



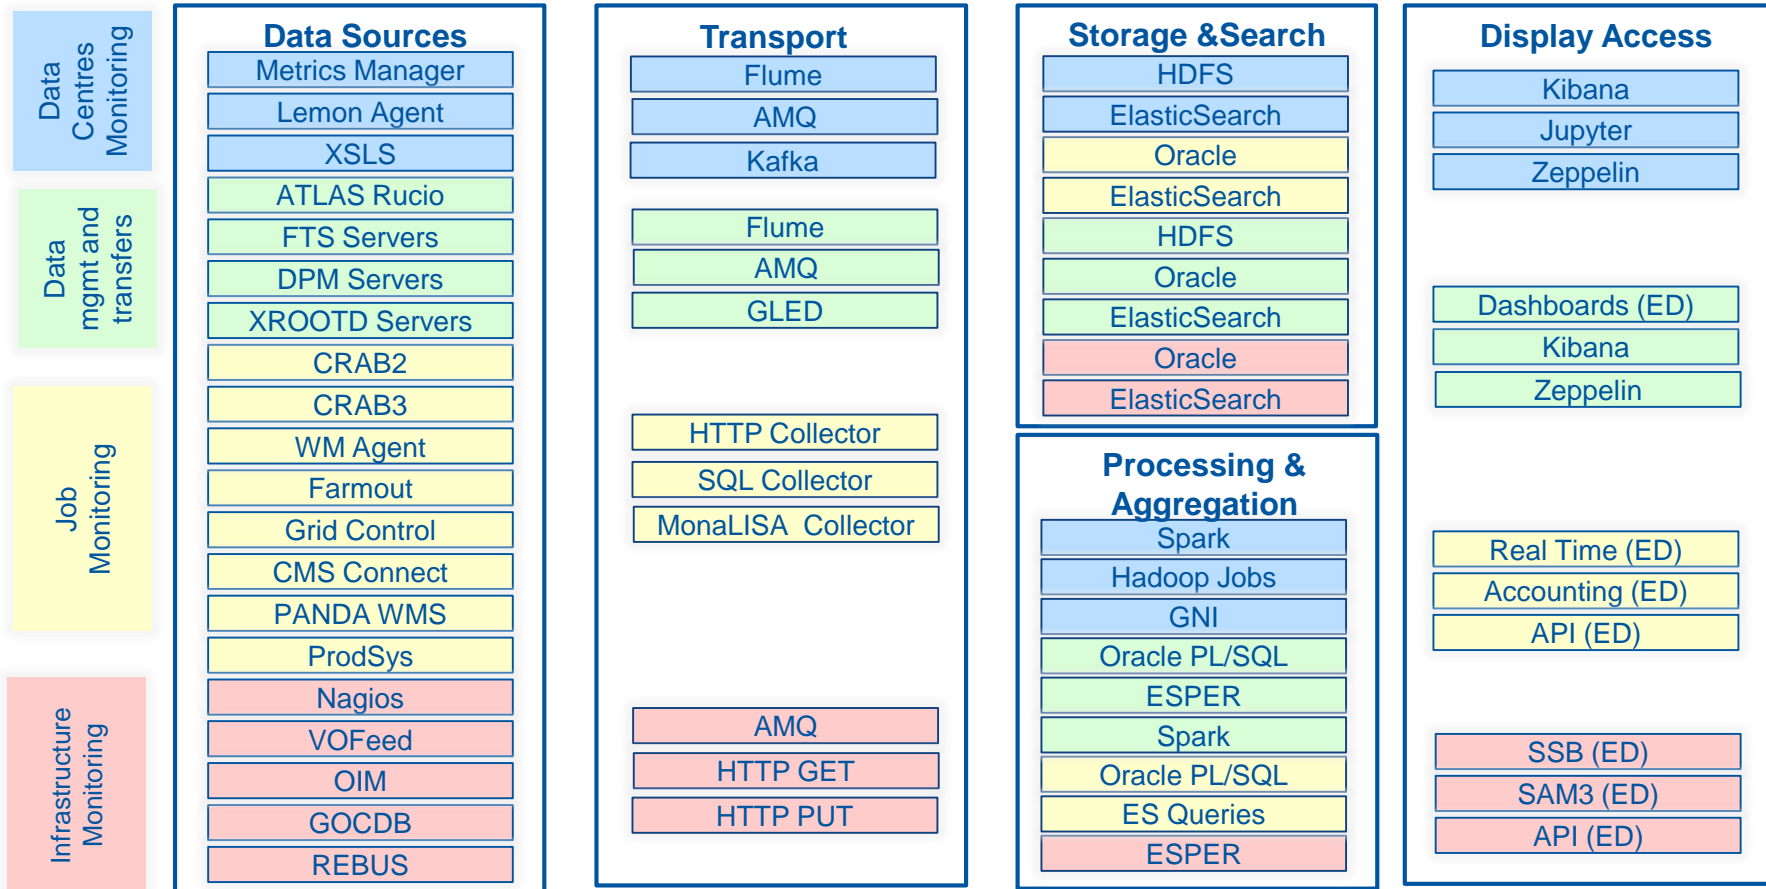
IT Monitoring Service

Update for Sites

Alberto AIMAR CERN-IT
for the MONIT Team

Architecture and Data Flow

Current Monitoring



Unified Monitoring

Data Sources

Metrics Manager

Lemon Agent

XSLs

ATLAS Rucio

FTS Servers

DPM Servers

XROOTD Servers

CRAB2

CRAB3

WM Agent

Farmout

Grid Control

CMS Connect

PANDA WMS

ProdSys

Nagios

VOFeed

OIM

GOCDB

REBUS

Transport

Flume

AMQ

Kafka

Storage & Search

Hadoop HDFS

ElasticSearch

InfluxDB

Processing & Aggregation

Spark

Hadoop Jobs

GNI

Data Access

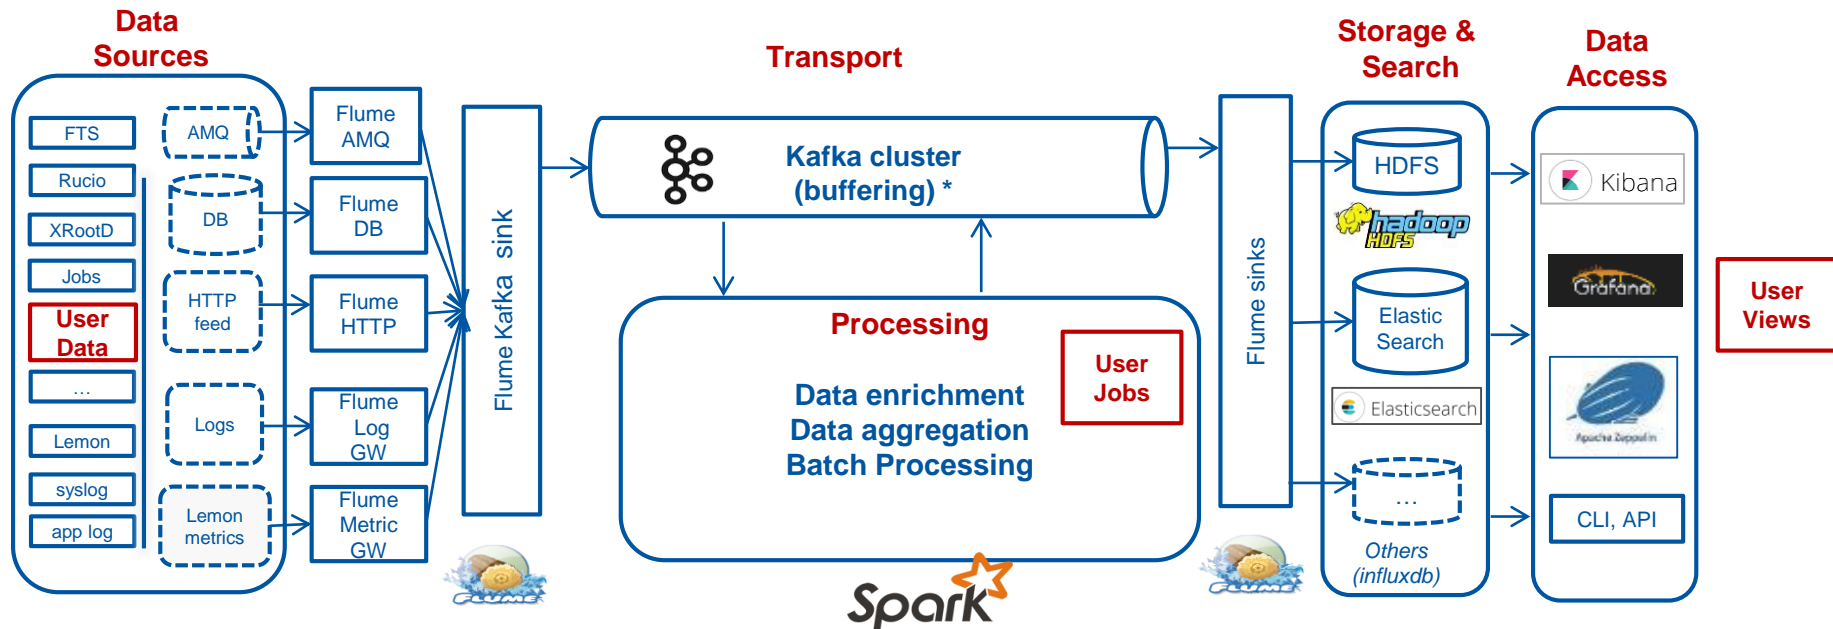
Kibana

Grafana

Jupyter (Swan)

