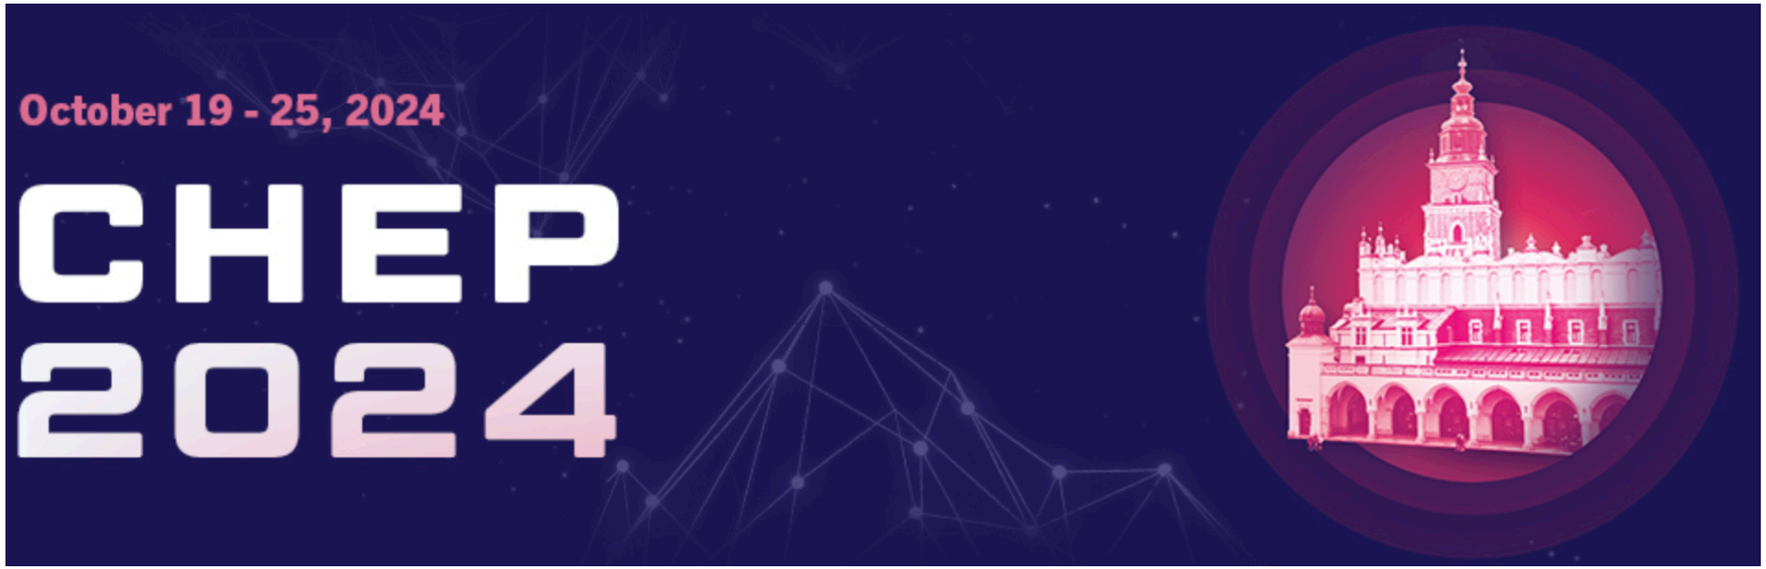

Status of DUNE Offline Computing

Michael Kirby on behalf of the DUNE Collaboration
27th Computing in High Energy and Nuclear Physics Conference
Oct 21, 2024 - Krakow, Poland

