

Machine Learning for ATLAS Data Management


Mario Lassnig

European Organization for Nuclear Research (CERN)

mario.lassnig@cern.ch

Data management responsibilities

Organise all of the experiment's data
Make data available to experiment's workflows



The software stack that implements ATLAS DDM is called *Rucio*
Built upon the experiences from LHC Run-1 DDM system *Don Quijote 2 (DQ2)*
Modularized – extensible and open architecture
Scalable by design – beyond LHC Run-2

Generic and adaptive system – Also in use by
Alpha Magnetic Spectrometer experiment



Data volume

Files organised in datasets and containers

600 million files

20 million datasets & containers

40 million archived datasets

File replicas on storage

650 million replicas

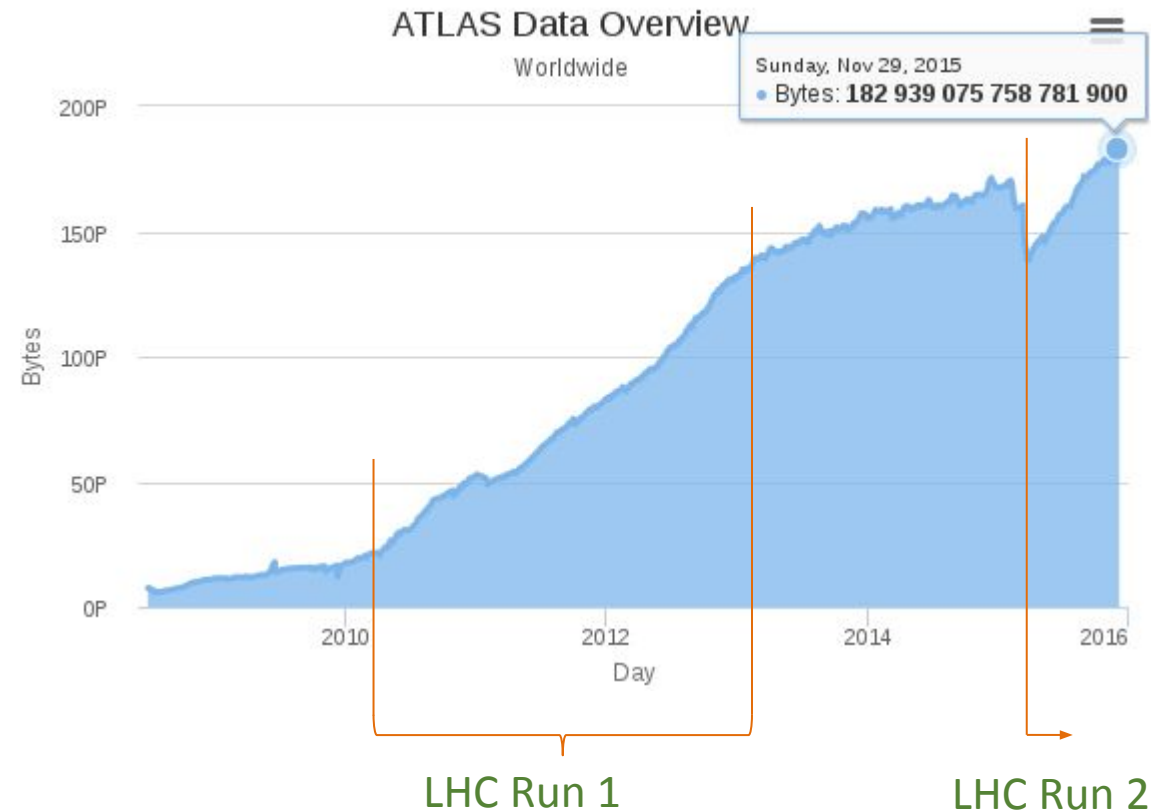
720 storage endpoints

145 data centres


Access control

9400 identities

5300 user accounts




Scheduled throughput



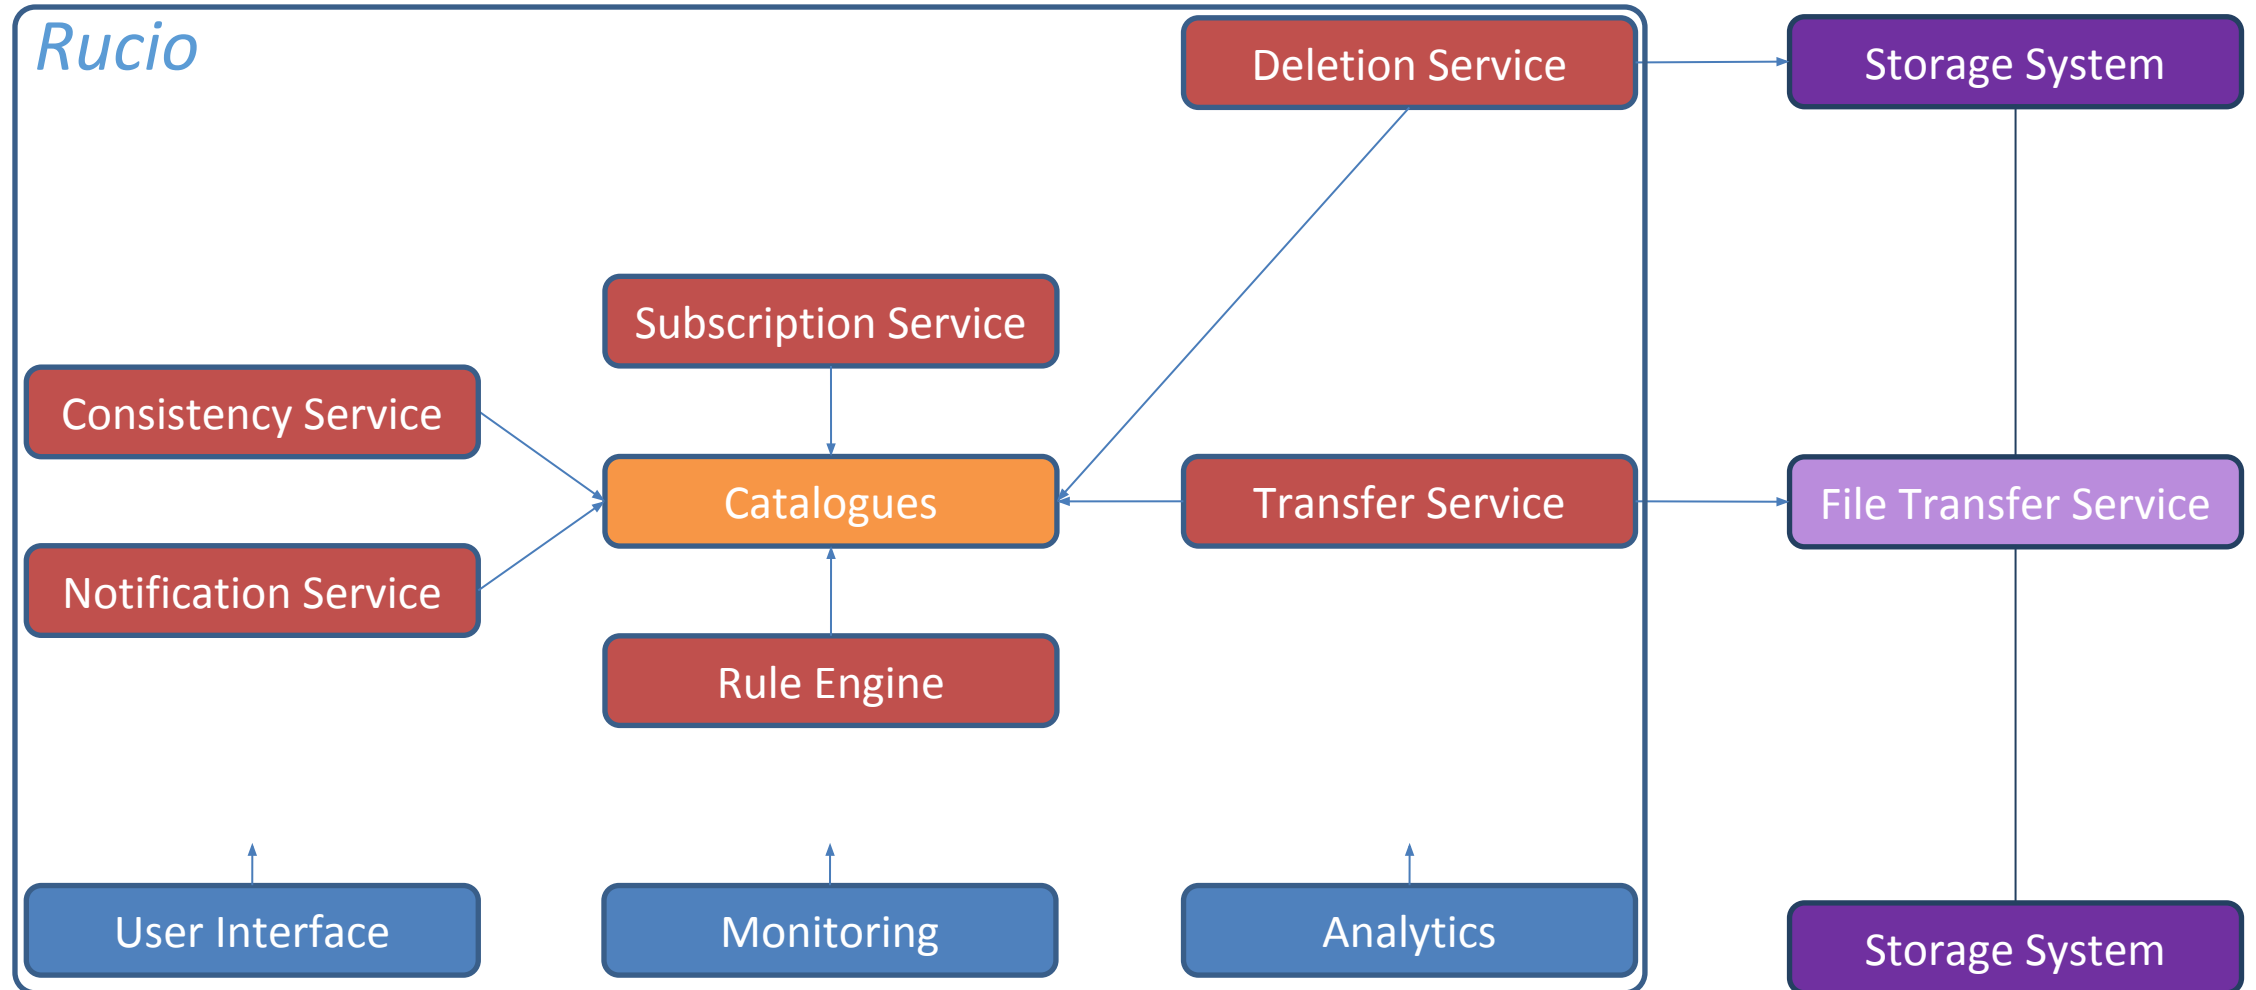
Throughput throttled on ISO/OSI Layer 7 (Application)
Limited storage capacity leads to dynamic data management
Constraint logic engine evaluates and processes data placement rules

