



Monitoring Update

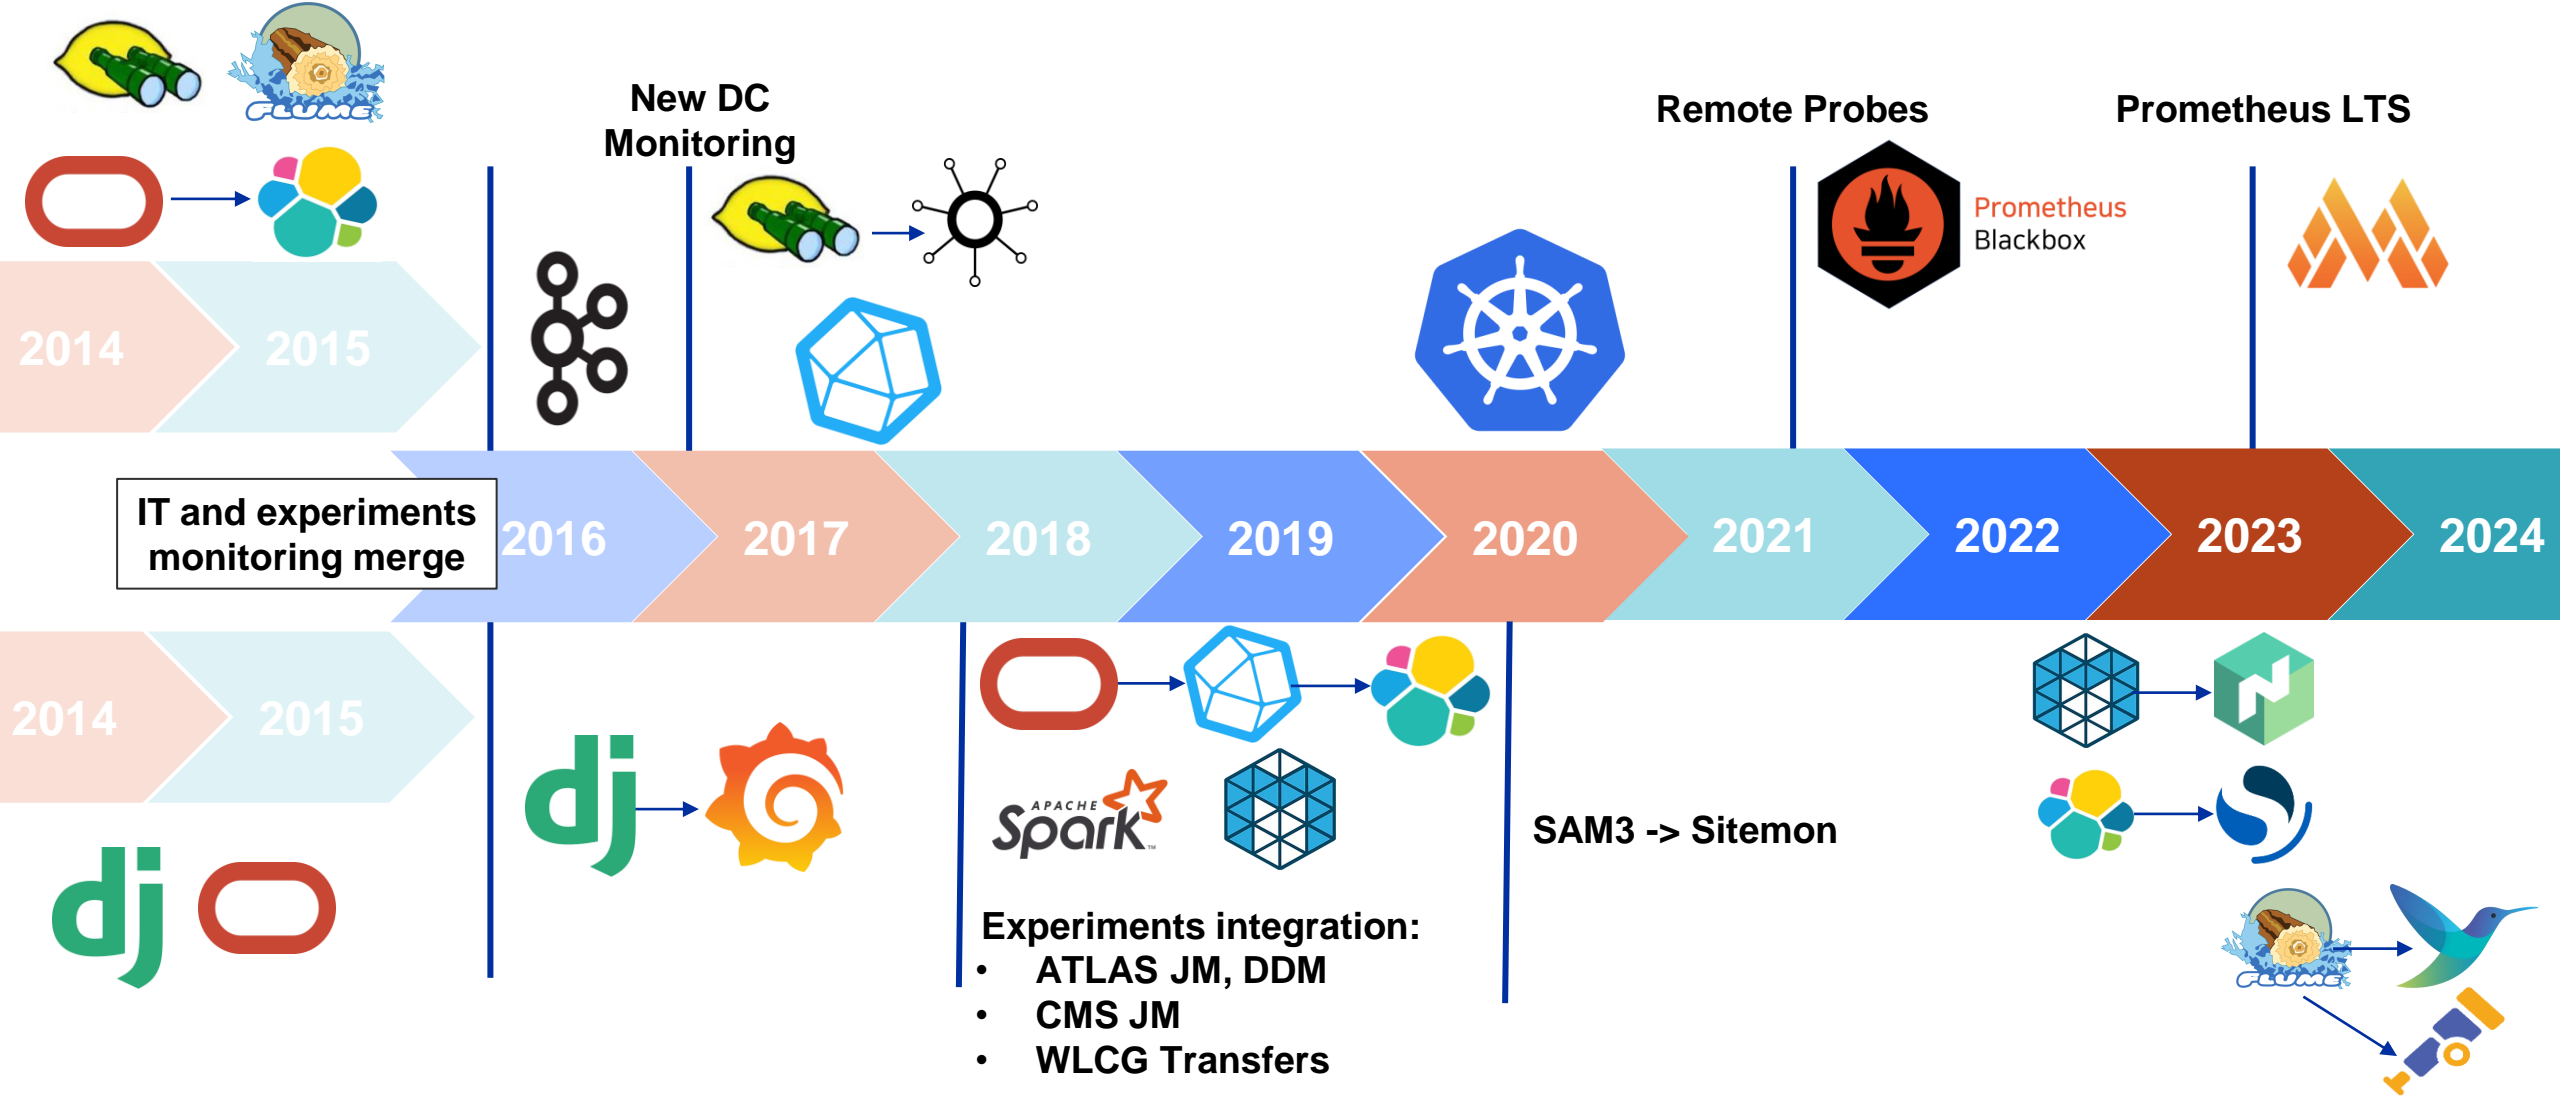
Borja Garrido Bear

13.09.2023

Index

- **Monitoring history overview**
- **Data Centre Monitoring**
 - Mimir pilot
 - OpenTelemetry ingestion layer
 - New "MONIT" bundle
- **WLCG Monitoring**
 - Site Network monitoring
 - New XRootD monitoring
- **Visualization**
 - Grafana v10

MONIT History Overview



Data Centre Monitoring

New services and planned evolution

Mimir Pilot Service (I)

Introduced in IT [ASDF meeting 24.08.2023](#)

Central place for Prometheus metrics aggregation and "Long term" archival

- **One single query point for multiple Prometheus deployments**
- **Support for High Availability**
- **Horizontally scalable**
- **Multi-tenancy out of the box**
- **Up to 40 days of data for the time being**
 - **Longer term options can be evaluated with users**



Mimir Pilot Service (II)

