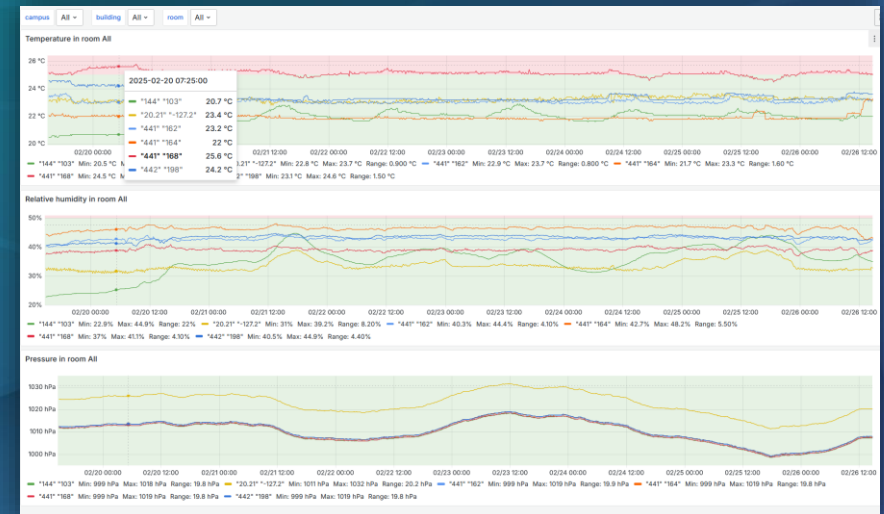
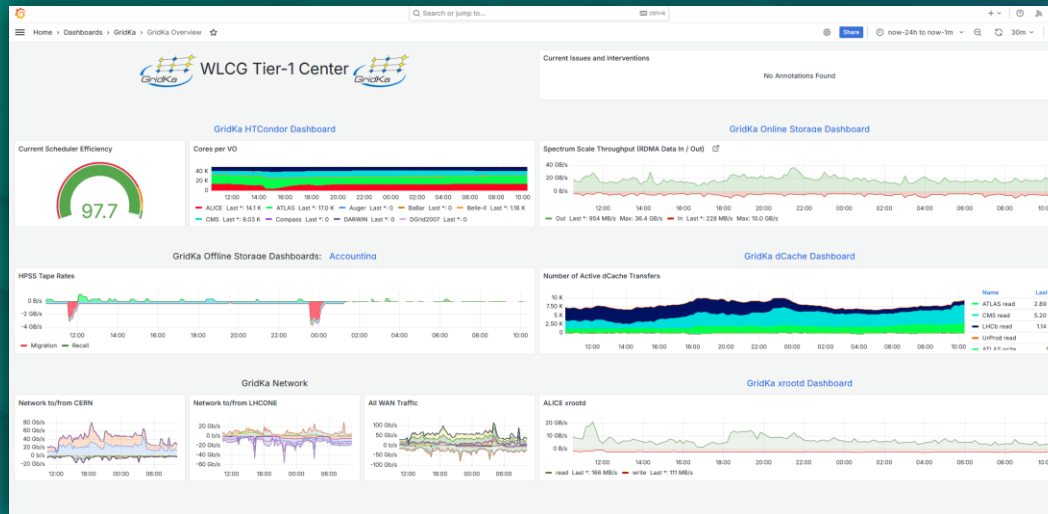



Infrastructure Monitoring for GridKa and beyond

Evelina Buttitta

Hepix Spring 2025 Workshop, 31 March - 4 April 2025, Lugano



Infrastructure resources



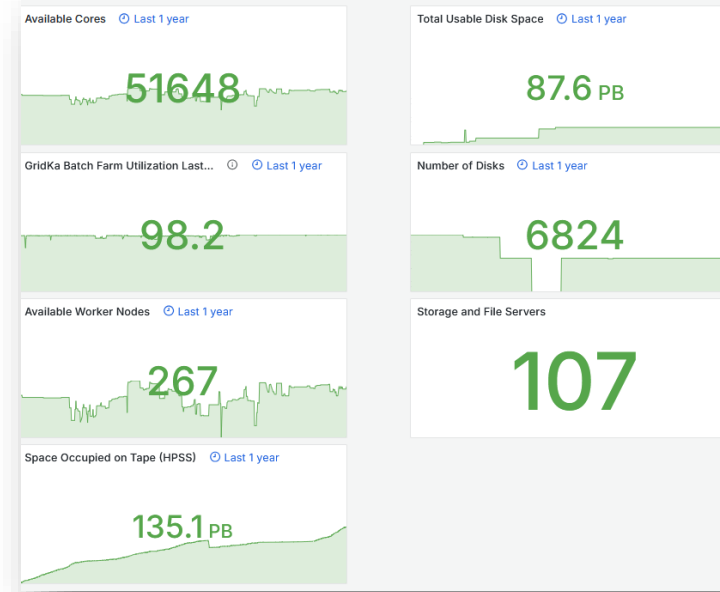
Online Storage





Offline Storage
(Tape)