Unscheduled throughput

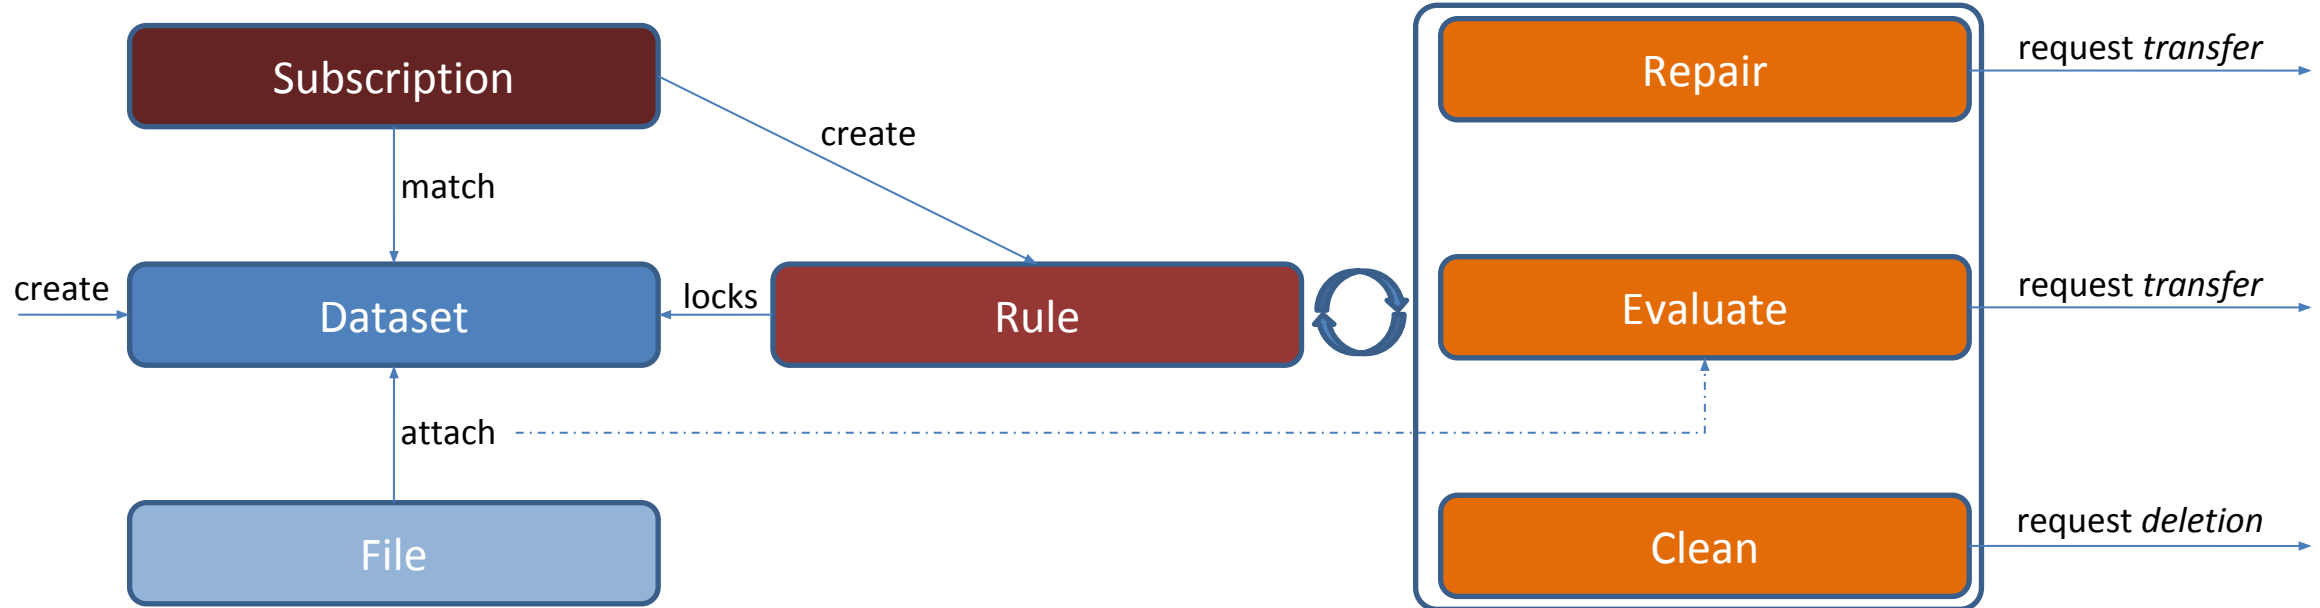


Jobs, CLIs, ... also access data – requires interaction with Rucio
Unpredictable interaction rate – potentially disruptive

High Level Architecture



Rule engine



Implicit data management – specify what, not how – use subscriptions & replication rules



