ATLAS Jobs data inside CNAF BDP infrastructure

Setup of a BDP cluster to study performances and error analysis of ATLAS Jobs logs

Whoami

- Giacomo Levrini,
- Contact: levrini@bo.infn.it
- Master Degree: Nuclear & Subnuclear Physics
- Thesis on hardware accelerators for ATLAS trigger
- Currently: PhD in Data Science and Computation
- Worked inside the Network group for ATLAS TDAQ
- Now working in collaboration with INFN-CNAF

INFN-CNAF

- **CNAF** (centro nazionale INFN “per la ricerca e lo sviluppo nelle tecnologie informatiche e telematiche”) is primary computational center for INFN
- Hosts **Tier-1** infrastructure for LHC data, providing support to the activity of storage, processing and data analysis.
- Hosts a Big Data Processing (**BDP**) infrastructure for managing and analysis of heterogeneous data.

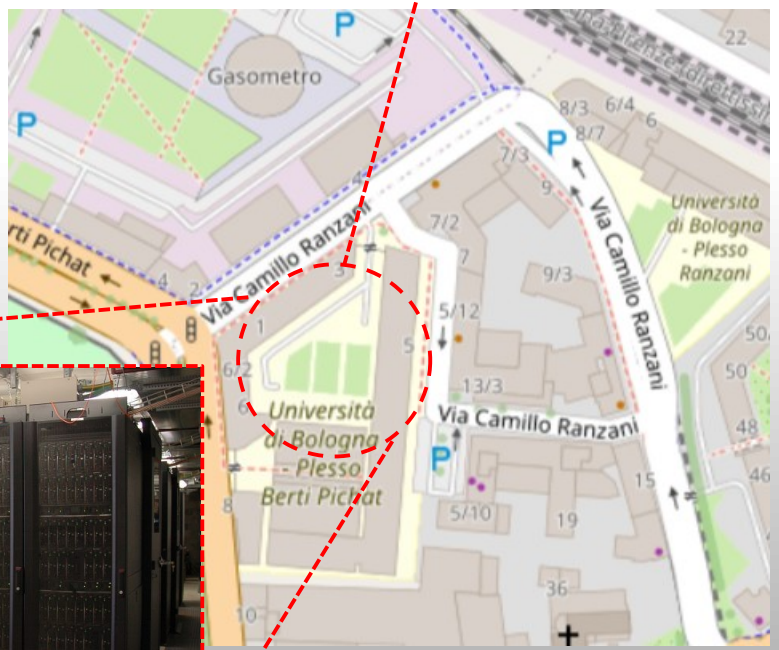

INFN = Istituto Nazionale di Fisica Nucleare

CNAF BDP Infrastructure

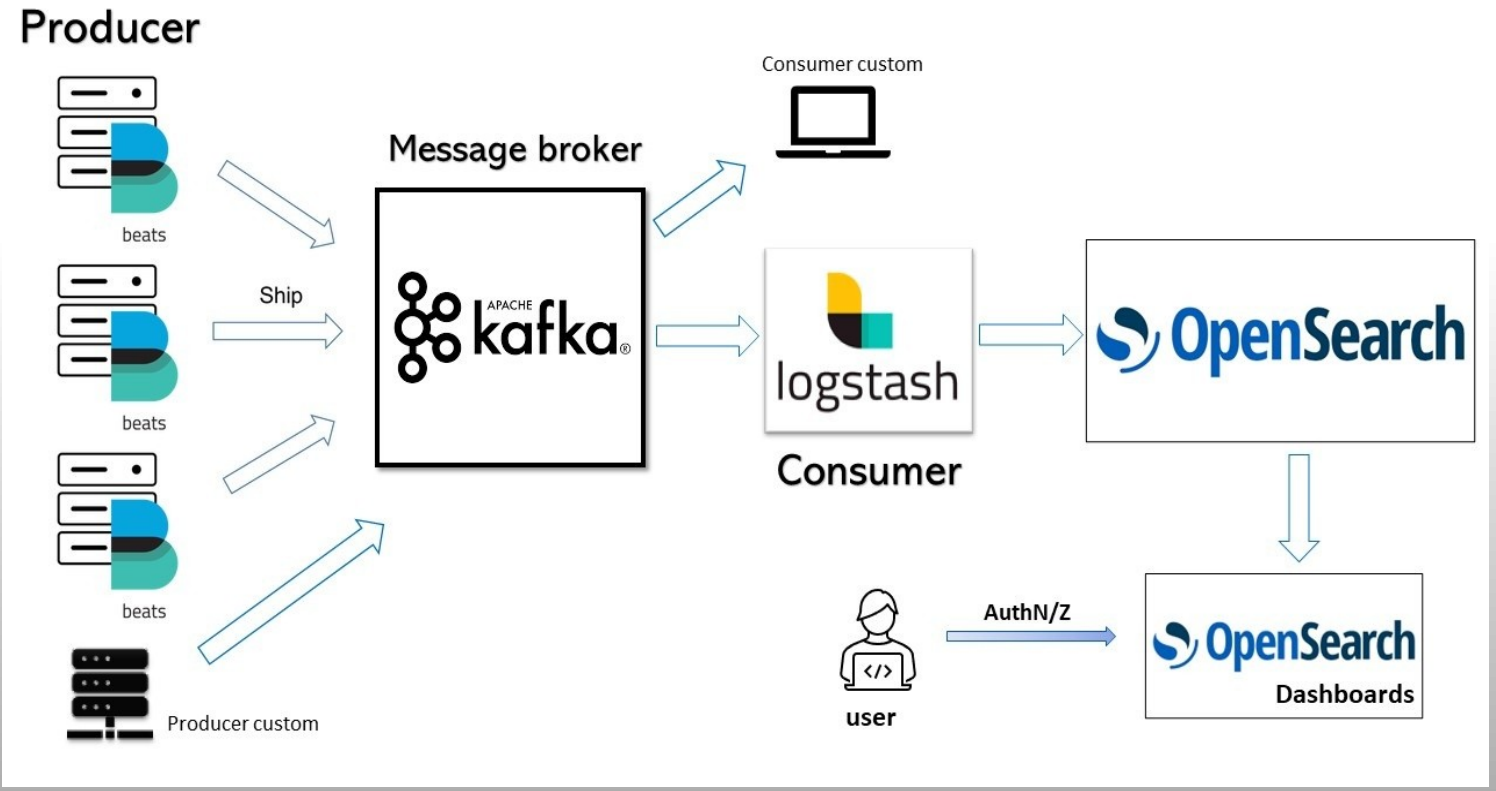

- **Filebeat** sends logs inside the BDP
- **Kafka** manages the data “topics” inside its machines
- **Logstash** reads the topics and sends them to **Opensearch** dashboards

