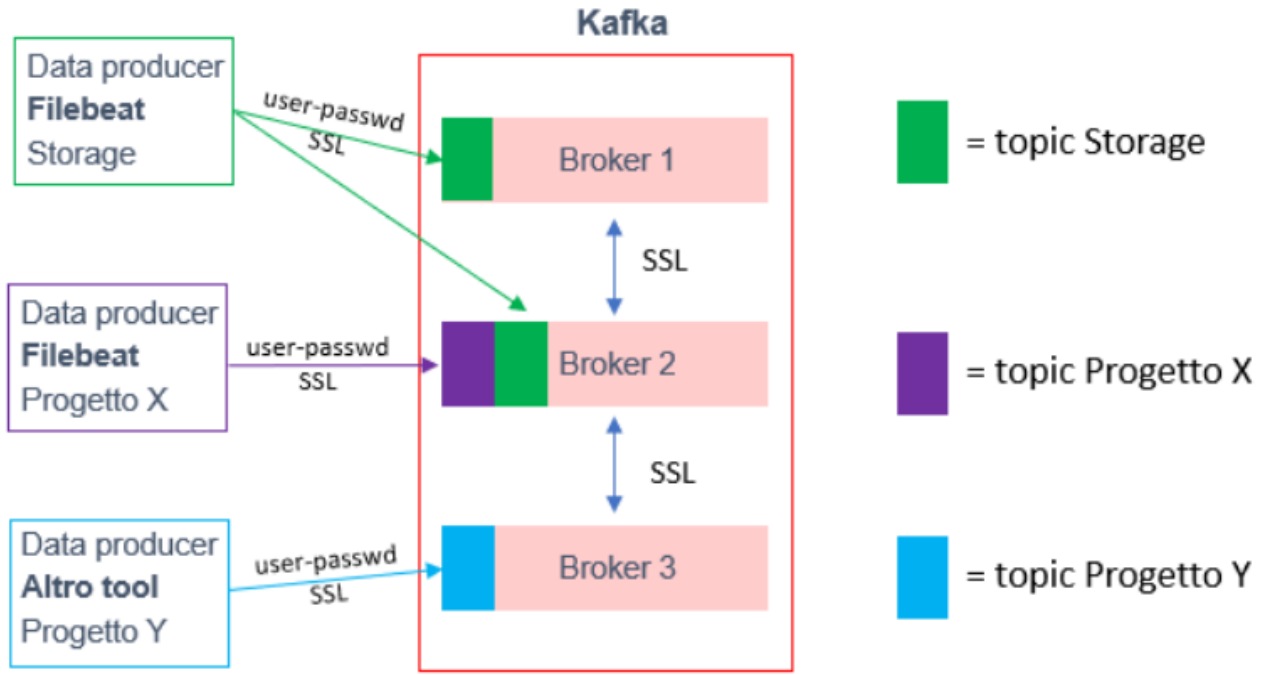

CNAF Big Data Platform

The save ATLAS computational job data and analysis of the jobs usage Big Data Platform of CNAF is the most important factor when trying to analyse the overall efficiency of the ATLAS computational job in the Italian big data center.