Single monitoring backend for different environments

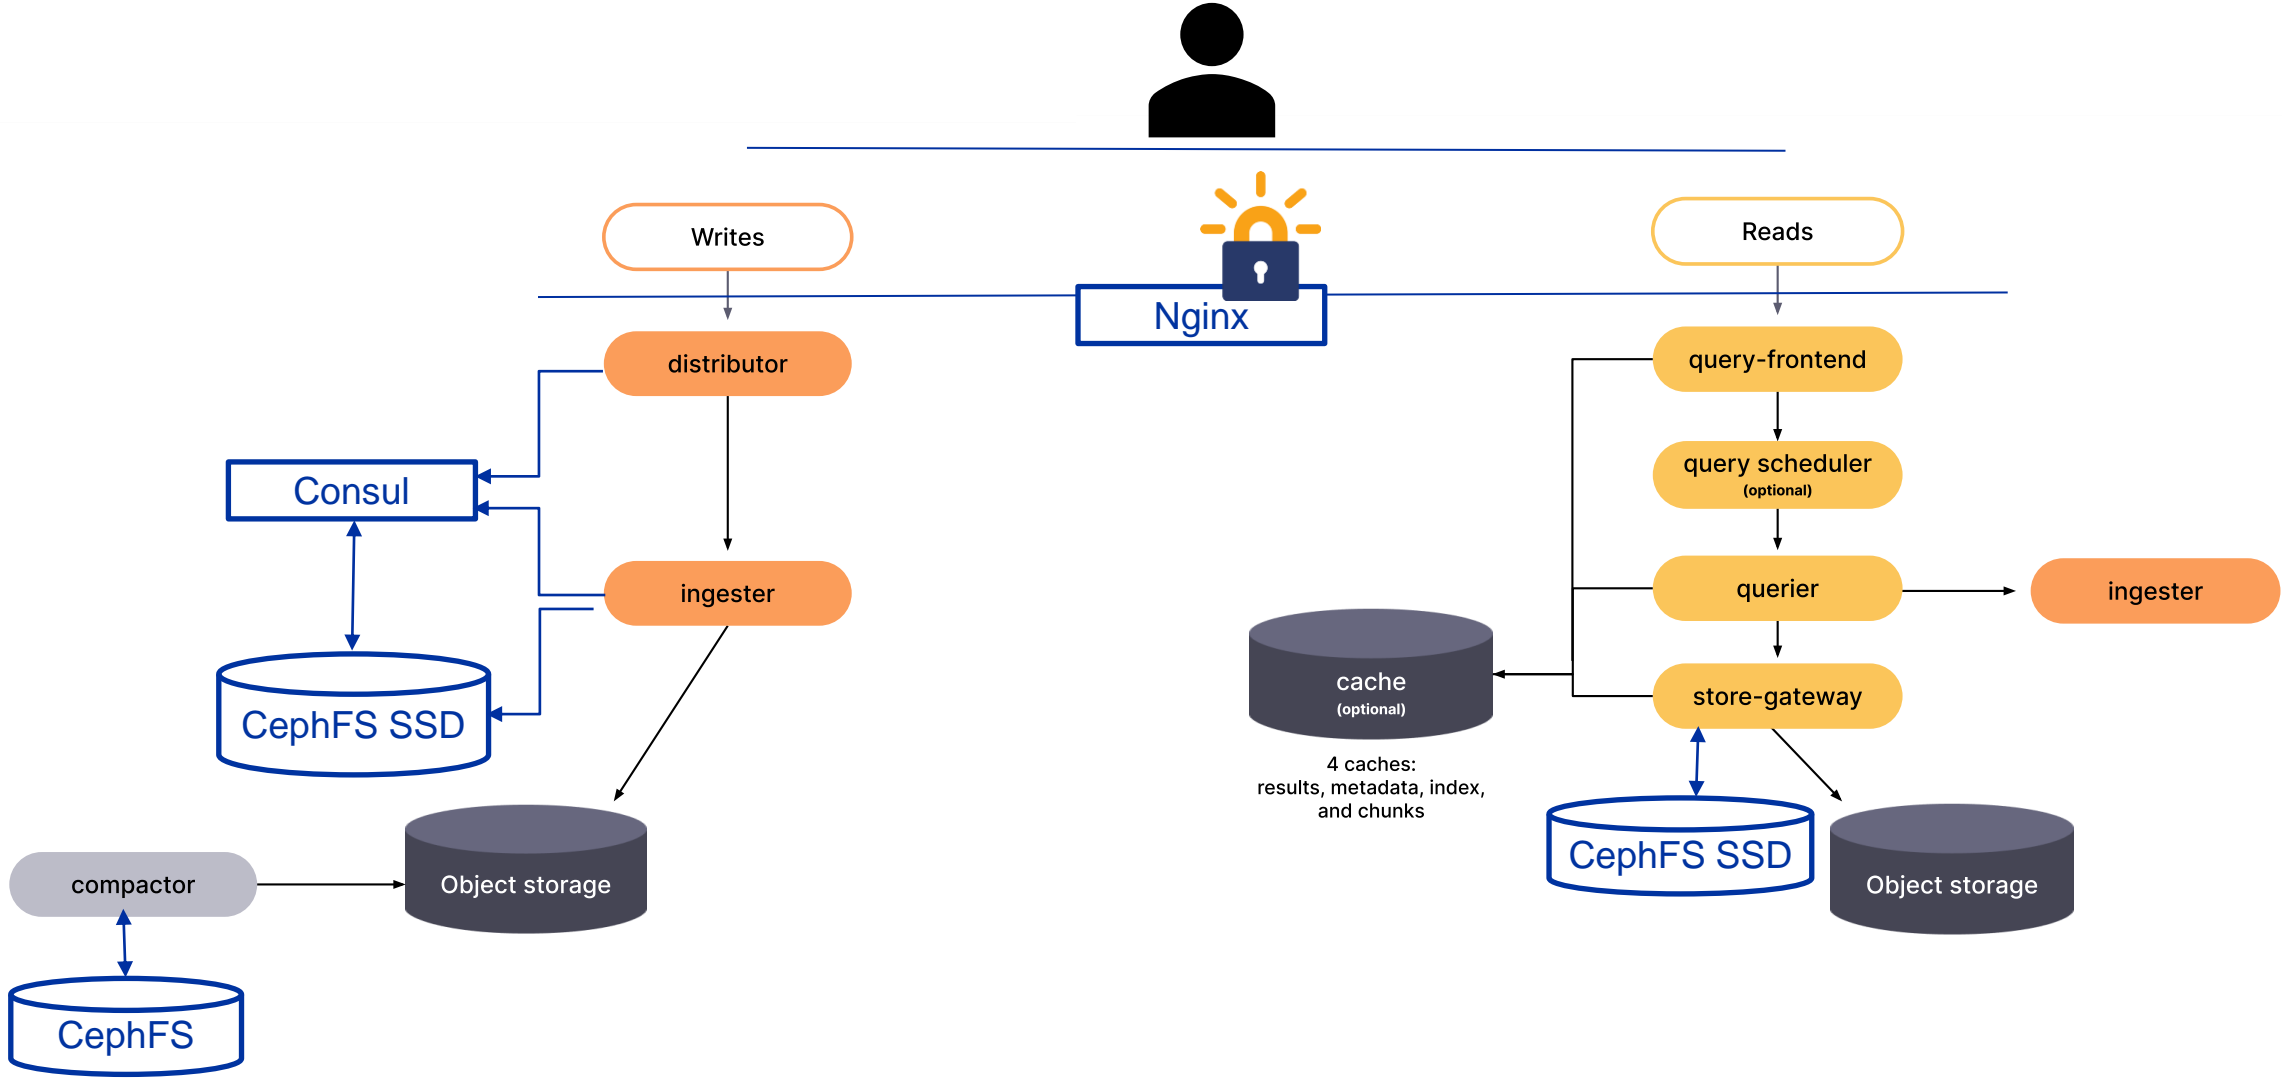
- Puppet managed hosts
- Windows infrastructure
- Kubernetes deployments

Scale deployment to handle 100M active series

- Should cover for DC monitoring needs, will be re-evaluated as required

Already available, usage [documentation](#)

