



Nikhef



# Testing framework and monitoring system for the ATLAS EventIndex

Elizaveta Cherepanova<sup>1\*</sup>, Fedor Prokoshin<sup>2</sup>,  
Miguel Villaplana Pérez<sup>3</sup>

1: University of Amsterdam/Nikhef, 2: JINR Dubna, 3: IFIC Valencia

**CHEP2023**

May 8-12, 2023

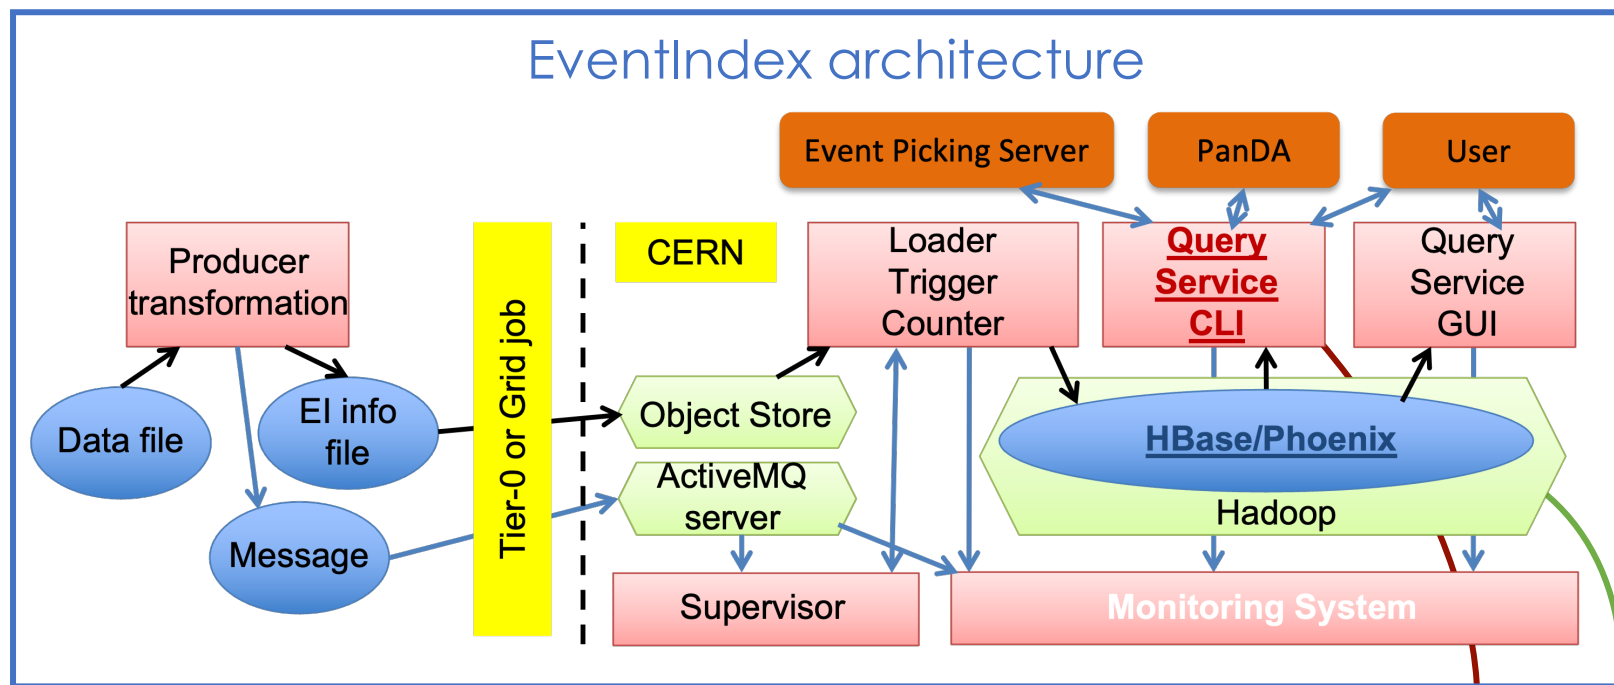
[\\*elizaveta.cherepanova@cern.ch](mailto:elizaveta.cherepanova@cern.ch)



The ATLAS EventIndex (EI) is the global catalogue of all events collected, processed or generated by the ATLAS experiment at the CERN LHC accelerator

The system provides:

- a way to collect and store event information using modern technologies
- various tools to access this information through command line and web services
- an indexing system that points to these events in millions of files scattered through a worldwide distributed computing system



More information in the posters:

- [Deployment and Operation of the ATLAS EventIndex for LHC Run 3](#)
- [Query Service for the new ATLAS EventIndex system](#)
- [HBase/Phoenix-based Data Collection and Storage for the ATLAS EventIndex](#)

The complexity of the EventIndex system requires the monitoring of its sub-systems and processes

## Producer:

- collects data about status and the performance of the system (~15k values a day)
  - scan of the logs and hdfs space
  - REST requests
  - analysis of the web-pages
- transfer data to the InfluxDB through the HTTP requests

