DUNE

Doug Benjamin (BNL) on behalf of the DUNE computing consortium

Data Challenge 2024 Workshop

9 Nov 2023

Acknowledgements

- Many Thanks to the following folks for providing information and slides for this talk:
- Steve Timm, Heidi Schellman, Mike Kirby, Andrew McNab
- Credit goes to them, and blame goes to me.

Introduction

Quick reminder about DUNE

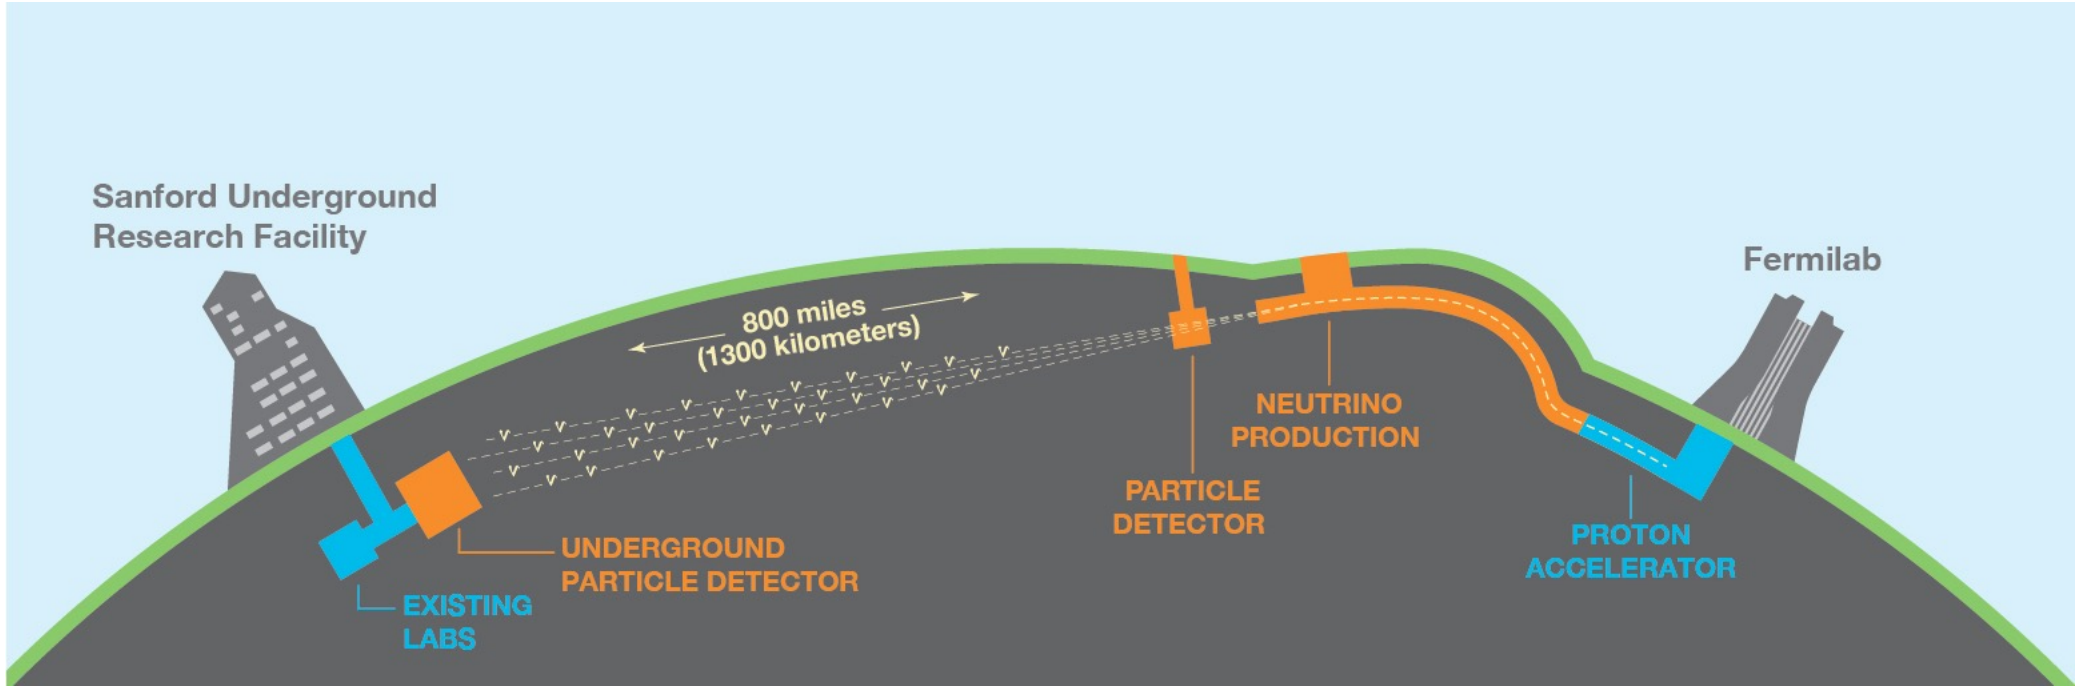

- neutrino experiment studying neutrino oscillation parameter (mass ordering, matter vs antimatter asymmetry, unitarity), proton decay, supernova neutrinos, and more.
- four very large LAr TPC (17 kT) at 4850 ft underground in Lead, SD (Homestake Mine)
- near detector onsite at Fermilab being designed (3 sub-detectors, two that move)
- two prototypes at CERN - (ProtoDUNE II Horizontal Drift - ProtoDUNE II Vertical Drift)