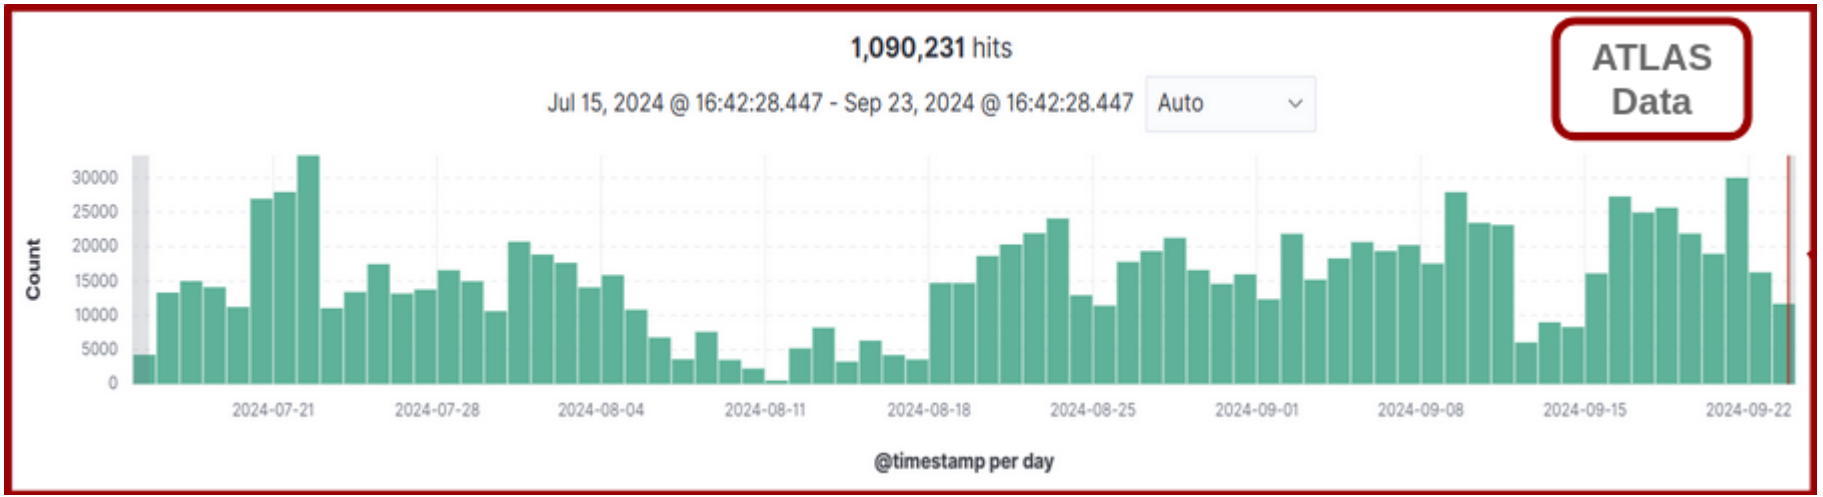

Jobs Distribution in CNAF VM 9

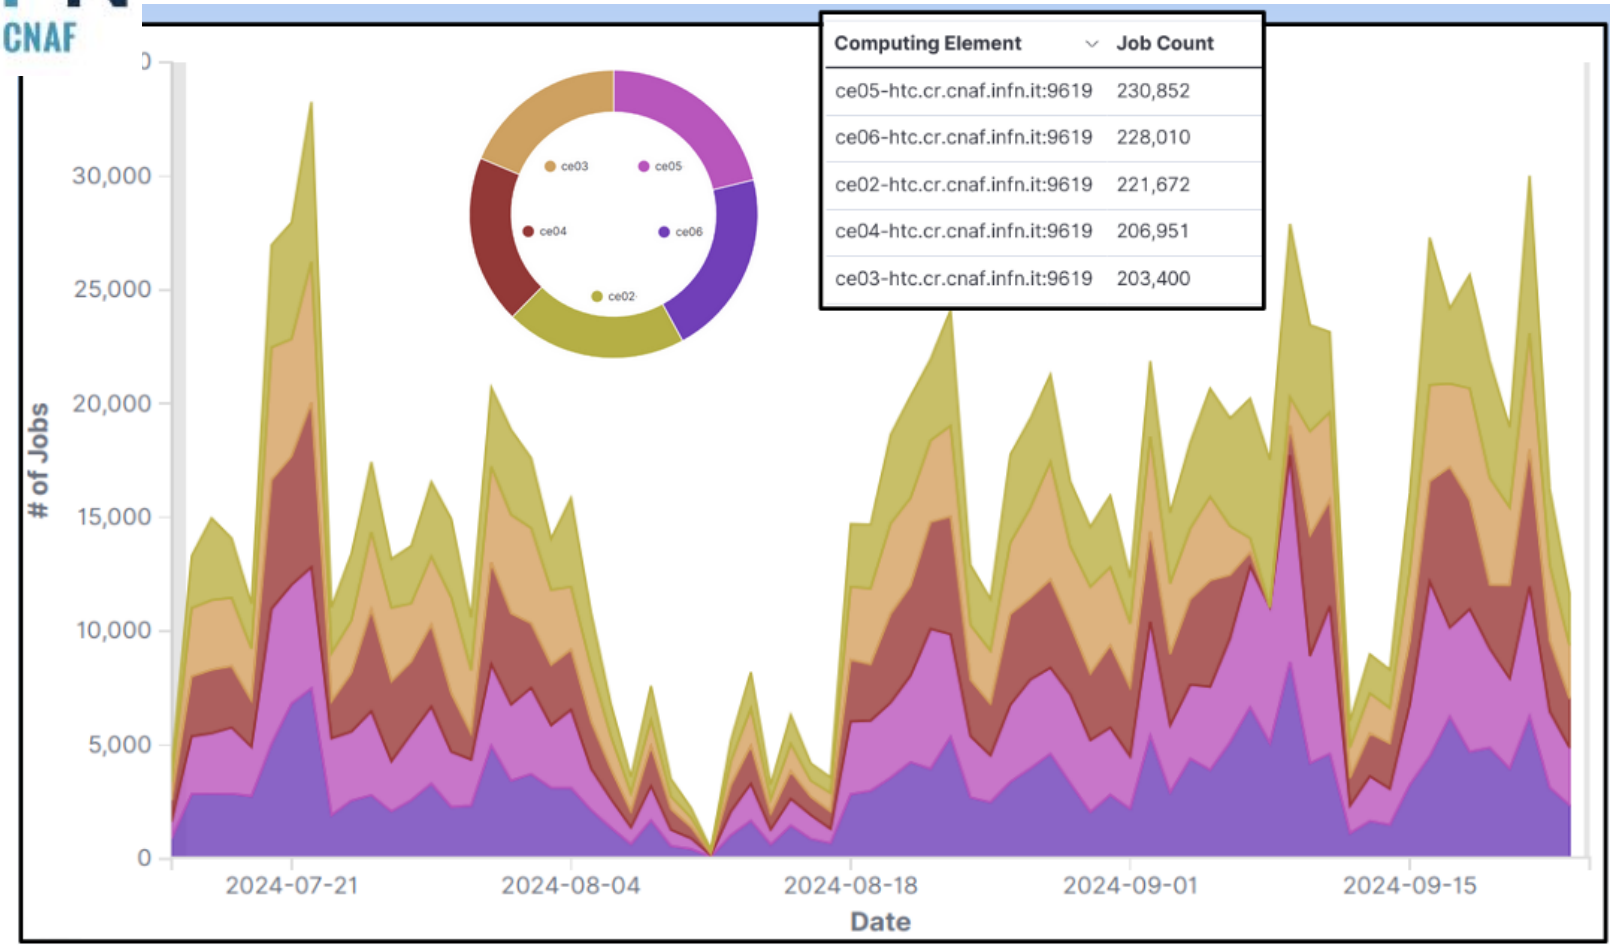

