# Experience with Rucio in the wider HEP community

Martin Barisits, Thomas Beermann, David Cameron, James Alexander Clark, Riccardo Di Maria, Gabriele Gaetano Fronzé, Ian Johnson, Mario Lassnig, Cédric Serfon, and Eric W Vaandering

RUCIO
SCIENTIFIC DATA MANAGEMENT

# Introduction

- Rucio provides a mature and modular scientific data management federation
  - **Seamless integration** of **scientific and commercial** storage and their network systems
  - Data is stored in **global single namespace** and can contain **any potential payload**
  - Facilities can be **distributed at multiple locations** belonging to **different administrative domains**
  - Designed with **more than a decade of operational experience** in very large-scale data management
- Rucio is free and open-source software licenced under Apache v2.0
- Principally developed by and for the ATLAS experiment, now used by many other communities
  - Some of them highlighted in this talk
- Community-driven development process works well
  - Software by- and for the community

# Rucio in the ATLAS EXPERIMENT

- Rucio has been working extremely well in 2020
  - 530+ Petabytes over 1+ billion files
  - Stable overall throughput in access, transfer, deletion
  - Central operations team has been busy
    - Data lifecycle management, space management
    - Commissioning CERN Tape Archive & Data Carousel
  - Participation in many R&D projects
    - TPC, QoS, AAI, Caching, NOTED, MAS, Popularity
    - Virtual Placement, Google Cloud, Amazon Cloud
- New deployment on Kubernetes
  - Rucio packaged via Docker and Helm, and deployed via Flux
  - Multiple clusters for production, plus integration
  - Fully monitored with Prometheus and Timber (ELK)
  - Request balancing with HAProxy and Loadbalancer Service
  - Workload balancing based on heartbeats

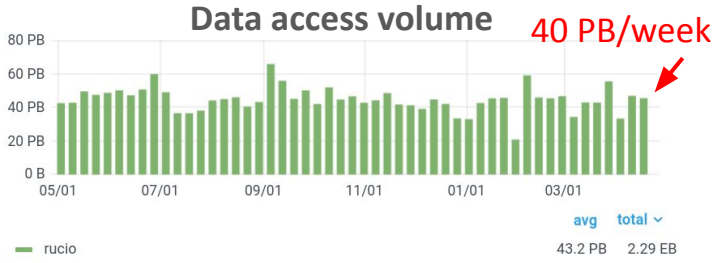


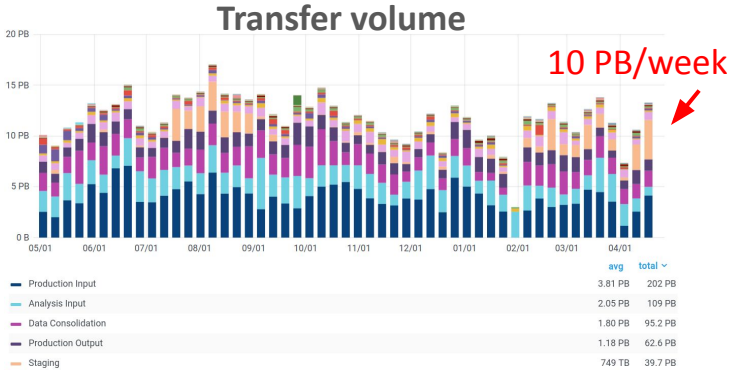

- Belle II recently moved to Rucio
- A lot of new developments done :
  - Integration with Dirac (see talk about Dirac/Rucio integration)
  - New features like chained subscription
- Since the migration, Rucio is running stable. The instance is the biggest one used PostGreSQL with 100M files
- Further detail in plenary talk

See Cedric's CHEP plenary talk from Tuesdayç

See Ruslan's CHEP talk just before this one :-)