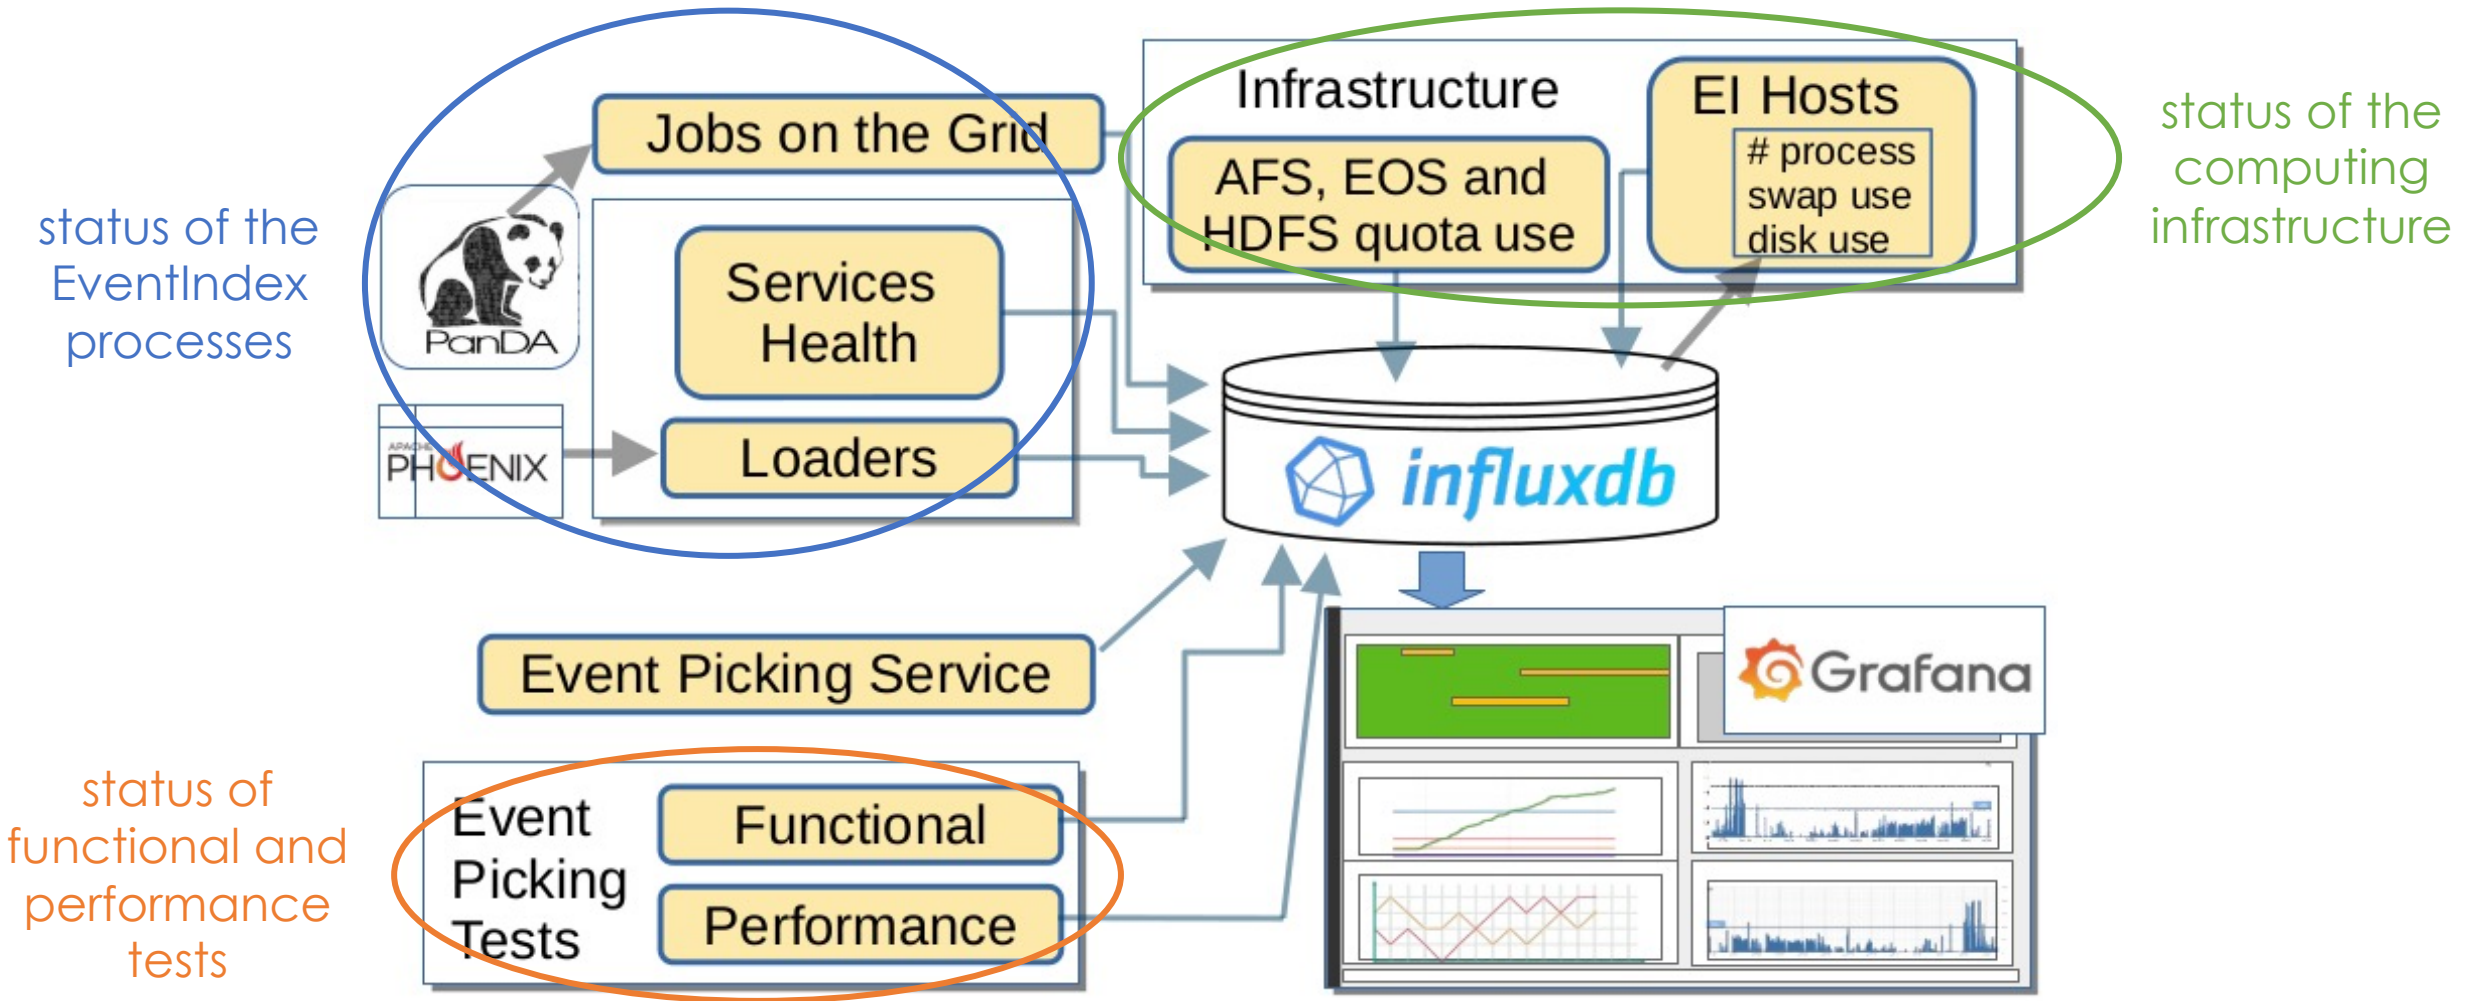


## Viewer:

- graphical presentation of the data – Grafana
- every EI sub-component has a status dashboard
- important parameters of each module have dedicated dashboards

- Scheduler (acrontab) controls data collection and processing
- The status of each EI service is also fed into the global ATLAS service monitoring view

Functional schema for the monitoring system of the EventIndex components





Test event picking in all newly-indexed samples both locally and via PanDA

Dataset indexation is triggered automatically when a dataset is marked as complete in AMI (the ATLAS metadata catalog)

The test runs **twice per week**: Monday and Thursday:

- scan EventIndex catalog and get a list of datasets indexed between 24 hours and 4 days ago
- for each new sample found
  - get reference to file (GUID) and the runNumber-eventNumber pair of a random event in the dataset
  - test event lookup **locally** using CLI and runNumber-eventNumber pair
    - If unsuccessful log sample name
    - If successful **send 2 GRID jobs**, which pick the given event