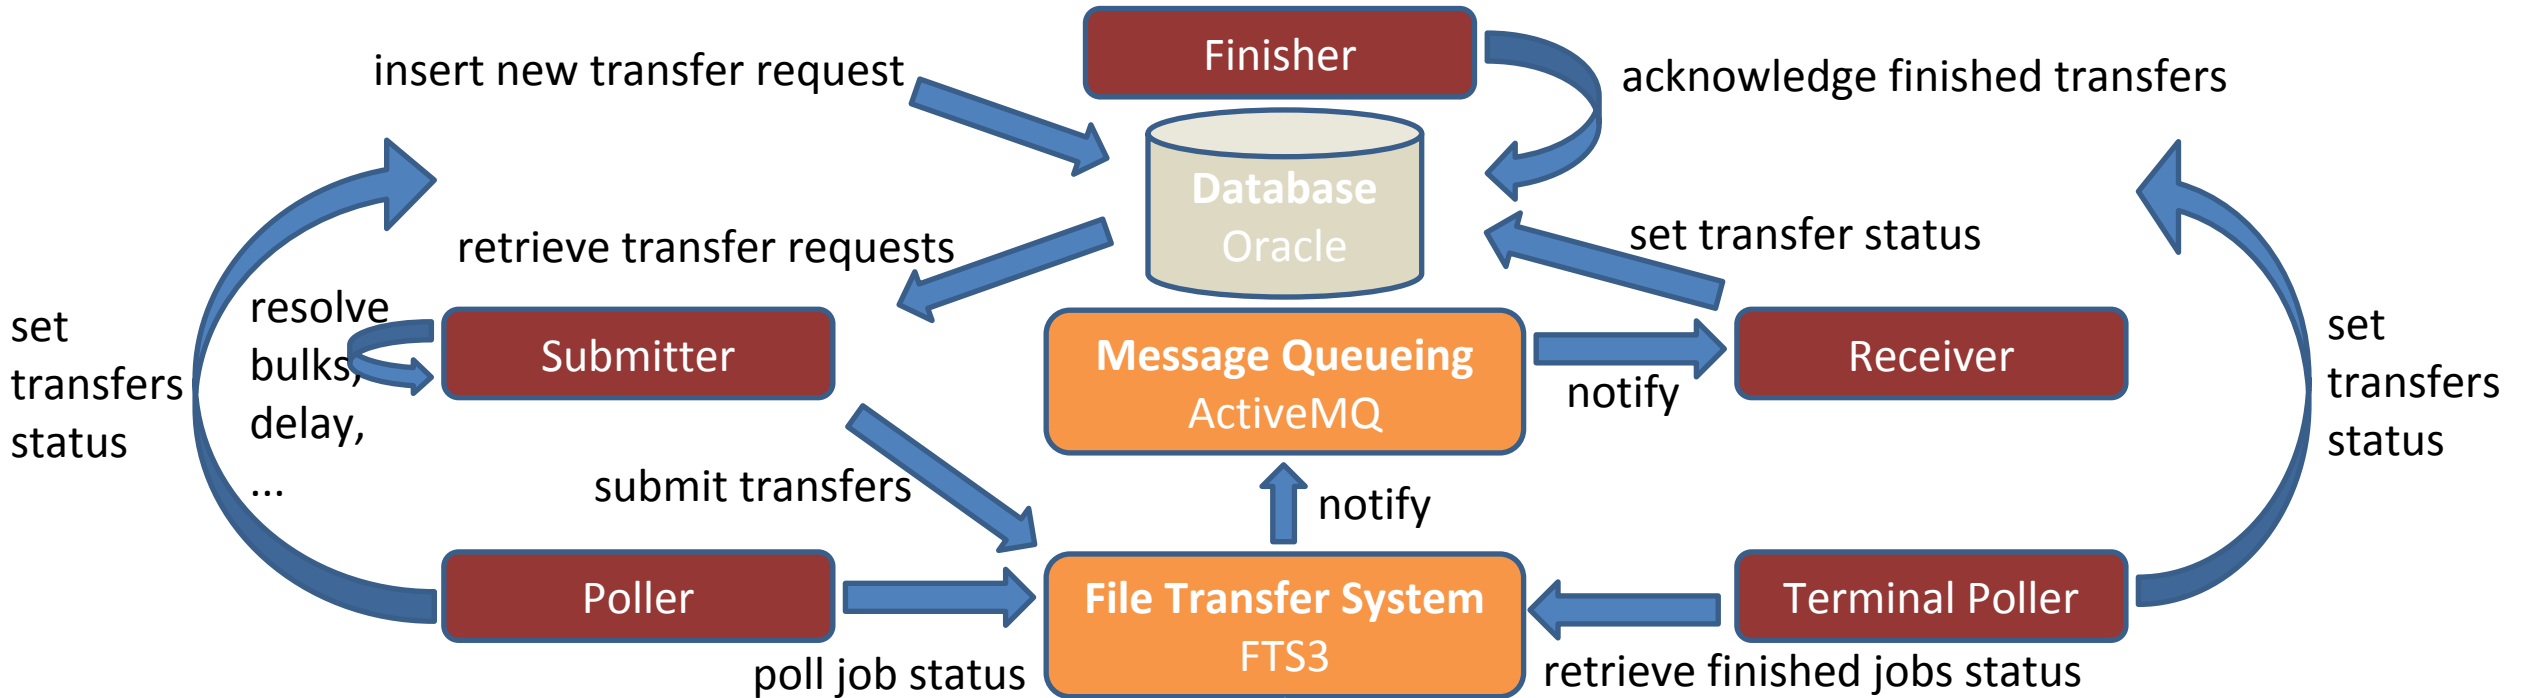
e.g., {mc15_13TeV*, copies: 2, type=tape} or {user.mlassnig.01, copies: 1, country=at}

Replication rules express interest in data – protect replicas from deletion

Smart potential destination selection – volume, distance, quota, grouping

Files without rules are eligible for deletion – kept as cache until space is needed