Number of files transferred daily since the migration to Rucio

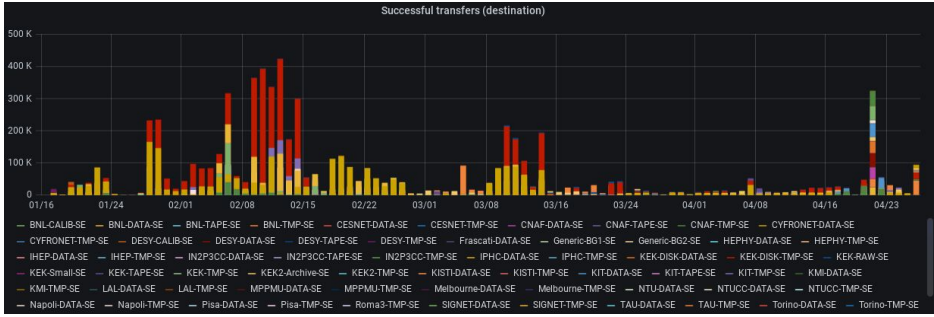


Evolution of data volume since the migration to Rucio

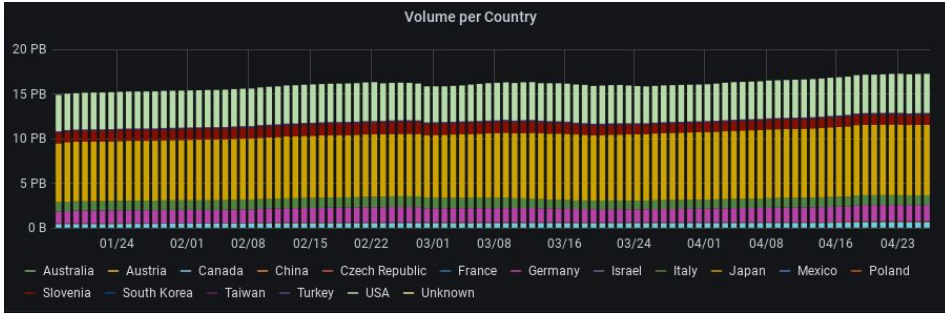

# CMS

CMS transitioned to Rucio in late 2020. Seamless transition with no downtime from PhEDEx to Rucio

Remaining PhEDEx movement is link testing, now transitioned

Still to work on

- File access traces
- Automated cleanup/repair
- Explore new functionality in workflow system enabled by Rucio