Optimized location of the Far Detector

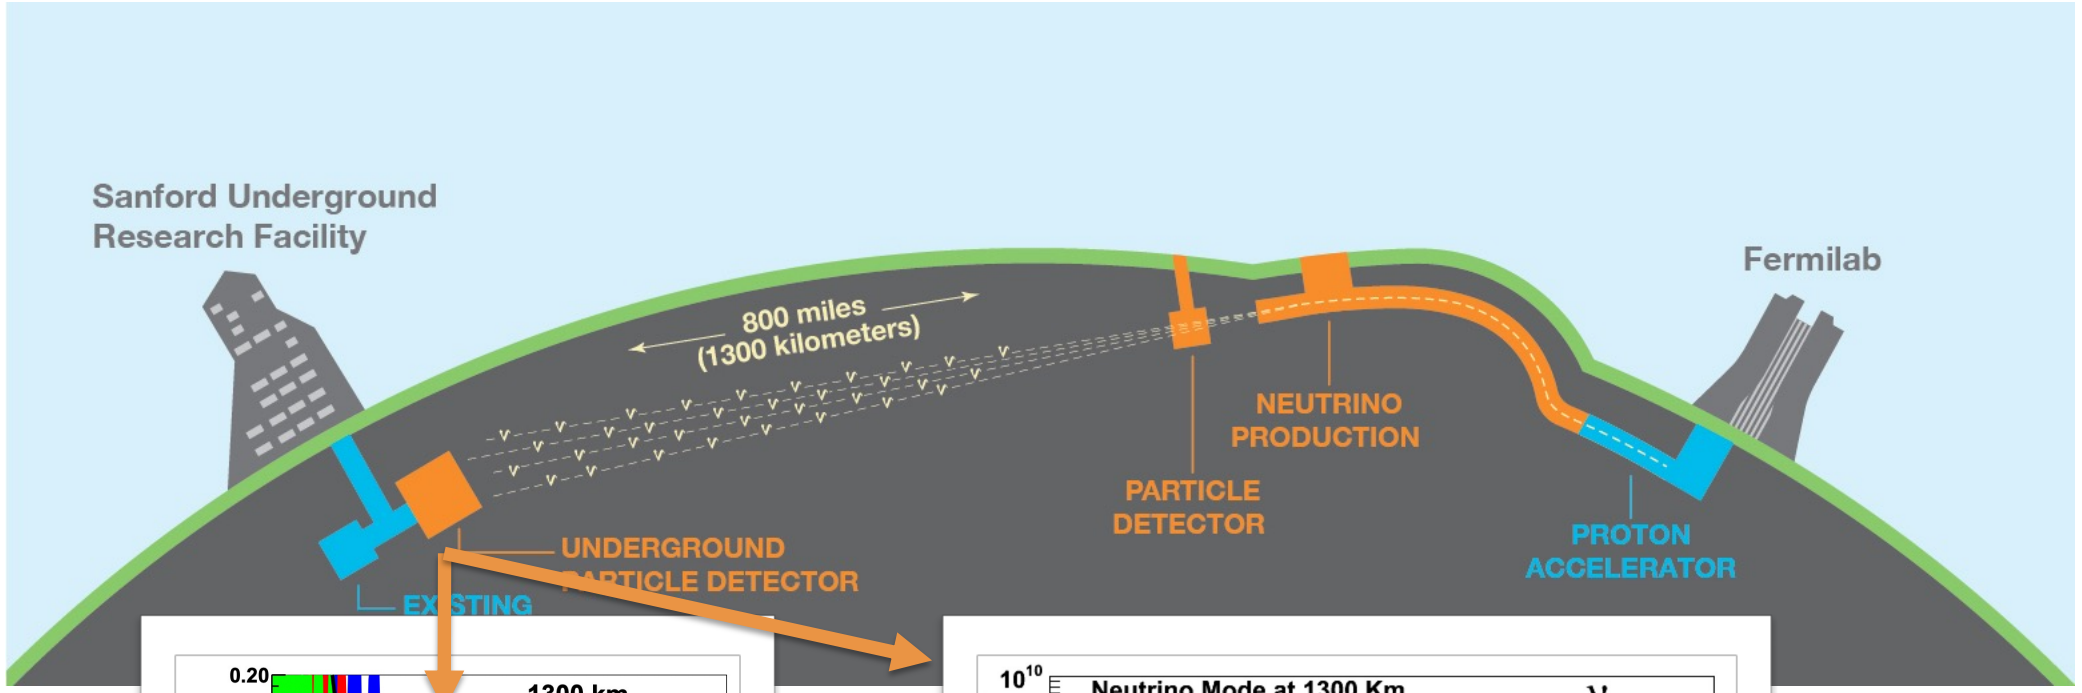

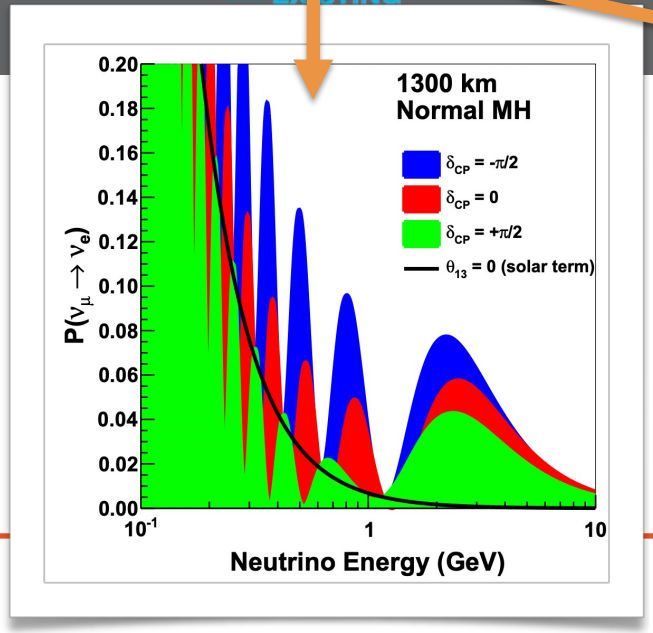

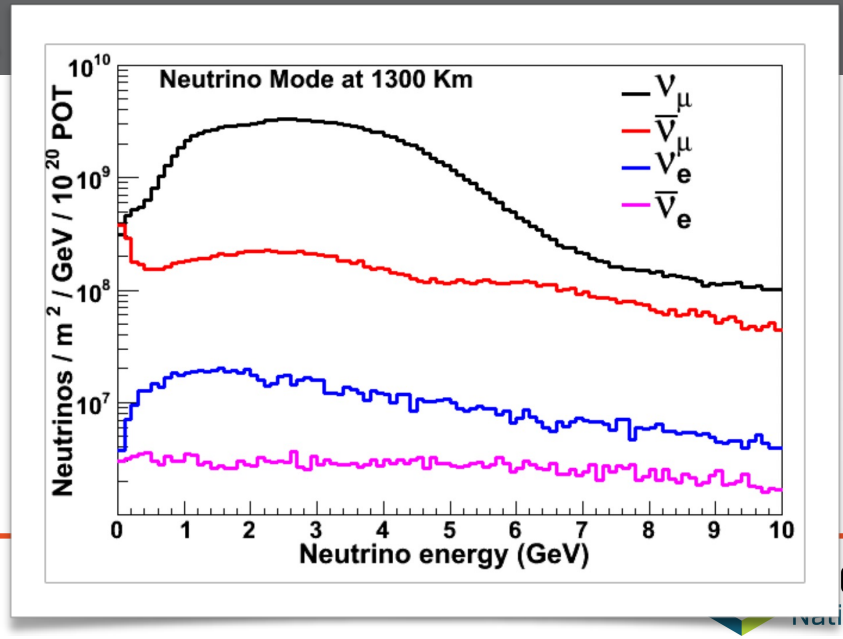

Supernova neutrino Burst

- A primary science driver “Detect and measure the ν_e flux from a core-collapse supernova within our galaxy” - DUNE Far Detector Technical Design Report
  - Sensitive to neutrinos from around 5 MeV to a few tens of MeV.
  - Good pointing resolution to a supernova (~ 5°)
- Network SURF to rest of world – 100 Gbe
- Need to transfer out 140 TB for precision ‘prompt’ processing