Infrastructure in details

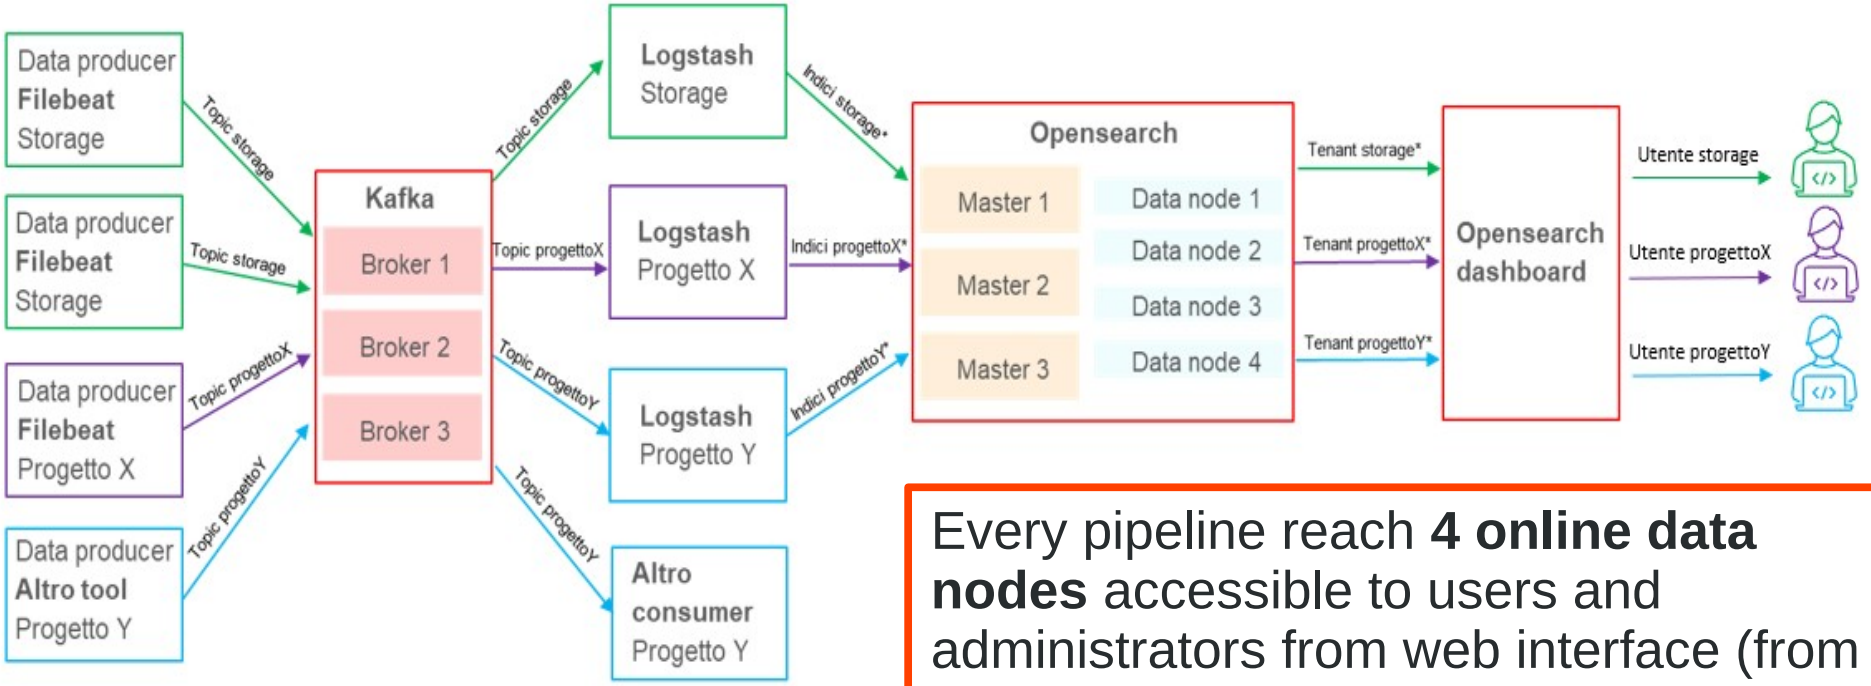

Every pipeline reach **4 online data nodes** accessible to users and administrators from web interface (from Opensearch)

OpenSearch Output Control

- Topic “atlas” is already available inside the BDP, data is flowing constantly (simulated simple messages to check the pipeline)