Zeppelin

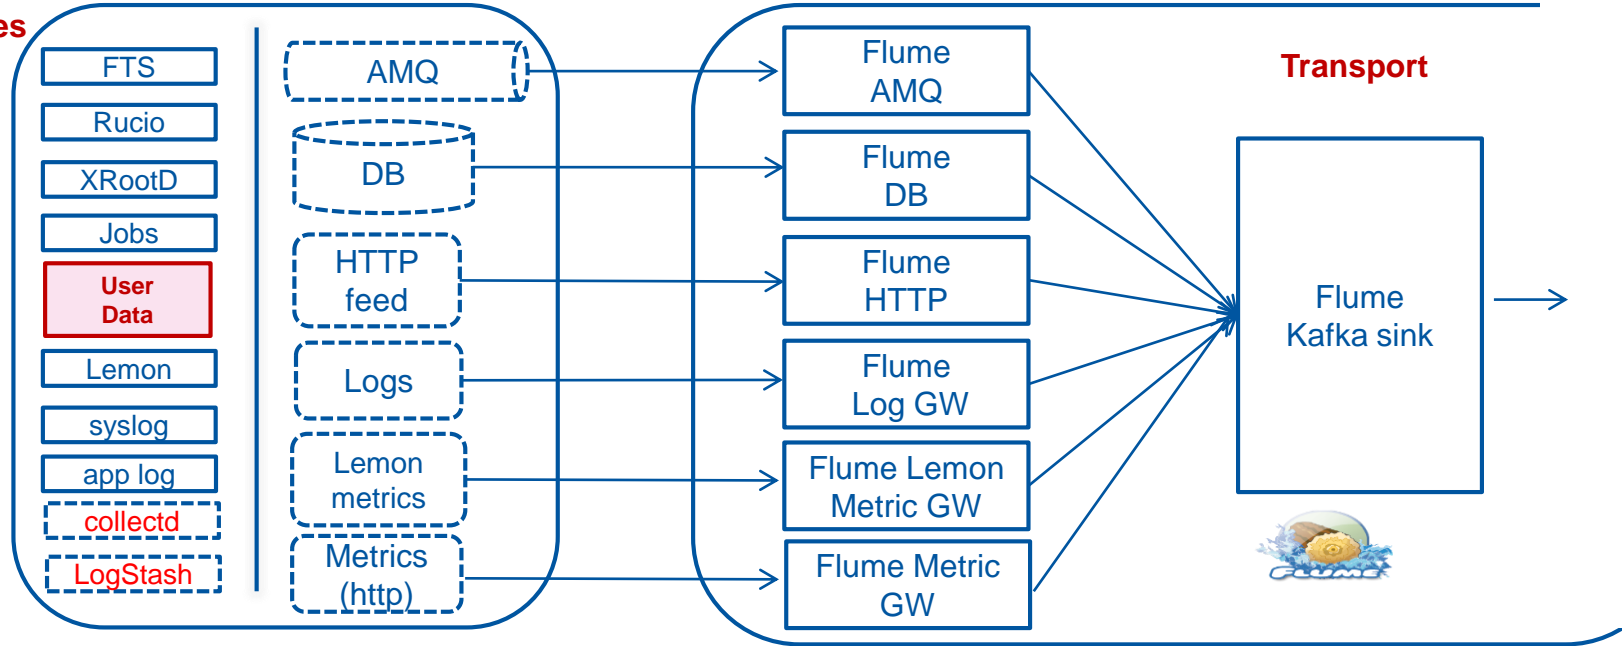
Unified Monitoring Architecture



- Data now 200 GB/day, 200M events/day. At scale 500 GB/day
- Current retention period 12 h, at scale 24 h
- Proved effective in several occasions

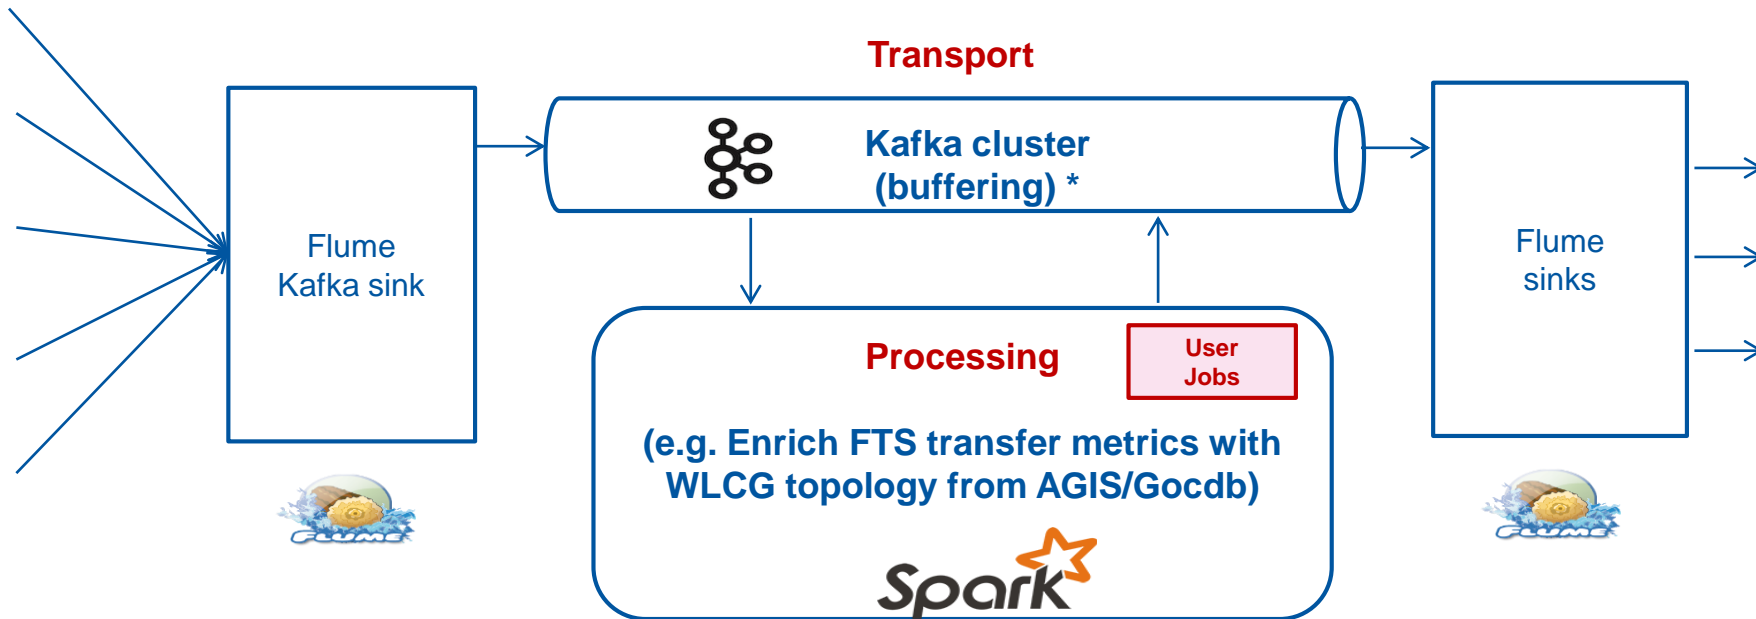
Unified Data Sources

Data
Sources



- Data is channeled via Flume, validated and modified if necessary
- Adding new Data Sources is documented and fairly simple

Unified Processing



- Data now 200 GB/day, 200M events/day. At scale 500 GB/day,
- Current retention period 12 h, at scale 24 h

Data Processing

Stream processing

Data enrichment

- Join information from several sources (e.g. WLCG topology)

Data aggregation

- Over time (e.g. summary statistics for a time bin)
- Over other dimensions (e.g. compute a cumulative metric for a set of machines hosting the same service)

Data correlation

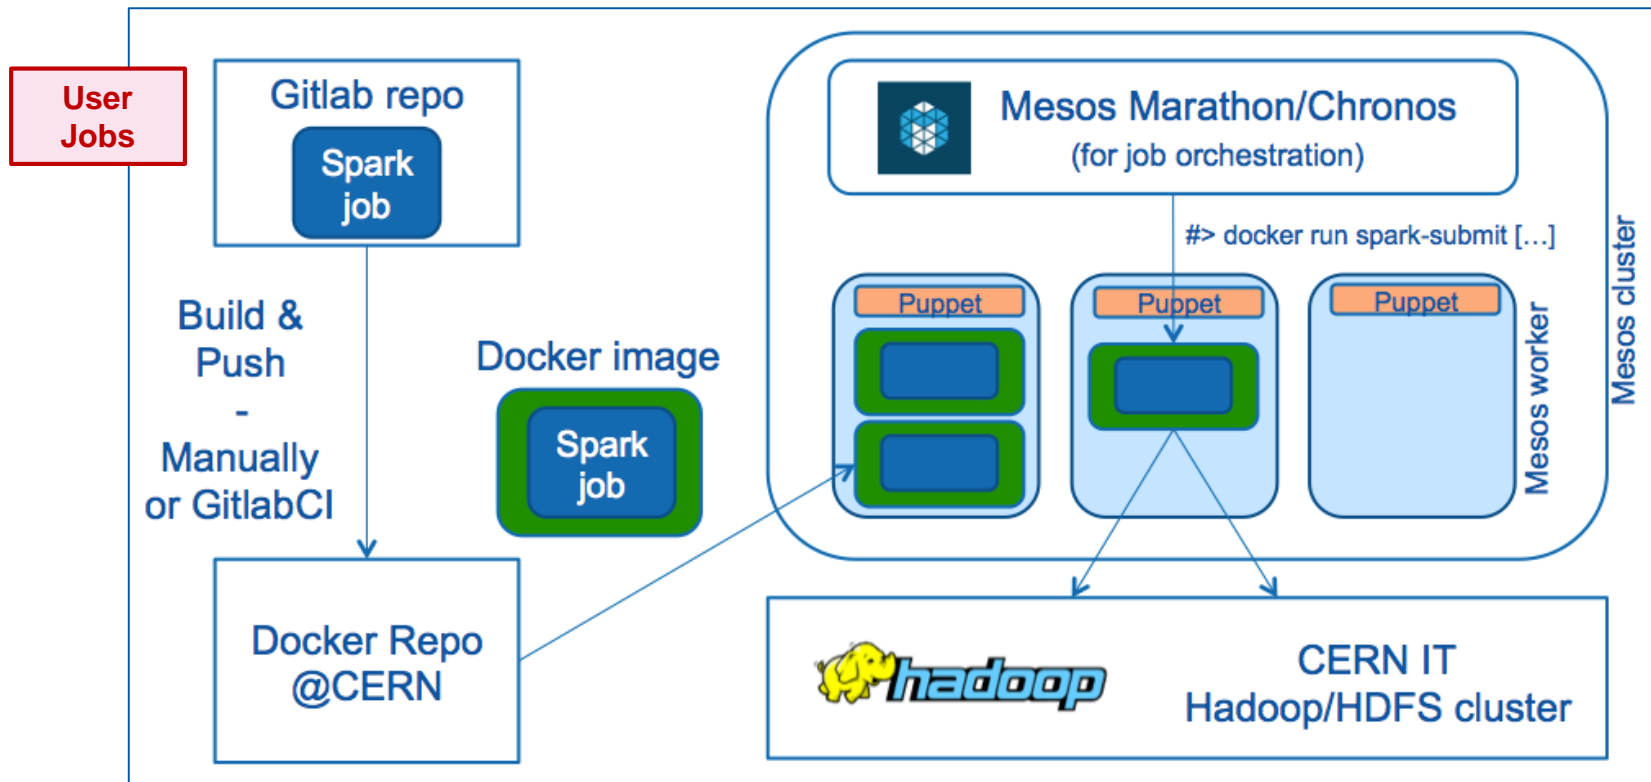
- Advanced Alarming: detect anomalies and failures correlating data from multiple sources (e.g. data centre topology-aware alarms)