Crab Nebula Supernova Remnant – James Webb Space Telescope

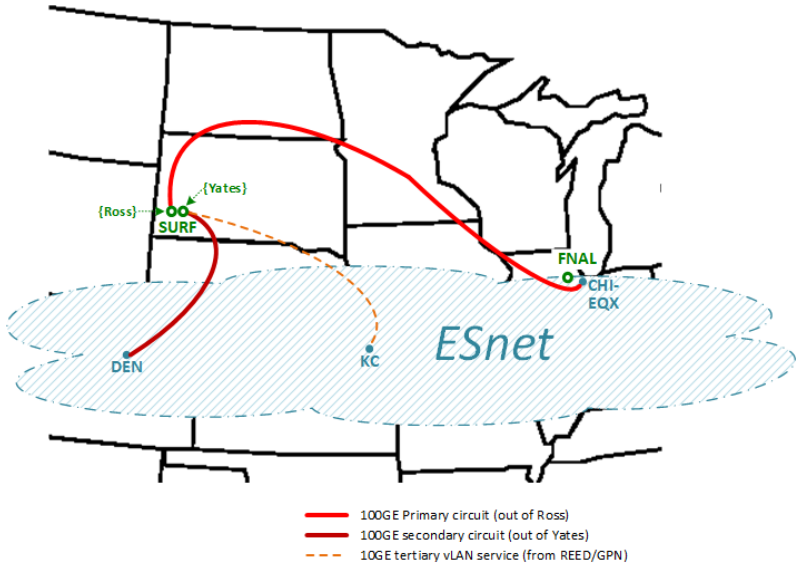

Brookhaven National Laboratory DUNE

Far Detector Dataflow and Trigger Records

- beam coincidence events are extremely important, but of limited total volume
  - ~1 Hz beam rate
    - majority of beam spills will not interact in FD
  - 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 - one FD module
  - work w/ trigger primitives for immediate optical follow up
  - transfer out 4 hours and process in 4 hours for precision optical observations
- DUNE requirement - less than 30 PB/year total to permanent storage from all active FDs

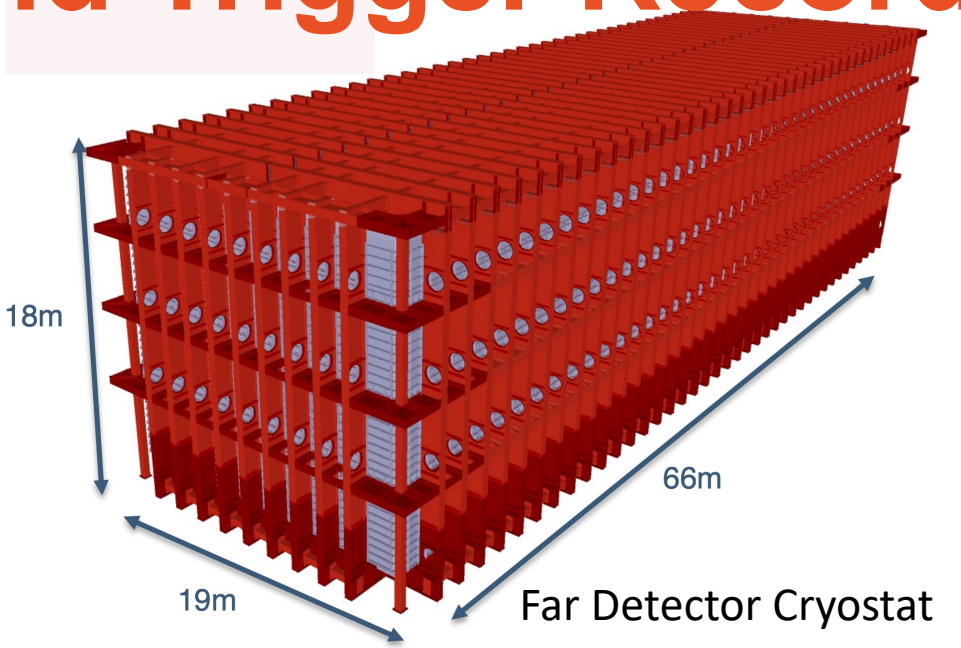

Far Detector Cryostat

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	$\leq$3.8 GB	35 TB/year
Calibrations	2/year	750 TB	1.5 PB/year
Total			9.4 PB/year

Recently published DUNE Computing CDR - https://arxiv.org/abs/2210.15665

DUNE Computing Resource Model

- less “tiered” than current WLCG model — > flatter model proposed by HSF DOMA working group
  - take advantage of existing WLCG sites that can add DUNE access
  - require reasonable minimum size - storage elements
  - allow for CPU only sites with **data streaming**
- collaborating institutions (or groups of institutions) provide significant disk resources (~1PB chunks)
- plan to use common tools for most services
- participation in the HSF process important to provide and integrate new solutions