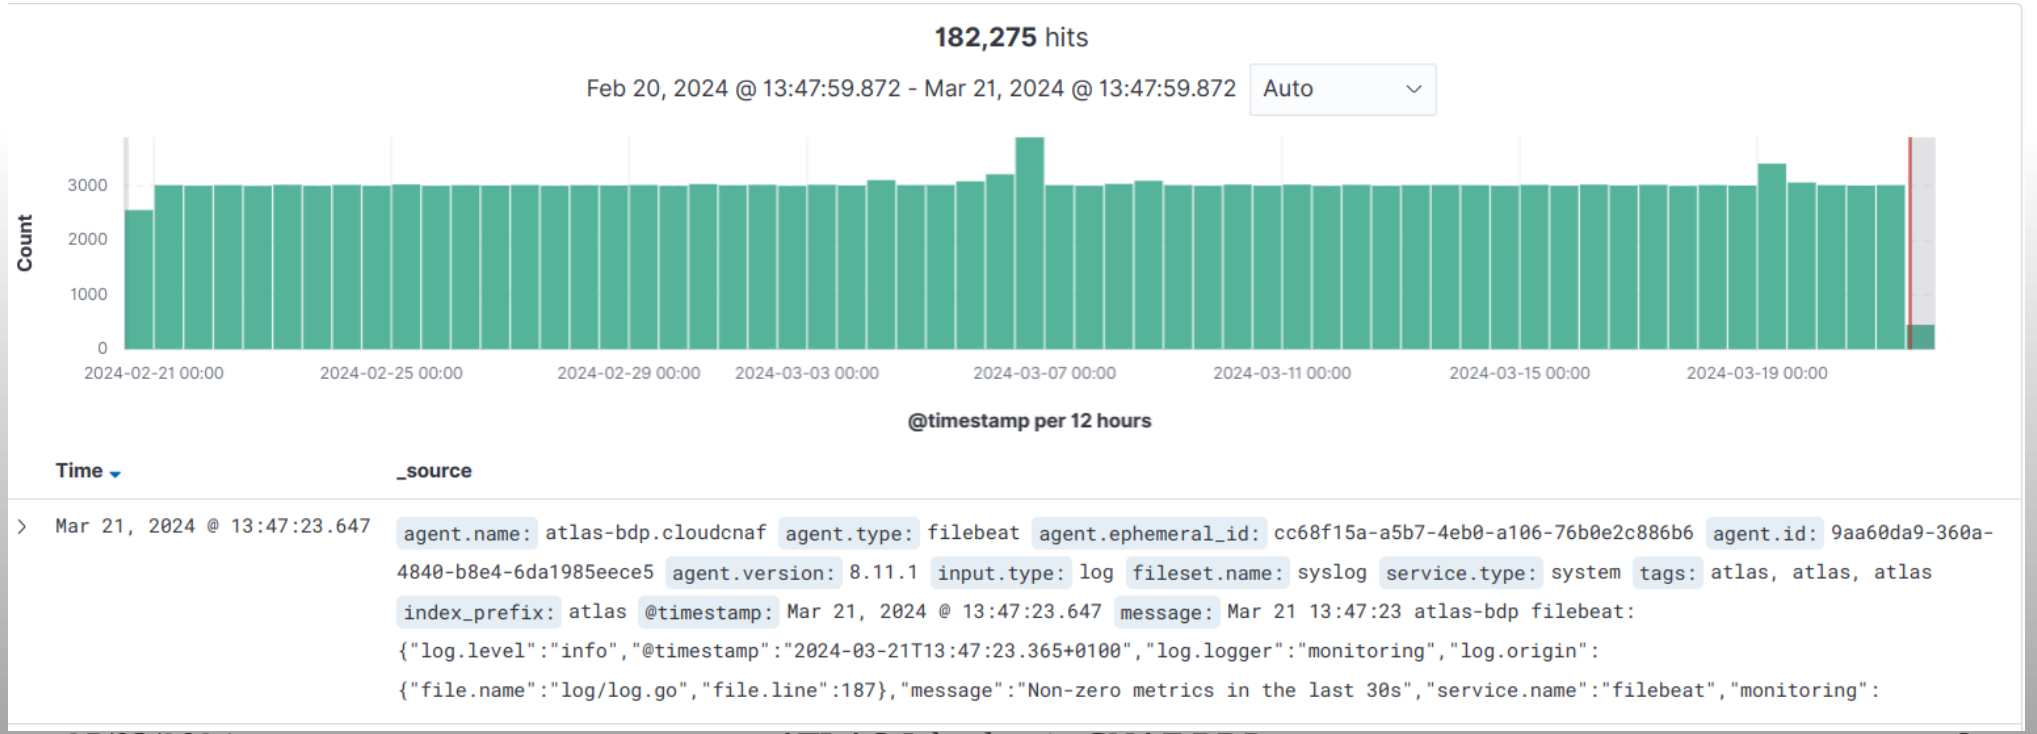

PanDa Jobs Data

Queue name	Queue type	Region	Status	Job type	Resource type	N running slots	N jobs total	% failure	assigned	activated	running	transferring	merging	finished	failed
ANALY_INFN-T1_GPU	analysis	IT	online	all	all	0	65	6.7	1	2	0	0	0	56	4
INFN-CNAF	unified	IT	online	all	all	6,662	18678	1.3	191	722	3694	626	4761	7287	93
INFN-CNAF_ARM	unified	IT	online	all	all	1,160	1178	0	21	362	145	0	0	460	0

- Access to PanDA (Production and Distributed Analysis) servers to collect logs from logs inside Tier-1 machine before they expire
-

Studies of interest

- Almost direct analysis on **failed jobs**
  - Investigate on possible reasons of the errors and failures
- BDP data recollection will allow delayed analysis on **computation performance** of both “long” period of time or specific machines
  - The cluster may keep data online up to 1 month, lately stored offline (still accessible)
  - From log information is possible to track the machines which handled a specific job (from job ID)

Conclusions and future works

- **Currently:** the BDP pipeline is set and functioning for the topic ATLAS, thus able to register properly data
  - Possibility to add in future more specific topics of ATLAS or new topics (in agreement with CNAF administrators)
- **Next step:** activate through ATLAS proxy the extraction of logs from completed jobs
  - Time periodic collection of logs, to be stored in a local machine at CNAF and sent to the BDP with a producer
- **Further step**: start to analyze error logs and machine performances

Thanks for the attention!

Backup – Opensearch indices search

atlas

Index ↓	Health	Managed by policy	Status	Total size	Size of primaries	Total documents
atlas-2024.03.21	Green	No	Open	16.9mb	9.9mb	3210
atlas-2024.03.20	Green	No	Open	32.8mb	16.4mb	6013
atlas-2024.03.19	Green	No	Open	34.2mb	17.1mb	6465
atlas-2024.03.18	Green	No	Open	32.6mb	16.3mb	6013
atlas-2024.03.17	Green	No	Open	32.7mb	16.3mb	6014
atlas-2024.03.16	Green	No	Open	32.8mb	16.4mb	6025
atlas-2024.03.15	Green	No	Open	32.7mb	16.3mb	6008