Batch processing

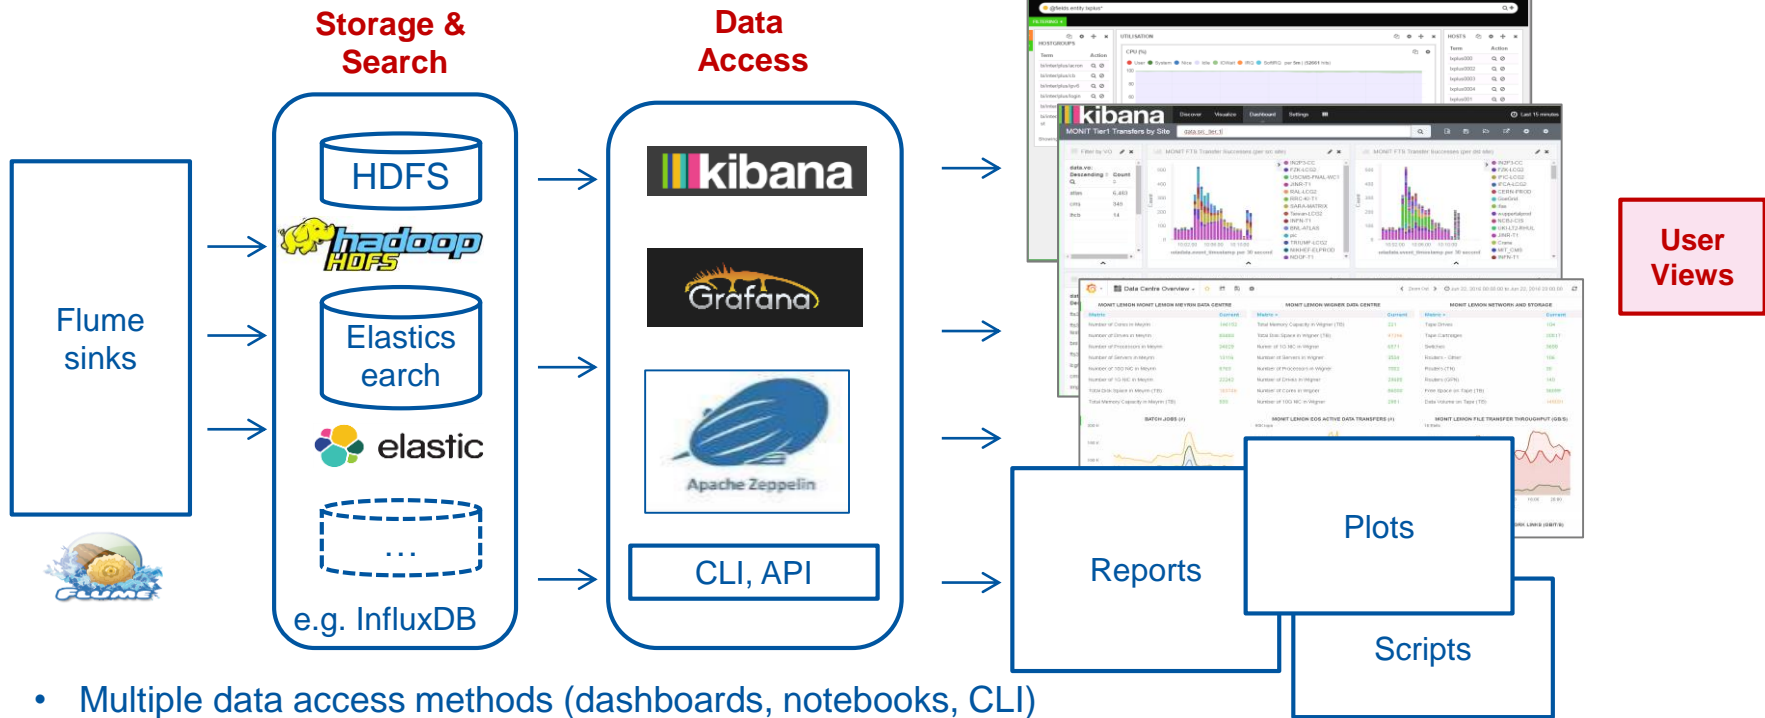
- Reprocessing, data compression, reports

Monitoring Processing Platform

Technologies: Reliable and scalable job execution (Spark), Job orchestration and scheduling (Marathon/Chronos), Lightweight and isolation deployment (Docker)



Unified Access



- Multiple data access methods (dashboards, notebooks, CLI)
- Mainstream and evolving technology

Data and Visualization

Data Storage and Search	
ElasticSearch	Short-term storage and index (months, depends on data)
InfluxDB	Short-term time series storage (months)
HDFS	Long-term archive (years)

Visualization	
Kibana	Data from ElasticSearch Dashboards and full search/filter/discovery of data
Grafana	Data from ElasticSearch, InfluxDB Dashboards optimized for time series plots
Zeppelin	Data from HDFS, ElasticSearch, InfluxDB Notebooks for analysis, reports and plots Native support for Spark
Swan (Jupyter)	Data from HDFS or ElasticSearch, InfluxDB Notebooks for analysis, reports and plots Integration with HEP toolbox (ROOT, CERNBOX, CVMFS, etc)

Status

Data Sources and Transport

- Moving all data via new transport (Flume, AMQ, Kafka)

Storage and Search

- Data in ES and HDFS, identical format

Processing

- Doing aggregation and processing via Spark

Display and reports

- Using only standard features of ES, Kibana, Spark, Hadoop
- Introduce notebooks (Zeppelin, Swan) and data discovery

Services Proposed

Monitor, collect, visualize, process, aggregate, alarm

- Metrics and Logs

Infrastructure operations and scale

Helping and supporting

- Interfacing new data sources
- Developing custom processing, aggregations, alarms
- Building dashboards and reports

Reference and Contact

Kibana Dashboards

monit.cern.ch

Feedback/Requests (SNOW)

cern.ch/monit-support

Early-Stage Documentation

cern.ch/monitdocs

MONIT Home - Dashboard

Discover Visualize Dashboard Settings

MONIT Home

MONIT Home: Welcome

Welcome to the MONIT dashboards portal. Here you can browse metrics and logs from the CERN data centres and the WLCG infrastructure. You can also find instructions on how to create custom dashboards.

MONIT Home: WLCG

Data Management Monitoring

- File Transfer Service
- XRootD (beta)
- WLCG Transfers (beta)
- ATLAS Data Management (beta)
- ATLAS DDM Queued Files (beta)

Job Monitoring

- ATLAS Real-Time (beta)
- ATLAS Accounting (beta)

Infrastructure Monitoring

- ETF Summary (beta)