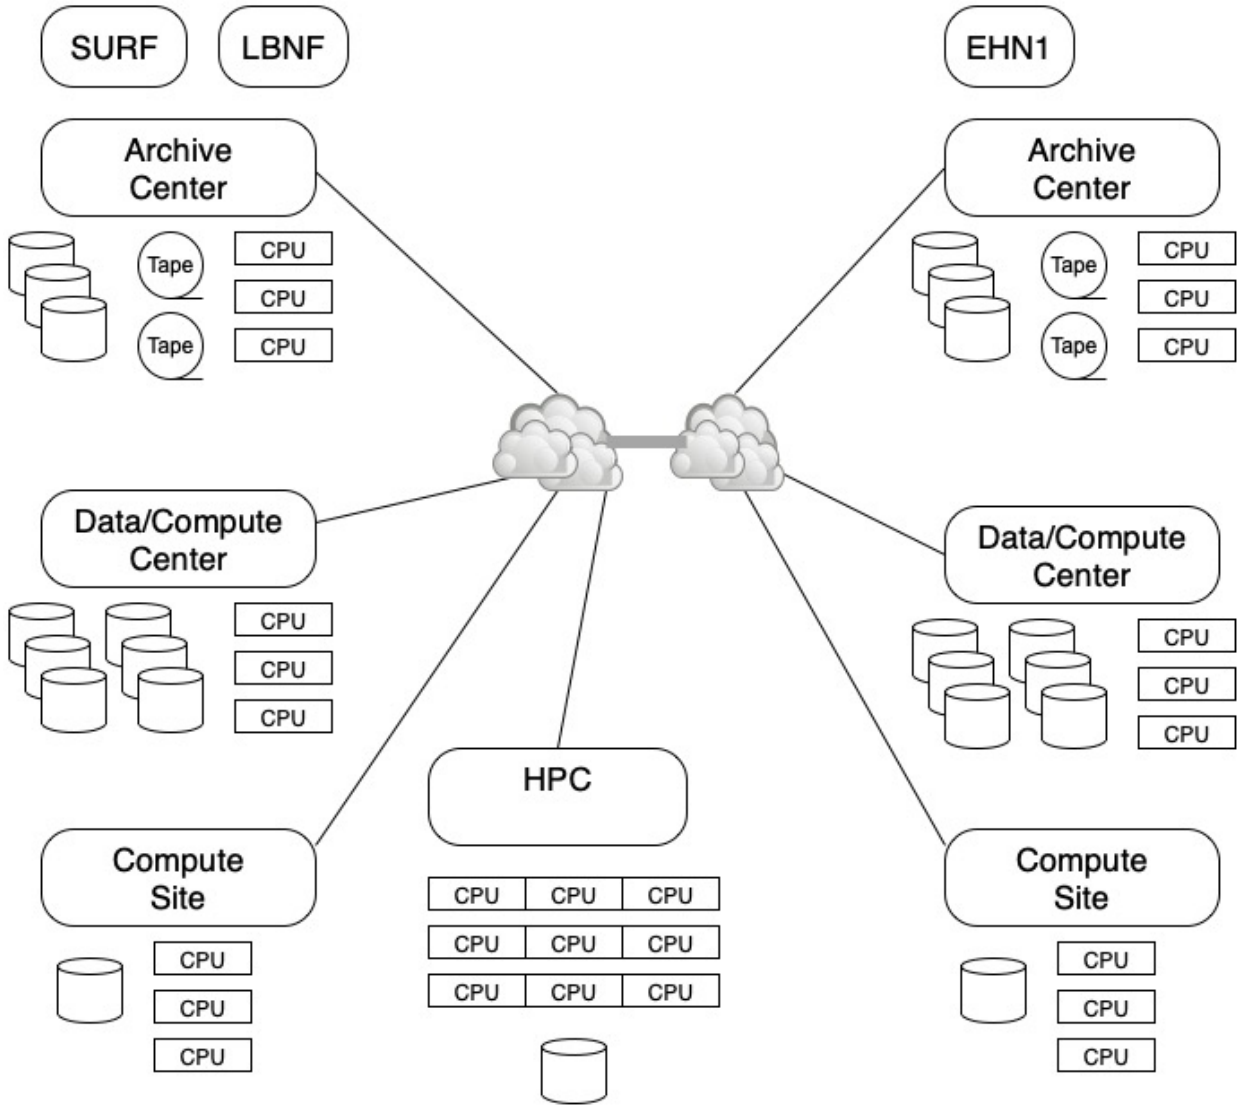

Brookhaven National Laboratory DUNE

CERN and Neutrino Platform currently hosting ProtoDUNE

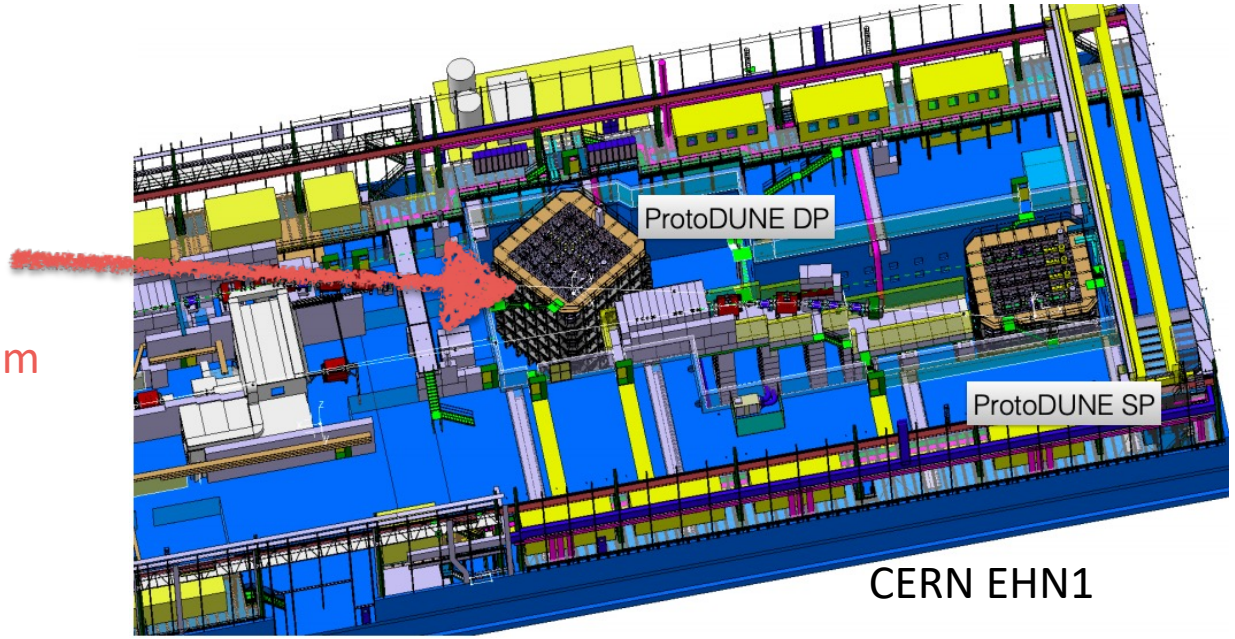

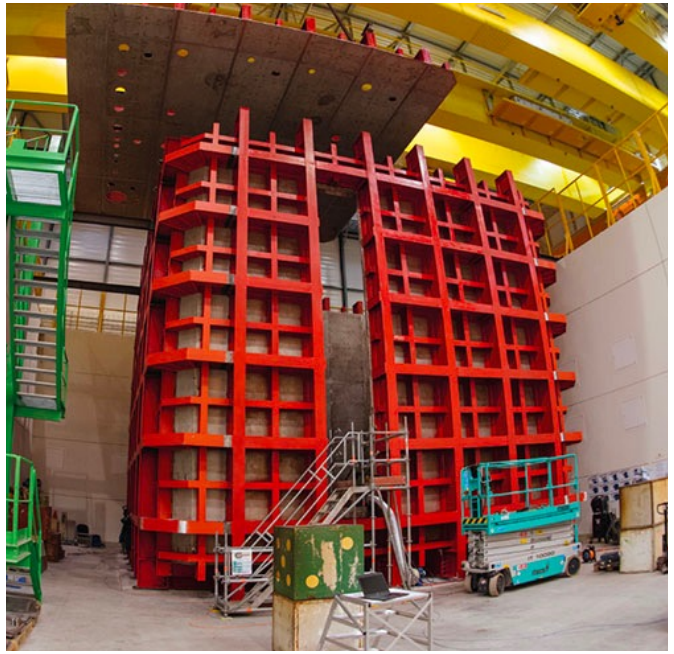

- ProtoDUNE Single Phase and Dual Phase
  - constructed and operated during the timeframe of 2018 - 2020
  - ProtoDUNE SP took 6 weeks of beam (~25 Hz)
  - invaluable information about performance, construction, and operations

- ProtoDUNE II currently under construction
  - Horizontal Drift and Vertical Drift
  - HD hoping for LAr filling 2024Q1
  - VD assembly underway for 2024 operations
  - Beam operations in 2024