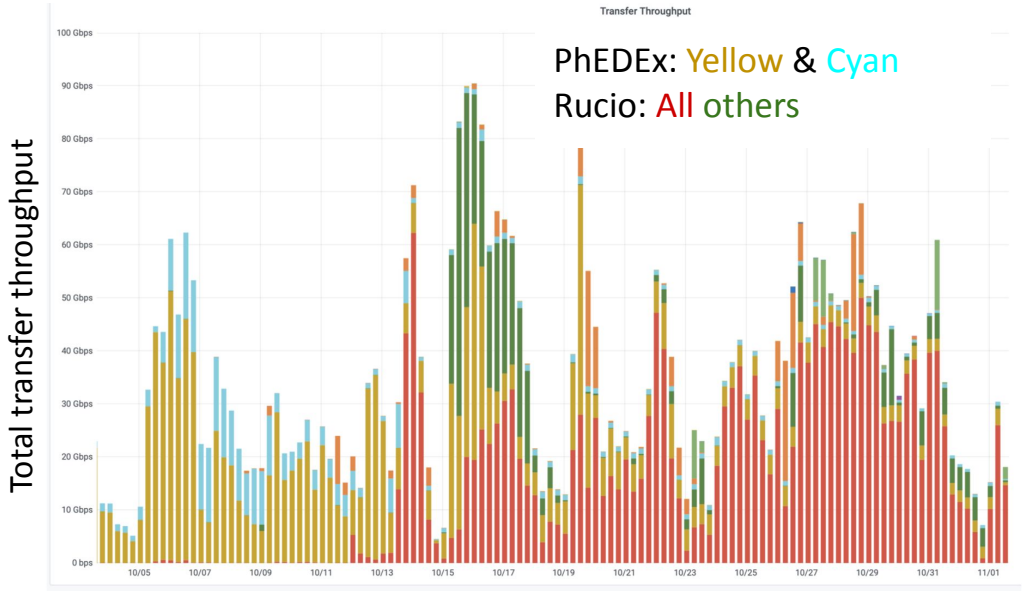

# ESCAPE

See Riccardo's CHEP talk from Wednesday

- Project goals:
  prototype an infrastructure adapted to exabyte-scale needs of large science projects; ensure sciences drive the development of EOSC; address FAIR data management principles.
- ESCAPE sciences at different phases of evolution:
  building up or consolidating computing model - needs of science projects drive services requirements;
  leveraging the existing expertise in WLCG on data organisation, management and access.
- Data Lake 24-hour Full Dress Rehearsal (workshop):
  24/7 availability to ESCAPE users even though not-production;
  9 sciences embarked FDR as specific testing-focused time-window → now, towards **Data Analysis Challenge 2021**.

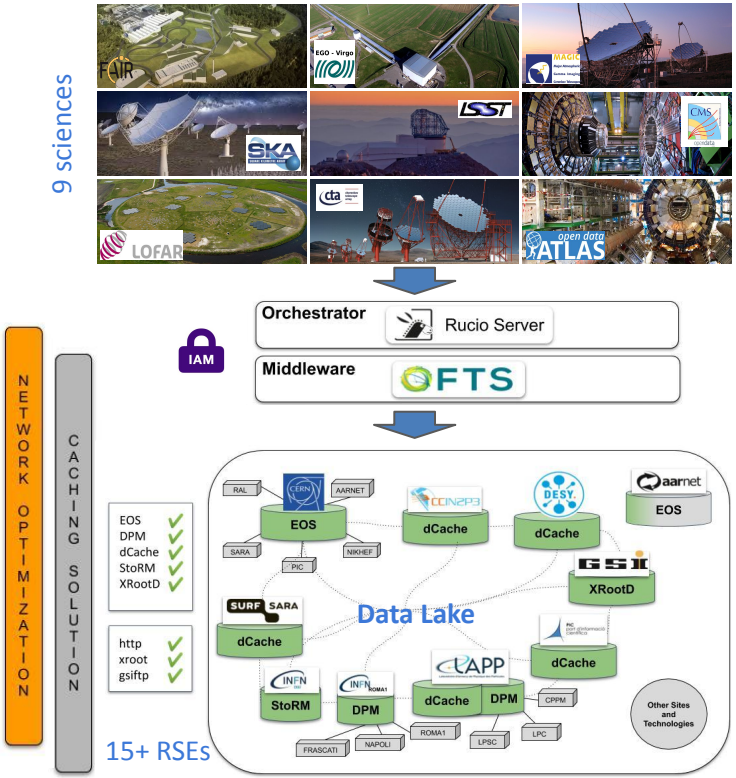

# Lightweight Dark Matter eXperiment

- LDMX is a small fixed-target dark matter experiment
- Currently running a pilot distributed computing projec to do detector simulations
  - Using Rucio for data management and ARC for workflow management
- Rucio used as replica catalog and metadata store
  - Taking advantage of Rucio's generic metadata system
- ~200TB and 1M files cataloged so far