MONIT Home: Data Centre

Overview

- Data Centre Overview

Host Metrics

- Host Metrics
- GNI Notifications

Service Metrics

- Service Metrics
- Service Availability

MONIT Home: User Info

Build your own Dashboard

- Naming Conventions
- Instructions
- Examples

Custom Dashboards

- Projects and Services

Support

- User Documentation
- Requests and Incidents

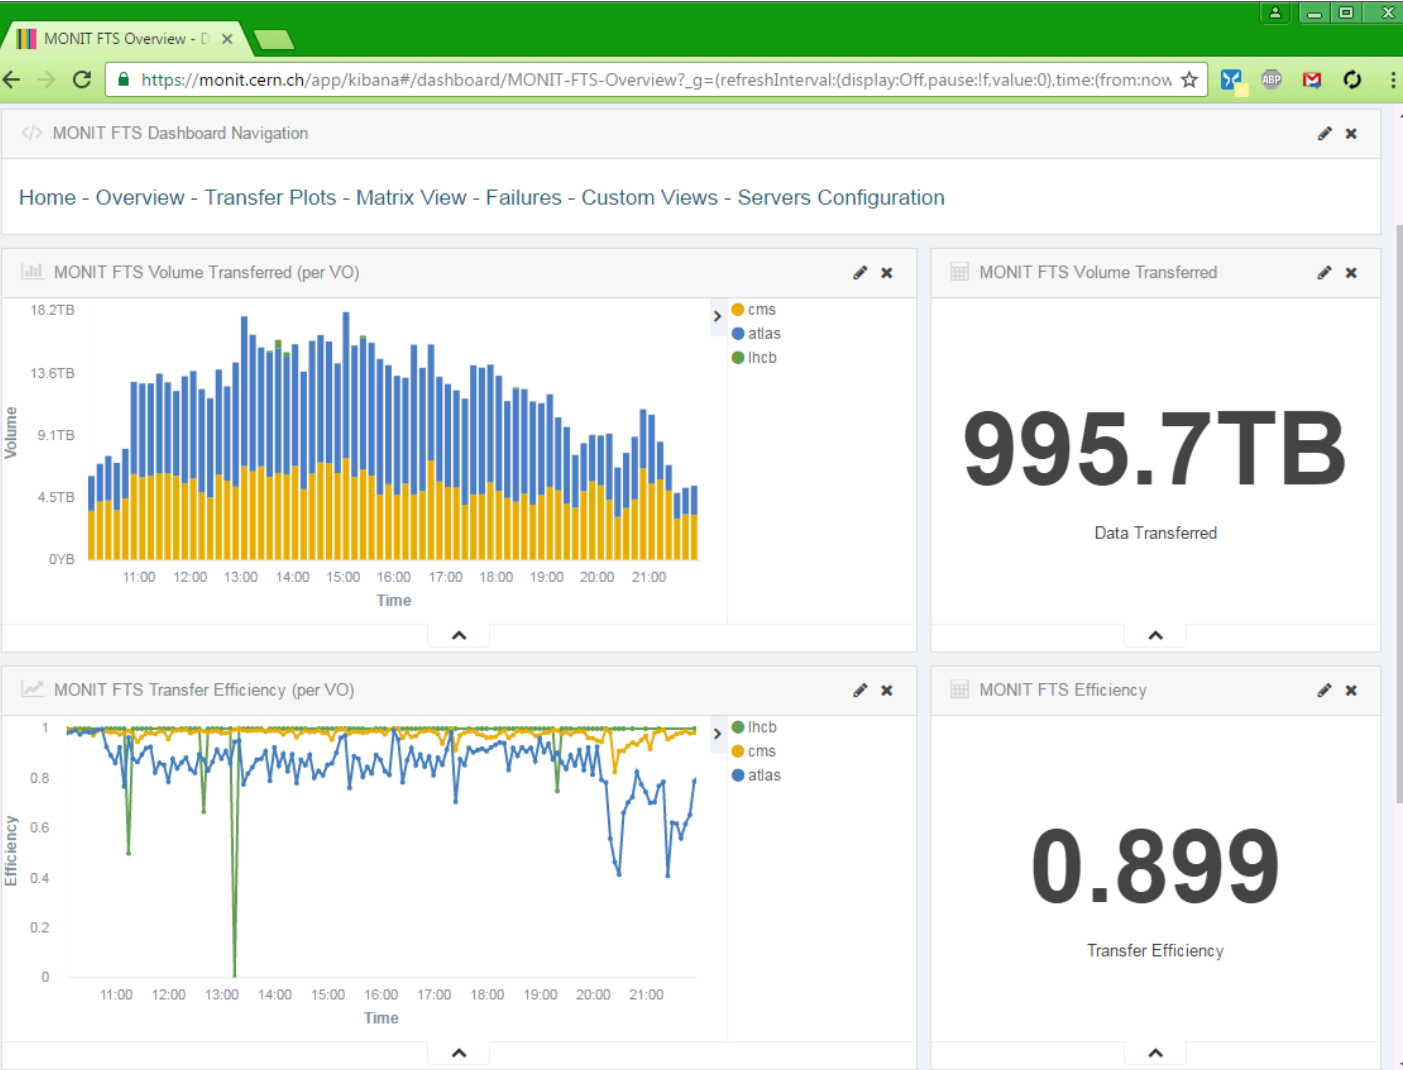
- Kibana homepage

Data Available

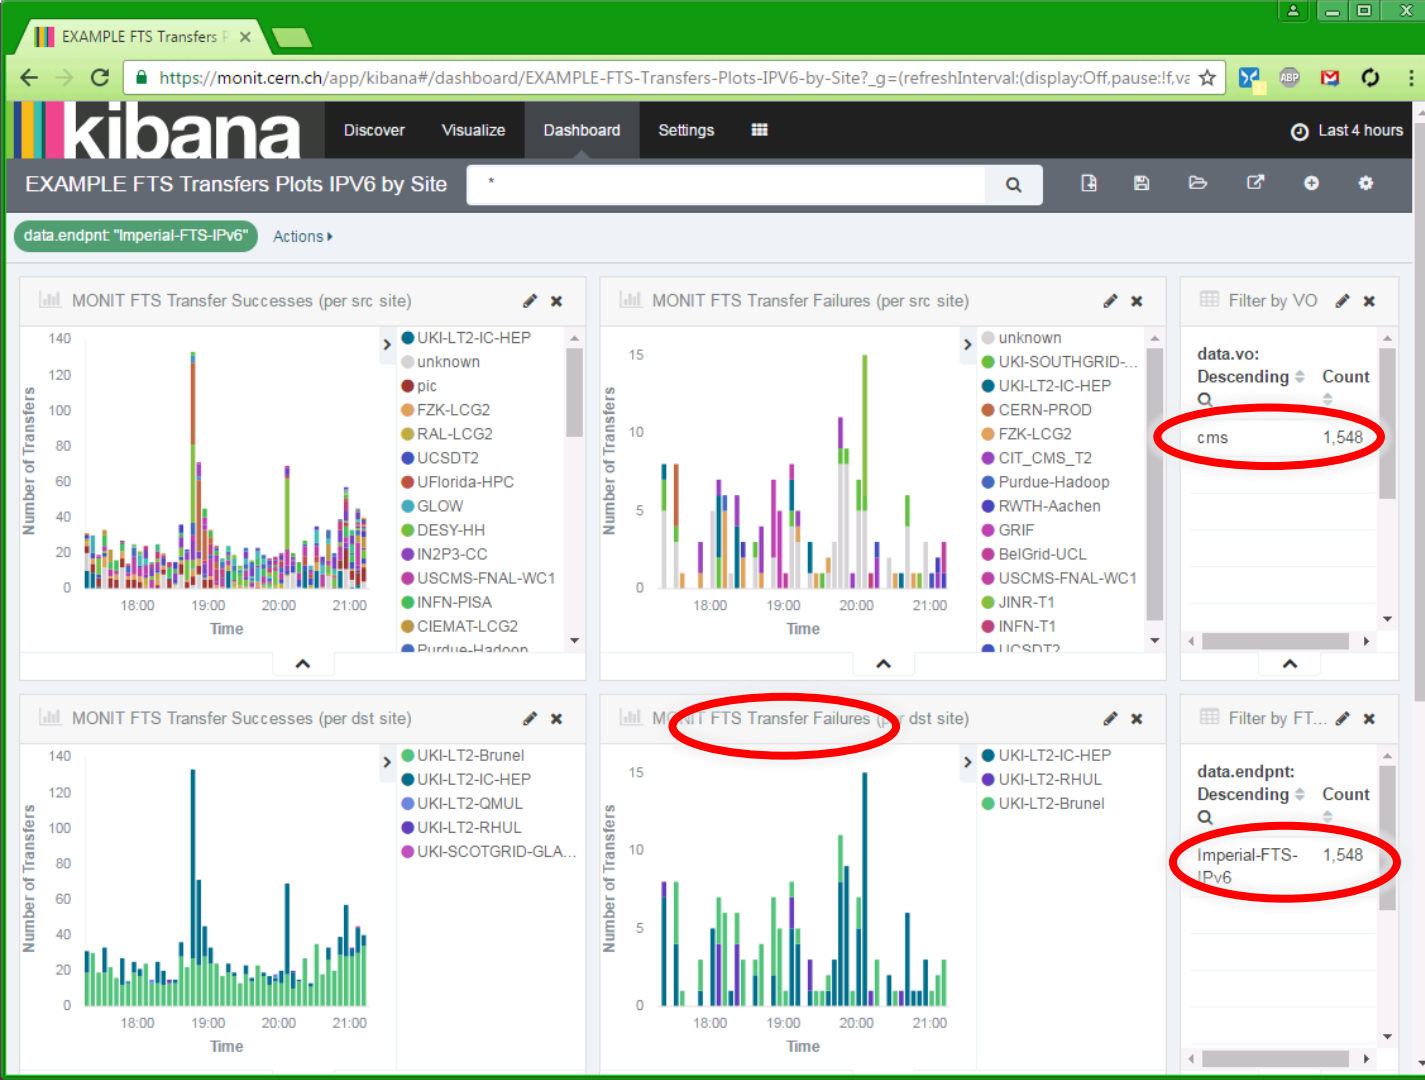
- FTS
- XRootD
- Job Monitoring Real-Time
- Job Monitoring Accounting
- SAM raw data
 - ETF
 - ALICE

