


DUNE Rucio Deployment and Issues

Steven Timm

For DUNE Data Management Group

Rucio Workshop, Sep 2021

Deep Underground Neutrino Experiment



DUNE Rucio Update

- As of the previous Rucio community workshop, March 2020, just before lockdown, DUNE was in the following situation:
 - 14 PB under Rucio management, 1190737 DID's, 2655705 replicas
 - 15 RSE's commissioned
- Much of the above data from ProtoDUNE @ CERN
- Electronics and coldbox testing now set to turn on @ CERN again this month, with ProtoDUNE beam data approx. 1 year from now.
- By the time ProtoDUNE II (a.k.a. “horizontal drift”) takes beam it is our goal to be in a native all-Rucio data movement system.
- Have added lots of monitoring thanks to students working on project.

Monitoring: Transfers


- [Rucio Kibana Monitoring](#). Shows **queued**, **failed**, **submitted**, **done**.

[rucio] Total Transfers


11,770
Total Transfers

82.082TB
Total Transferred Bytes

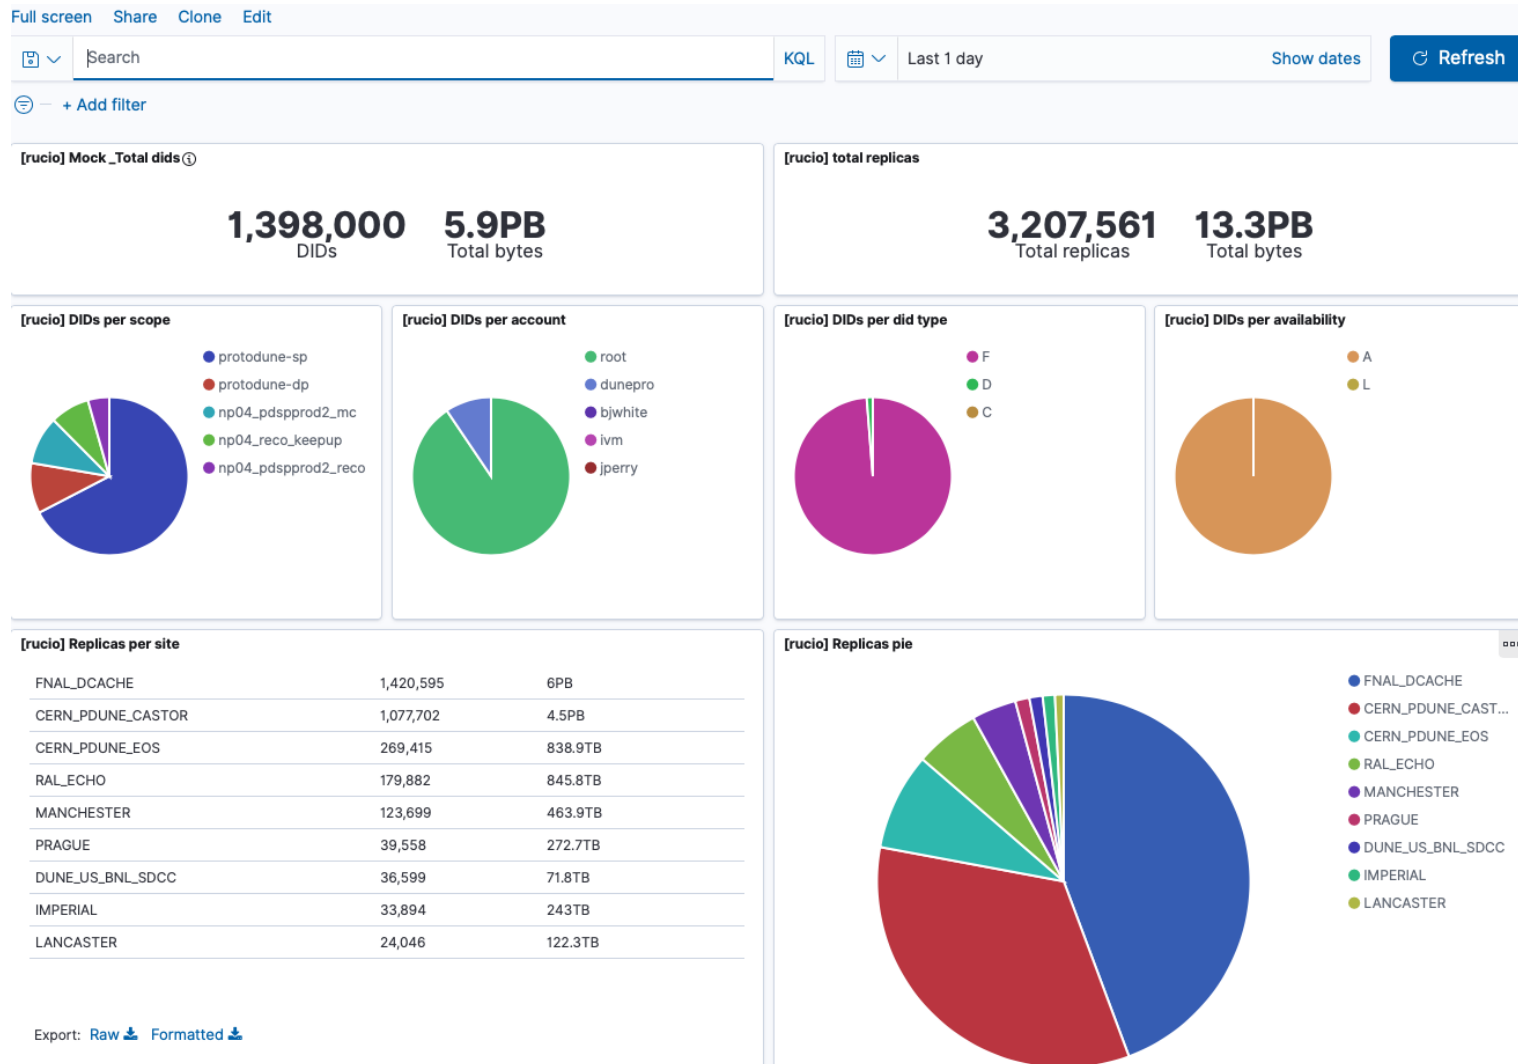
[rucio] Transfer Efficiency



[rucio] Transfers by Events




Rucio total monitoring




Rucio country/scope/container monitoring


Rucio Usage by Country



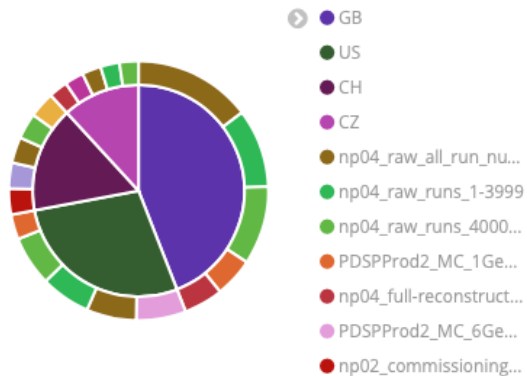
Rucio Usage by Scopes




Rucio Usage by Containers




Rucio Country by Containers



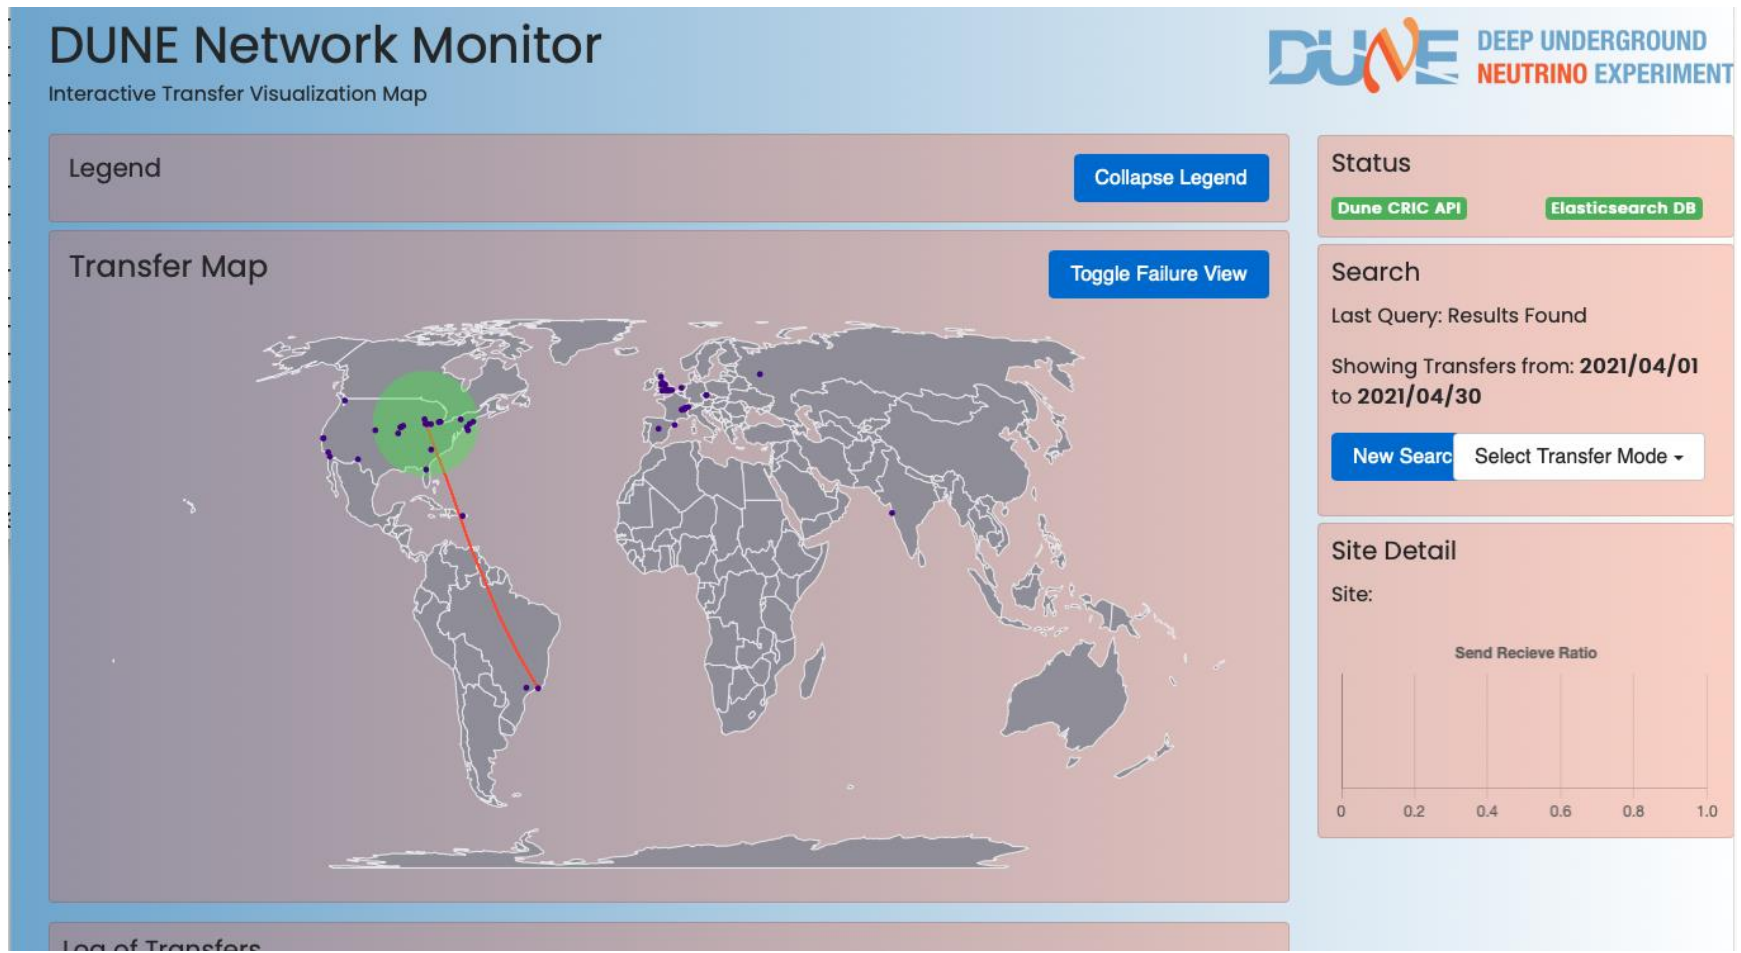
Rucio Containers by Scope



Rucio Scopes by Country



Monitoring added:



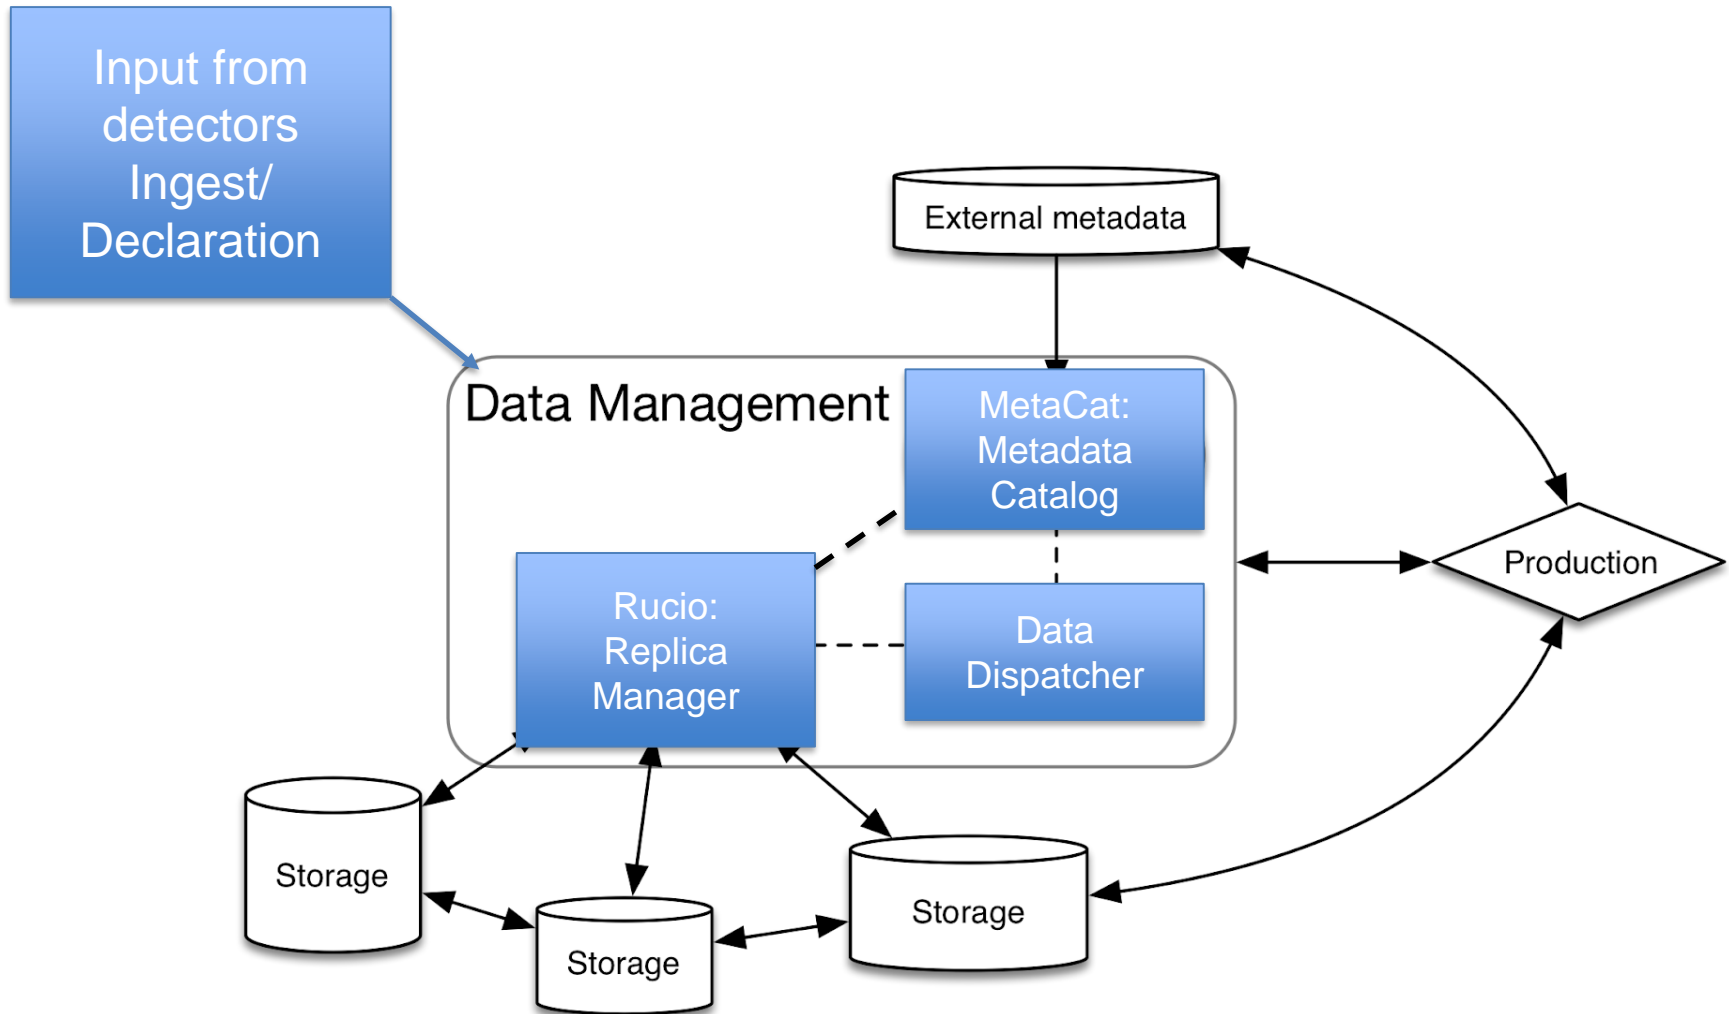
Challenges

- Fermilab runs OKD (Open Kubernetes Distribution) containerized platform
 - Open Source upstream product of RedHat OpenShift
 - Not the same as vanilla Kubernetes, Helm charts have to be modified.
- Made OKD-based containerized version be production early in 2021.
 - Version 1.22.x not stable in OKD. Lots of out of memory errors.
 - (10 file declares or a bunch of rucio list-parent-did enough to knock it over)
 - One nasty mode it falls over and then fails to restart
 - Hoping that 1.26.x does better.
 - If not, will request conventional non-containerized deployment.