Brookhaven National Laboratory

DUNE

Quick Reminder about DUNE

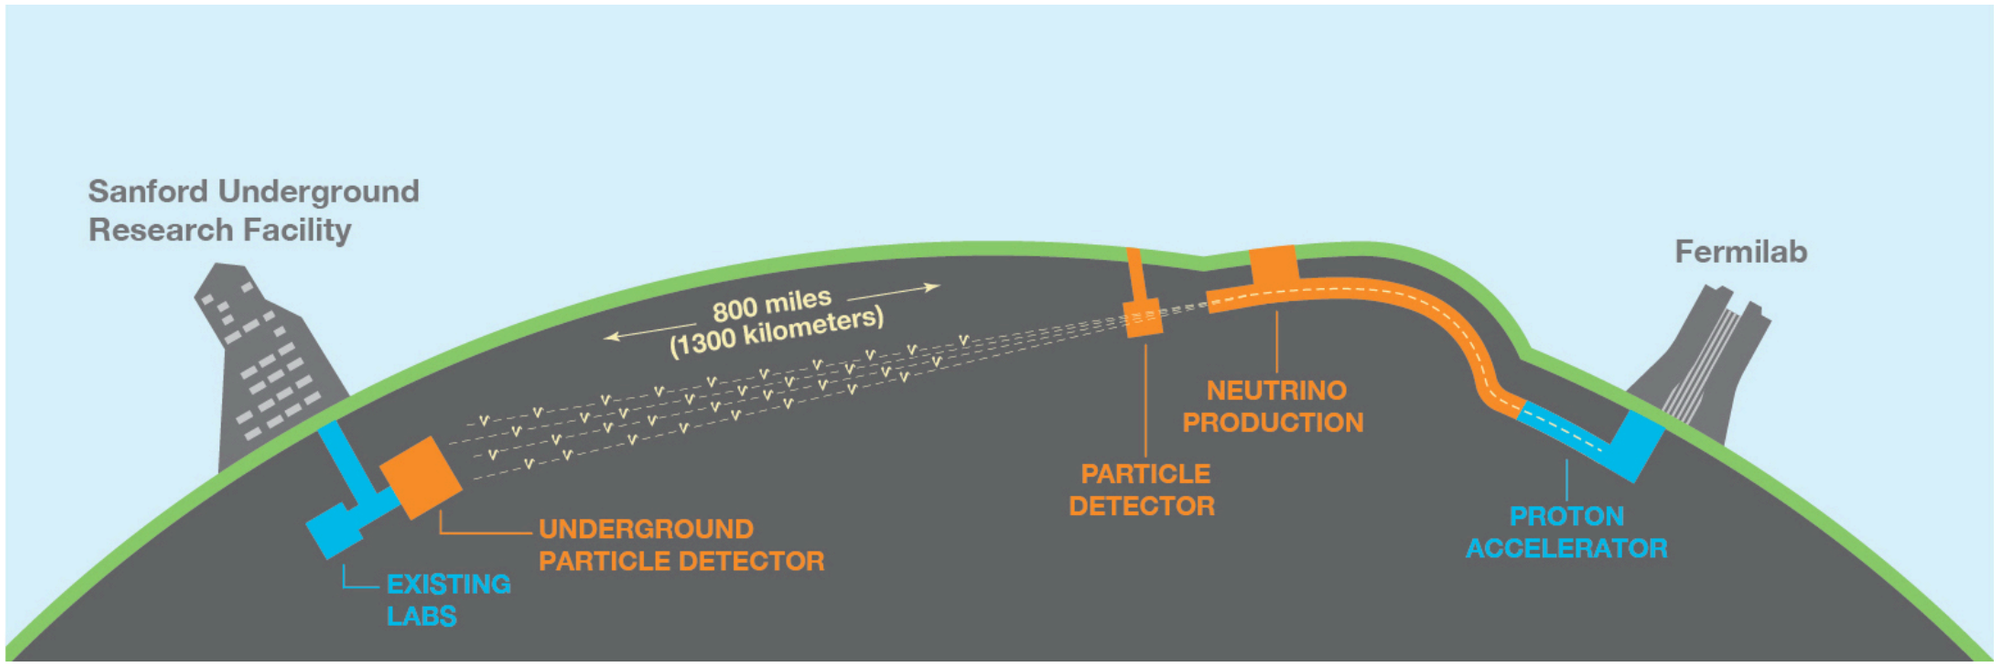

- neutrino experiment studying neutrino oscillation parameter (mass ordering, matter vs antimatter asymmetry, unitarity), proton decay, supernova neutrinos, and more.
- four very large LAr TPC (17 kT) at 4850 ft underground in Lead, SD (Homestake Mine)
- near detector onsite at Fermilab being designed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