GridKa Network



GridKa Batch Farm

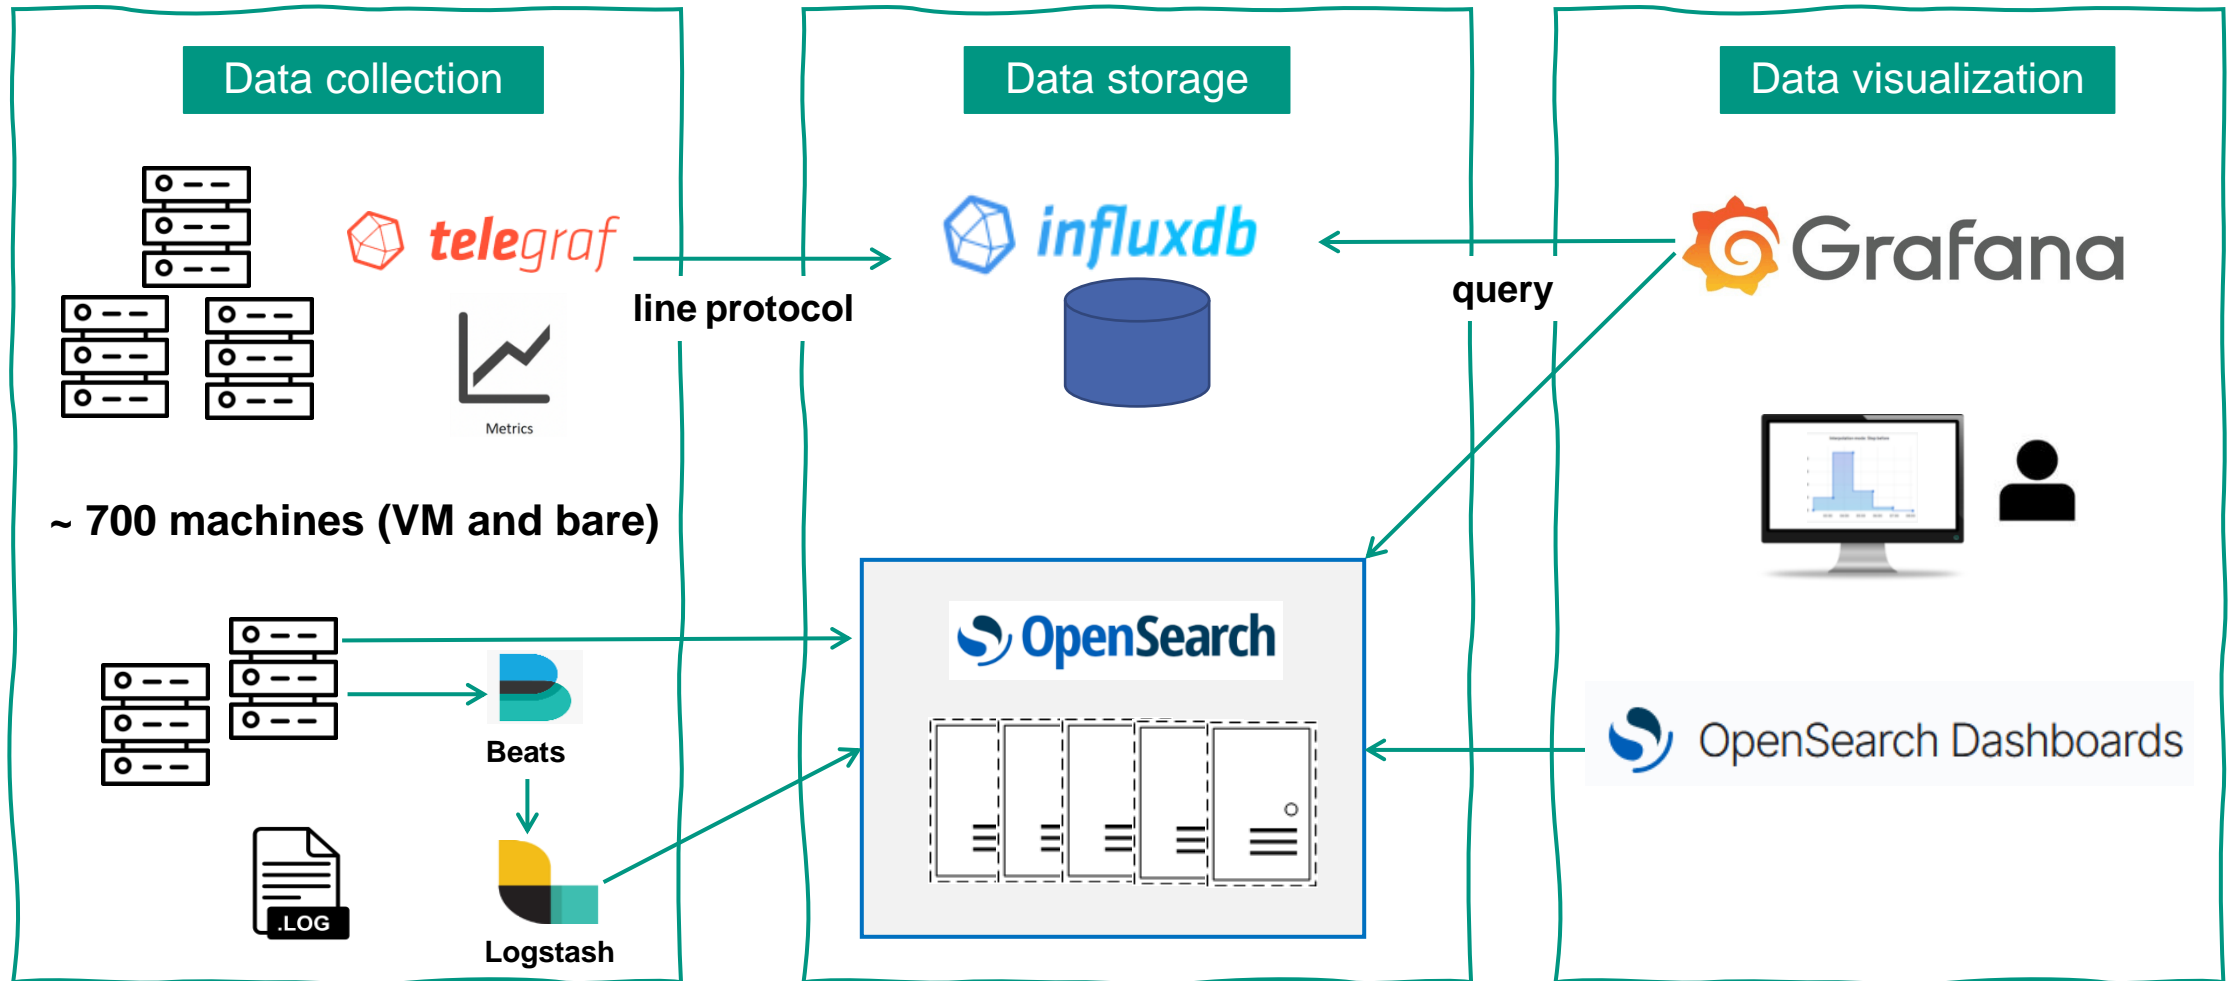


BWDATA
ARCHIV

Motivations

- Constantly **monitor health and availability** of the services
 - Give users a fast overview of services/applications status
 - **Detect and report issues** fast and in detail
 - Centralize, search, analyze a large volume of data (e.g. logs)
- The infrastructure monitoring provides insights and visibility into the health and status of all data center **by tracking specific metrics and logs in real time**