      - Using the runNumber-eventNumber pair
      - Using the runNumber-eventNumber pair and the GUID
  - Send the complete list of samples that failed local event lookup to EventIndex Ops mailing list

- Check test jobs' status in PanDA:

Task ID	Request ID	Task name	N files total	N files done	N files failed	%	Status (JEDI)	Duration, days	Task logged status	Jobs failure, %
32889647	706168	group.proj-evind.EPRAW.NG.v0.data22_13p6TeV.00440613.physics_ZeroBias.merge.AOD.f1321_m2153.20230327/	1	1	0	100	done	0.04	-	0
32889646	706167	group.proj-evind.EPRAW.YG.v0.data22_13p6TeV.00440613.physics_ZeroBias.merge.AOD.f1321_m2153.20230327/	1	1	0	100	done	0.13	-	0
32889645	706166	group.proj-evind.EPRAW.NG.v0.data22_13p6TeV.00440613.physics_Main.merge.AOD.f1321_m2153.20230327/	1	1	0	100	done	0.16	-	16.67
32889644	706165	group.proj-evind.EPRAW.YG.v0.data22_13p6TeV.00440613.physics_Main.merge.AOD.f1321_m2153.20230327/	1	1	0	100	done	0.03	-	0
32889643	706156	group.proj-evind.EPRAW.NG.v0.data22_13p6TeV.00440613.physics_Main.deriv.DAOD_PHYS.f1321_m2153_p5442.20230327/	1	1	0	100	done	0.16	-	0

