Scale Tests of the New DUNE Data Pipeline

Steven Timm, Fermilab,
Wenlong Yuan, Univ. of Edinburgh,
Douglas Benjamin, Brookhaven National Lab.
For the DUNE Collaboration

CHEP 2023. 11 May 2023

DUNE Far Detector Data Flow

- Beam coincidence events are extremely important, but of limited total volume - ~1 Hz beam rate (most don't trigger)
  - active online trigger in development
  - Region-of-Interest within module
  - online compression and zero-suppression being considered
- Solar neutrino triggered events
- Cosmic ray events and calibrations
- Supernova readout events
  - ~140 TB in 100 seconds per module
  - transfer out in 4 hours and process in 4 hours to point for optical follow up
- DUNE requirement - less than 30 PB/year total to permanent storage from all active FDs (~1 GB/s)
- More details in Michael Kirby talk at this same conference

Far Detector Cryostat

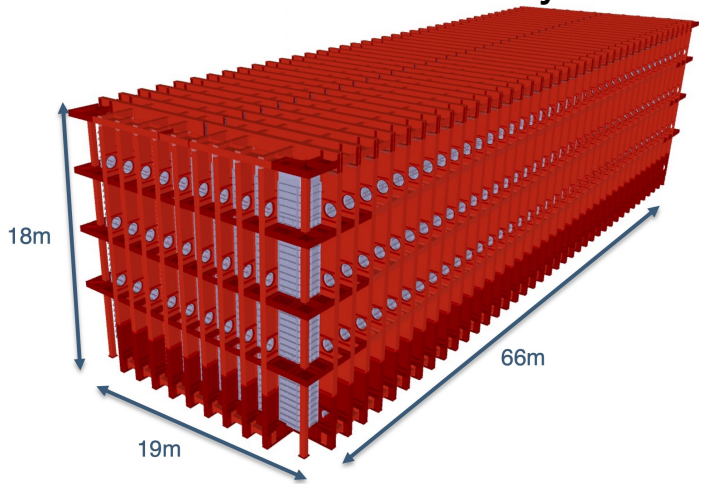

Data volumes from DUNE Computing CDR

Process	Rate/module	size/instance	size/module/year
Beam event	41/day	3.8 GB	30 TB/year
Cosmic rays	4,500/day	3.8 GB	6.2 PB/year
Supernova trigger	1/month	140 TB	1.7 PB/year
Solar neutrinos	10,000/year	≤3.8 GB	35 TB/year
Calibrations	2/year	750 TB	1.5 PB/year
Total			9.4 PB/year

CERN and Neutrino Platform currently hosting ProtoDUNE

- ProtoDUNE Single Phase and Dual Phase
  - constructed and operated 2018 - 2020
  - ProtoDUNE SP took 6 weeks of beam (~25 Hz)
  - invaluable information about performance, construction, and operations
  - 5% of size of eventual far detector

CERN EHN1

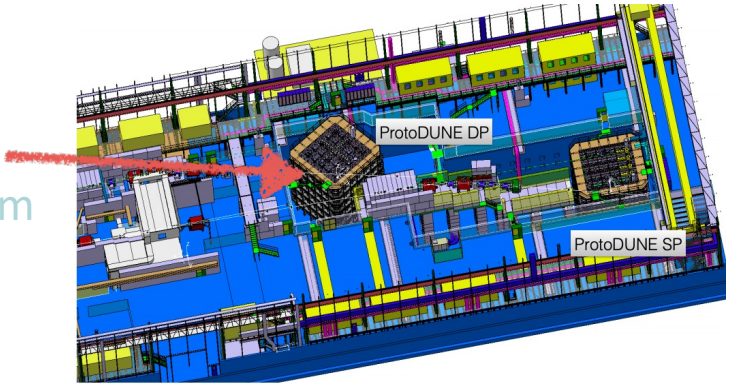

- ProtoDUNE II currently under construction
  - Horizontal Drift and Vertical Drift
  - HD anticipate filling LAr 2023Q4
  - VD assembly will proceed in parallel and operate after HD operations
  - Beam operations in 2023/2024
  - Charged particle beam data rate from ProtoDUNE (180MB/event), 4.5GB/s peak rate
  - Comparable to expected DC data rate from far detector (1GB/s)