Examples
Documentation

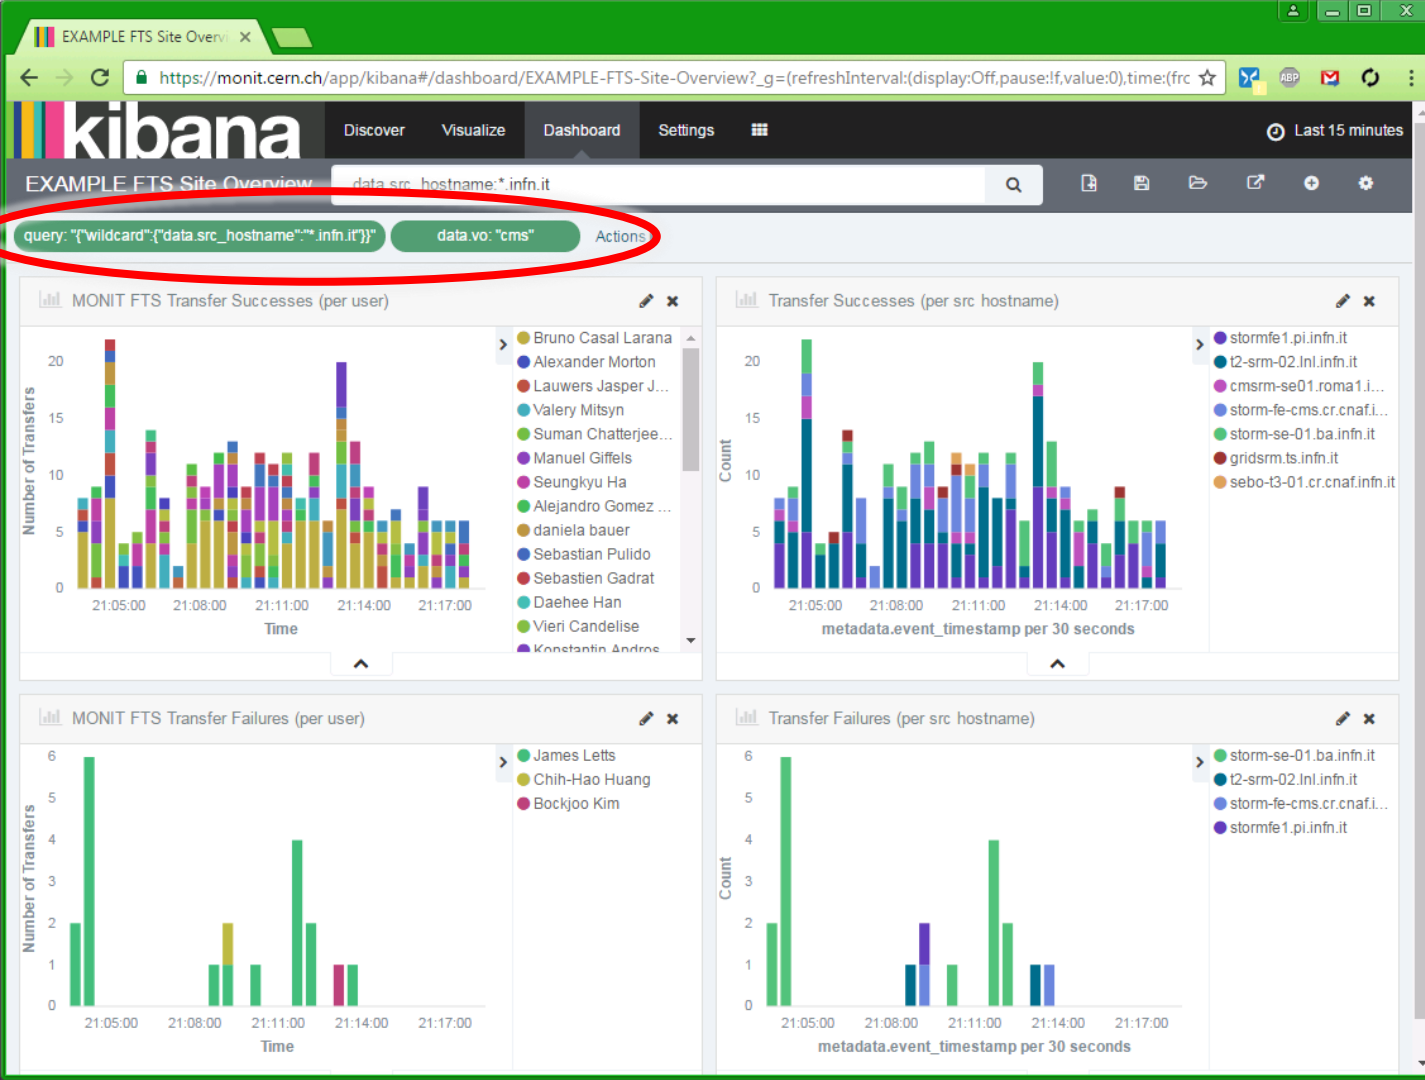
Link to Custom Projects



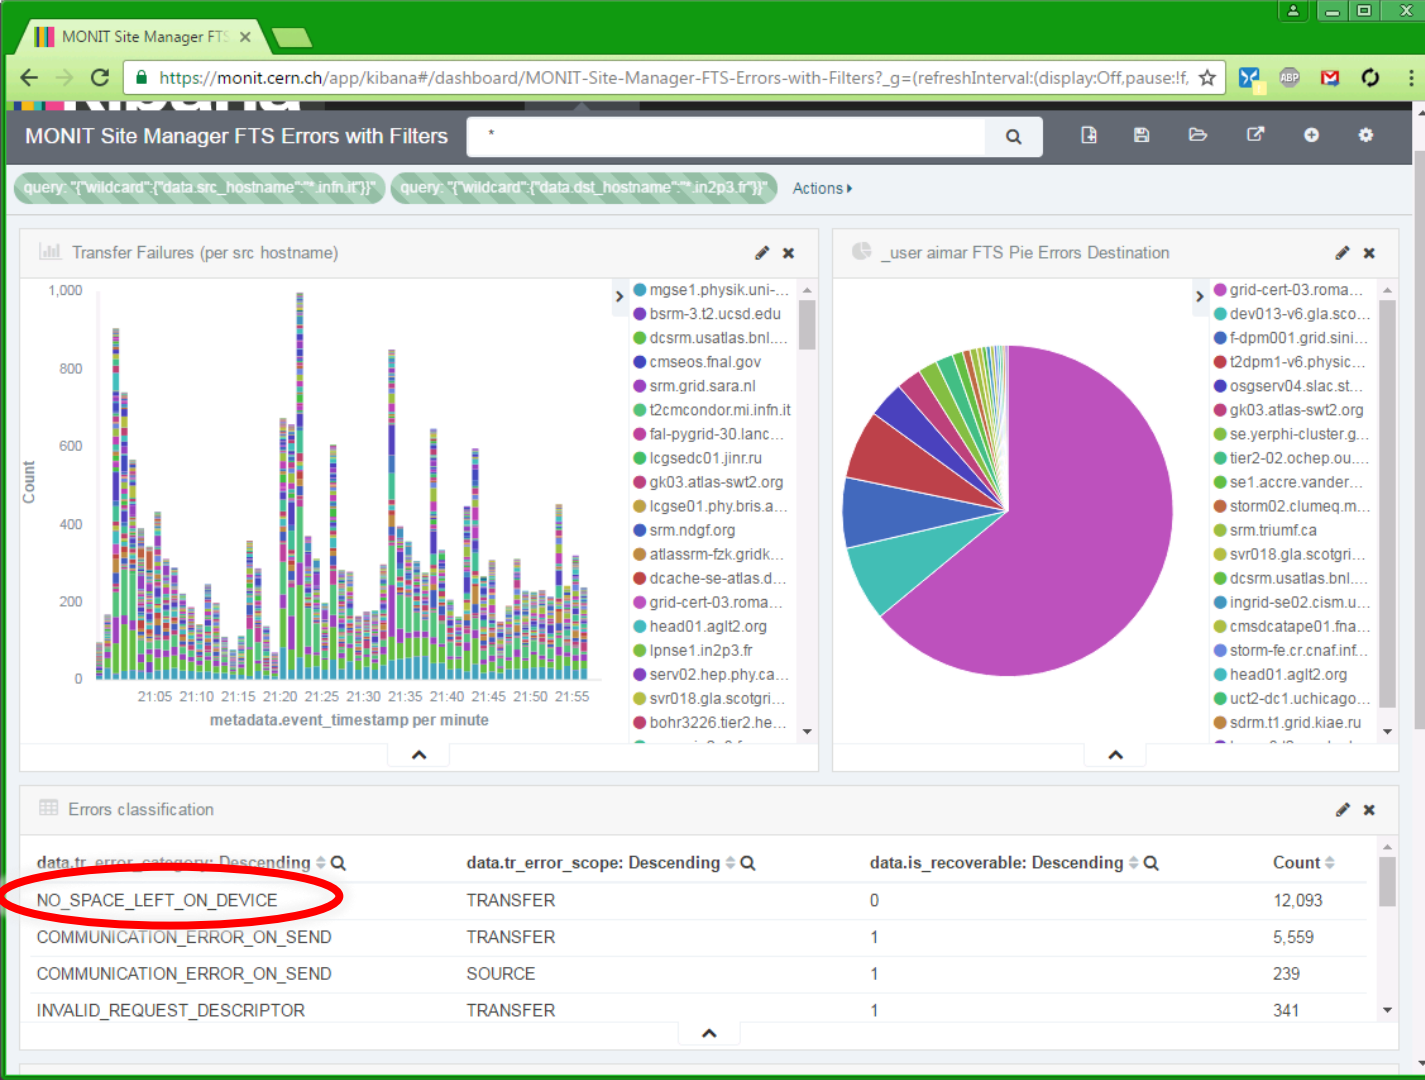
Transfers Overview



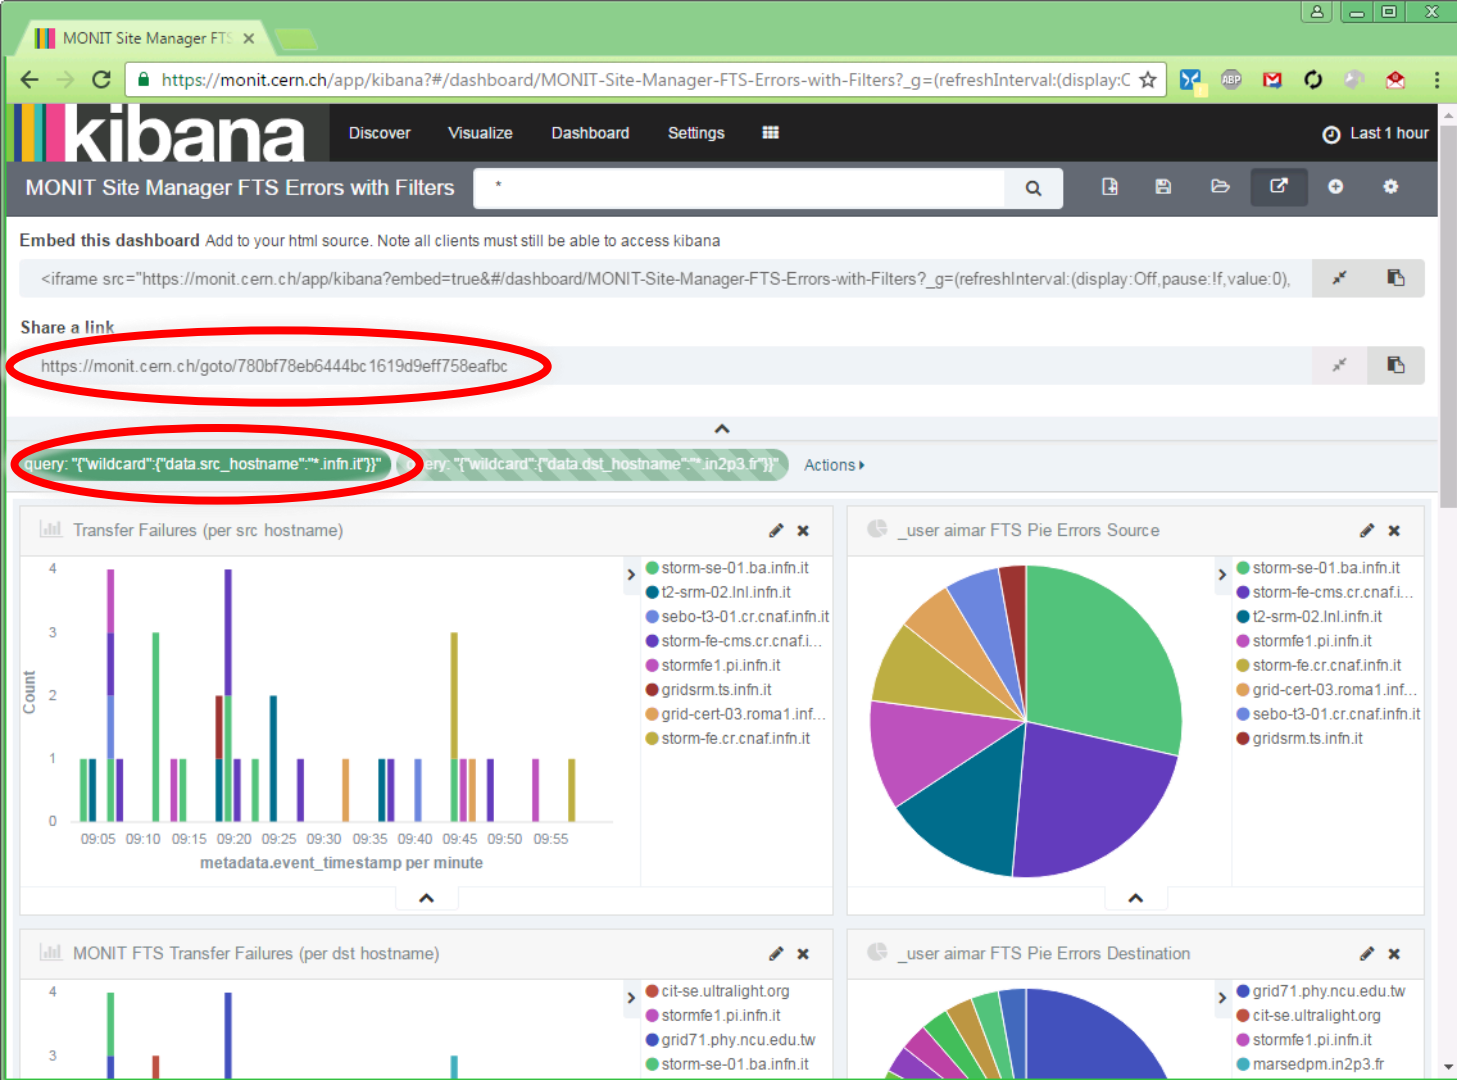
Transfer failures over IPV6



Success and failures from a src_site or a domain