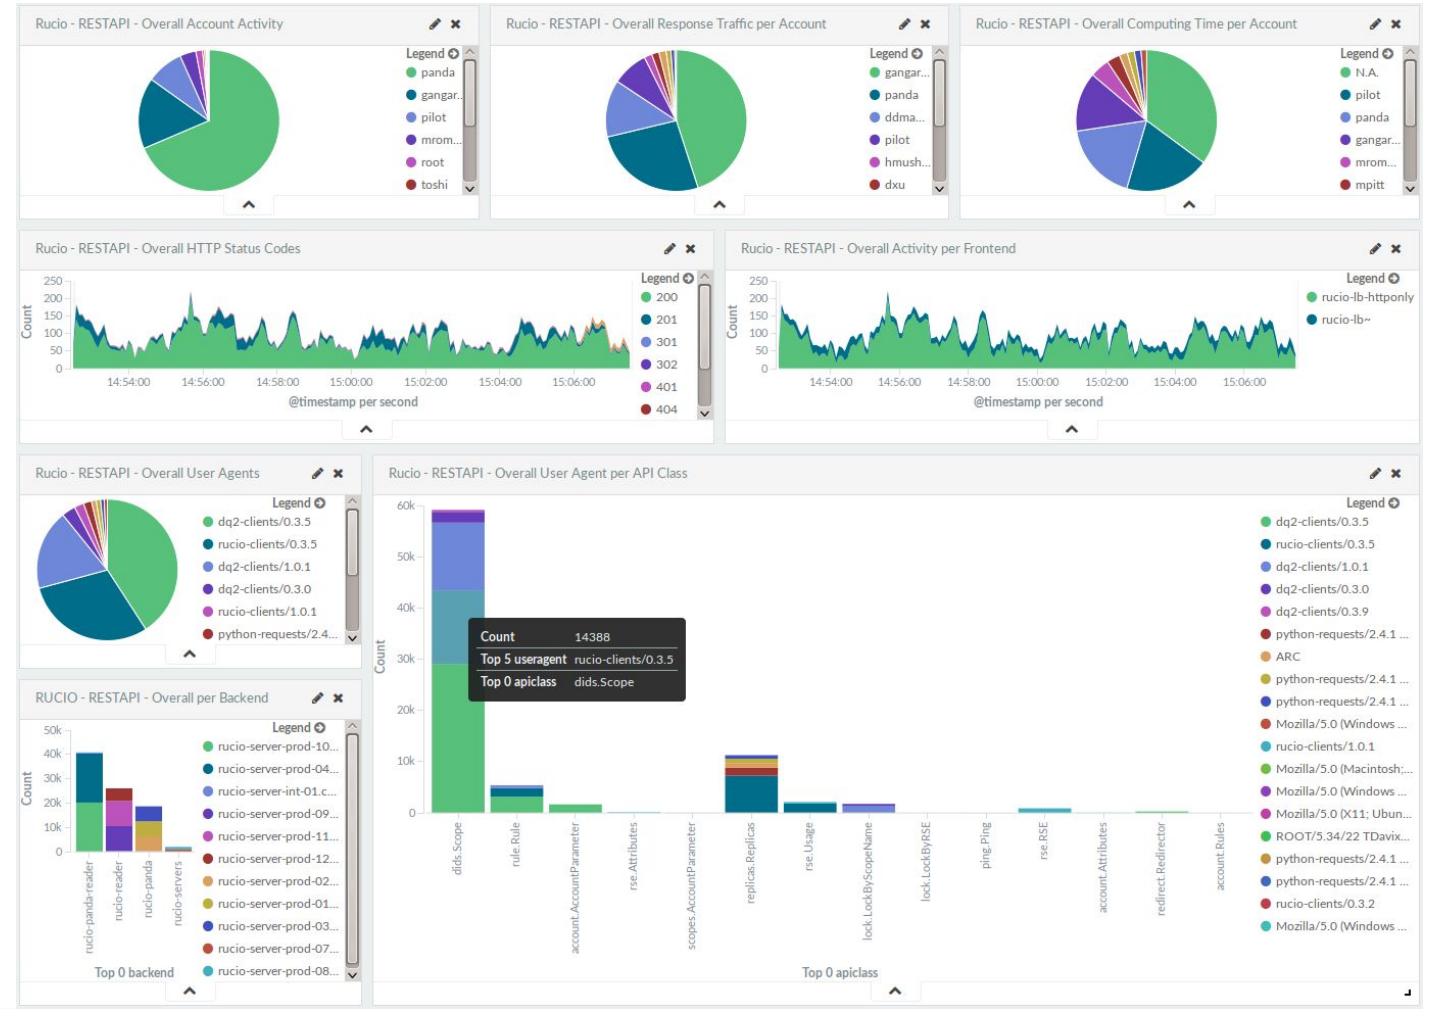
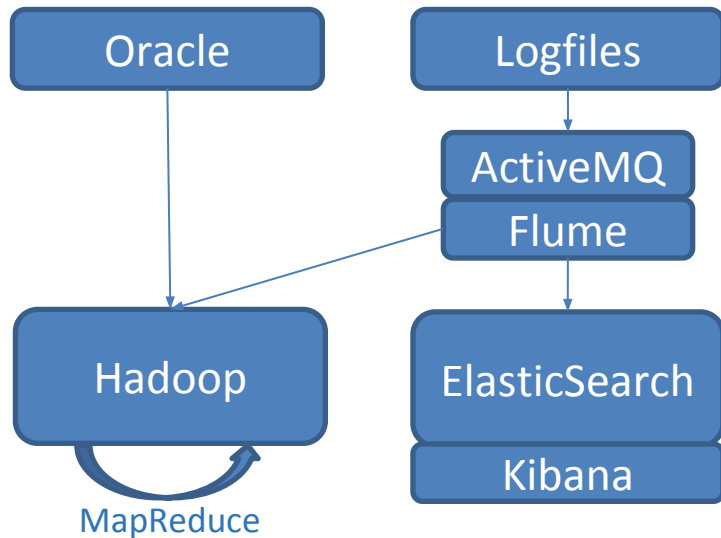
Scheduled transfers