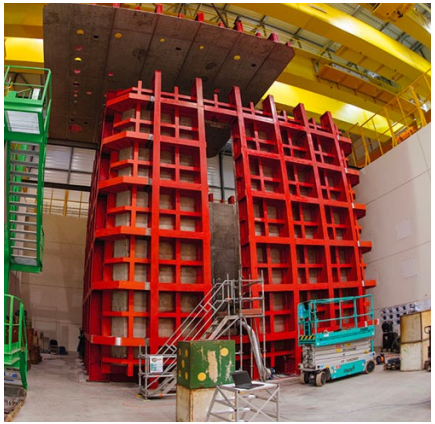

New DUNE Data Management System

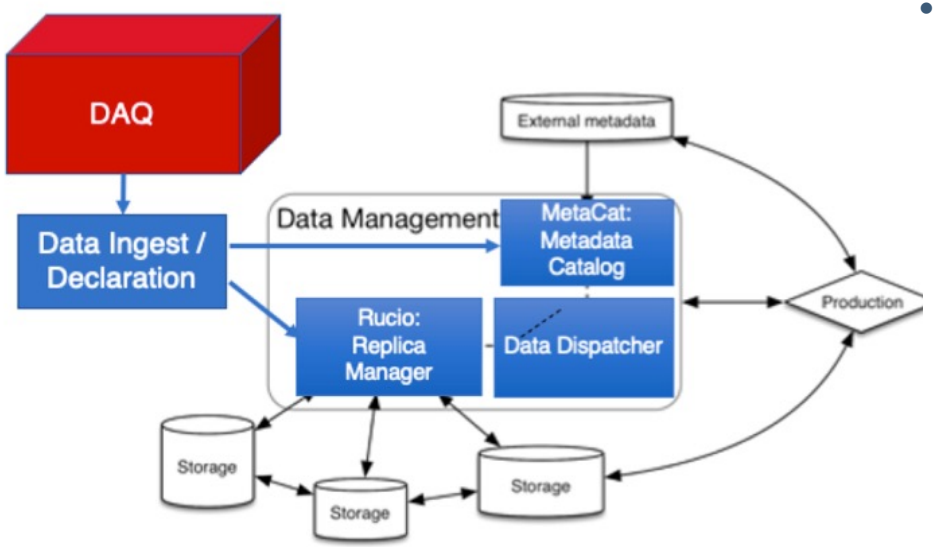

Other experiments already using these new data management tools

- Fermilab replacing legacy unified data management / workflow system
  - Rucio for file replica cataloging and management
    - Well-known tool adopted by most big experiments in HEP community
    - Uses FTS3 for file transfers
  - MetaCat—new metadata service
    - Allows for hierarchical metadata and extensible queries against other databases.
  - Data Dispatcher—simple projects and data file delivery
- In addition JustIN workflow system also developed by UK groups.

DUNE Data Challenge 4 (Summer 2022)