DUNE data pipelines

ProtoDUNE Run 2 and DUNE Data Pipeline Diagram

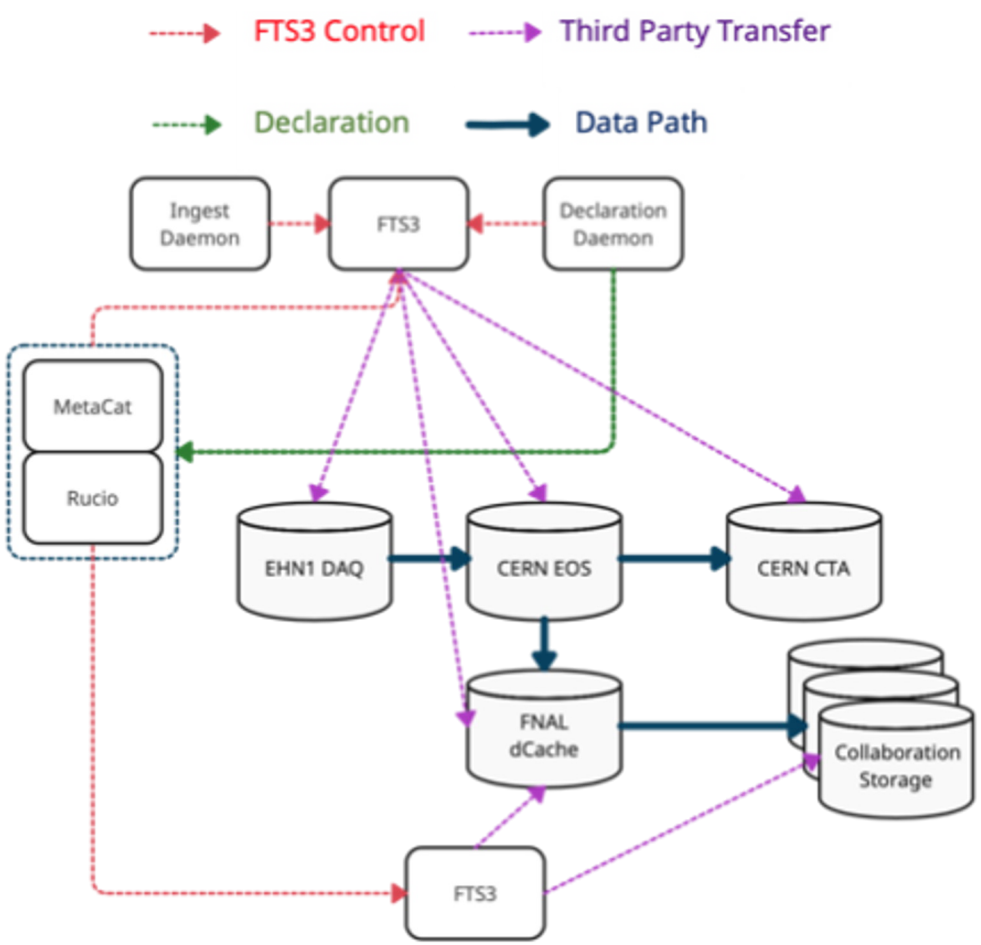

- Ingest Daemon and Declaration Daemon
  - Ingest daemon brings files from experimental systems to dropbox
    - Can operate without connection to Fermilab.
    - All transfers done via FTS3
  - Declaration daemon declares them to MetaCat and Rucio and makes rules to get them to the final destinations.
  - being adapted for 2x2 test beam at FNAL
- 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.

Brookhaven National Laboratory DUNE

DUNE Workflow System

- Workflow system - **justIN** (developed in UK):
  - justIN designed according to DUNE requirements to manage processing campaigns for users, working groups, and the production team
  - Matches workflows to sites “near” unprocessed files
  - Uses a “just in time” philosophy (late binding)

- Data flows with justIN
  - Data is streamed into the compute nodes from “nearby” RSE (either at same site or across WAN from a close site)
  - Rucio upload used to store output at “nearby” RSE
    - Output RSE does not need to be same as input RSE
  - Rucio/FTS - used to transfer data back to FNAL

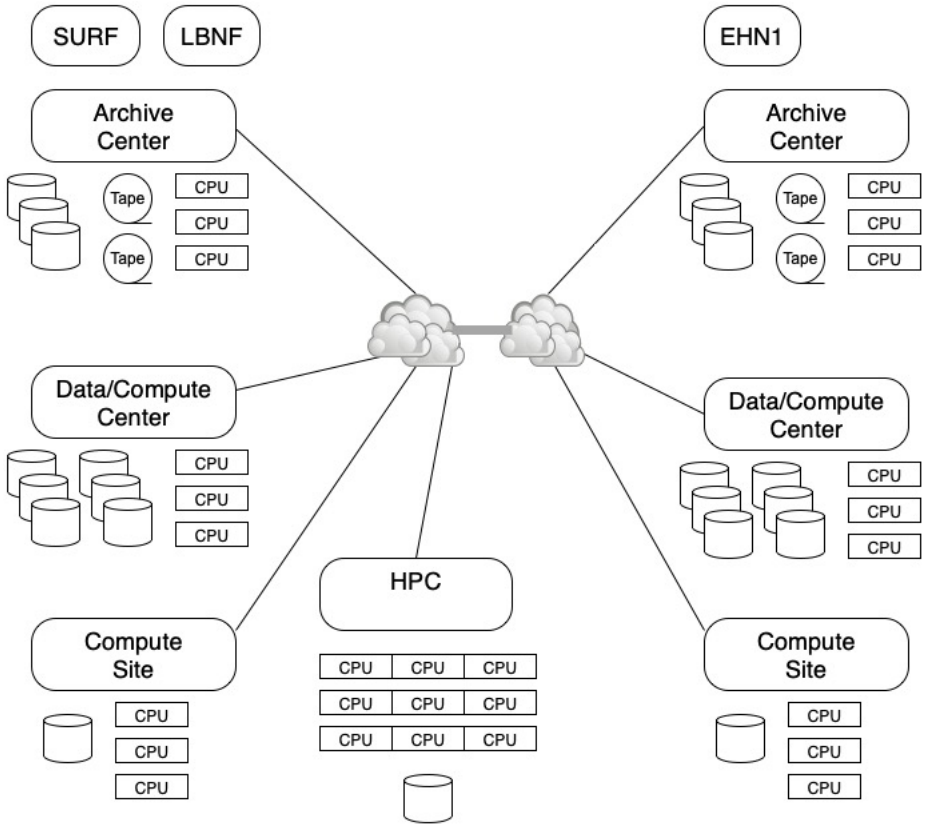

DUNE Computing (flat) model
Prioritizes local file access when possible
Will use “nearby” sites to access data as needed