- In contact with PanDA developers to improve error logs

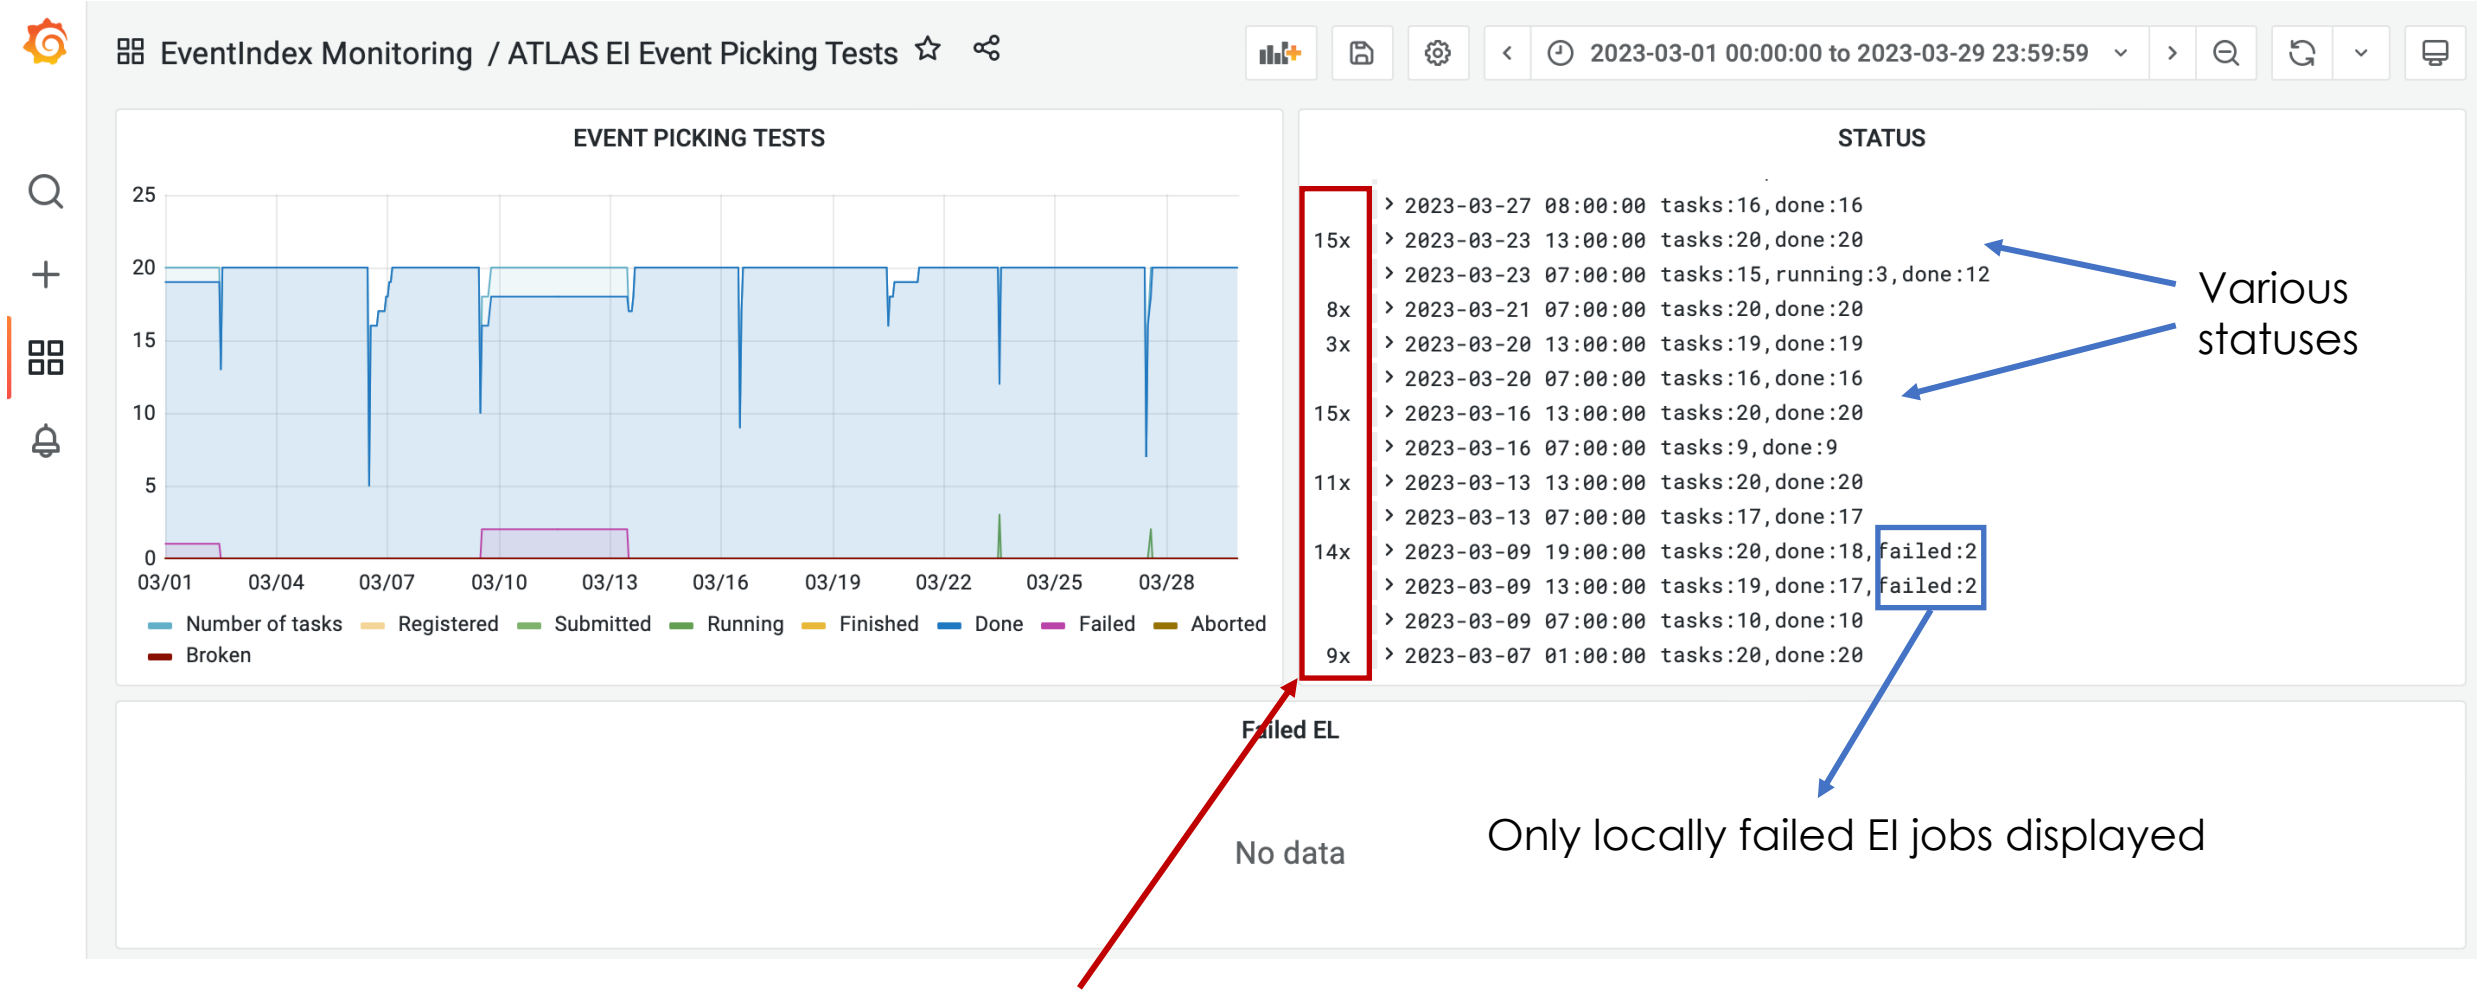
### Failed to list log files

```
Something went wrong while the log file downloading. [guid=690f4553-896c-4f17-90f0-8f23d7543dfb, scope=group.proj-evind, lfn=group.proj-evind.EPRAW.YG.v0.data22_13p6TeV.00440613.physics_BphysDelayed.merge.AOD.f1321_m2153.20230309.log.32713754.000003.log.tgz]
File download failed with command [export RUCIO_ACCOUNT=atlpn; export X509_USER_PROXY=/data/atlpn/x509up_u25606; export ATLAS_LOCAL_ROOT_BASE=/cvmfs/atlas.cern.ch/repo/ATLASLocalRootBase; source ${ATLAS_LOCAL_ROOT_BASE}/user/atlasLocalSetup.sh -3;source $ATLAS_LOCAL_ROOT_BASE/packageSetups/localSetup.sh rucio; rucio download --dir=/opt/prod/pythonpath/core/media/filebrowser/690f4553-896c-4f17-90f0-8f23d7543dfb group.proj-evind:group.proj-evind.EPRAW.YG.v0.data22_13p6TeV.00440613.physics_BphysDelayed.merge.AOD.f1321_m2153.20230309.log.32713754.000003.log.tgz]. Output: [[7mlsetup [0m lsetup <tool1> [ <tool2> ...] (see lsetup -h):
[7m lsetup asetup [0m (or asetup) to setup an Athena release
[7m lsetup astyle [0m ATLAS style macros
[7m lsetup atlantis [0m Atlantis: event display
[7m lsetup eiclient [0m Event Index
[7m lsetup emi [0m EMI: grid middleware user interface
[7m lsetup ganga [0m Ganga: job definition and management client
[7m lsetup lcgenv [0m lcgenv: setup tools from cvmfs SFT repository
[7m lsetup panda [0m Panda: Production AND Distributed Analysis
[7m lsetup pyami [0m pyAMI: ATLAS Metadata Interface python client
[7m lsetup root [0m ROOT data processing framework
[7m lsetup rucio [0m distributed data management system client
[7m lsetup scikit [0m python data analysis ecosystem
[7m lsetup viewc [0m Set up a full LCG release
```