Far Detector Design and Data Volume

- beam coincidence events are extremely important - large when compared with information density - ~4 GB/drift window
- cosmic ray events and calibrations projected to generate largest total volume
- supernova readout events
  - ~140 TB in 100 seconds - one FD module
  - work w/ trigger primitives for optical follow up (5-10 minutes timescale)
  - transfer out 4 hours and process in 4 hours for timely reporting of the neutrino spectrum
- DUNE requirement - less than 30 PB/year total to archival storage from all active FDs

Far Detector Cryostat

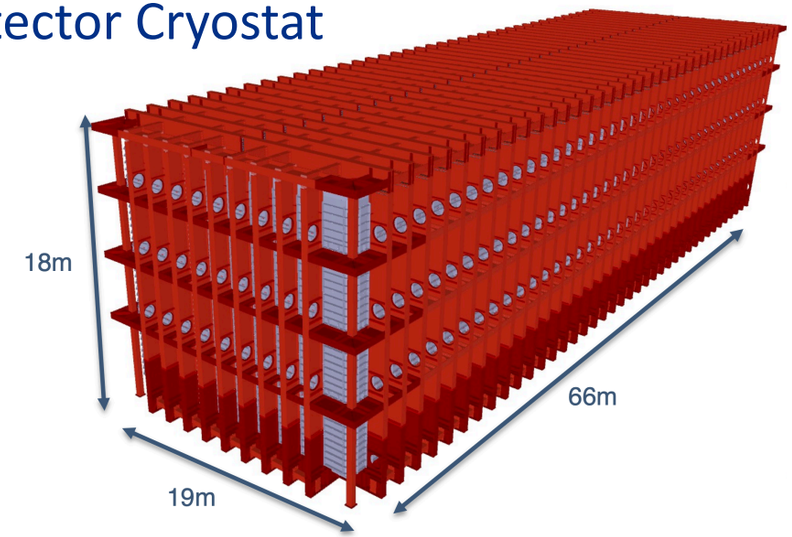

Process	Rate/module	size/instance	size/module/year
Beam event	41/day	3.8 GB	30 TB/year
Cosmic rays	4,500/day	3.8 GB	6.2 PB/year
Supernova trigger	1/month	140 TB	1.7 PB/year
Solar neutrinos	10,000/year	$\leq$3.8 GB	35 TB/year
Calibrations	2/year	750 TB	1.5 PB/year
Total			9.4 PB/year

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

Near Detector Design and Data Flow

- designed to constrain neutrino beam flux, precisely measure cross sections, minimize detector response uncertainties
- three subdetectors: ND-LAr, Muon Spectrometer, and SAND
- intense neutrino beam dictate different detector design >15 overlapping v-interactions
- physics and detector design precipitate move towards GPU focused software for reco and sim
- ongoing efforts on scalability of ND software
  - redesigning data structures to reduce GPU memory requirements
  - identifying bottlenecks and organizing effort
  - important to decrease the workflow footprint for optimal use of GPU resources at facilities such as NERSC and other HPC resources