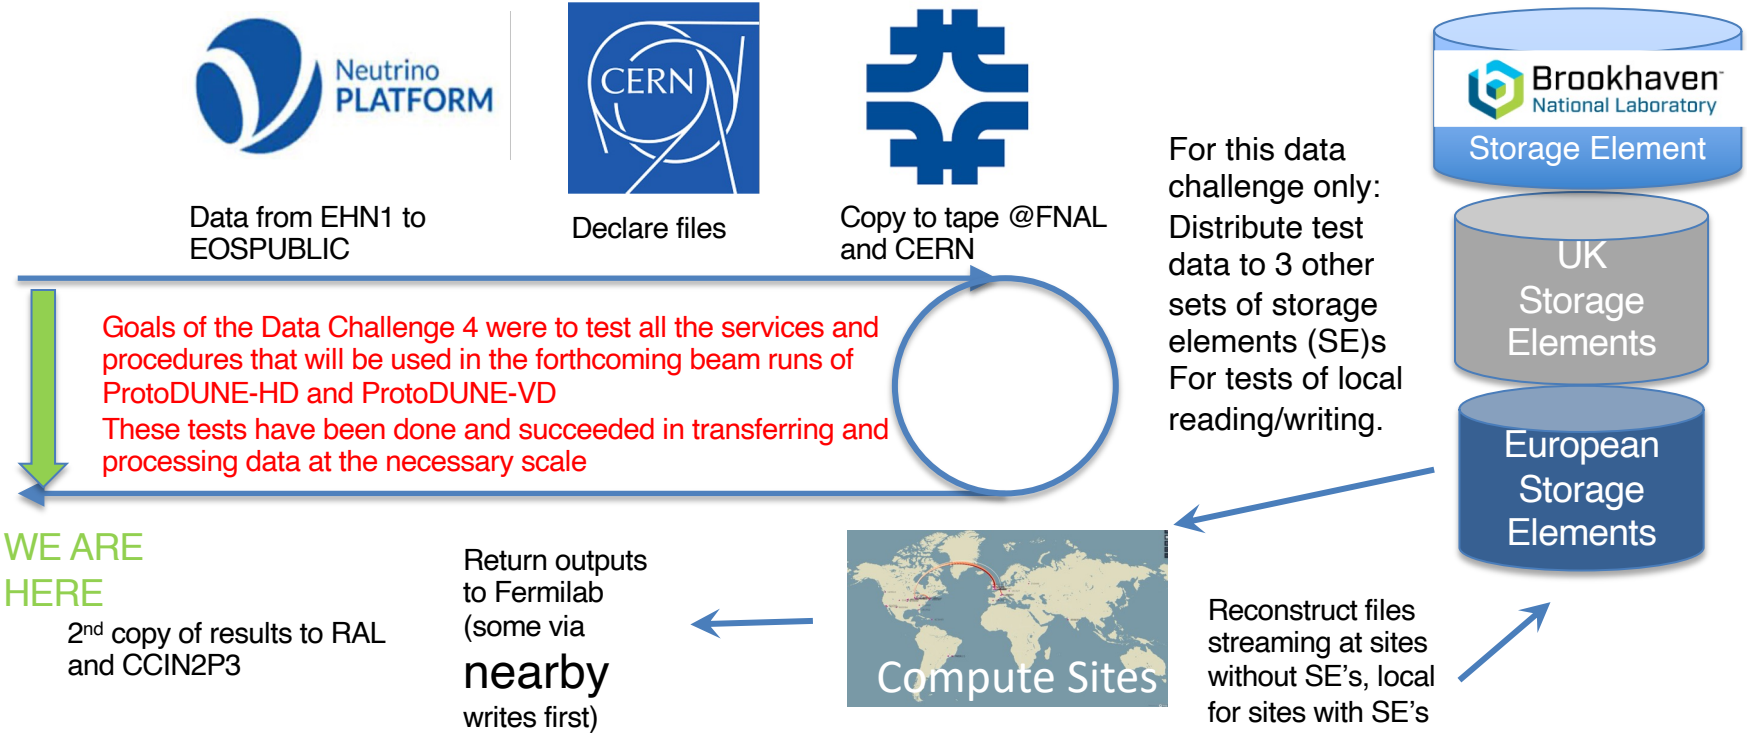

Data Pipeline Diagram

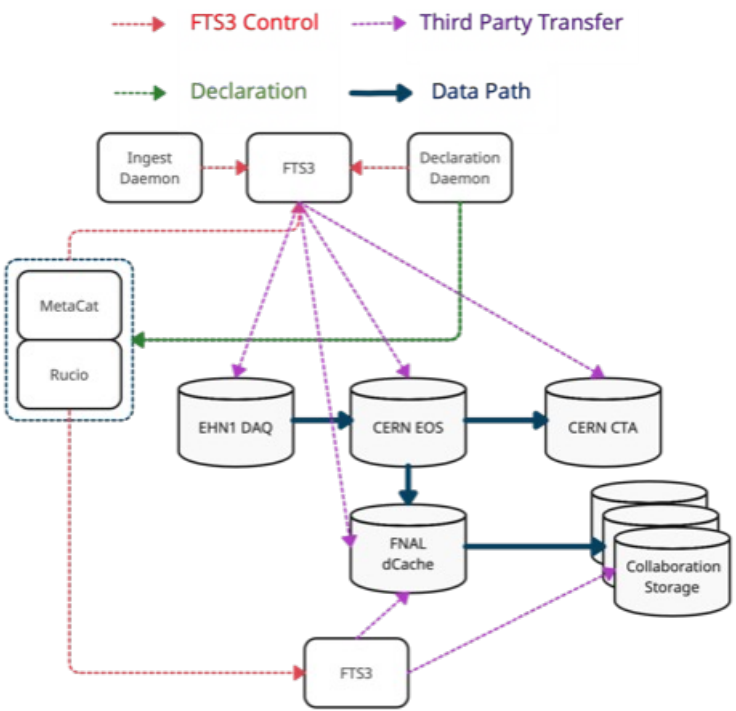

- Legacy file transfer system replaced by Ingest Daemon and Declaration Daemon
  - Ingest daemon brings files from experimental systems to dropbox
    - Can operate without connection to Fermilab.
    - All transfers done via FTS3
  - Declaration daemon declares them to MetaCat and Rucio and makes rules to get them to final destinations.
- 2 copies of raw data on tape
- 1 copy of sim/reco on tape
- 2 copies of sim/reco on disk distributed across global storage elements.

Summary of Data Challenge 4 Data Pipeline Phase

- Ran 5 days, July 11-15, Midnight UTC to Midnight UTC
- Copied a set of files from a sample area, renamed them with a timestamp and generated metadata for them.
  - 3 samples: Vertical Drift Cold Box top electronics (3GB), Vertical Drift Cold Box bottom electronics (4GB). Horizontal Drift Monte Carlo(1.4GB)
- Generated 250GB 4x an hour on 4 machines.
- Ingest daemon copied it to CERN EOS
- Declaration daemon declared it to MetaCat, Rucio and SAM
- Generated ~500TB of data total
- Near end of the challenge reached a peak of 3.6GB/s=28.8Gbit/s
- Once pipeline was done we had 2 copies of the data set, one at CERN and one at Fermilab
- For testing of locality reading, fanned it out: 1 copy split among 5 SE in the UK, 1 copy split between 3 SE in Europe, 1 full copy at BNL. So 5 copies of test data set total.