Jobs Status

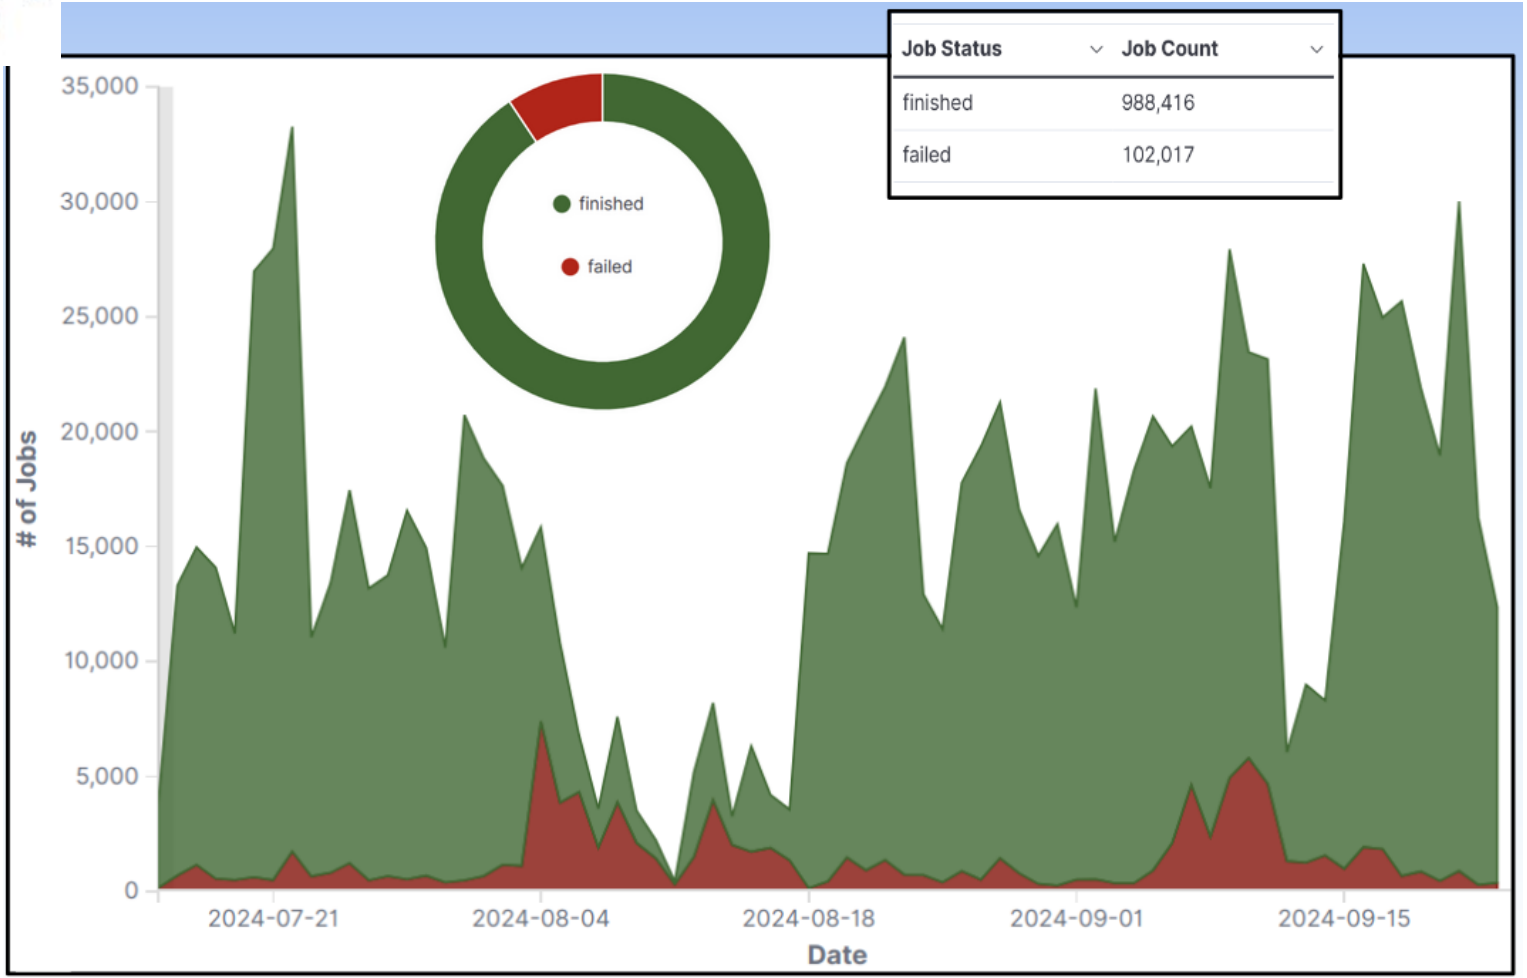
Job Status
Job Count
finished
988,416
failed
102,017
finished
failed
35,000
30,000
25,000
20,000
15,000
10,000
5,000
0
# of Jobs
2024-07-21
2024-08-04
2024-08-18
2024-09-01
2024-09-15
Date