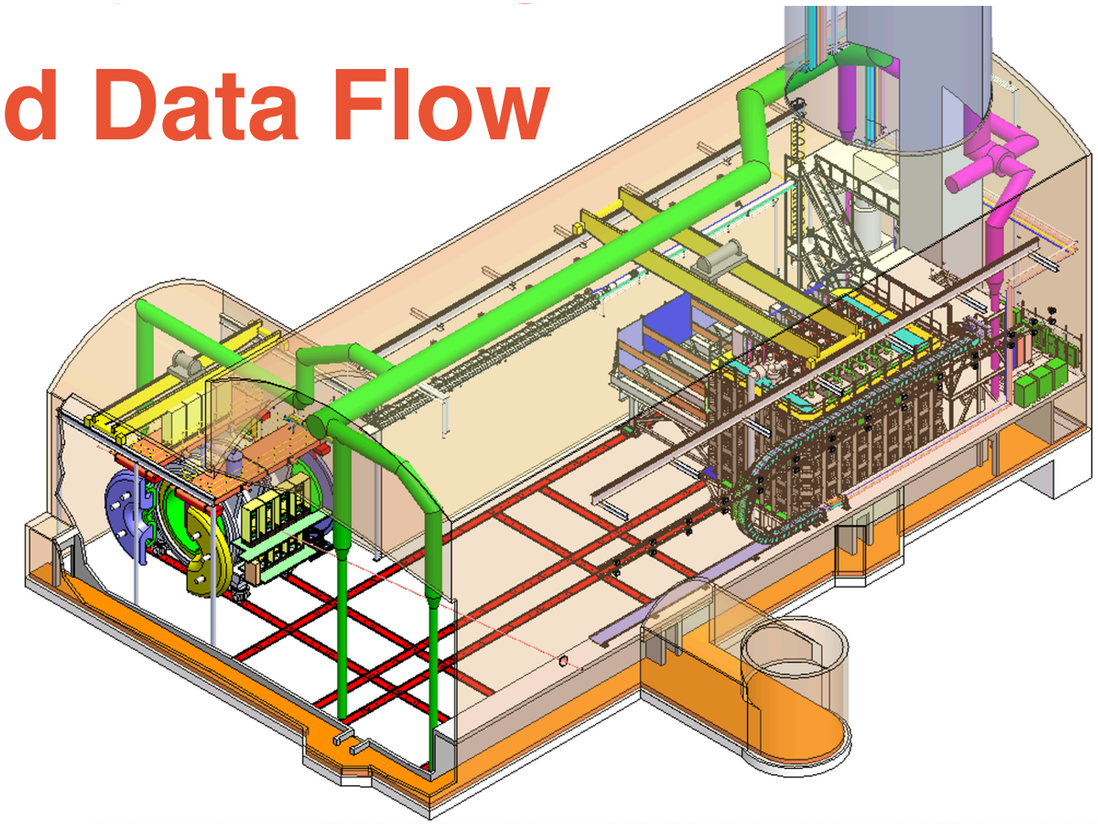

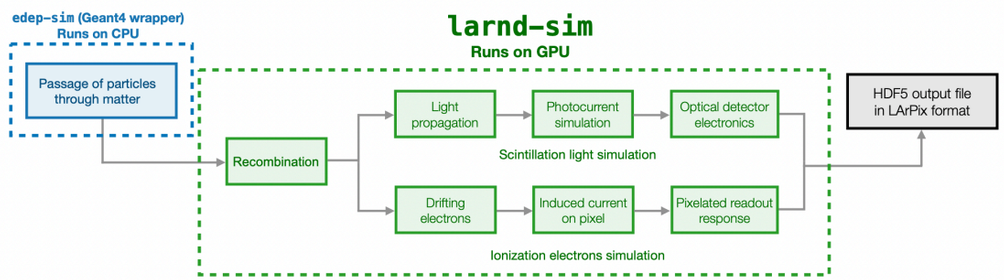

GPU based simulation of ND - LAr
2023 JINST 18 P04034

CERN Neutrino Platform hosting
ProtoDUNE II - Horizontal Drift (NP04) and Vertical Drift (NP02)

- ProtoDUNE-HD currently operating
  - Horizontal Drift completed beam run - 10 weeks of beam from the SPS
  - tertiary beam [1-7 GeV], both polarities
  - protons, kaons, pions for cross section measurements
- ProtoDUNE -VD soon to operate
  - Vertical Drift preparing fill later this year
  - beam operations for VD in 2025
  - anticipate 8 weeks of beam
- DAQ provides trigger generation and continuous stream of trigger primitives - simplified “hits” from TPC