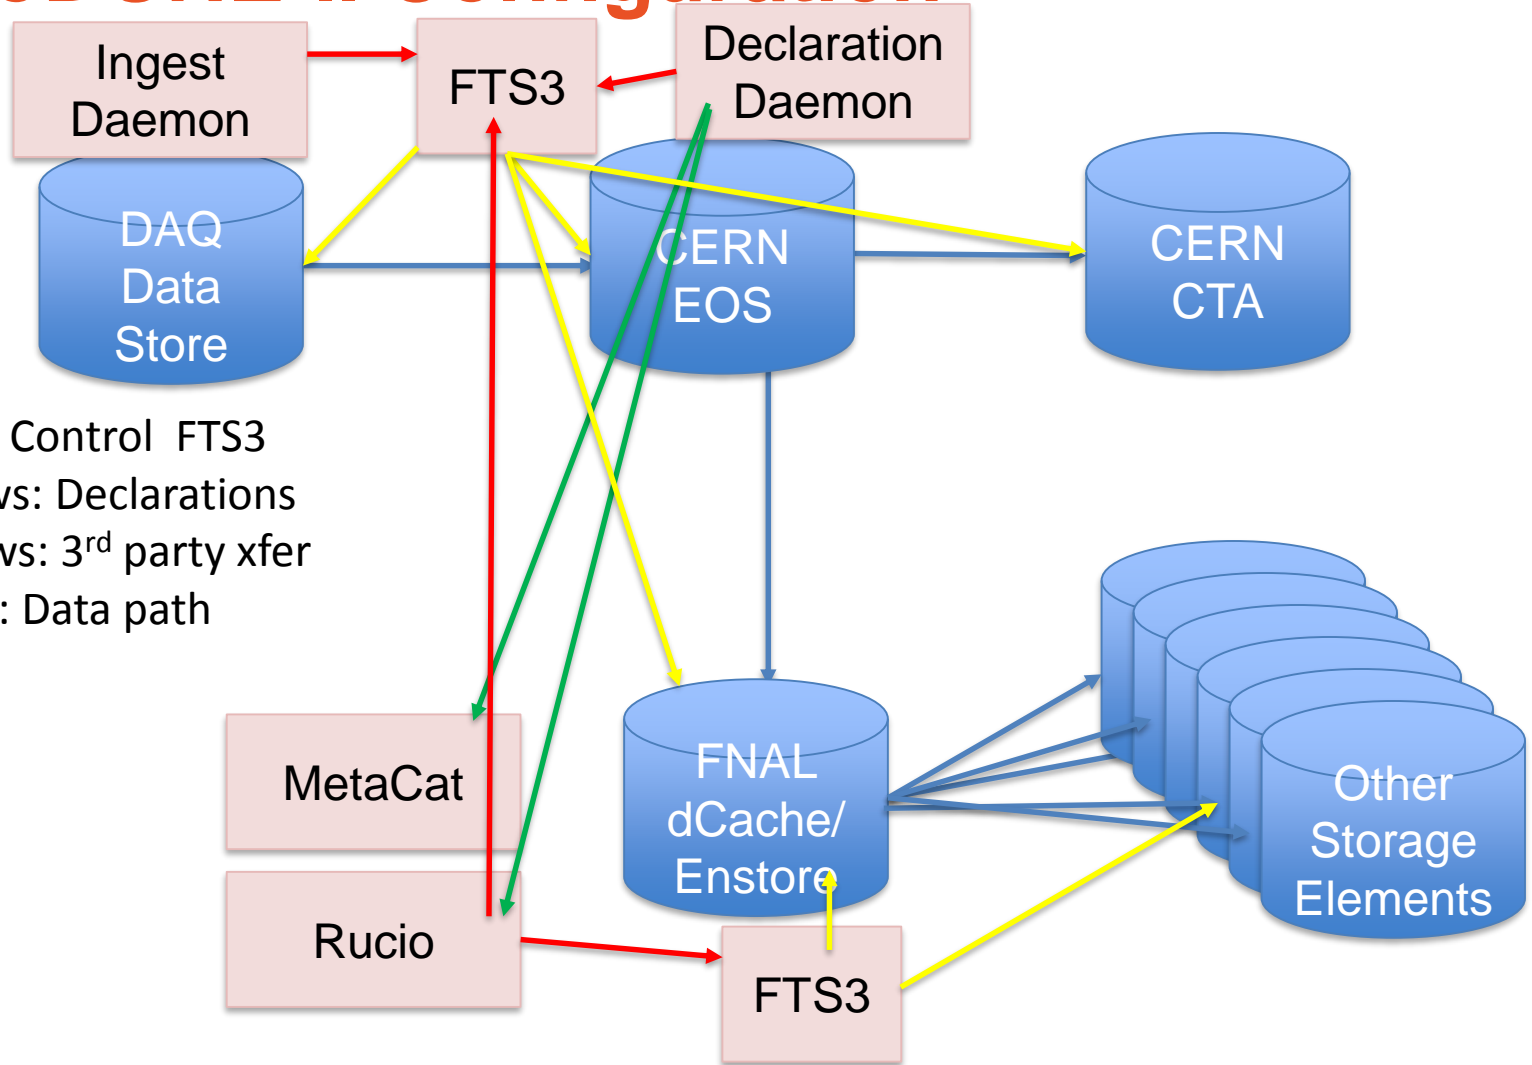
Features

- In progress:
 - Mapping to non-deterministic paths based on metadata
 - Lightweight Rucio client
 - Require file exist in external “MetaCat” metadata server before being declared to Rucio
- New: Quality of Service:
 - On combination disk/tape storage Element (Fermilab public dCache)
 - Need way to make a Rucio rule and have the file staged from tape to disk, or vice versa
 - Have heard this is being worked on for other experiments. Status?

DUNE Data Management System



ProtoDUNE-II Configuration



Red arrows: Control FTS3
Green arrows: Declarations
Yellow arrows: 3rd party xfer
Blue arrows: Data path

Status

- Getting all of our existing files to be known by Rucio is now on our critical path to deployment.
- Once that is done we still have a Data Dispatcher to write to be able to completely get rid of the SAM system at Fermilab.
- And a Data Ingest Daemon and a Data Declaration Daemon to run at the remote sites and initiate the Rucio / FTS3 transfers.
- And then do it all again with JWT Tokens.