# Monitoring the functional tests



- Overall view of test jobs' status over time in Grafana



How many times this status was read by Grafana



## Data for the tests:

- Run2 LHC data (2015-2018) ~50000 events
  - runNumber-eventNumber lists for several datasets in each year
- Two combinations of event lists:
  - Lists with events from different datasets (mixed): data15, data17
  - Lists with events from datasets with 1 million events (1M): data15, data18
- Lists are fixed OR generated randomly before running the query

## Query:

- Event Lookup
  - Fast search of events using CLI
  - Return events GUIDs, full dataset names, data types (RAW, AOD, DAOD)



With randomly generated event lists (running every **1 hour** in acrontab)

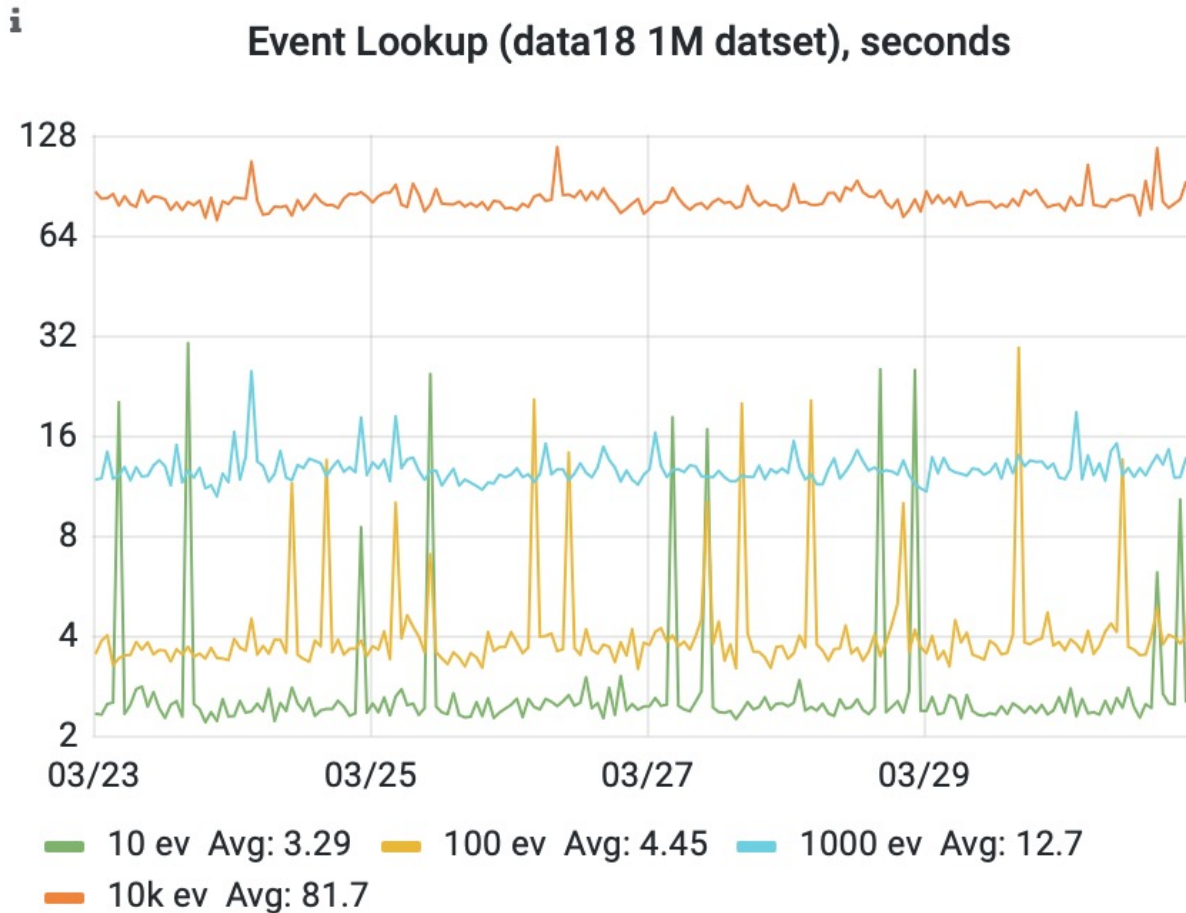
dataset	nevents	Avg time, s
data15 (1M)	1000	18.4
data15 (mix)	1000	23.6
data17 (mix)	1000	36.7
data18 (1M)	10	3.29
data18 (1M)	100	4.45
data18 (1M)	1 000	12.7
data18 (1M)	10 000	81.7