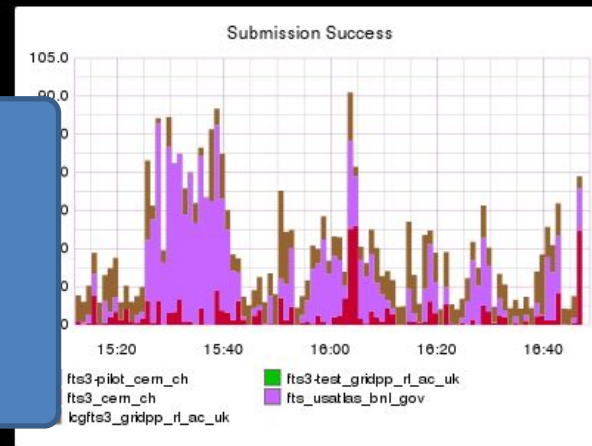
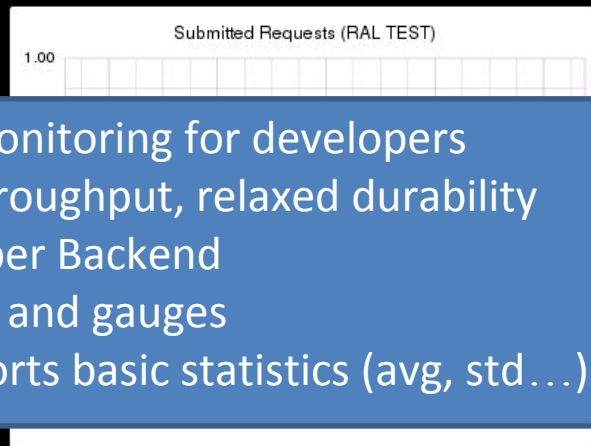
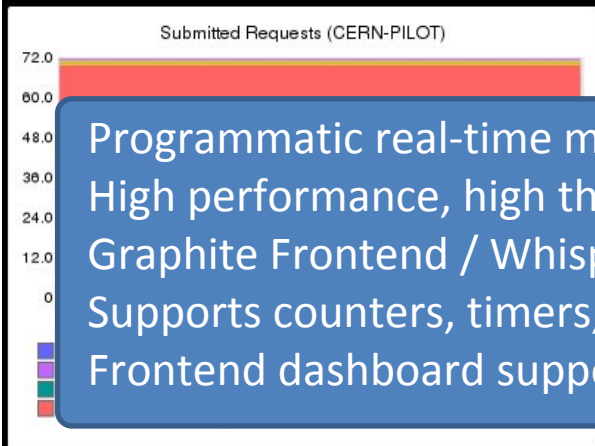


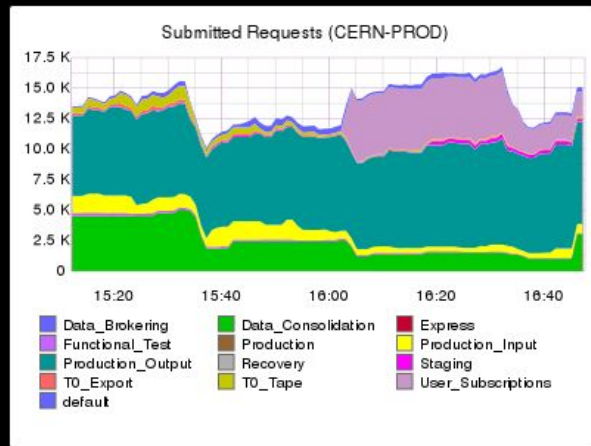
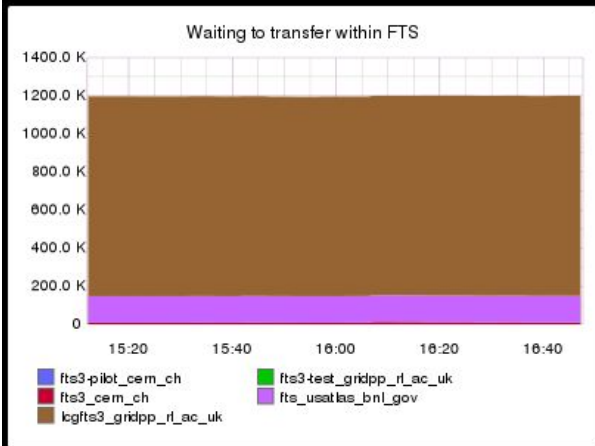


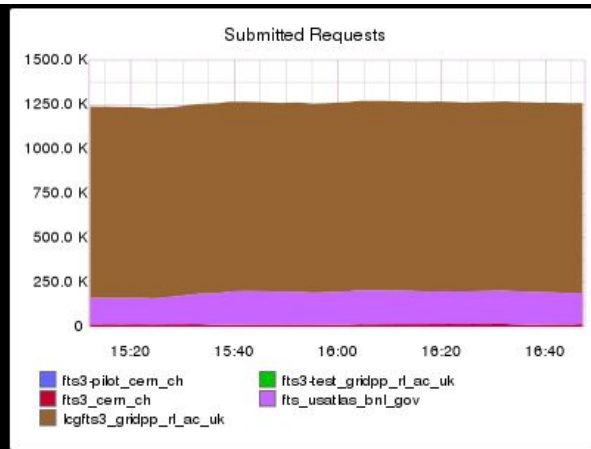



Data sources

From what can we learn?