Security and authorization

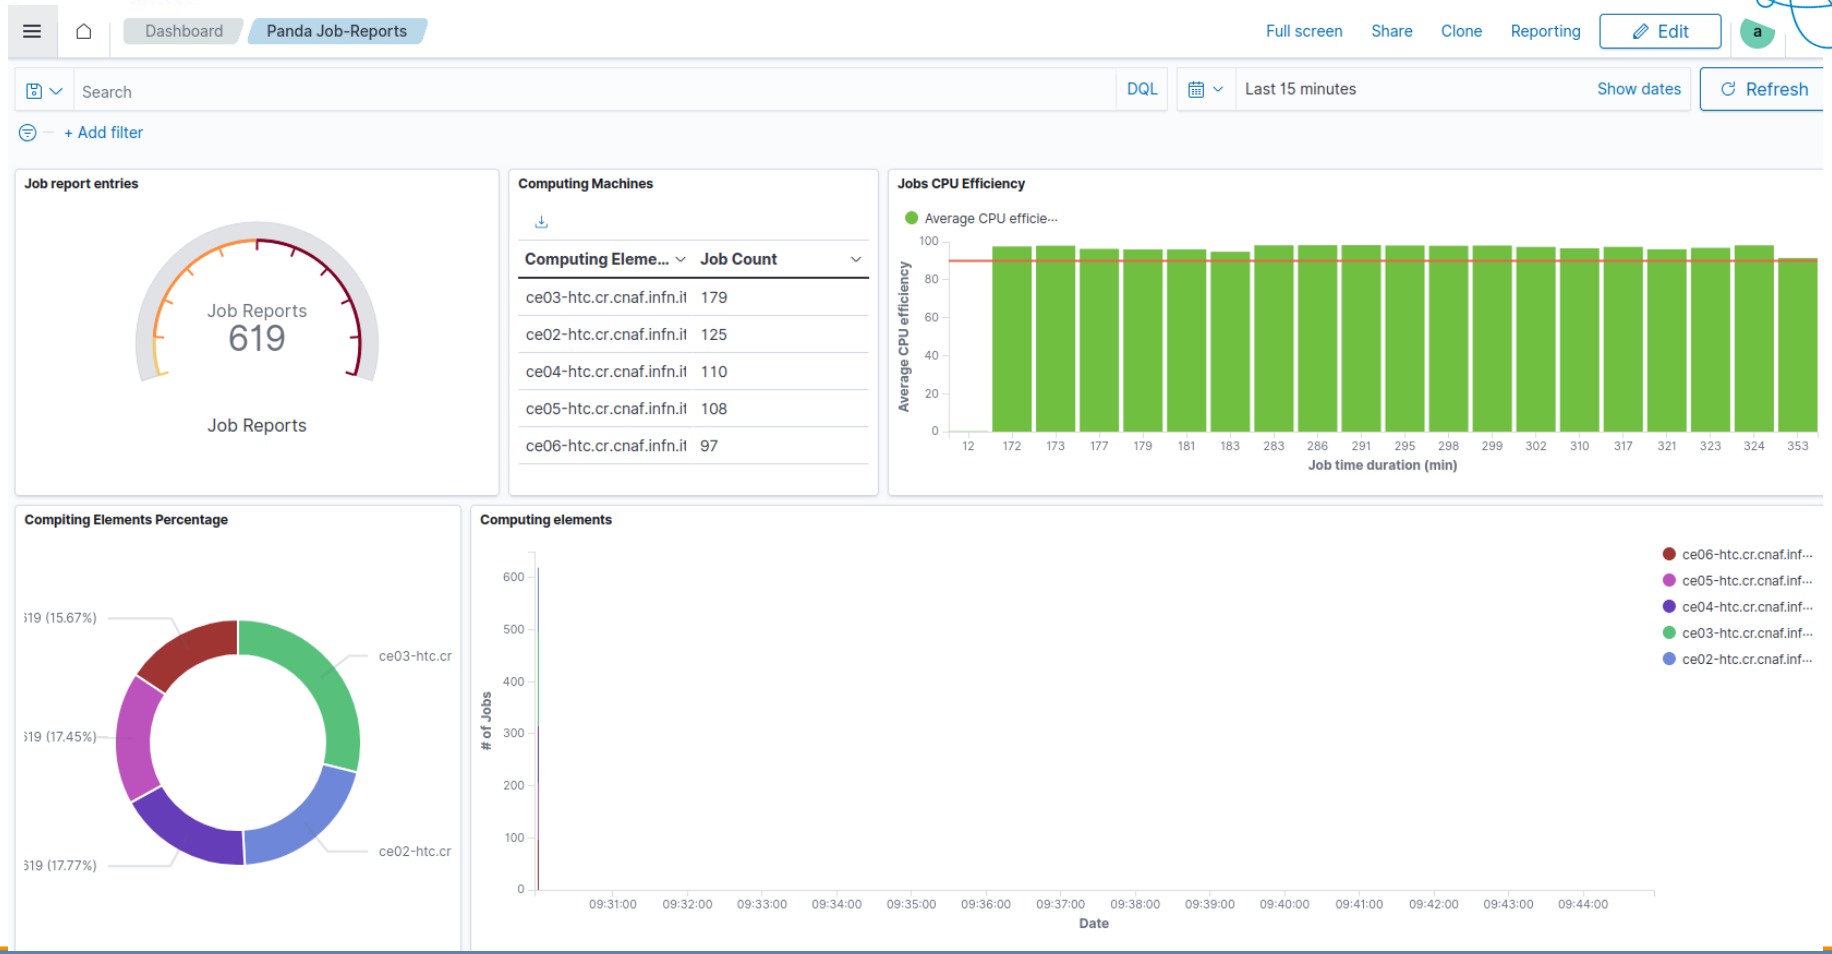

Error Messages

Different type of errors have been written inside the logs. Data collection will proceed along to have more entries with different errors.