Mimir Pilot Service: Architecture




Collection layer

OpenTelemetry Collection layer

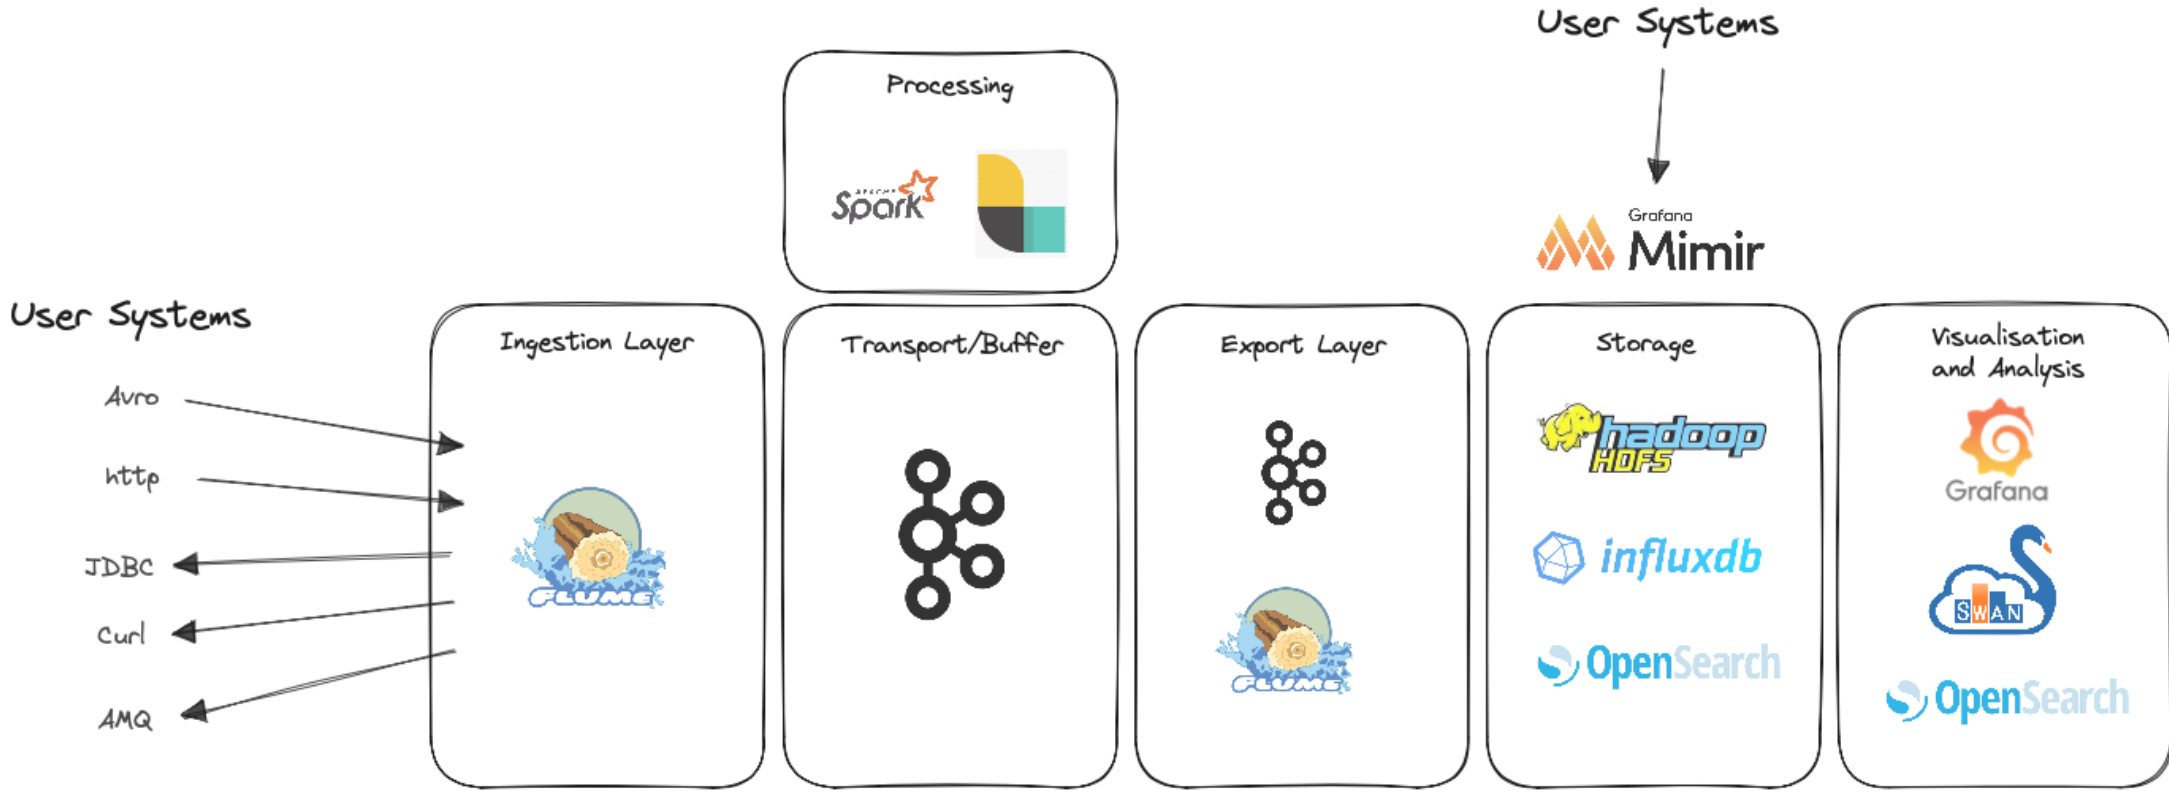
- **Single ingestion layer for metrics, logs and traces using OTLP protocol**