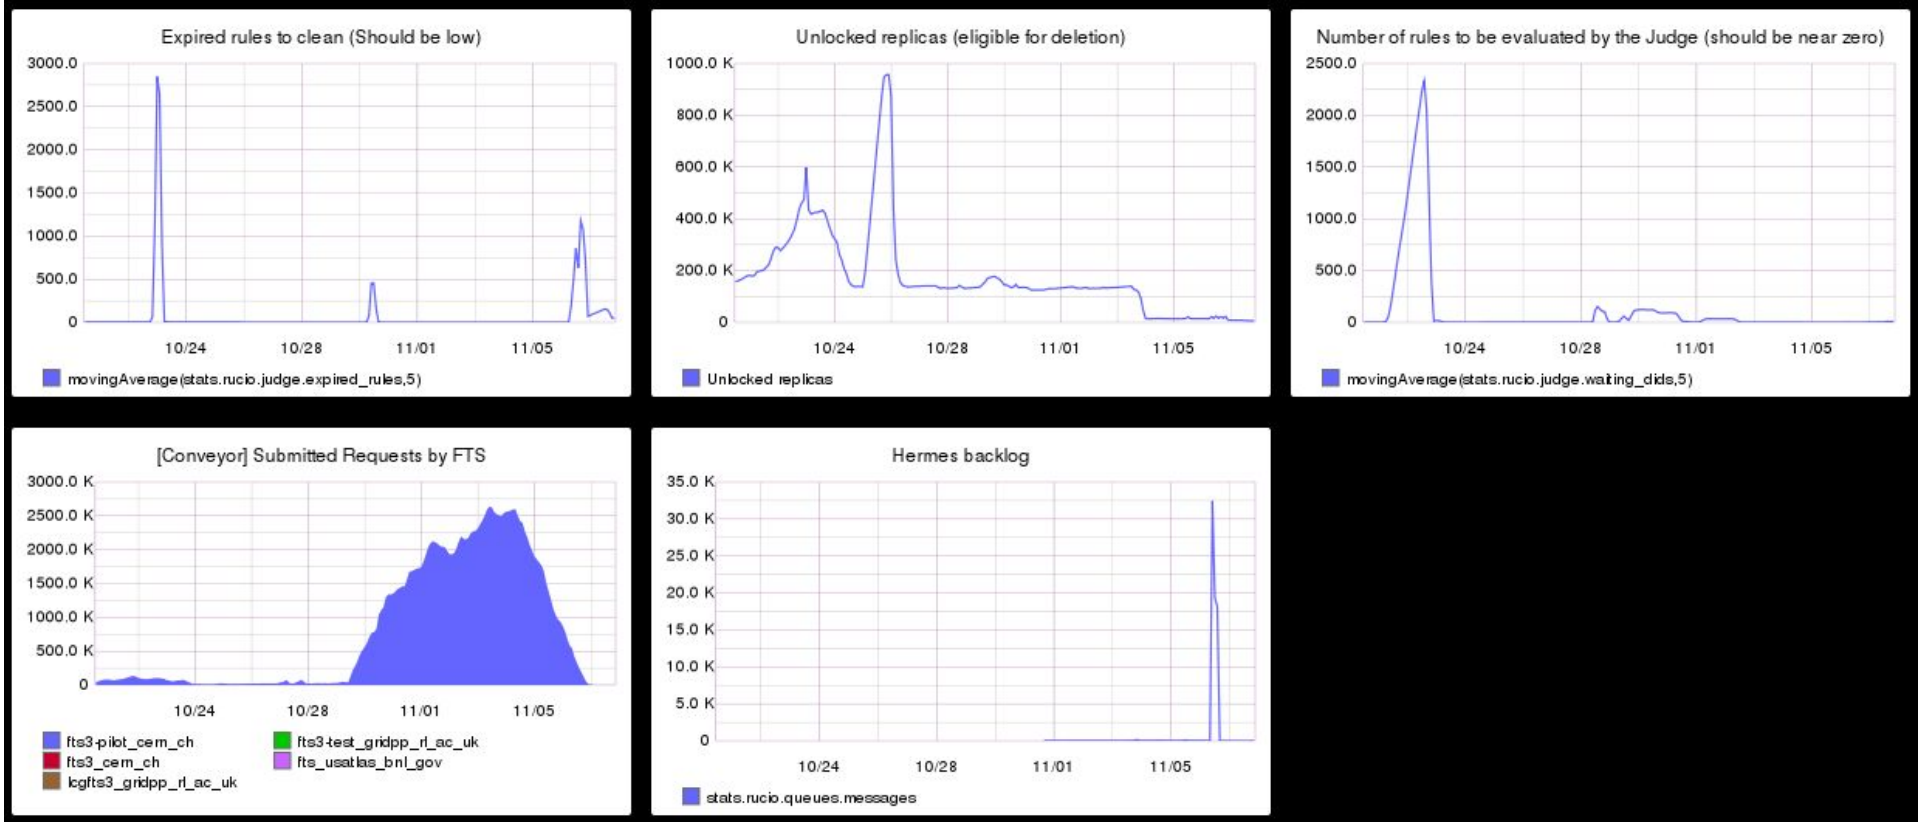
We differentiate
offline sources
online sources