Brookhaven National Laboratory DUNE

ProtoDUNE HD/VD data flows

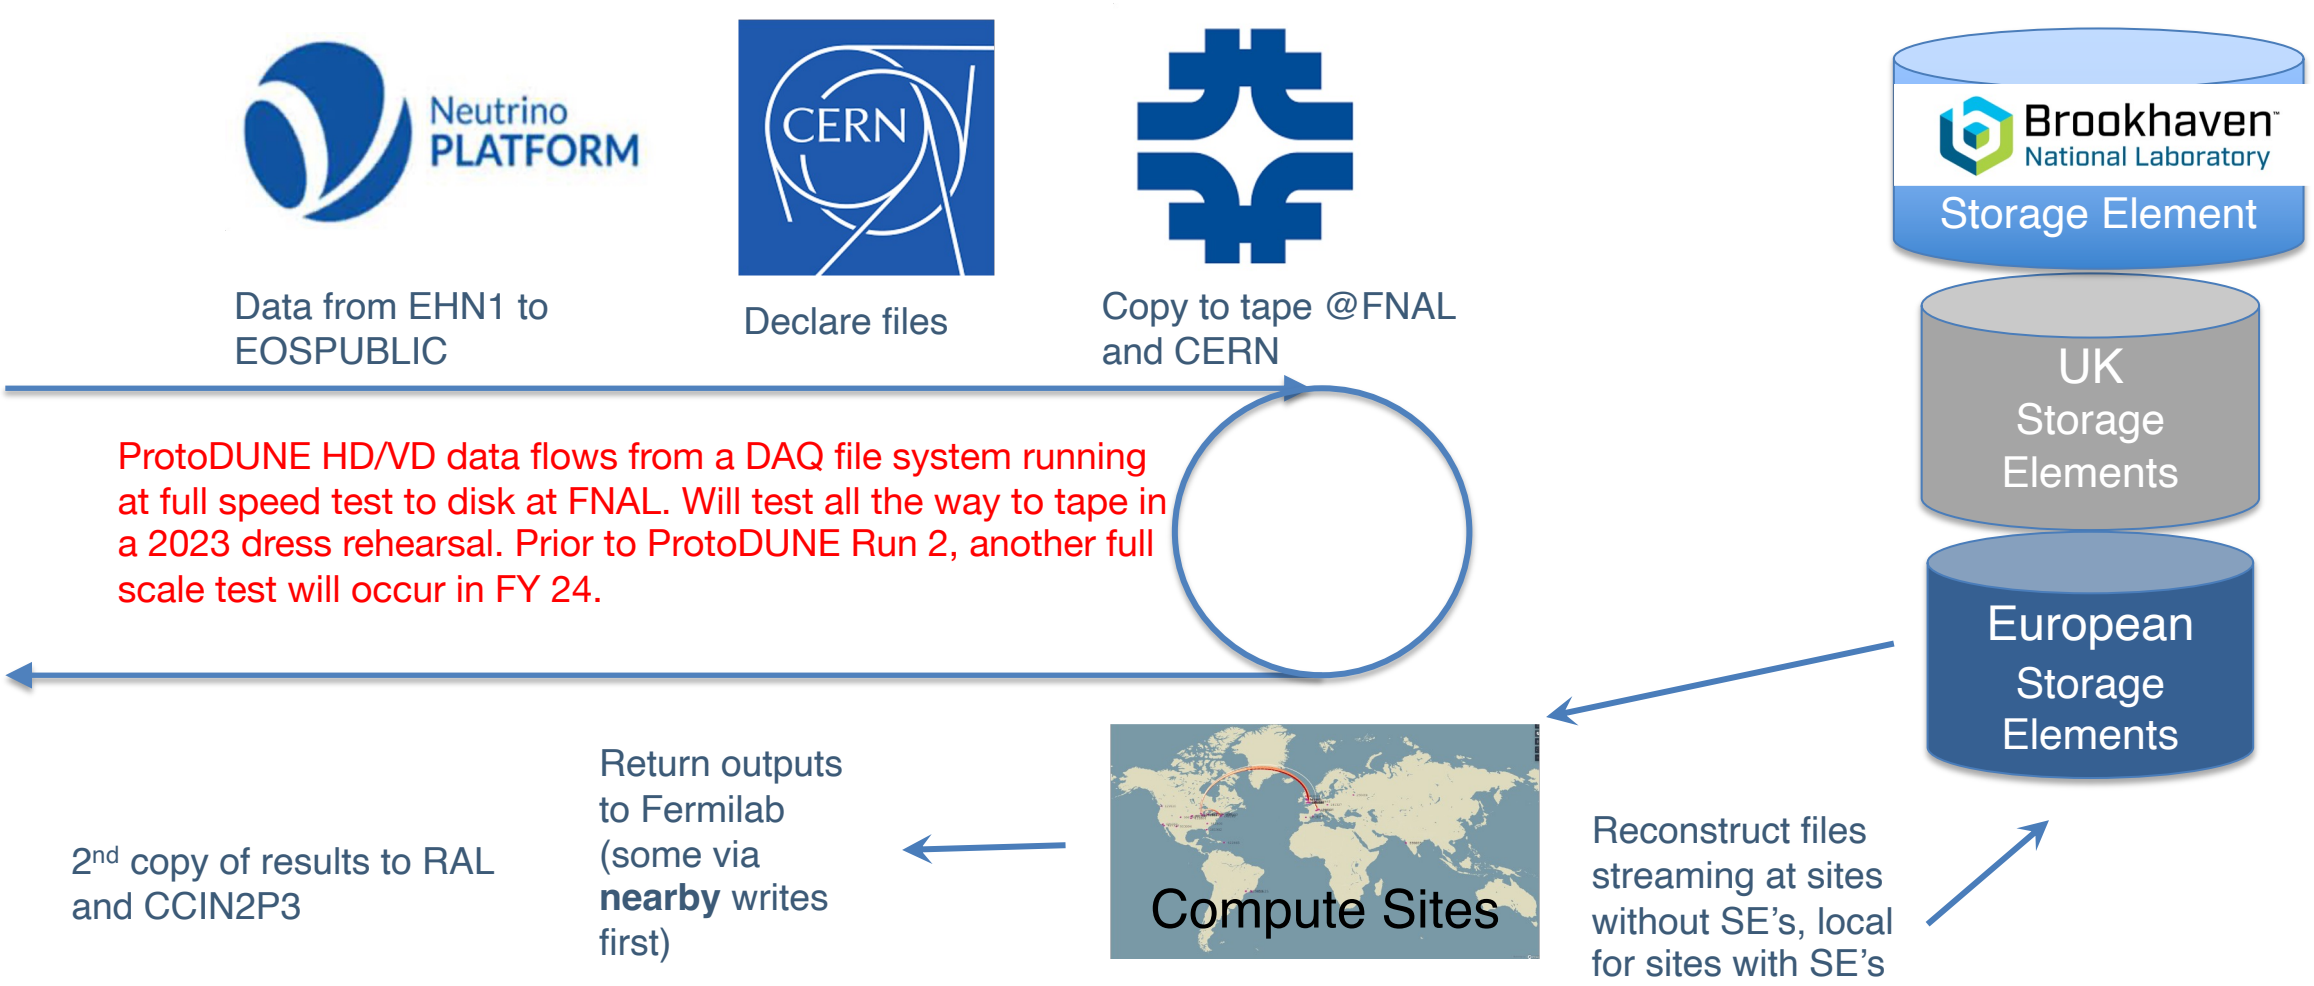

DUNE Phase I data flows from SURF

SURF

Declare files

Copy to tape @FNAL and CERN

Brookhaven National Laboratory Storage Element

UK Storage Elements

European Storage Elements

This data pipeline will be well tested before beam comes to both the DUNE Far Detector and Near Detector. Two copies of raw data will be on tape. Transient copies of data files for efficient processing will be on both in the US and non-US sites.

Reconstruct files streaming at sites without SE's, local for sites with SE's

Return outputs to Fermilab (some via **nearby** writes first)

2nd copy of results to other sites as dictated by DUNE policy.