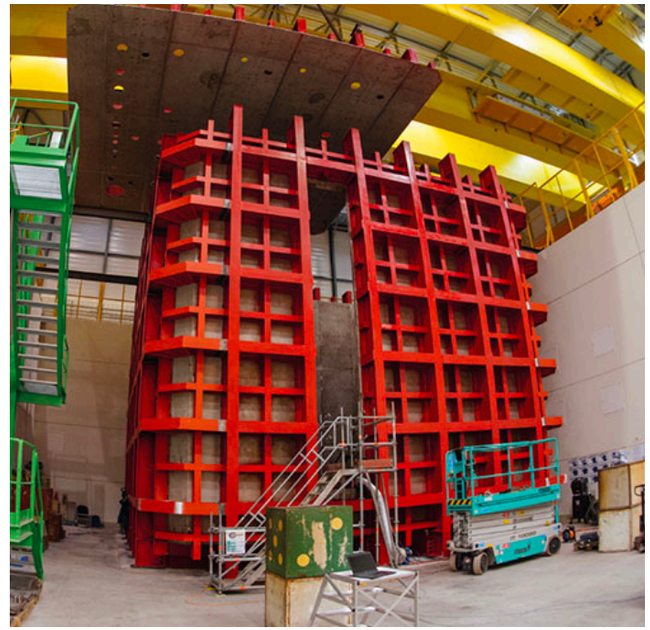

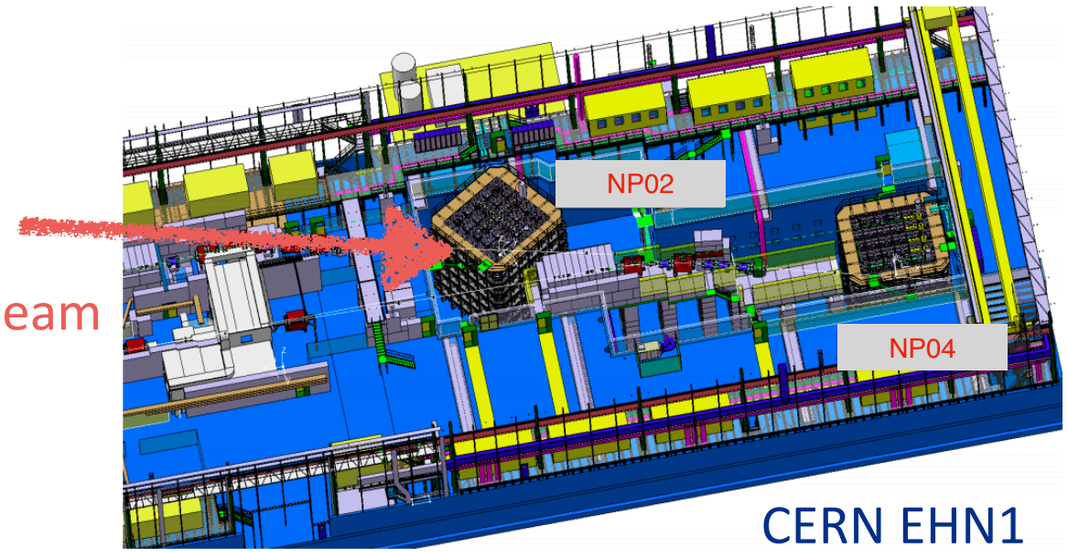

Data Management and Data Pipeline from NP04 (PDHD)

- pipeline utilizes sequences of tools (FTS3, Metacat, Rucio, and custom ingest/declaration daemons)
- early estimates anticipated 2-3PB of beam data to be written
- Aug 10-11 consecutive days 100+ TB
- wrote 4.5PB of data
- moved about 2PB of protodune-hd raw data to other disk sites in one week to accommodate additional data taking
- files moved to BNL, PIC, NIKHEF, FNAL, PRAGUE, RAL-PP, SURFSARA
- details on development by Wenlong Yuan Wed talk: DUNE Rucio development and monitoring