Monitoring architecture



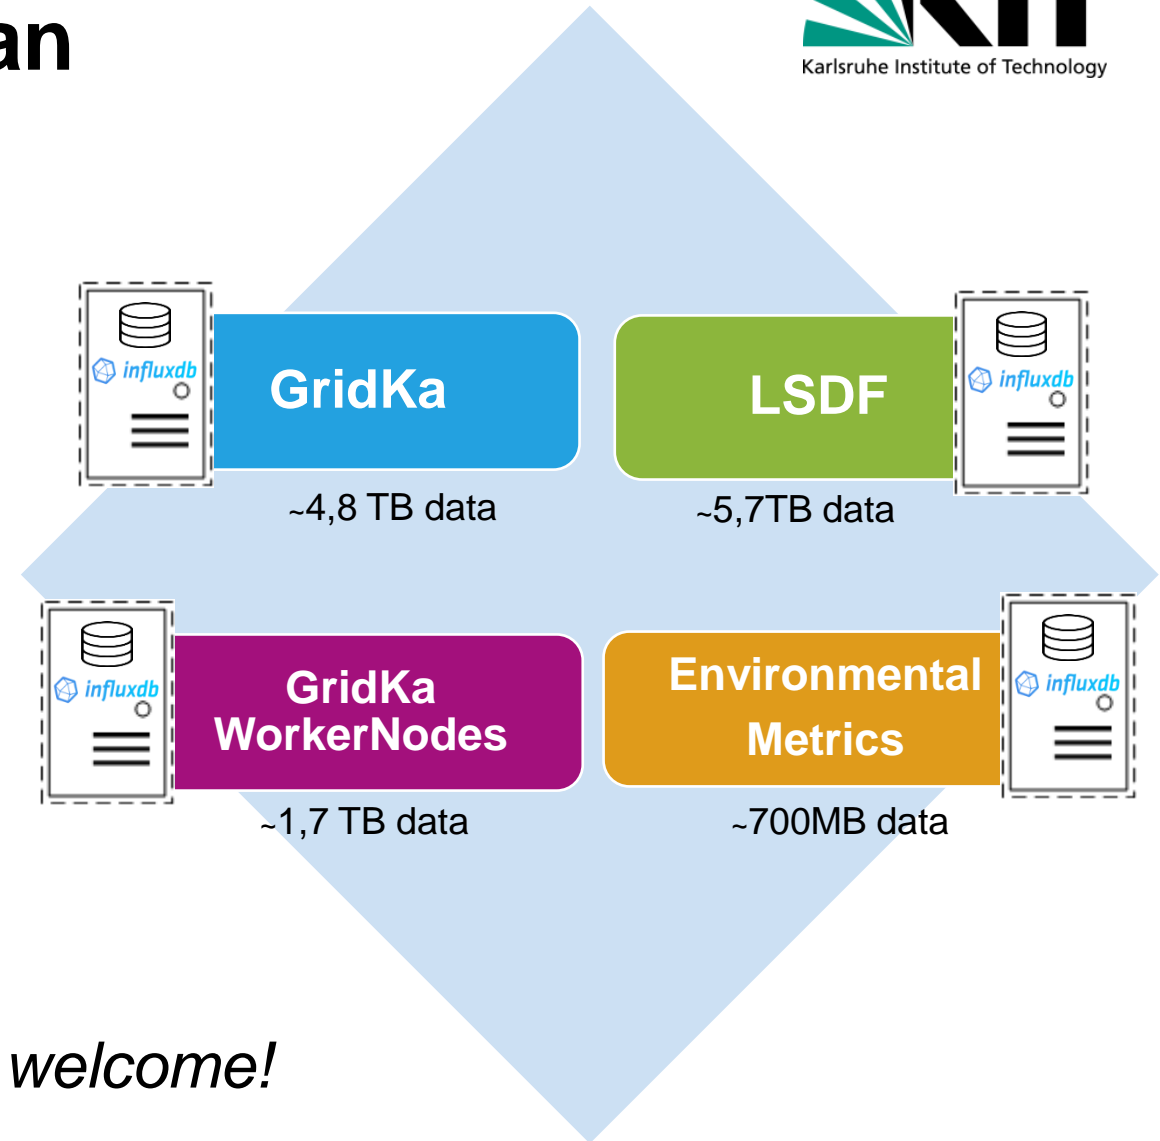
Collected metrics and logs

- Host Metrics (Disk, CPU, Memory, Network...)
- GridKa Experiments (Data Transfer, Disk usage per experiment)
- HTCondor (Running jobs info)
- Storage:
 - dCache, XRootD (IO Metrics and Accounting)
 - Storage Scale (GPFS) (IO Metrics)
 - Disk Systems (controller, disk, raid status and IO metrics)
 - Tape Monitoring
 - LSDF Online Storage
- Environmental Metrics (PDUs, humidity, temperature)