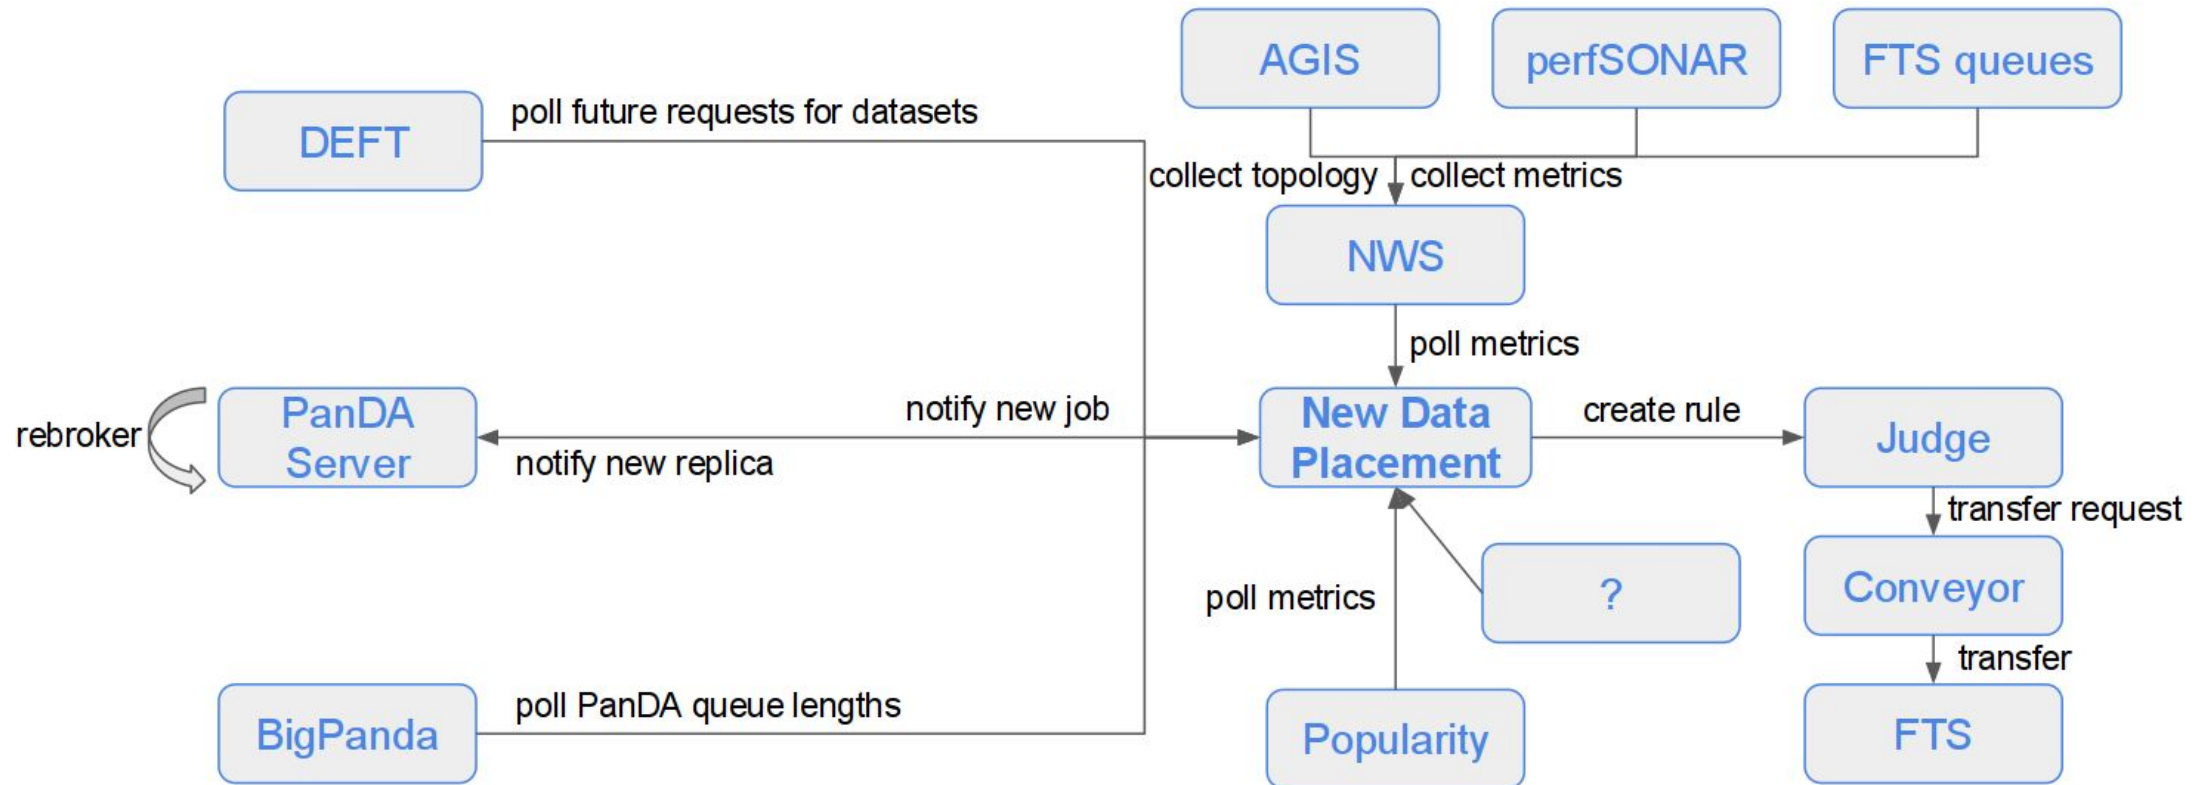
Programmatic real-time monitoring for developers
 High performance, high throughput, relaxed durability
 Graphite Frontend / Whisper Backend
 Supports counters, timers, and gauges
 Frontend dashboard supports basic statistics (avg, std...)

Anomaly detection




Things don't always go well — we need shifter interventions
Error-prone and costly — write up of post-mortem analysis for computing management

Dynamic data placement



Studied well in literature — Has optimal solution if full knowledge is available (topology & workload)
We have proposed a new dynamic data placement system — currently prototyping implementation



www.cern.ch