Monitoring tools installed in all the Data Centre nodes

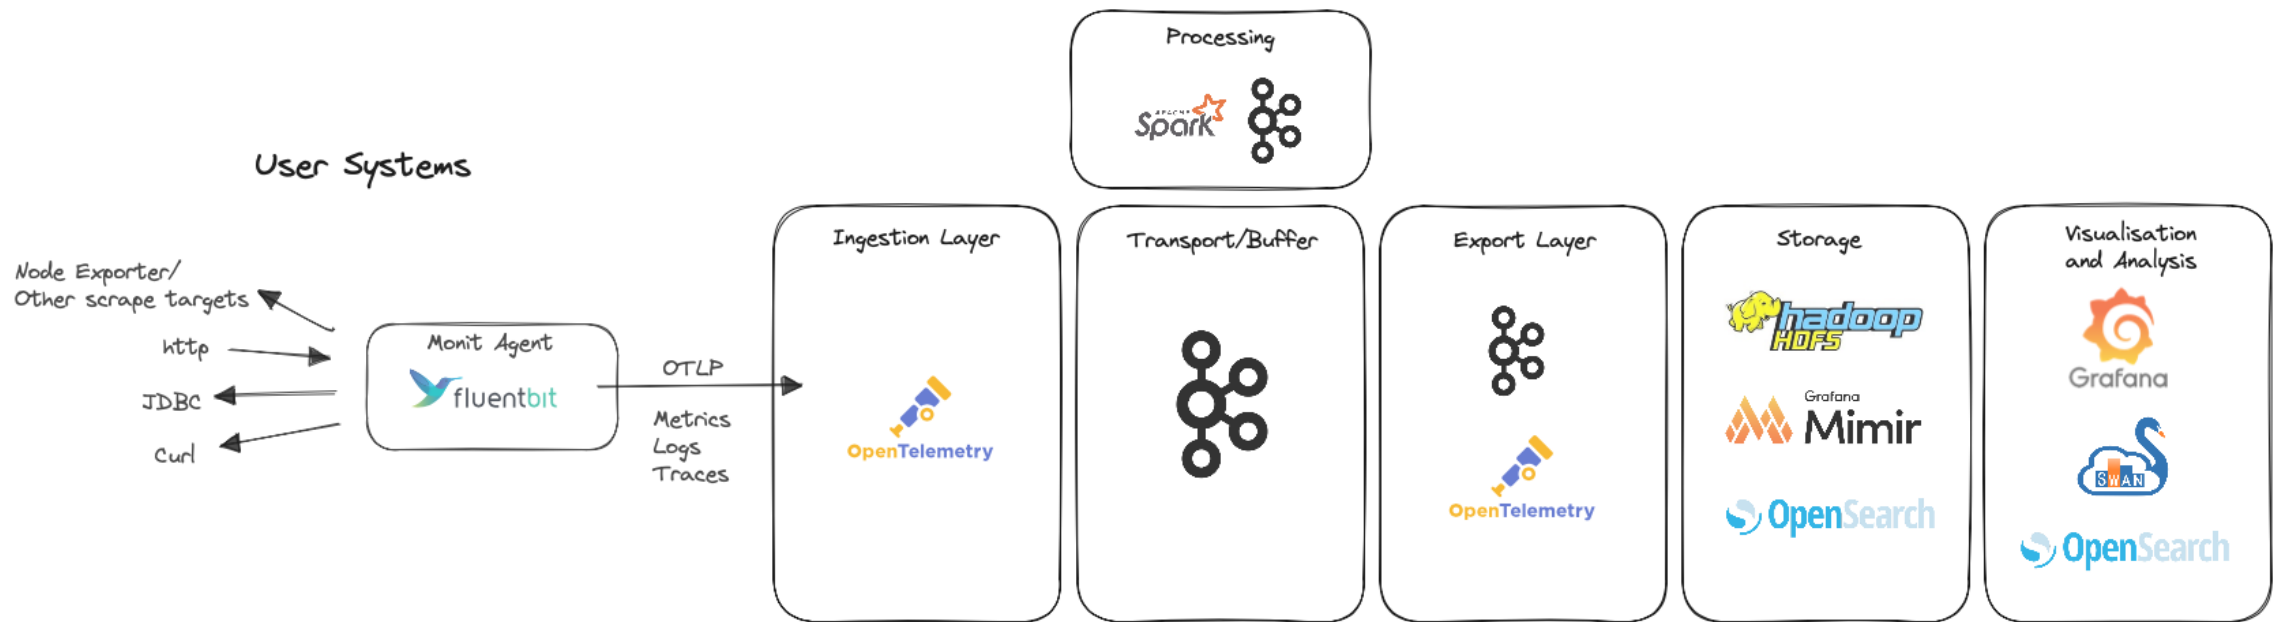
- **Current bundle:**
 - Collectd as collector agent, Apache Flume as forwarding agent
- **New bundle:**
 - Collectd + Node exporter as collector agents
 - Collectd is kept to provide compatibility for services that invested in their own plugins
 - Fluentbit as forwarding agent
 - Lighter than Apache Flume, OTLP support



Current Monitoring Infrastructure (simplified)



Intended Monitoring Infrastructure (simplified)



WLCG Monitoring

New services and planned evolution

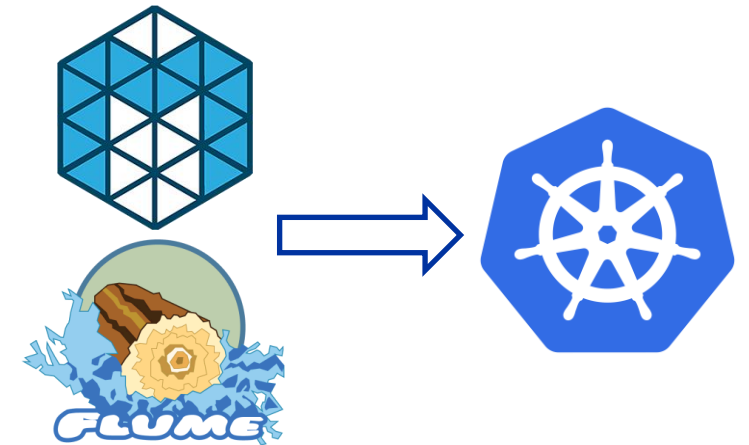
MONIT infrastructure related changes

Processing layer

- **Marathon/Chronos (Mesos) infrastructure decommissioned**
 - Job scheduling moved to Nomad
 - Spark Jobs switched to execute in cluster mode
 - Hadoop Analytix infrastructure for the compute
 - Spark Jobs upgraded to Spark 3

Collection layer

- **Apache Flume scheduler to Kubernetes**
 - Affects several WLCG flows
 - CRIC data, Downtimes, Sitemon profiles, Site network
 - Already moved Site network, rest will follow in the short term