- Logs of GridKa and LSDF storage, HTCondor, Tape HPSS, GridKa networks, dCache, XRootD

InfluxDB current and future plan

- Running 4 InfluxDB instances OSS v2.7

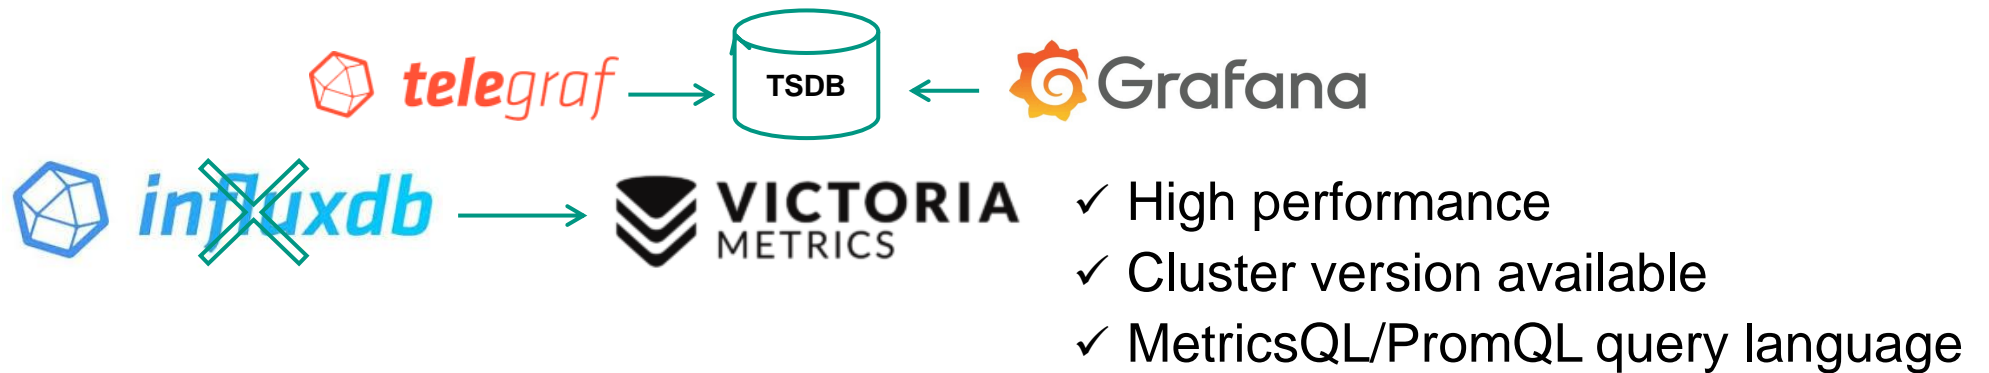
- Test the new version *InfluxDB 3 Core* with new capabilities and enhanced performances
- Migrate InfluxDB to other timeseries databases like Grafana solutions (e.g. Mimir) or VictoriaMetrics



Suggestions and sharing of experiences are welcome!