FTS3 monitoring: CERN->FNAL link

3.6 GB/s !

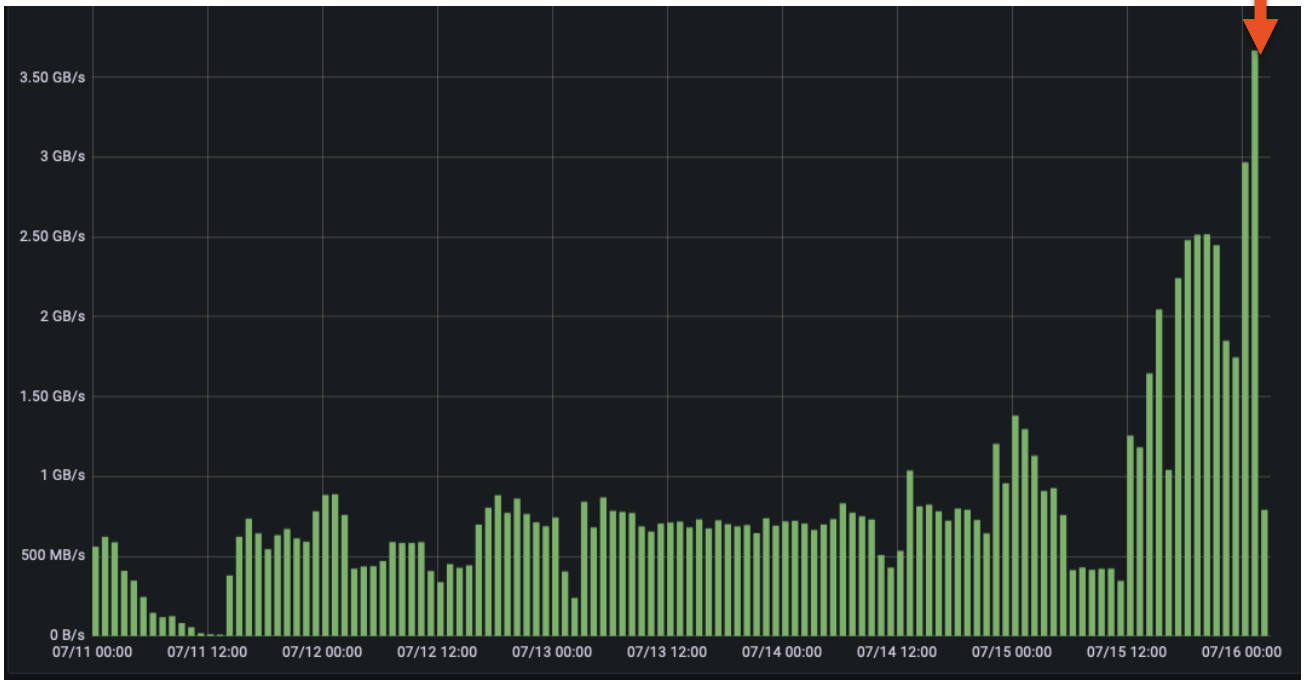

07/11 00:00

07/16 00:00

https 3rd party transfers eospublic.cern.ch -> fndca1.fnal.gov

Big improvements during run:

- Ingest Daemon ->100 threads
- Speed up metadata script
- Increase cores of declaration daemon from 2 to 8
- Increase limit of simultaneous transfers in FTS3 from 143->500

Fermilab DUNE

CERN–FNAL traffic

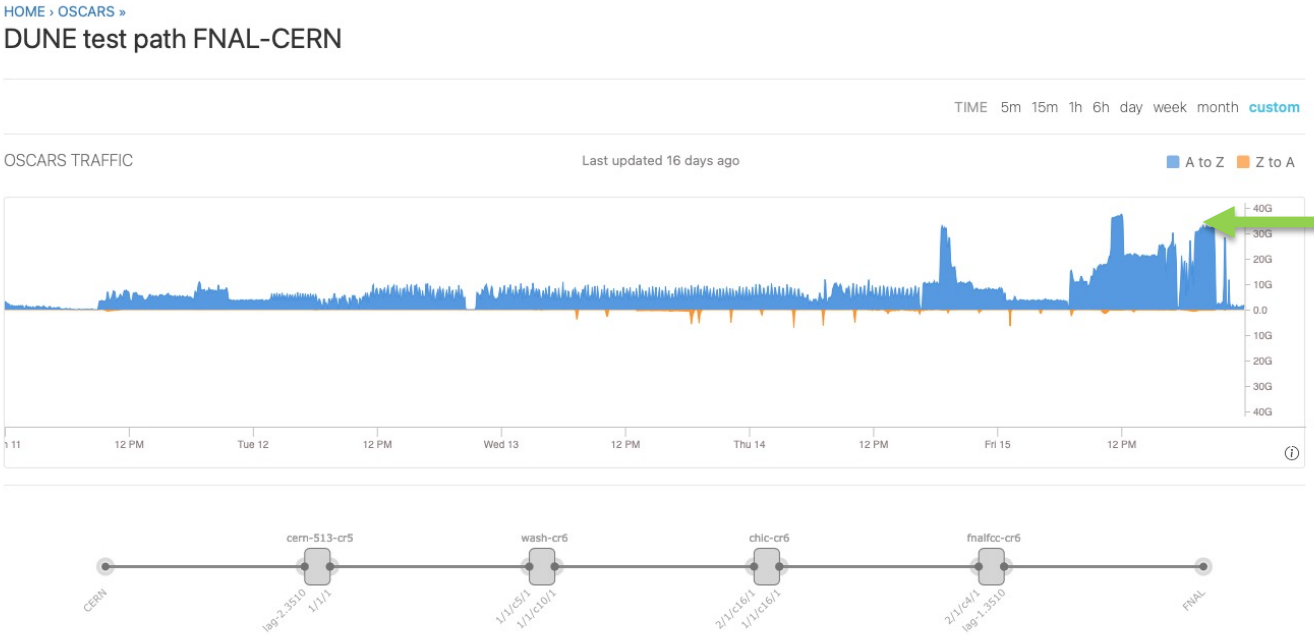

28.8GBit/s
=3.6GB/s

Dedicated OSCARS link between CERN and Fermilab. Blue is CERN->Fermilab, Orange is Fermilab->CERN
We had enough network and transfer bandwidth to do the data challenge and still take raw data on the cold box.