See Lene Kristian's CHEP talk from this session

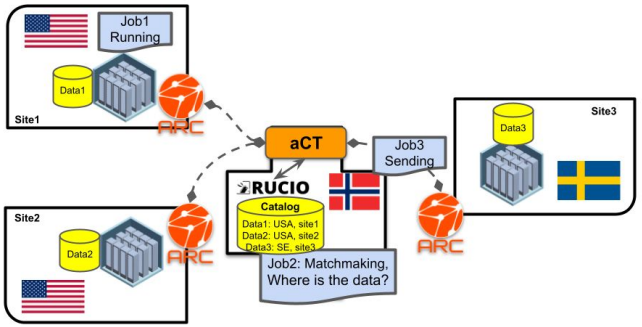


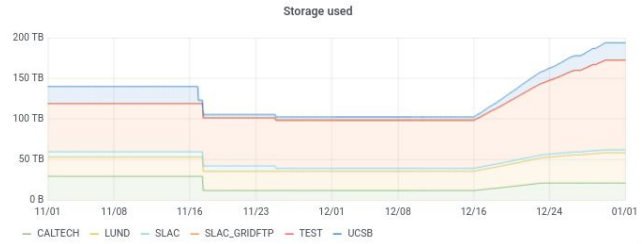

- LIGO-Virgo archival data management:
  - Replication of aggregated instrument data from instrument sites to central archive & CVMFS origin(s) for offline analysis
- Phasing-out legacy bulk data management system (LDR), in favour of Rucio
- Stable k8s deployment on [Nautilus/PRP]
- Rucio now managing all new offline datasets
- Future plans:
  - Complete LDR ⟶ Rucio transition
  - User-facing native Rucio data discovery based on Rucio's inequality filters (see next slide)

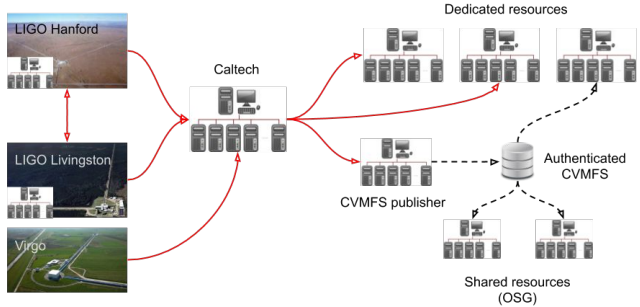


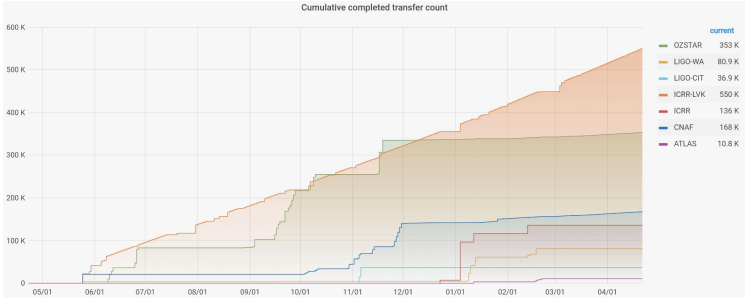


1 year of stable operation

- LIGO-Virgo's architecture required some specific development:
  - per-RSE checksum algorithms availability (merged)
  - metadata based inequality filters (in progress)
- Management of derived data & analysis output via Rucio is being evaluated
- The Rucio FUSE-POSIX mount utility can aggregate multiple Rucio servers at client level
- Analysis jobs outputs and byproducts might be handled by satellite Rucio deployments, separated from the bulk data distribution one