With fixed event lists (running every **2 hours** in acrontab)

dataset	nevents	Avg time, s
data17 (1M)	1000	6.97
data17 (mix)	1000	34.0

\*Green ones are displayed  
8-12 May 2023 - CHEP2023

2023/03/23 – 2023/03/30



Response times of the event lookup queries for data18

\*peaks in the response times are due to interference with other activities on the HBase cluster

With randomly generated event lists  
(running every **1 hour**)


dataset	nevents	Avg time, s
data15 (1M)	1000	18.4
data15 (mix)	1000	23.6
data17 (mix)	1000	36.7
data18 (1M)	10	3.29
data18 (1M)	100	4.45
data18 (1M)	1 000	12.7
data18 (1M)	10 000	81.7

With fixed event lists (running every  
**2 hours**)

dataset	nevents	Avg time, s
data17 (1M)	1000	6.97
data17 (mix)	1000	34.0

\*Green ones are displayed  
8-12 May 2023 - CHEP2023

2023/03/23 – 2023/03/30



Response times of the event lookup queries

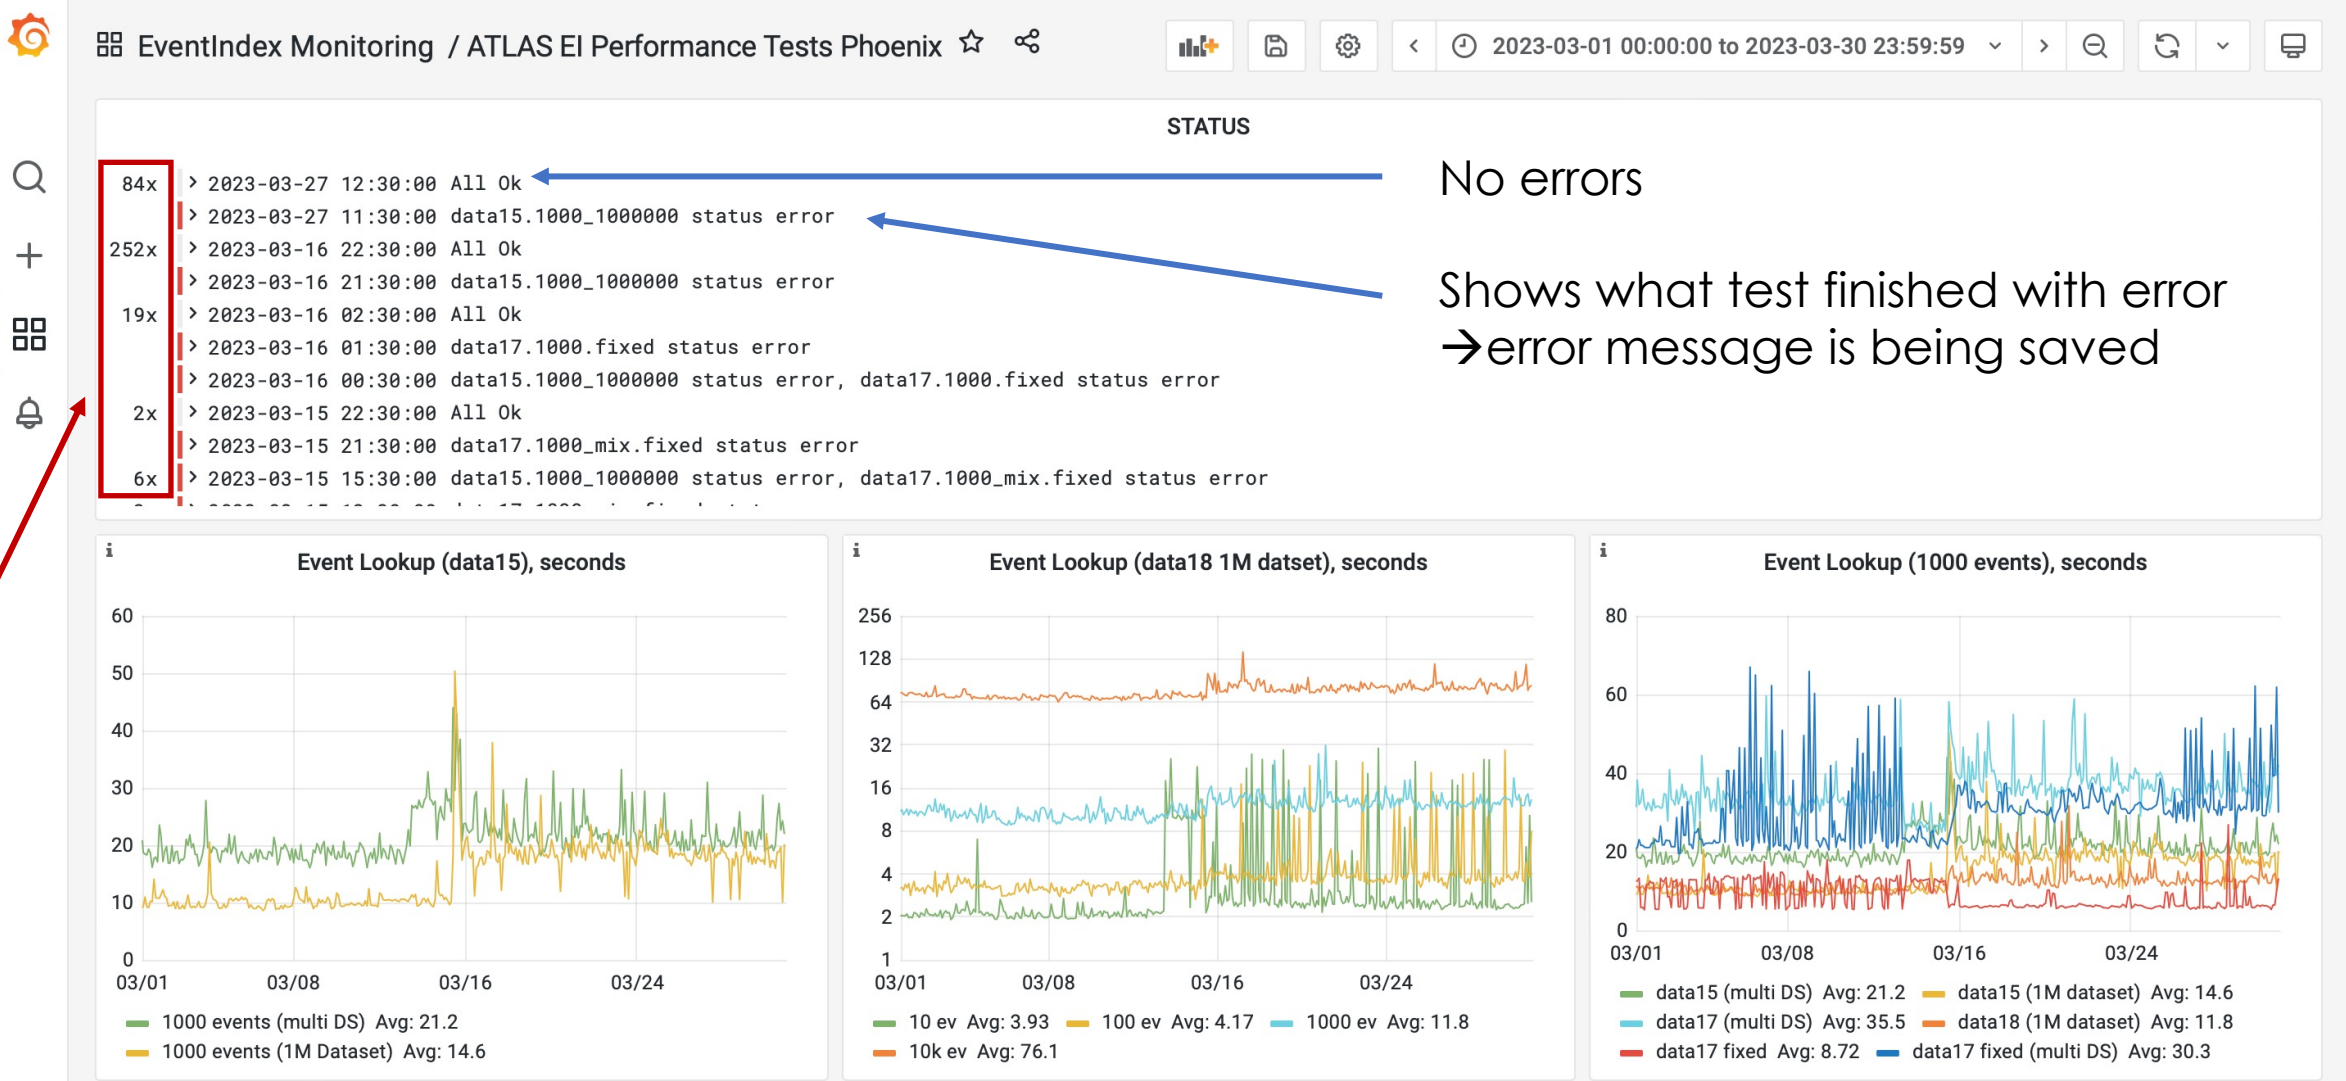
\*peaks in the response times are due to interference with other activities on the HBase cluster

# Monitoring the performance tests



- Test results have been written to JSON files → sent to Grafana

How many times this status was read by Grafana





- The Monitoring System provides a good representation of the status of all EventIndex sub-systems and major parameters
- Functional and performance tests control the well-being and the responsiveness of the EventIndex tool
  - Run regularly
  - Run locally and remotely (on GRID)
  - Quick indicators of technical problems, not necessarily caused by EventIndex malfunctions
  - Enable performance studies and system optimization

Thank you for your attention!