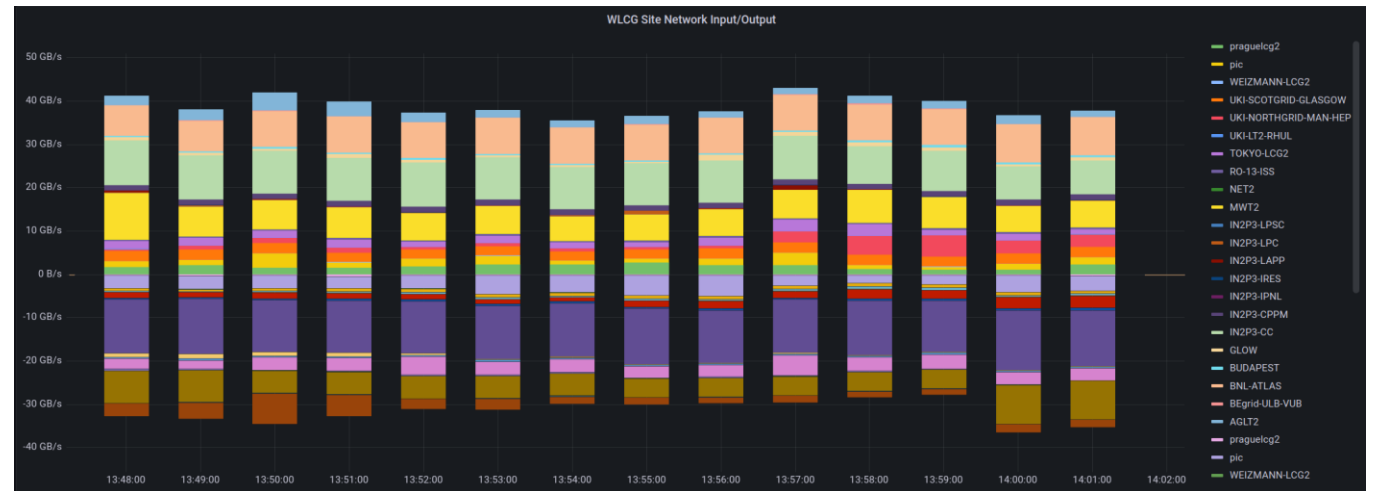


WLCG Site Network Monitoring

Work done as part of the WLCG Monitoring TaskForce

Based on pull model of JSON files provided by sites

- Cronjob running in Kubernetes
 - Updates available URLs from CRIC
 - Queries sites for information once per minute
- Data stored for one year
- Dashboard available in [Grafana](#)



New XRootD Monitoring (XRootD servers)

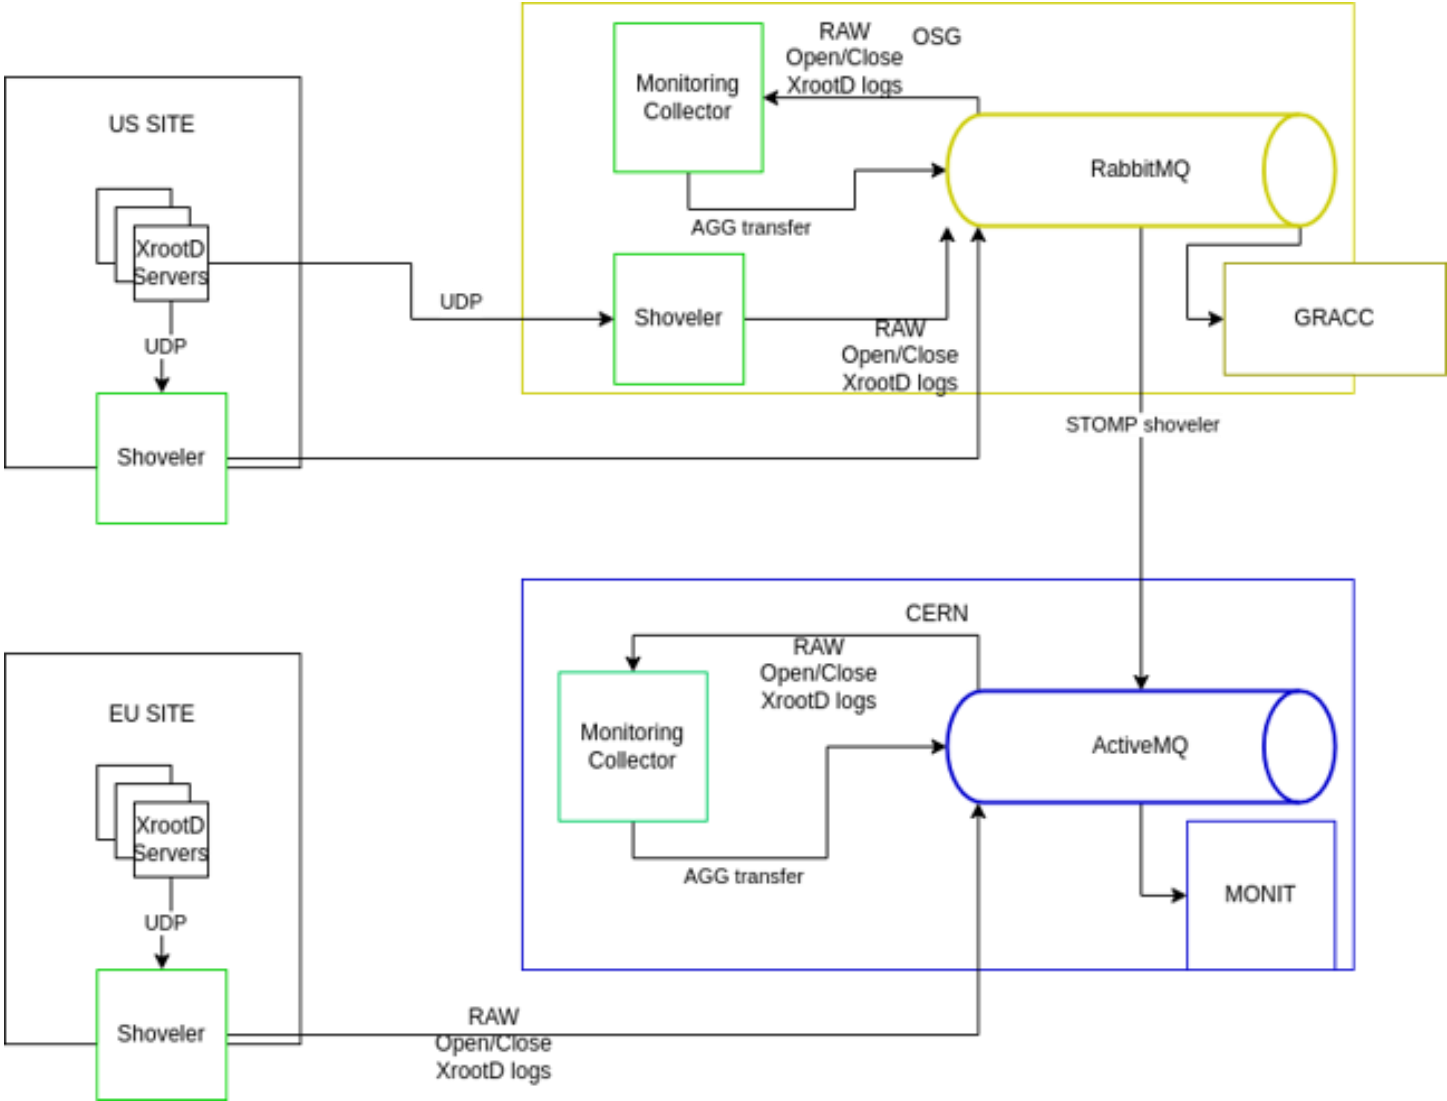
Work done as part of the WLCG Monitoring TaskForce

Last update in [CHEP 2023](#)