Investigate and filter for a specific error

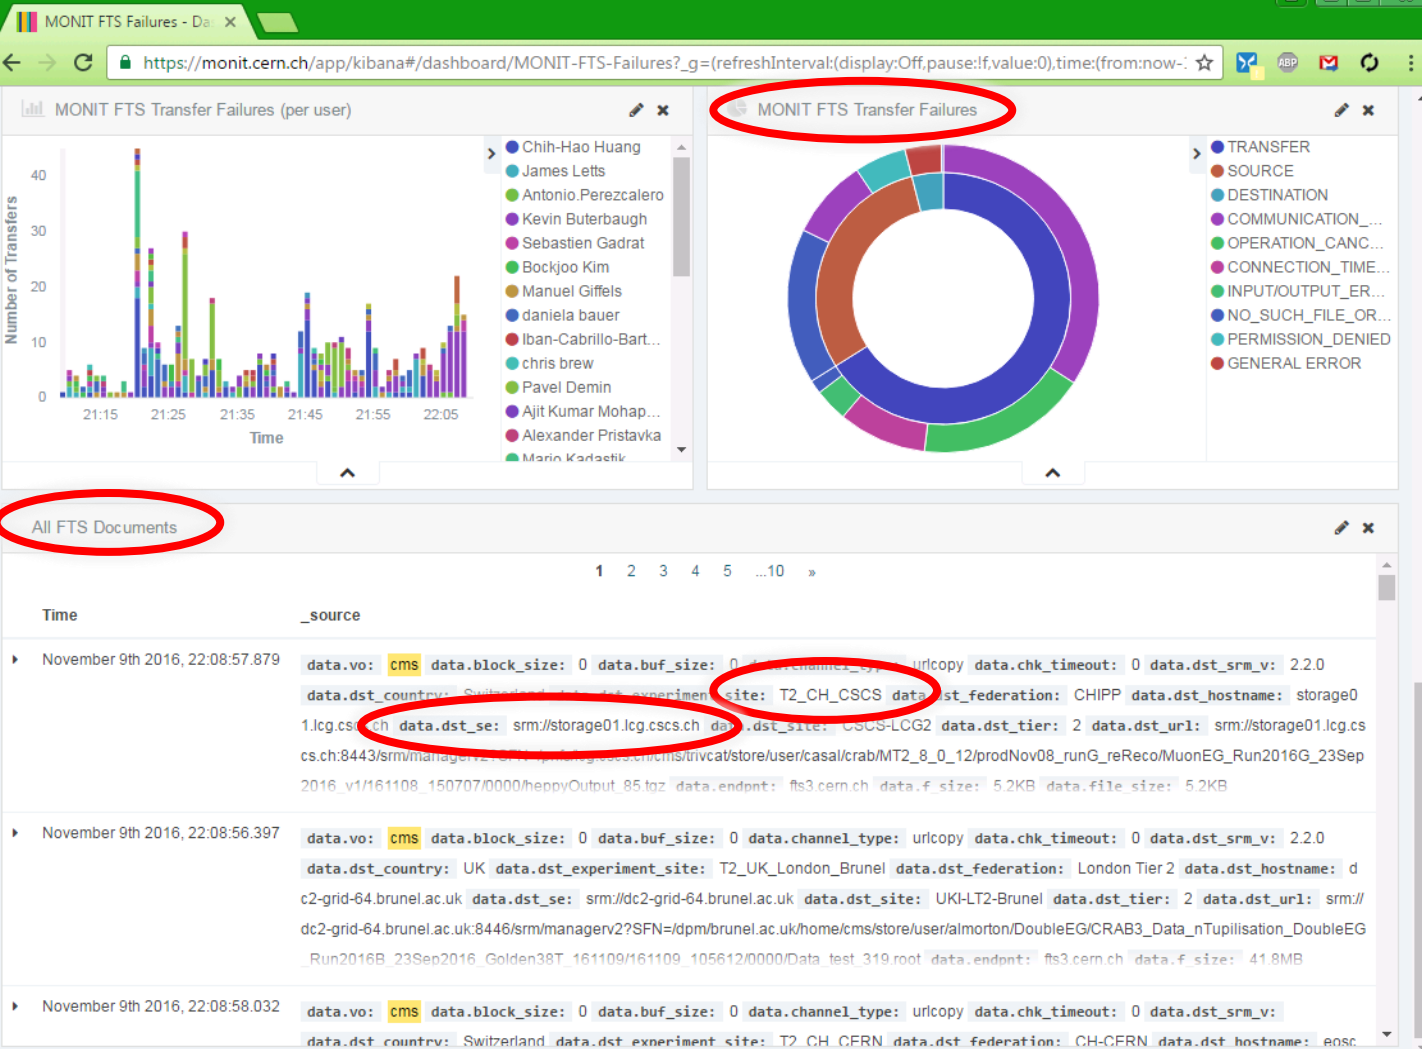


FTS Errors

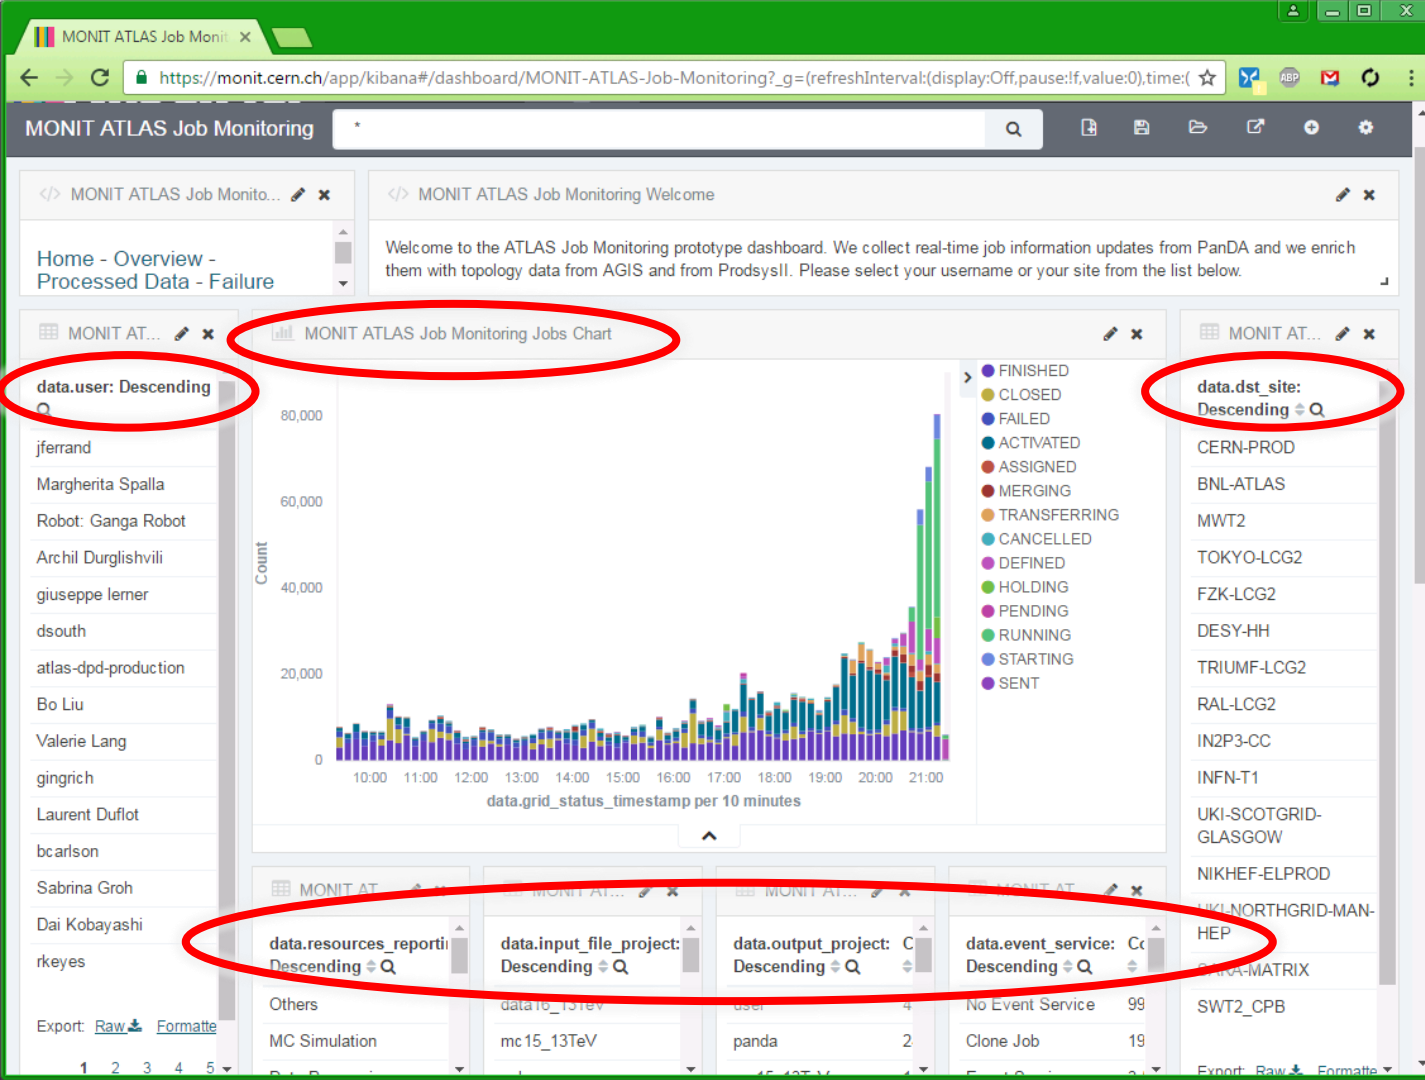
Select by a domain

Short link

GGUS ticket or email



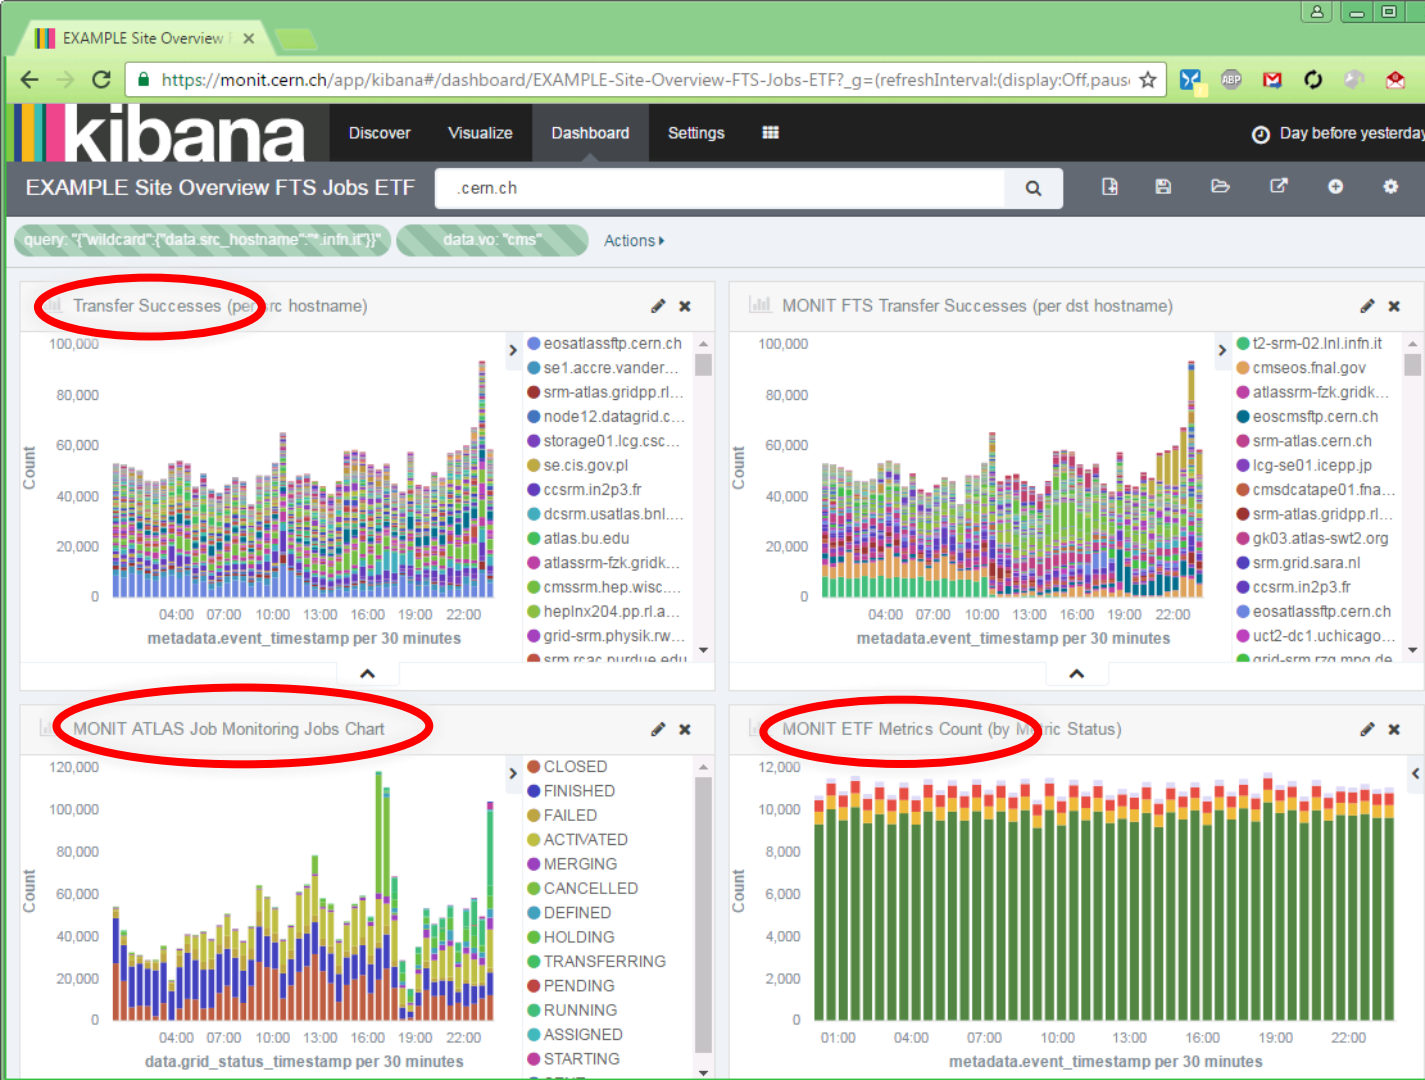
Filter and then
access the raw
data and logs