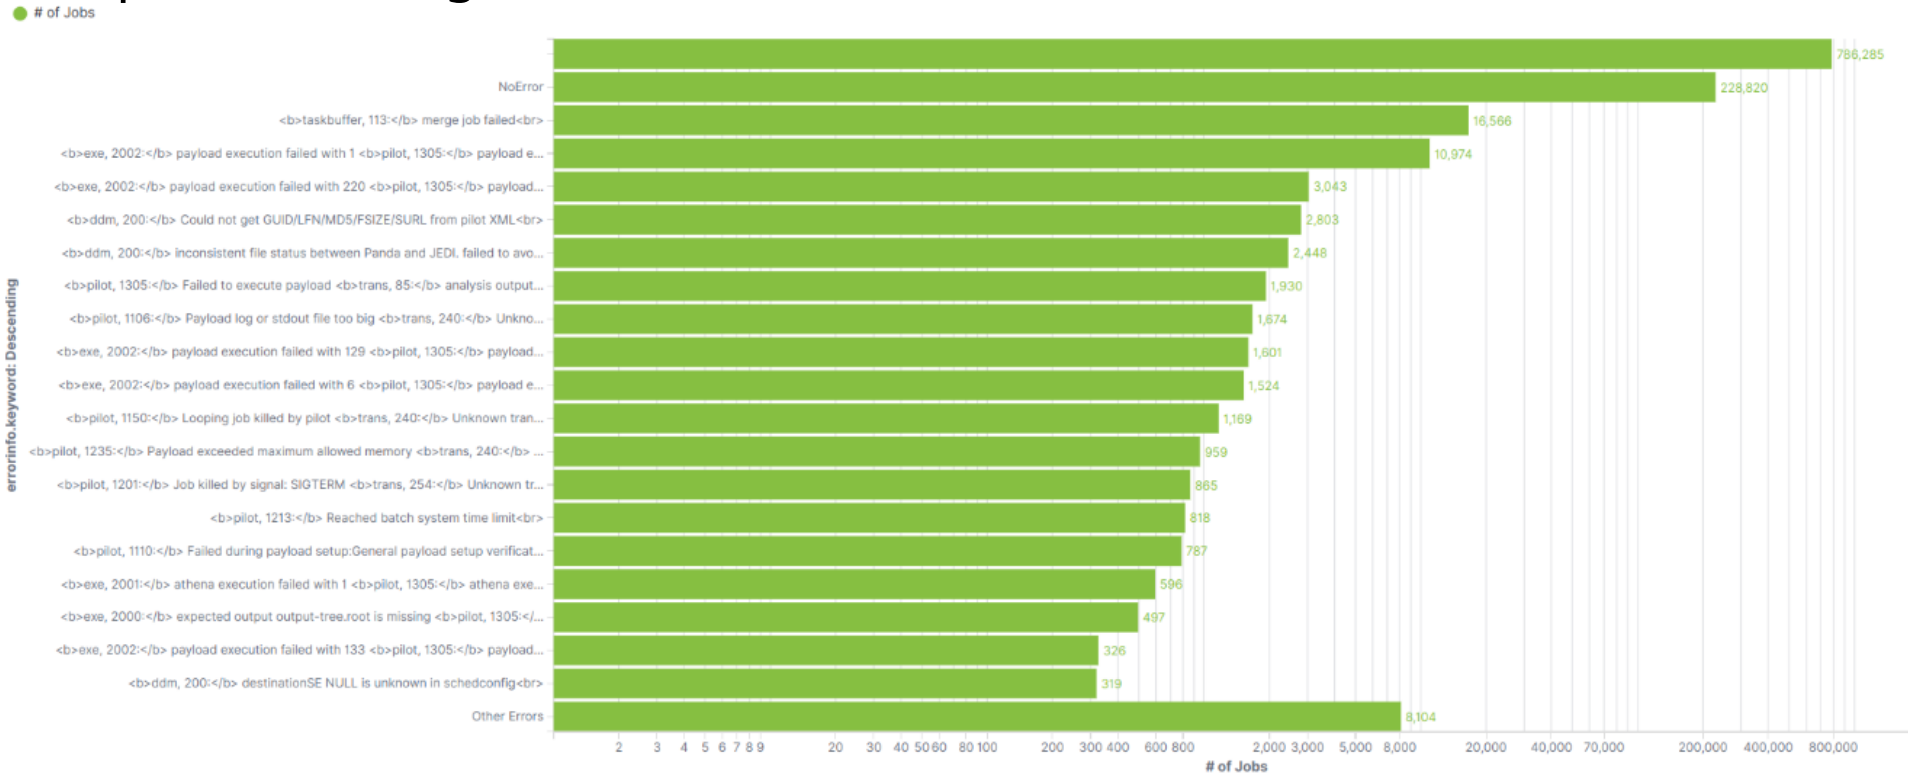

Security and authorization

Users and Admins access to Opensearch is granted via many authorization options:

- Opensearch Dashboard or API;
- Double authentication method :
  - user and password;
  - local IAM CNAF and AAI INFN (iam.cnaf.infn.it);
- Authorization with IAM group (associated with permissions):
  - User IAM groups are automatically mapped to OpenSearch backend role;
  - Every User assigned to a given project (topic) may read or write inside the branch of the corresponding assigned topic, granting access to indices inside OpenSearch.

Next steps

Machine Learning