VictoriaMetrics as alternative TSDB

- Testing InfluxDB v2 migration to VictoriaMetrics



Challenges

- New data model for the metrics
- Import of large historical dataset from InfluxDB v2
- Translation of many InfluxQL/Flux queries in Grafana Dashboards to MetricsQL

Grafana

<https://grafana-sdm.scc.kit.edu>

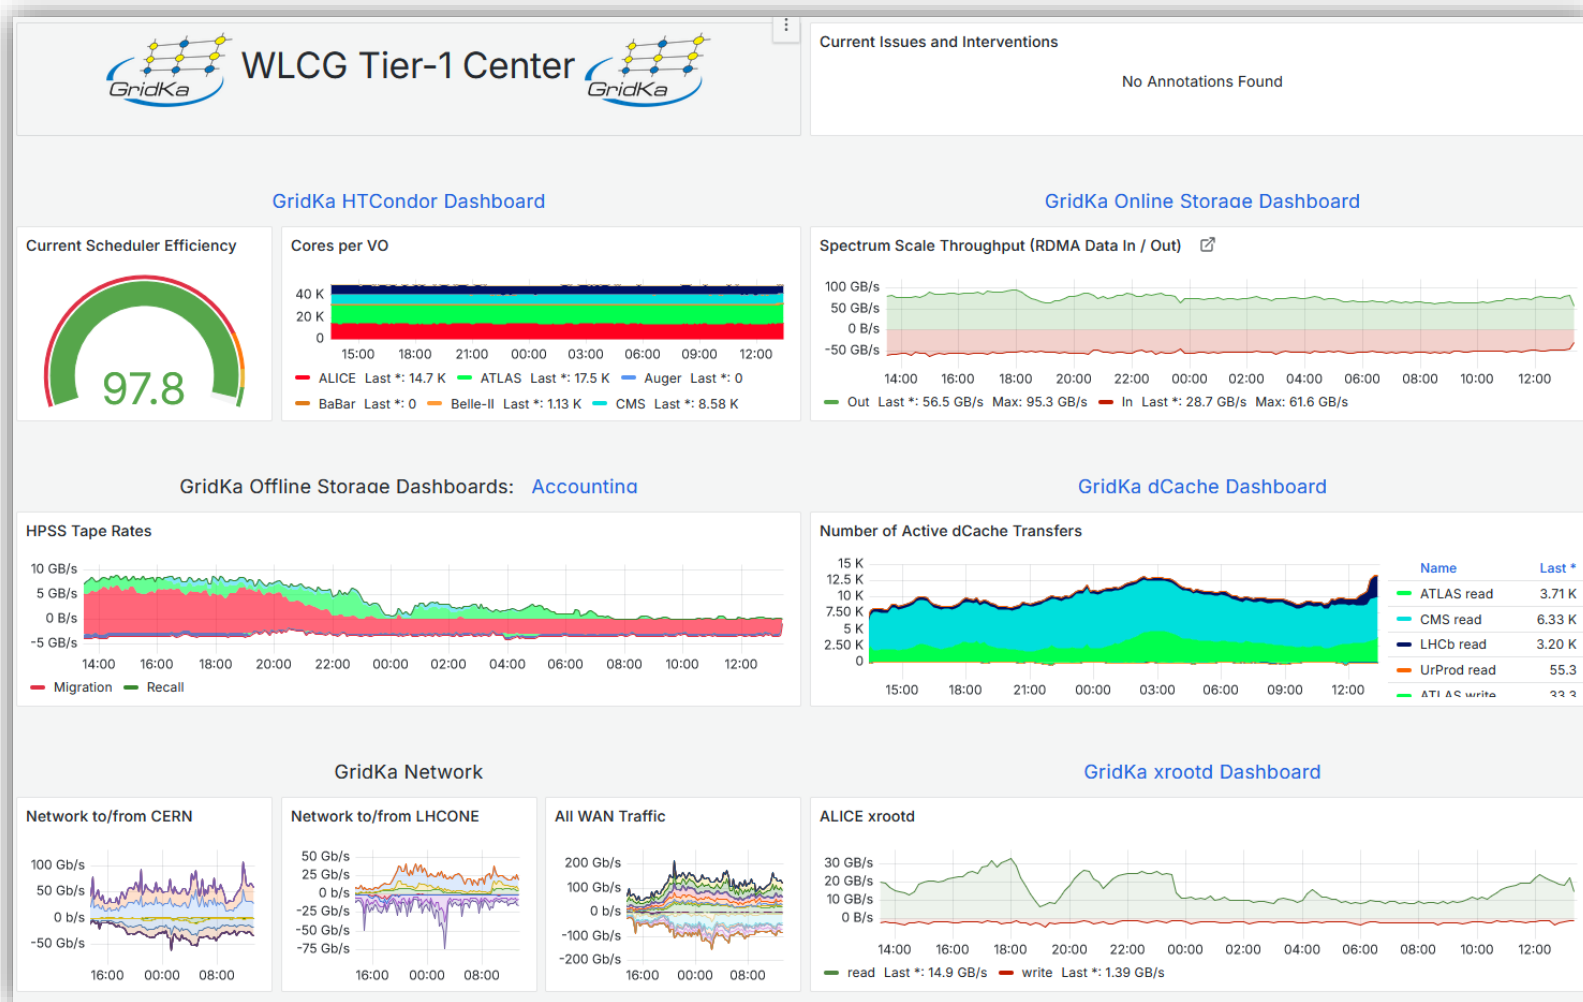


Auth/Authz via
KIT OpenID

~ 60
dashboards

~ 40
active users

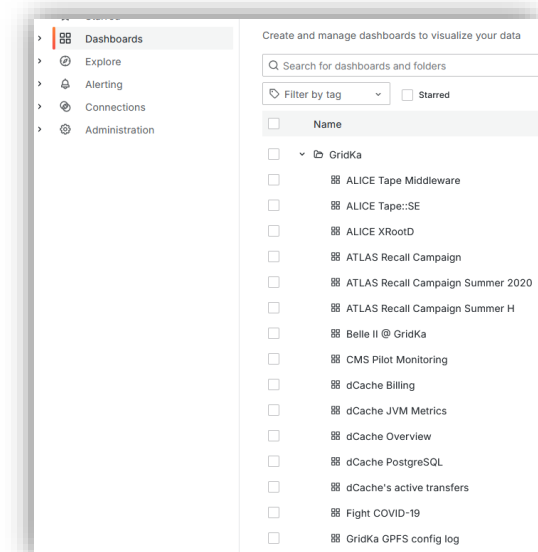