Job Monitoring

Real Time and Accounting

Filter by Site, User, cores, etc.

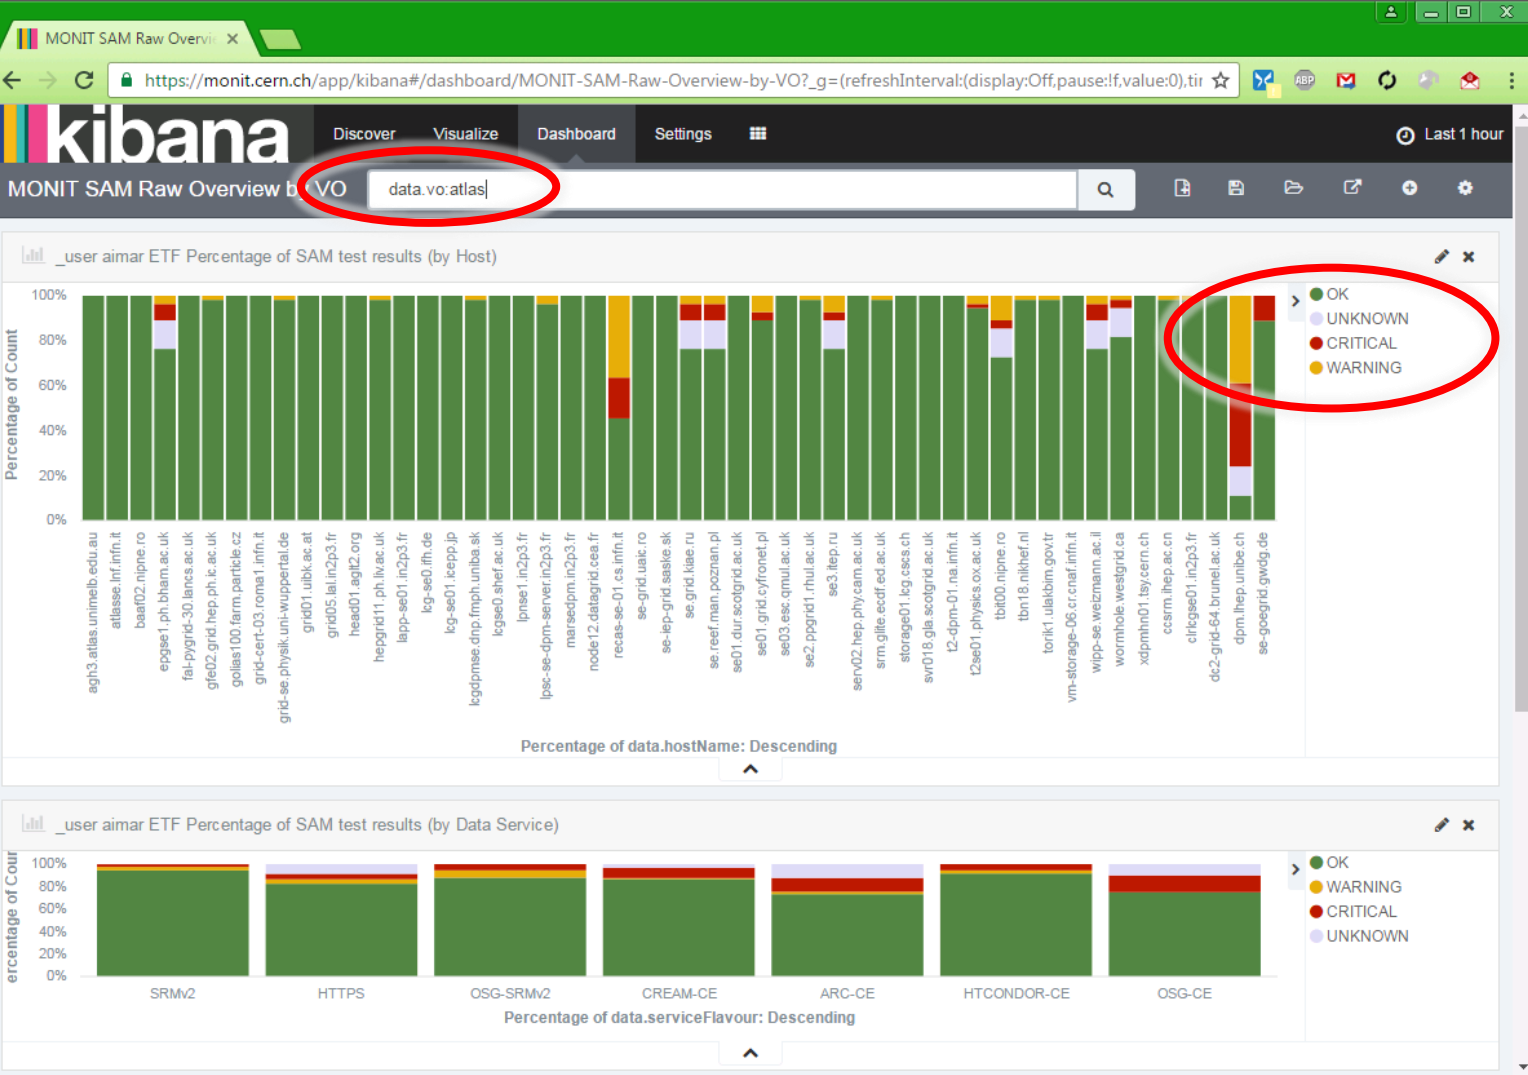


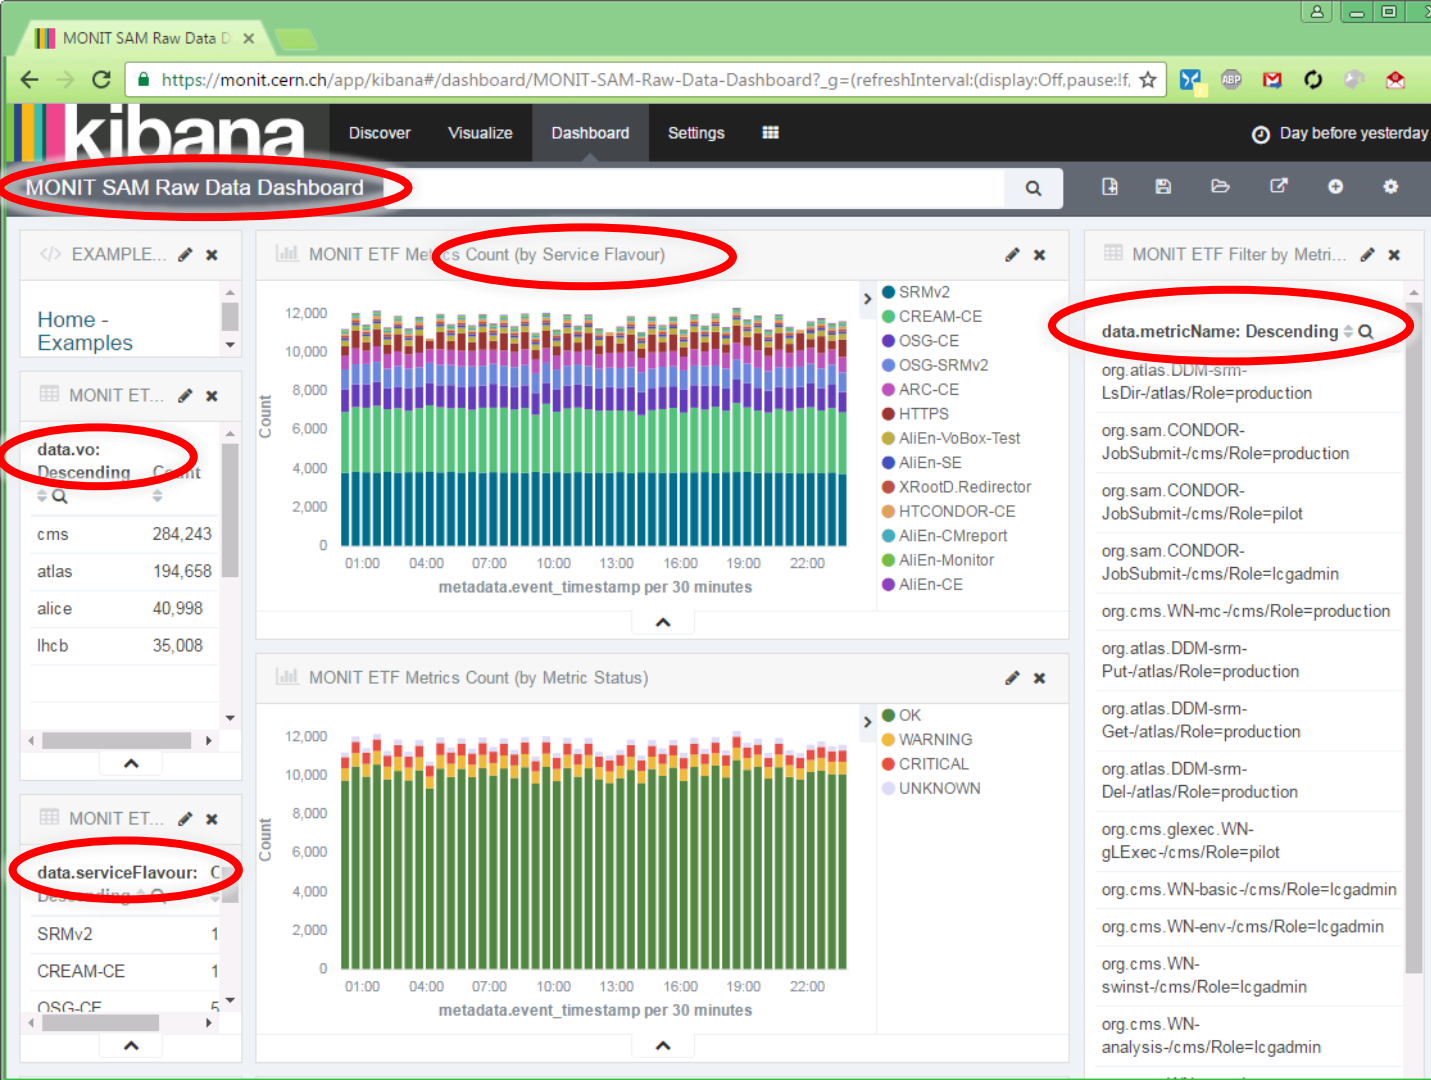
Site Overview
showing

- FTS
- Jobs
- SAM Raw

cern.ch hosts

SAM Raw Data





SAM Raw Data Overview

VO
Hosts
Metrics
Flavors

Discovering Data | Monit

monitdocs.web.cern.ch/monitdocs/discovering-data.html

Monitoring Service

1. Scope and Mandate

2. Service Overview

2.1. Status and Plans

2.2. Architecture

3. Data Access

3.1. Data Schema

3.2. Monit Kibana Portal

3.3. Discovering Data

3.4. Creating Dashboards

3.5. Processing Platform

3.6. Naming Conventions

4. Data Producers

4.1. Data Centre Metrics

4.2. WLCG Metrics

4.3. Service Logs

Published with GitBook

ES Documents and Fields

You will then be able to see all fields in each document and the more popular fields.

Discover - Kibana

monitdocs.web.cern.ch/app/kibana#/discover?_g=(refreshInterval:(display:Off,pauseIf:(value:0)),time:(from:now-15m,mode:quick))

July 15th 2016, 09:53:55.981

monit_prod_vofeed_*
monit_prod_xrootd_*
monit_prod_xrootd_enr_gled_*

Selected Fields
? _source
Available Fields

Popular
data.dst_url
data.f_size
data.job_id
data.src_hostname
data.src_se
data.src_url
data.throughput
data.timestamp_chk_src...
data.timestamp_chk_src...
data.transfer_duration
_id
_index
_score
_type
data.block_size
data.buf_size
data.channel_type

data.block_size: 0 data.buf_size: 0 data.channel_type: urlcopy data.chk_timeout: 1,800
data.dst_srm_v: 2.2.0 data.dst_country: Germany data.dst_experiment_site: FZK-LCG2
data.dst_federation: DE-KIT data.dst_hostname: atlasrm-fzk.gridka.de data.dst_se: srm://atlasrm-fzk.gridka.de data.dst_site: FZK-LCG2 data.dst_tier: 1 data.dst_url: srm://atlasrm-fzk.gridka.de:8443/srm/managerv2?SFN=/pnfs/gridka.de/atlas/disk-only/atlasdstatdisk-fzkiop.dat10-77ev/a6/9b/91stc_08004793_000

Link to /monit_prod_fts_enr_complete_v007-2016-07-15/enr_complete/AVxtjGTDqas_QRv7jv8

table	JSON
_id	AVxtjGTDqas_QRv7jv8
_index	monit_prod_fts_enr_complete_v007-2016-07-15
_score	
_type	enr_complete
data.block_size	0
data.buf_size	0
data.channel_type	urlcopy
data.chk_timeout	1,800
data.dst_country	Germany
data.dst_experiment_site	FZK-LCG2
data.dst_federation	DE-KIT
data.dst_hostname	atlasrm-fzk.gridka.de
data.dst_se	srm://atlasrm-fzk.gridka.de
data.dst_site	FZK-LCG2
data.dst_srm_v	2.2.0
data.dst_tier	1
data.dst_url	srm://atlasrm-fzk.gridka.de:8443/srm/managerv2?SFN=/pnfs/gridka.de/atlas/disk-only/atlasdstatdisk-fzkiop.dat10-77ev/a6/9b/91stc_08004793_000

These fields are those used for all the searches, visualizations and the filters in the Kibana dashboards in this portal. See some how they are used in the [Kibana Examples](#).

Access to raw data select by any of the fields

Important to know the structure of the data

27

MONIT SAM Raw Site View

Discover Visualize Dashboard Settings

MONIT SAM Raw Site View data.hostName:* .cern.ch

data.serviceFlavour: "OSG" data.metricStatus: "CRITICAL" data.hostName: "ce503.cern.ch" Actions

EXAMPLE... MONIT ETF Percentage of SAM test results (by Data Serv... MONIT ETF Percentage of SAM test results (by Host)

Home - Examples

MONIT ET...

data.vo: Descending Count

lhcb 10

cms 2

alice 1

Percentage of Count

Percentage of data.serviceFlavour: Descending

Percentage of data.hostName: Descending

MONIT ETF All Data Details

Time data.detailsData

December 1st 2016, 10:04:31.000 CRITICAL: Storage Path[gsiftp://eoscmsftp.cern.ch/eos/cms/store/unmerged/SAM/testSRM/SAM-eoscmsftp.cern.ch/cg-util]-ok;Storage Path[gsiftp://eoscmsftp.cern.ch/eos/cms/tier0/store/unmerged/SAM/testSRM/SAM-eoscmsftp.cern.ch/cg-util]-critical;\nTesting from: etf-18.cern.ch\nDN: /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=sciaba/CN=430796/CN=Andrea Sciaba/CN=1109873700/CN=731975511/CN=1054687011/CN=1461344890\nVOMS FQANs: /cms/Role=production/Capability=NULL, /cms/TEAM/Role=NULL/Capability=NULL, /cms/Role=NULL/Capability=NULL, /cms/TEAM/Role=NULL/Capability=NULL\nngfal2 2.9.3\nUsing ngfal2 listdir() \nParameters: \n surds: /osiftp://eoscmsftp.cern.ch/eos/cms/store/unmerged/SAM/testSRM/SAM-eoscmsftp.cern.ch/cg-util

December 1st 2016, 09:59:07.000 \nInfo : Processing MACHINEFEATURES\nInfo : MACHINEFEATURES=/var/run/wlcc-mjf-host-features\nError : Key total_cpu absent (or empty)\nInfo : Key hs06 found with value 82.2725\nInfo : Key grace_secs absent (or empty)\nInfo : Key shutdowntime absent (or empty)\nInfo : Processing JOBFEATURES\nInfo :