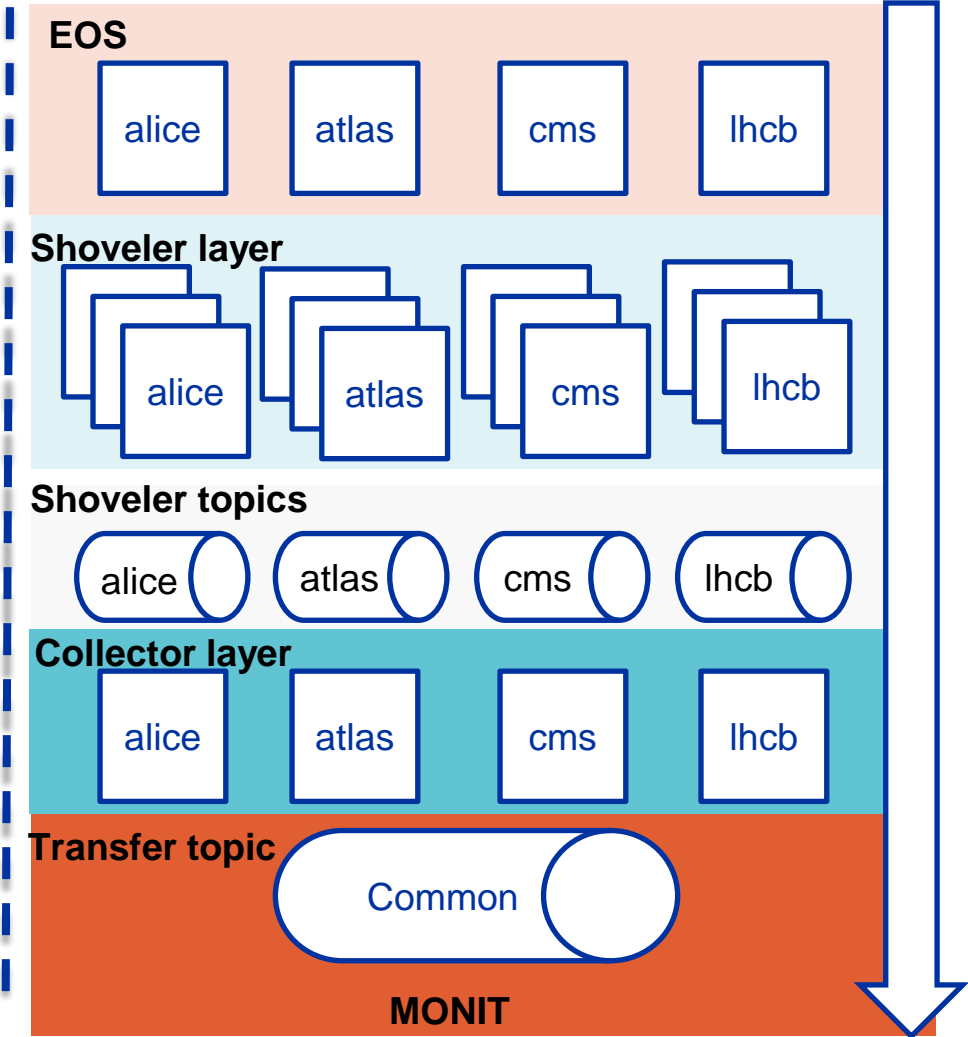
New temporary deployment at CERN

- **Focusing on being able to monitor CERN EOS servers**
 - Isolated flow per VO, dedicated shovelers, topics and collectors
 - Allows for us to hardcode a VO (needed for data enrichment)
- **Already started configuring some servers to validate data**
 - This includes topology enrichment process and aggregation

New XRootD Monitoring (XRootD servers) (II)



CERN Temporary deployment



New XRootD Monitoring (dCache + XRootD)