Distributed processing: justIN in Data Challenge 4

Plots of sites where processing occurred

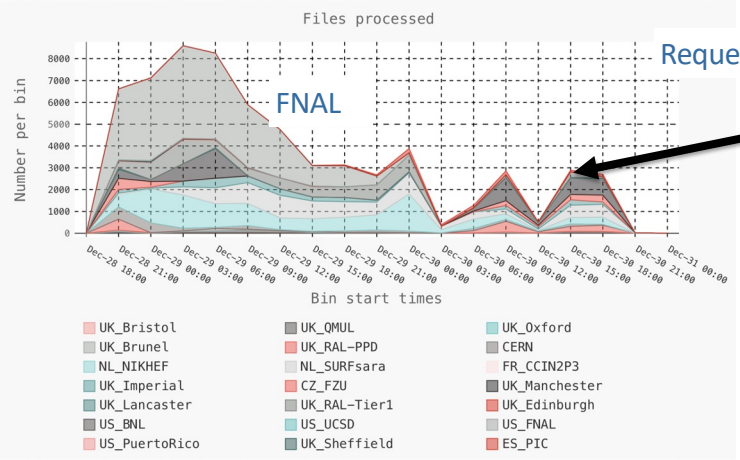

Not much data at Manchester (VD coldbox top)

Lots of data at Manchester! (VD coldbox bottom)

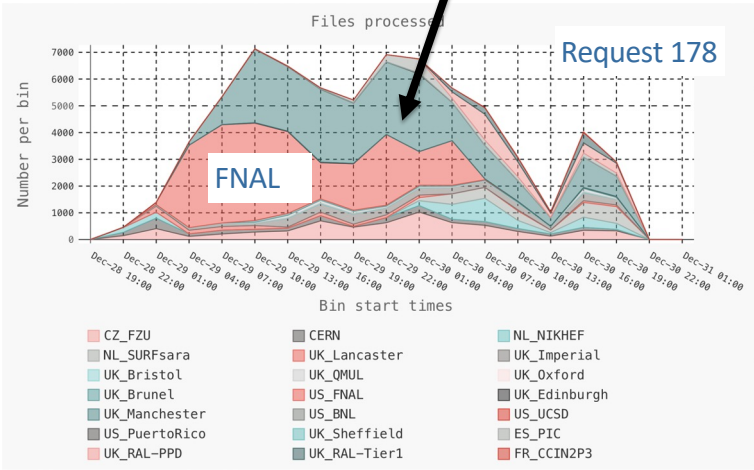

Live monitoring shows how jobs processing the files follow the data to suitable sites, using replica locations originally from Rucio

Conclusions

- We have tested that the full pipeline from experimental hall to CERN to Fermilab is functional and stable and can run at rates expected in the forthcoming ProtoDUNE II beam run (3.6GB/s)
- The software deployed and tested has been deployed in production now as our production data pipeline.
- Using this software to build a second pipeline for the DUNE Near Detector Liquid Argon Cube test beam this summer.
- Converting output of offline production to use this pipeline as well.
- Expect to use a very similar software stack when data starts to flow from SURF (1GB/s expected).

Acknowledgments

- Storage dept. at Fermilab for operating Rucio, MetaCat, and Data Dispatcher servers as well as writing MetaCat, Data Dispatcher, and the ingest and declaration daemons (Igor Mandrichenko, Brandon White, Dennis Lee)
- Jake Calcutt for Data Dispatcher testing
- Andrew McNab, Chris Brew, Raja Nandakumar for justIN design, development and testing.
- CERN IT staff for FTS3, EOS, CTA support.
- DUNE Data management team, S. Timm, D. Benjamin, W. Yuan.
- Rucio developers
- This research is based upon work supported by the US Department of Energy, Office of Science, Office of High Energy Physics.

Brookhaven National Laboratory

U.S. DEPARTMENT OF ENERGY | Office of Science

Fermilab DUNE