The The main goal for the future evolution of the data ingestion inside the BDP is to start to use a ML plugins in OpenSearch, whose scope will be to investigate failed jobs, categorizing the reason of the failures and eventually seeking for anomalies inside the machines. We would like to implement and train ML algorithm, whose final scope will be to read and categorize the whole data throughput coming from PanDa.

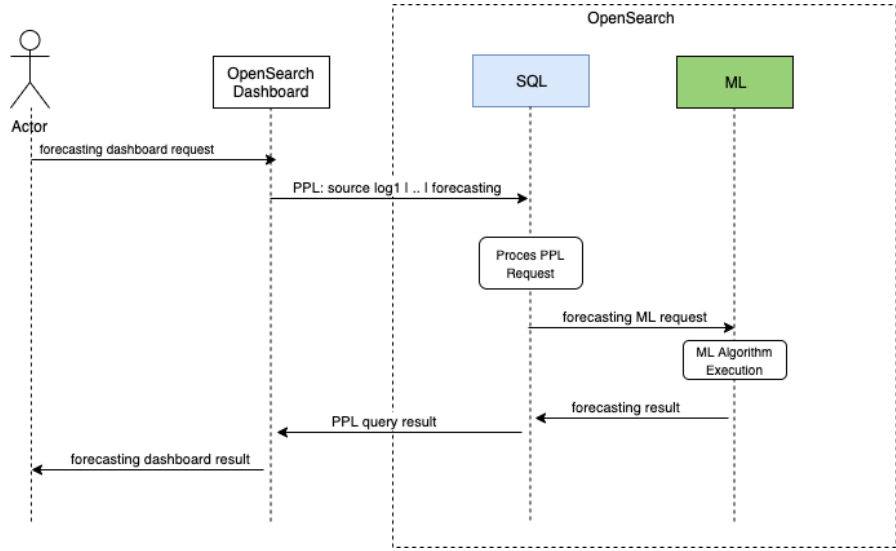

Thank you!

Contacts:

Giacomo Levrini, glevrini@bo.infn.it

Aksieniia Shtimmerman, ashtimmerman@cnaf.infn.it

Antonio Falabella, antonio.falabella@cnaf.infn.it

Enrico Fattibene, enrico.fattibene@cnaf.infn.it

Diego Michelotto, diego.michelotto@cnaf.infn.it

Giusy Sergi, giusy.sergi@infn.cnaf.it

Questions?