DUNE - Data Challenge 24 tests

DUNE Involvement in WLCG Data Challenge 24

"FD" Raw Data
archival storage

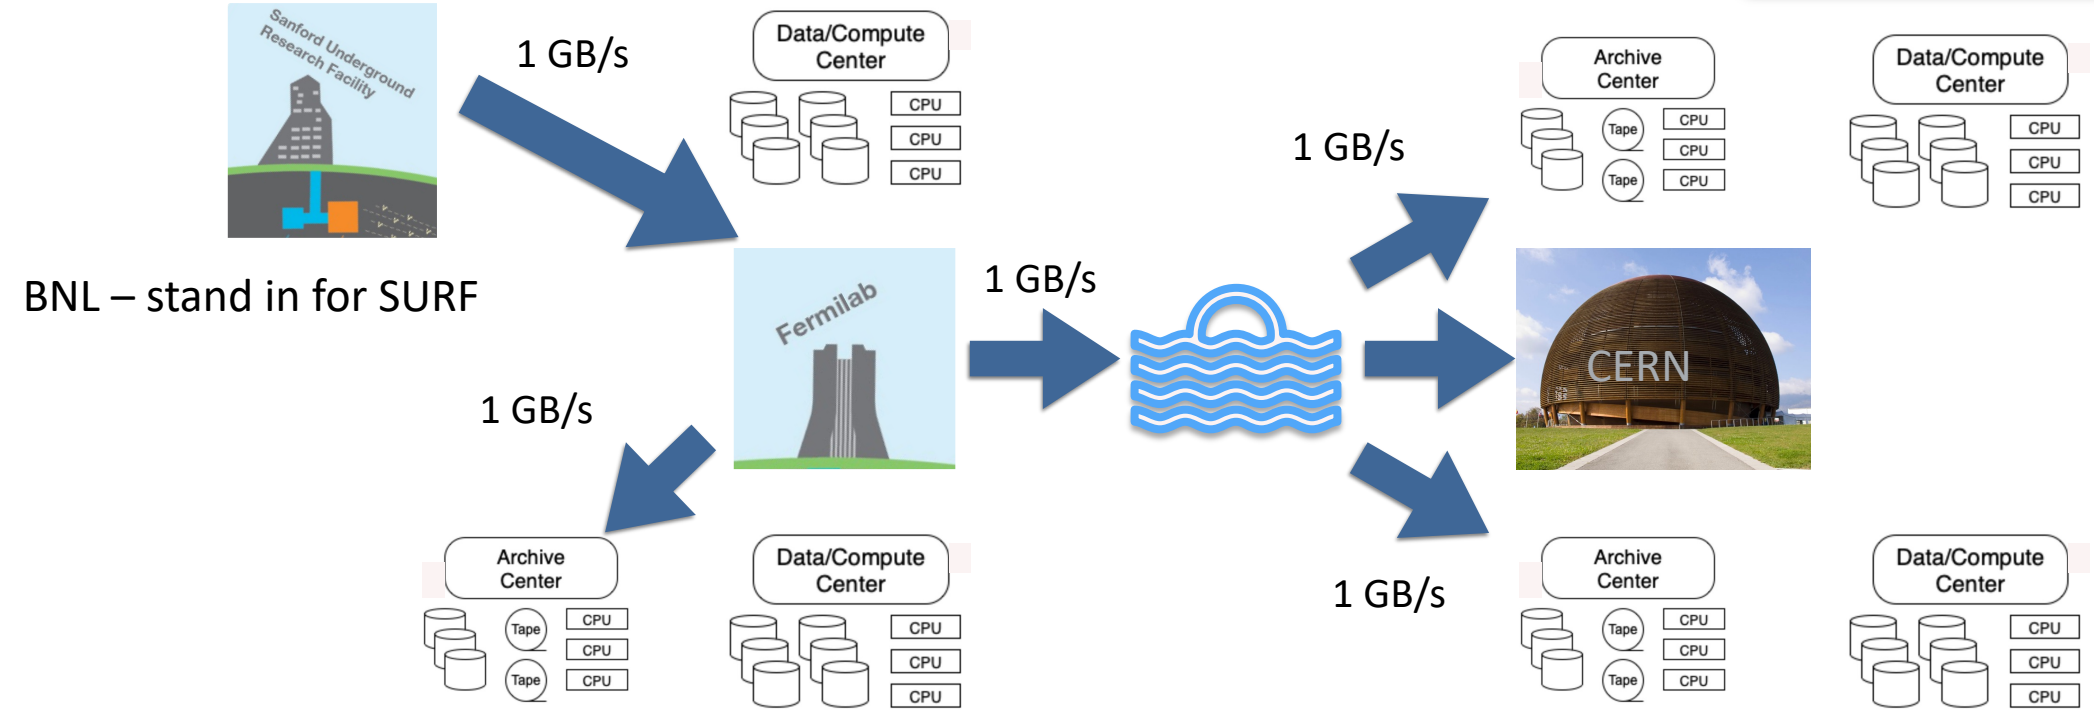

not to scale, not a technical design, it's just a cartoon

- Simulate the archival of 25% of the raw data rate from the Far Detector
  - translates to 1 GB/s from SURF to FNAL
  - replicate that "FD" raw data to archival storage facilities around the world
  - replicate the "FD" raw data to disk storage elements around the world for prompt access from compute elements
- Both job submission and RSE to RSE w/ token authentication/authorization