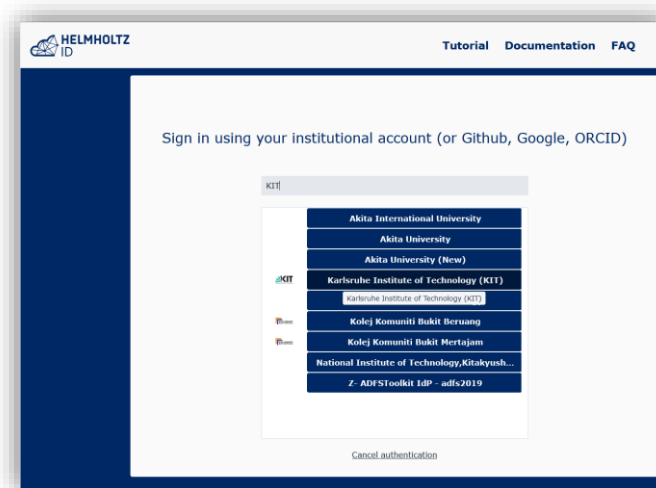

~ 20
provisioned
datasources



Grafana open instance

- The problem: share specific dashboards with external (no KIT) users
- What we have tried:
 - Enable anonymous access: there are security implications we have experienced
 - Use the *Public dashboard* feature: some inconvenient limitations
- Our solution: run a separate Grafana instance
 - Require auth, but allow auth from all of edugain (via Helmholtz-ID AAI)