Select the errors (CRITICAL)

See all details of the errors

The screenshot shows a web browser window with the URL `monitdocs.web.cern.ch/monitdocs/data_access.html`. The page has a light green header and a sidebar on the left. The sidebar contains a navigation menu with the following items: Monitoring Service, 1. Scope and Mandate, 2. Service Overview, 2.1. Status and Plans, 2.2. Architecture, 3. Data Producers, 3.1. Data Centre Metrics, 3.2. WLCG Metrics, 3.3. Service Logs, 3.4. Service Metrics, 4. Data Access (highlighted in blue), 4.1. Data Schema, 4.2. Kibana Dashboards, 4.3. Grafana Dashboards, 4.4. Zeppelin Notebooks, 4.5. Spark Jobs, 4.6. Naming Conventions, and Published with GitBook. The main content area has a white background and a top navigation bar with a hamburger menu, search icon, and social media icons. The main heading is **Data Access**. Below it, a paragraph states: "Data collected by the monitoring infrastructure can be accessed in different ways." The next heading is **Dashboards**. Below it, a paragraph states: "The main entry point is: <https://Vmonit.cern.ch>". This is followed by another paragraph: "This service includes two main types of dashboards:" and a bulleted list:

- public views: listed in the homepage
- private views: shared with a dedicated links

. A horizontal separator line follows. Below it, a paragraph states: "Use the articles in the left panel learn about:" and a bulleted list:

- [Monit Kibana portal](#)
- [Discovering data in the portal](#)
- [Naming conventions to follow](#)
- [Examples of Kibana dashboards](#)

. Another horizontal separator line follows. At the bottom, a paragraph states: "Support: To report any request or incident please refer to this support page: <http://cern.ch/monit-support>".

User Documentations

cern.ch/monitdocs

Common fields

Monitoring Service

1. Scope and Mandate
2. Service Overview
 - 2.1. Status and Plans
 - 2.2. Architecture
3. Data Producers
 - 3.1. Data Centre Metrics
 - 3.2. WLCG Metrics
 - 3.3. Service Logs
 - 3.4. Service Metrics
4. Data Access
 - 4.1. Data Schema
 - 4.2. Kibana Dashboards
 - 4.3. Grafana Dashboards
 - 4.4. Zeppelin Notebooks
 - 4.5. Spark Jobs
 - 4.6. Naming Conventions

Published with GitBook

- metadata.producer: data producer identifier
- metadata.type_prefix: data type (e.g. raw, agg)
- metadata.type: data producer category
- metadata.hostname: host producing the data
- metadata.event_timestamp: time the document was produced
- metadata.flume_timestamp: time the document arrived to flume
- metadata.kafka_timestamp: time the document arrived to kafka

For the detailed schema of each producer please use the "discovery" tab in [Monit](#).

WLCG Metrics

Due to the diversity of WLCG producers, several common fields are defined across all data sets:

- data.vo
- data.user
- data.src_country / data.dst_country
- data.src_site / data.dst_site
- data.src_experiment_site / data.dst_experiment_site
- data.src_federation / data.dst_federation
- data.src_tier / data.dst_tier

For data management related producers, the following additional fields are also defined:

- data.throughput
- data.operation_time

Naming Conventions | M x

← → ↺

monitdocs.web.cern.ch/monitdocs/dashboards_naming.html

🔍 ☆

Monitoring Service

1. Scope and Mandate

2. Service Overview

2.1. Status and Plans

2.2. Architecture

3. Data Producers

3.1. Data Centre Metrics

3.2. WLCG Metrics

3.3. Service Logs

3.4. Service Metrics

4. Data Access

4.1. Data Schema

4.2. Kibana Dashboards

4.3. Grafana Dashboards

4.4. Zeppelin Notebooks

4.5. Spark Jobs

4.6. Naming Conventions

Published with GitBook

Managed by MONIT team

Official

`MONIT title` (e.g. MONIT FTS Throughput)

Examples

`EXAMPLE title` (e.g. EXAMPLE FTS Transfers Plots)

Managed by MONIT users

Users

`_user account title` (e.g. `_user johndoe` FTS Error Investigation)

Projects

`_project projectname title` (e.g. `_project` FTS C5 Report)

VOs

`_vo voname title` (e.g. `_vo CMS` FTS Error Investigation)

Sites

`_site sitename title` (e.g. `_site IN2P3` Site Overview or `_site INFN` Global Overview)

Naming Conventions

Reference and Contact

Kibana Dashboards

monit.cern.ch

Feedback/Requests (SNOW)

cern.ch/monit-support

Early-Stage Documentation

cern.ch/monitdocs