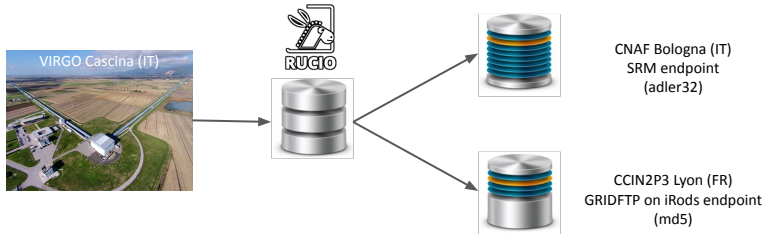


Example of per-RSE checksumming definition

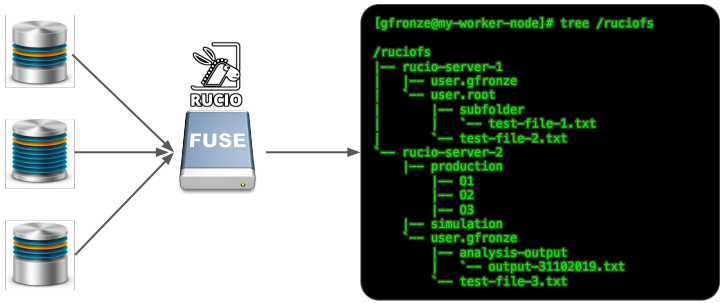


POSIX-like operations with multiple RUCIO servers

# Folding @ Home

- Citizen science project to research protein folding characteristics
- F@H currently manages ~1PB of data across 35 servers in ~10 million files
  - They expect an order of magnitude increase by next year
  - Their data management operations are completely manual: `ssh, rsync, s3cmd`
  - They have HPC centres (e.g., Hartree HPC in the UK) and several cloud storage
- After discussions with CERN COVID Team, they got interested in Rucio
  - Can potentially solve all their data management needs
  - We offered to host an instance at CERN for demonstration purposes
- Currently evaluating the deployment procedure and running first tests
  - Running xrootd servers with root x509 third-party-copy across two RSEs
  - Fully integrated with Rucio Folding instance @CERN and FTS @CERN
  - Demonstrated first third-party-transfers and client upload/downloads
- Exposed a lot of "grid"-cruft that makes it difficult to deploy our LHC environment tools for others

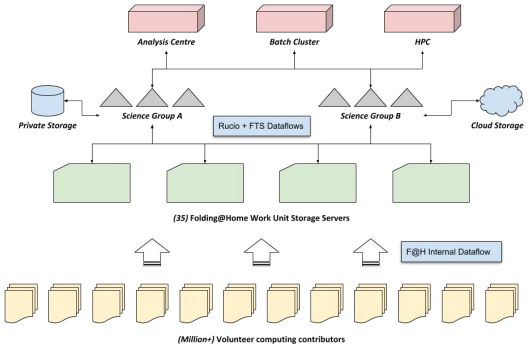

# Multi-VO Rucio

The concept - provide an easier way for small or new VOs to use Rucio without needing their own installation:

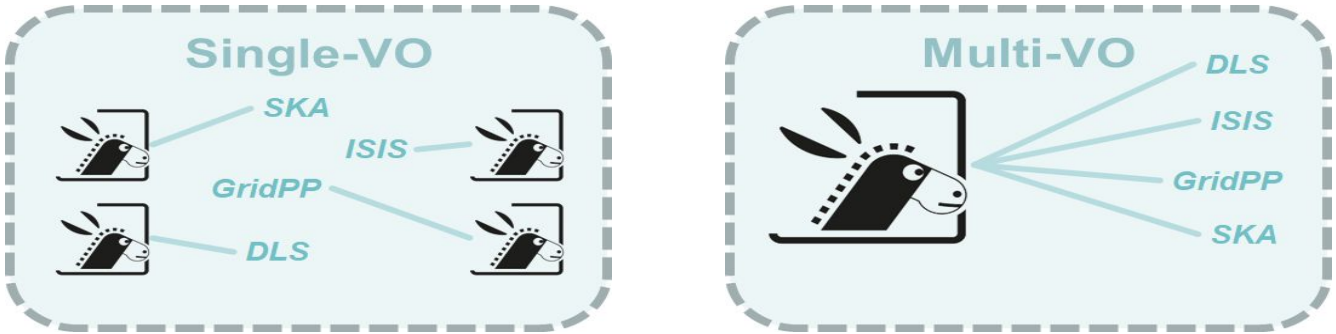


The implementation provides isolation between all VOs using the instance, and is under user trials at RAL.

# More information

**4th Rucio community workshop**
Sep 28 - Oct 01, 2021
https://indico.cern.ch/e/RUCIO2021

| | |
|---|---|
| Website | http://rucio.cern.ch |
| Documentation | https://rucio.cern.ch/documentation |
| Repository | https://github.com/rucio/ |
| Images | https://hub.docker.com/r/rucio/ |
| Online support | https://rucio.slack.com/messages/#support/ |
| Developer contact | rucio-dev@cern.ch |
| Publications | https://rucio.cern.ch/publications.html |
| Twitter | https://twitter.com/RucioData |