<https://grafana-sdm-open.scc.kit.edu>



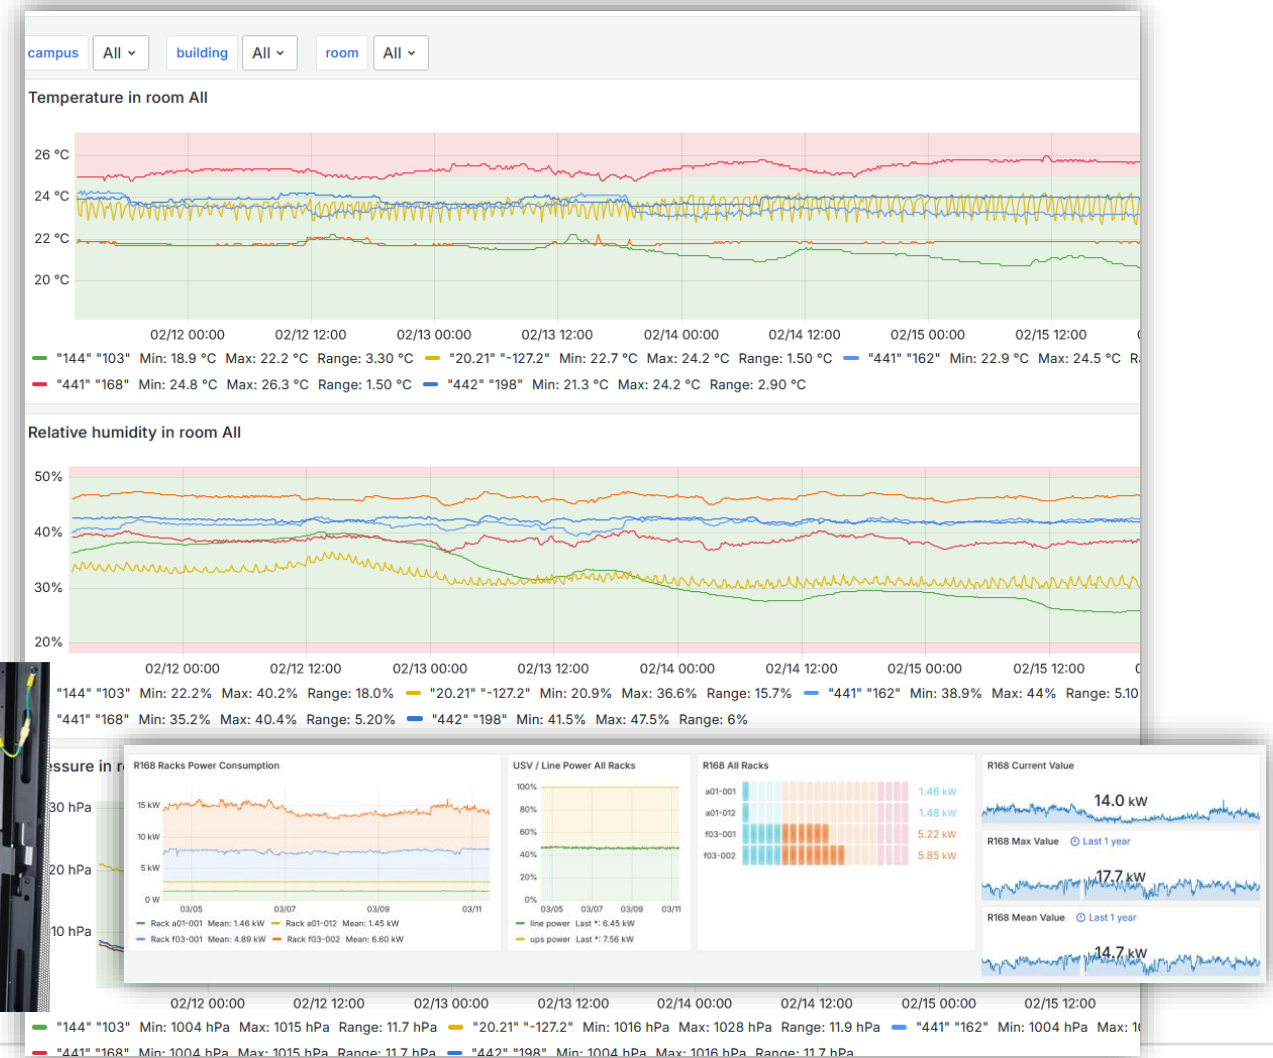
Grafana Datacenter Metrics

<https://datacenter-metrics.scc.kit.edu>

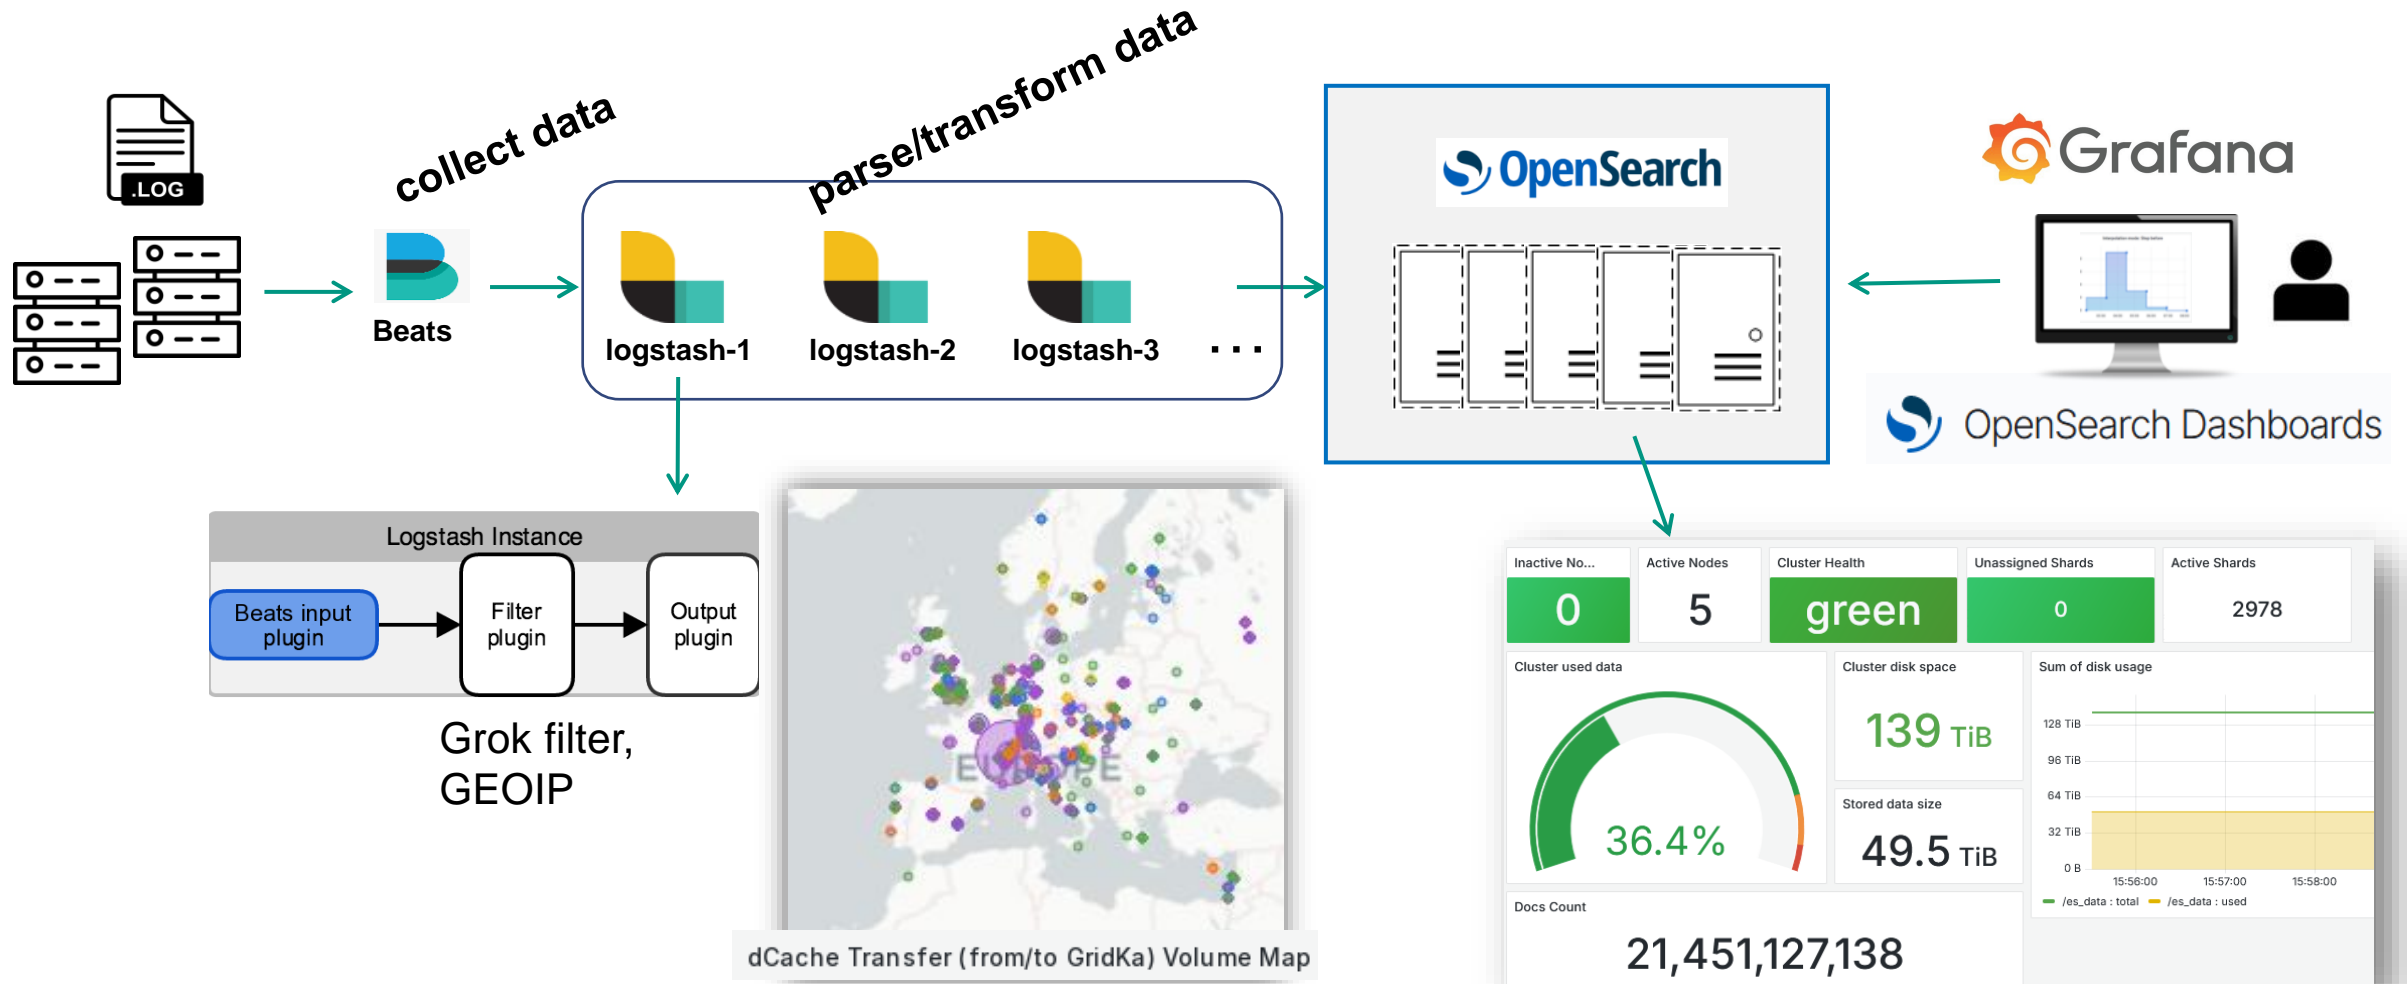


Visualization of environmental metrics:

- Sensors metrics (temperature, humidity, pressure) from all server rooms
- Water cooling temperature
- Power consumption (UPS, Power distribution units (PDU))



Monitoring with OpenSearch, Beats and Logstash



Summary and outlook

- The infrastructure monitoring plays a fundamental role in the operations of the whole infrastructure
- It provides a real time (and historical) overview of all applications and ensures efficiency and reliability of services
- Future of InfluxDB v2.x:
 - Update to InfluxDB v3
 - Replace InfluxDB with other timeseries databases (e.g. VictoriaMetrics)?
- On-going research about a standard and consistent observability framework for collecting and storing telemetry data (logs, metrics, traces), e.g. OpenTelemetry