Brookhaven National Laboratory DUNE

DUNE Involvement in WLCG Data Challenge 24

"FD" Raw Data
raw processing

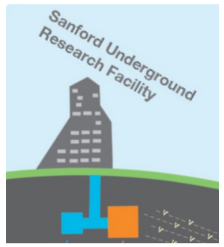

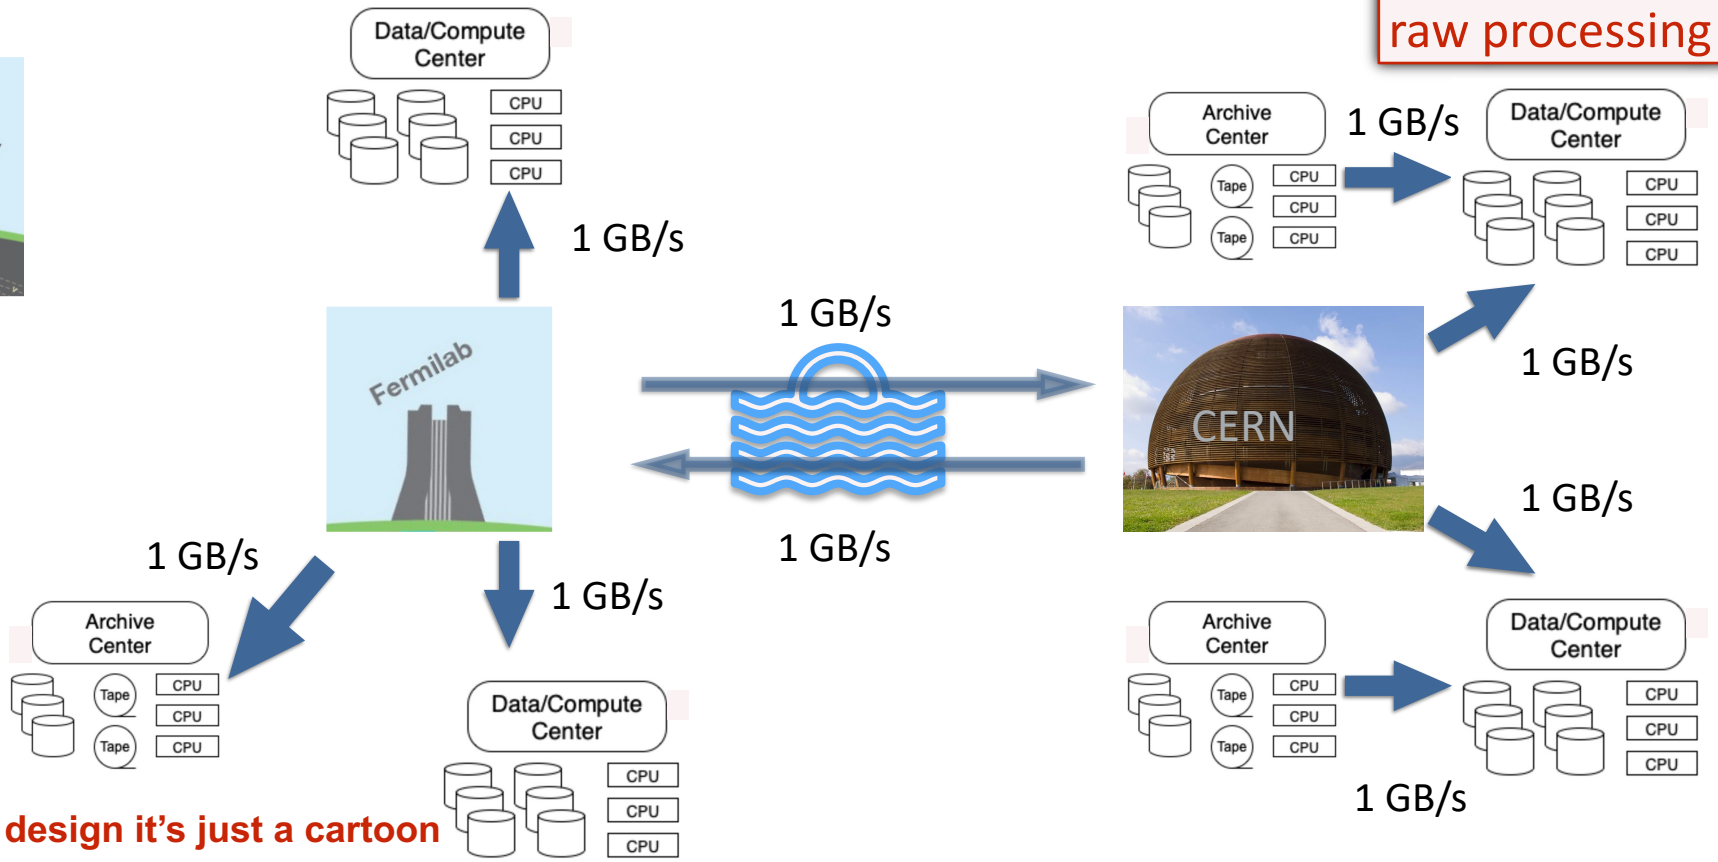

not to scale, not a technical design it's just a cartoon

- Maintain continuous processing workload at distributed sites commensurate with 25% "FD" raw data rate
  - utilize compute elements across the WLCG and OSG
  - match the locality of jobs with locality of data at nearby RSEs
- **Both job submission and RSE to RSE w/ token authentication/authorization**

Brookhaven National Laboratory DUNE

DUNE Involvement in WLCG Data Challenge 24

SuperNova Raw Data
rapid transfer & processing

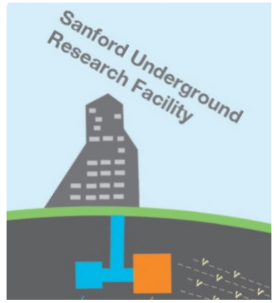

3.5 GB/s

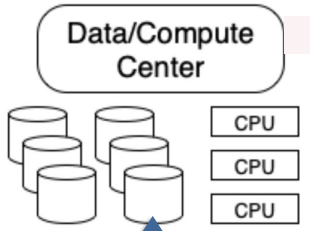

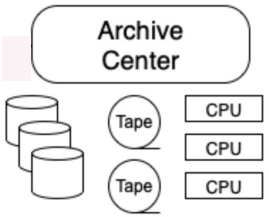

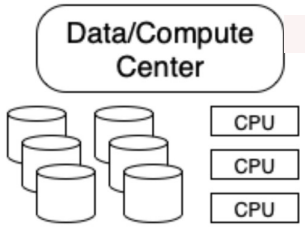

BNL – stand in for SURF

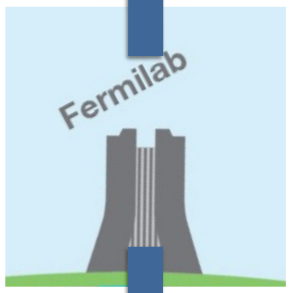

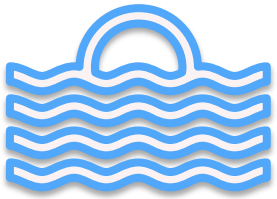

3.5 GB/s

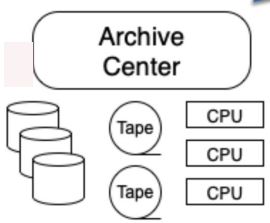

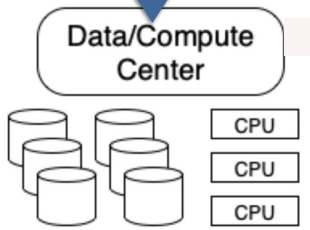

HPC - NERSC

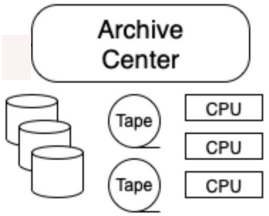

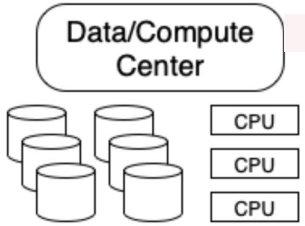

not to scale, not a technical design
it's just a cartoon

Note: maximum network rate from SURF –
100 Gbits/sec (12.5 GB/s)

Brookhaven National Laboratory DUNE

Current Status

- Next Week DUNE will conduct a Dress Rehearsal at head of ProtoDUNE HD/VD data taking in the CERN beam.
  - Many of the components and data flows needed by the DUNE DC24 activities will be test at a higher scale (~ 4 GB/s)
    - Only the FNAL to CERN data flow will not be tested (ProtoDUNE HD/VD data goes from CERN to FNAL)
- FTS load generator tool kit developed by Hiro Ito (BNL) used to test transfers BNL to FNAL (~ 2GB/s)
  - New additional dCache door being deployed at BNL in Nov. to increase Rate to > 4 GB/s
- Need to test FNAL to CERN data flow
- Still to test FNAL to NERSC data flow @ 3.5 GB/s at part of supernova burst test
  - In Nov will deploy xrootd base RSE at NERSC for DUNE
- Token testing and deployment still outstanding
  - We have token issuer know to work with key storage elements
  - Our token issuer (Cilogon) is not IAM based

Brookhaven National Laboratory DUNE

Conclusions

- DUNE looks forward to participating in the WLCG Data Challenge 24 activity.
- Plan to run our tests concurrently with peak LHC experiments tests
  - excellent coordination will be needed especially with time zone difference Europe and US
- Most data flows will be tested this month at levels at or above DC24 requirements
  - remaining dataflows will be tested in December.

- Token deployment still an outstanding issue
  - Need to verify that all DUNE Storage Elements can use DUNE tokens
  - readiness for Rucio/Fts to use DUNE tokens

Brookhaven National Laboratory DUNE