Work done as part of the WLCG Monitoring TaskForce

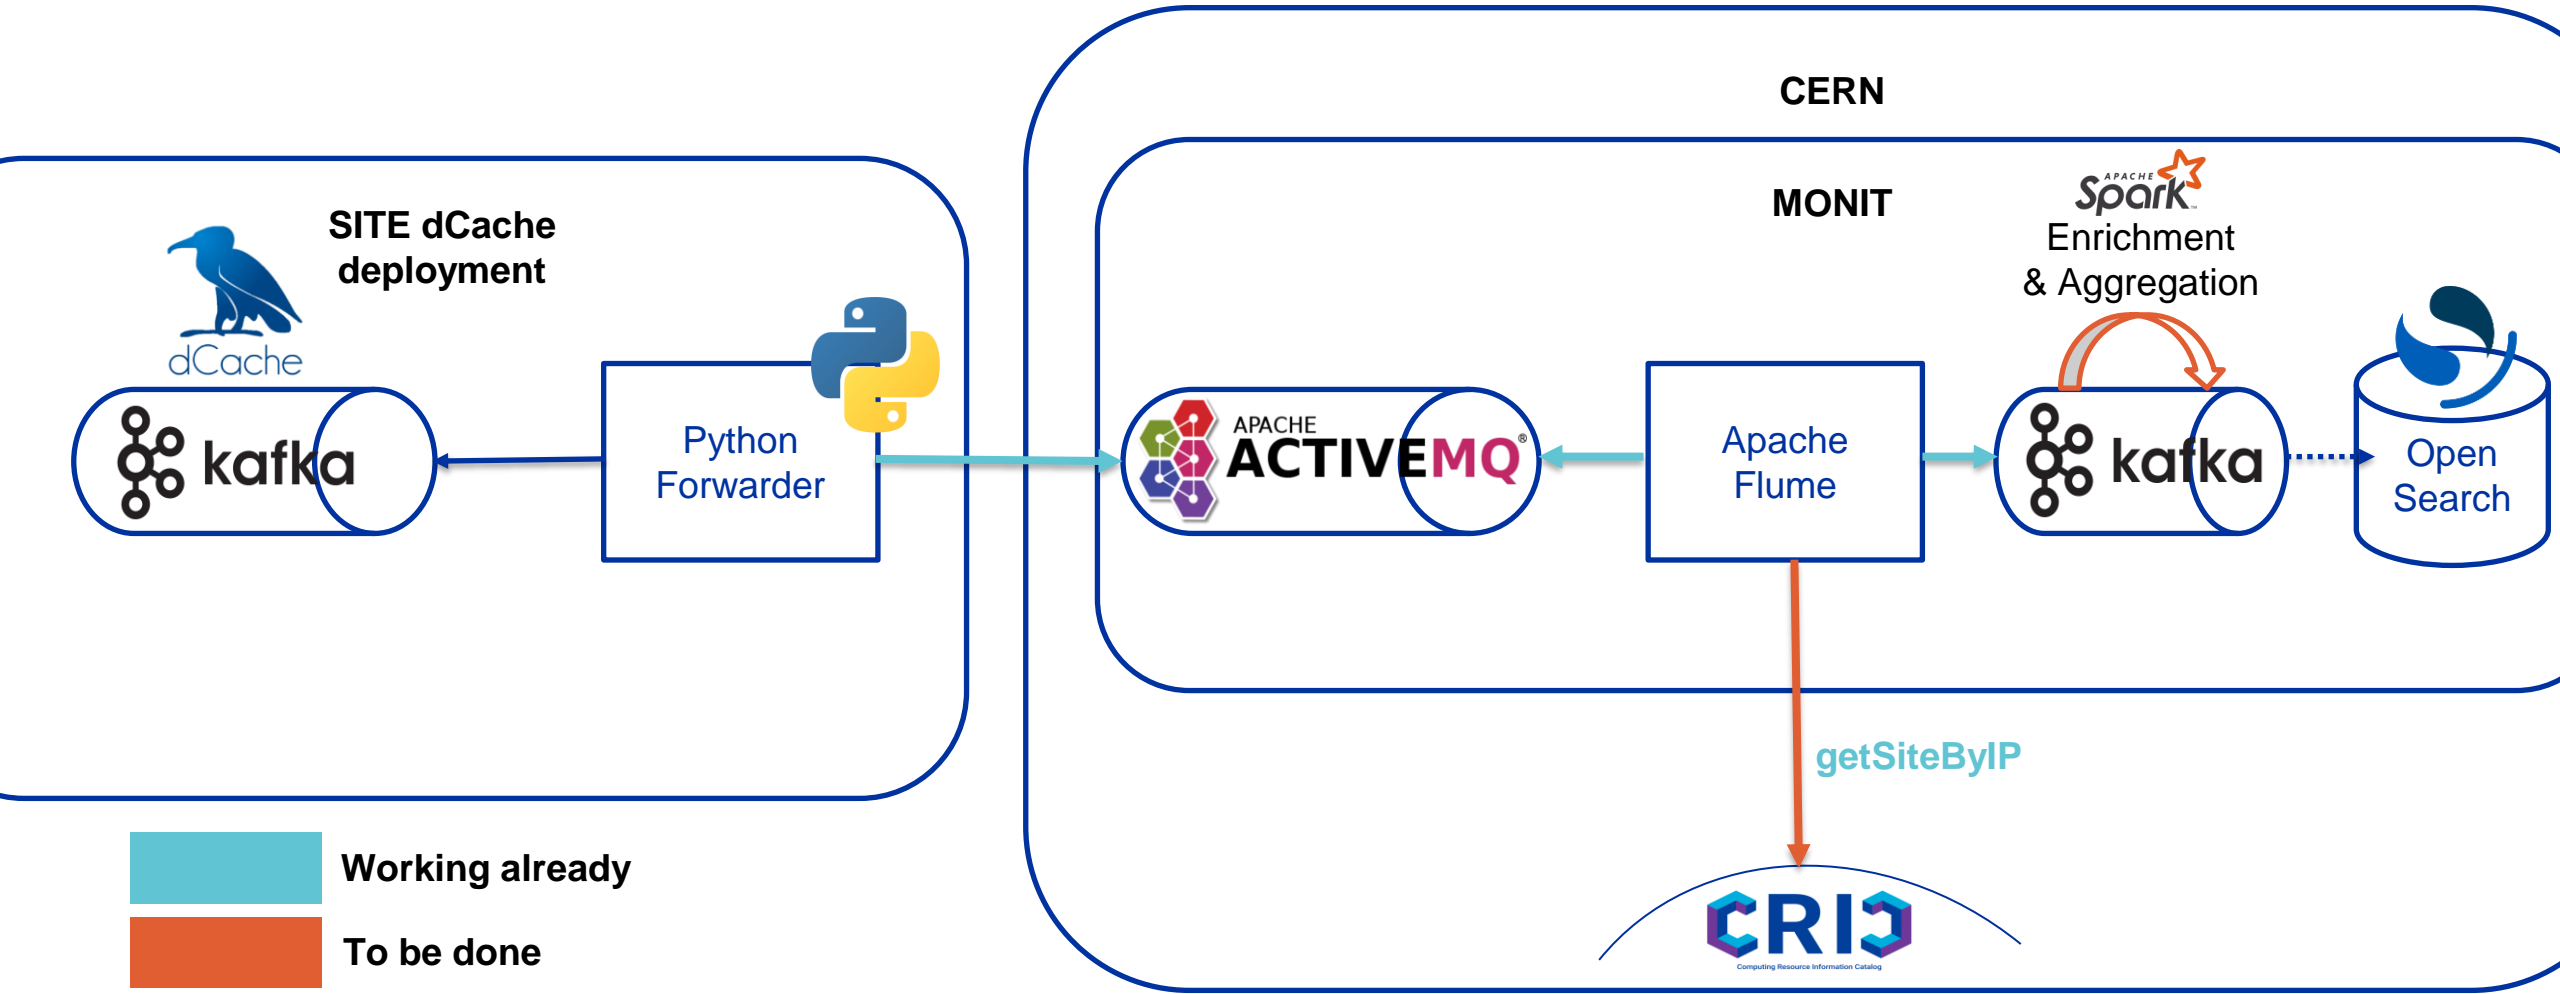
Prototype developed by dCache developers

- **Based on a Python script that connects Kafka to ActiveMQ**
 - Messages information to be compatible with WLCG standards
 - Will ship as part of the next "Golden" release (9.2 end of September)
 - We will include FNAL in the testing phase

Some extra work to be done on MONIT side

- **Resolve IP to site using CRIC provided API (Thanks!)**
- **Enrichment and aggregation**

New XRootD Monitoring (XRootD servers) (II)



Visualizations

Grafana v10

Update plan

31 August

10 October

[OTG0078726]

30 October

[OTG0079523]

4 December



Grafana 10 brings many new features



Upgrade process organised in two steps

- User Interface upgrade
- Unified Alerting



Migration process is starting NOW

Visit Grafana NEW

<https://monit-grafana-new.cern.ch>



Grafana 8 will run out of security release support in near future

Thank you !

Q & A

SNOW: [Monitoring Service](#)

Mattermost: [MONIT](#)

Docs: <https://monit-docs.web.cern.ch/>



home.cern