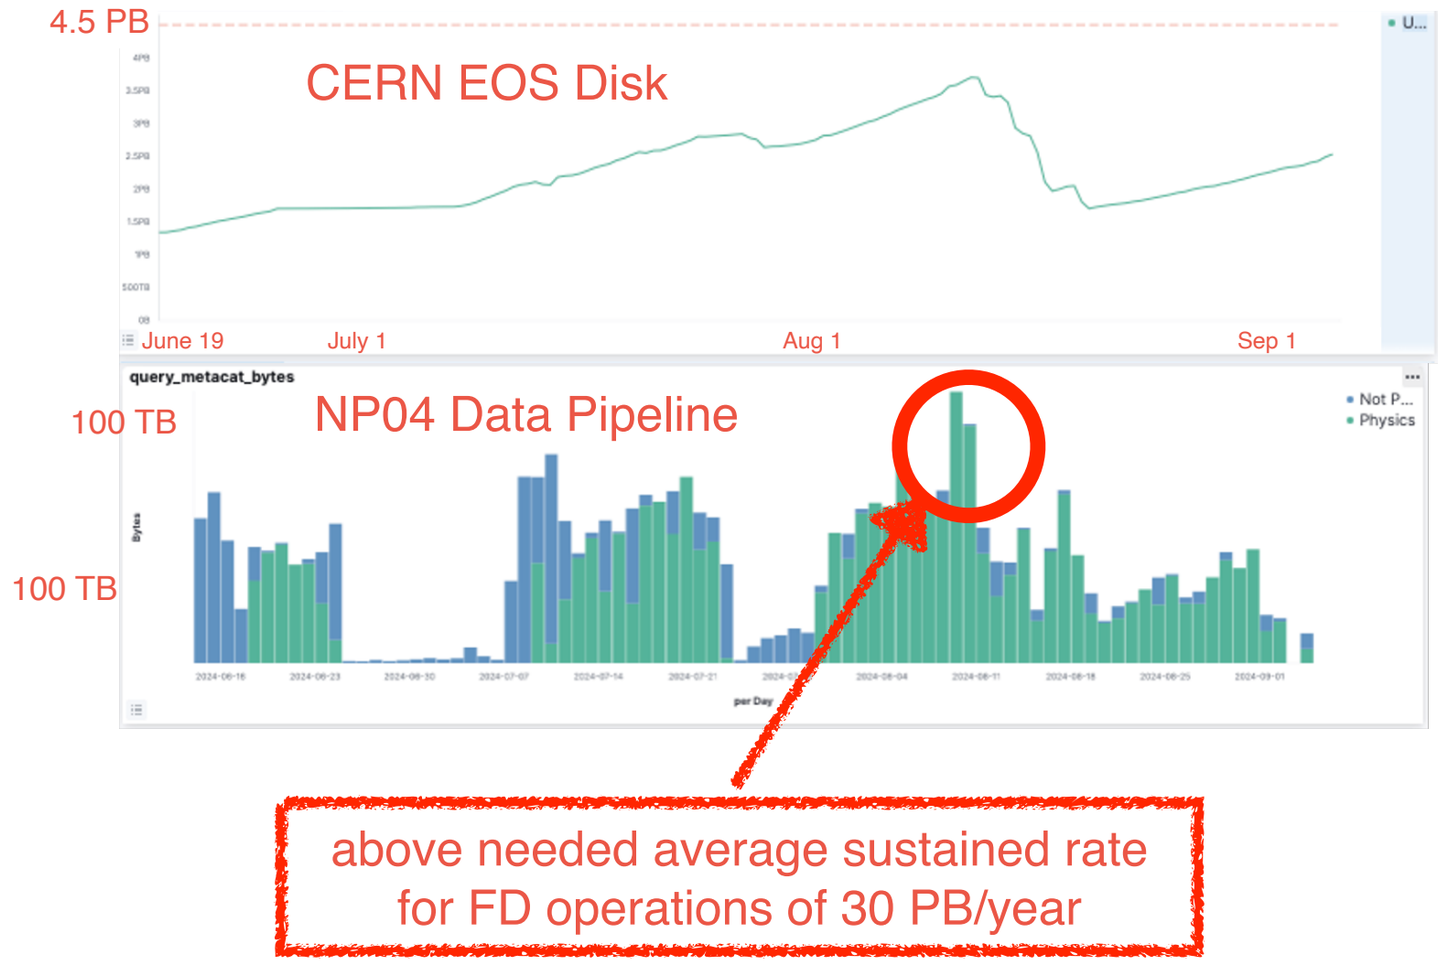

Production Operations

- responsible for generation of both prototype data and simulation for detector design/ sensitivity studies
- considerable challenges with large memory consumption and data volume movement
  - FD beam neutrino samples
  - FD low-energy generation
  - ND simulation with overlay
- successfully utilizing compute sites and RSEs from around the world

DUNE Production by Country and Site 2024

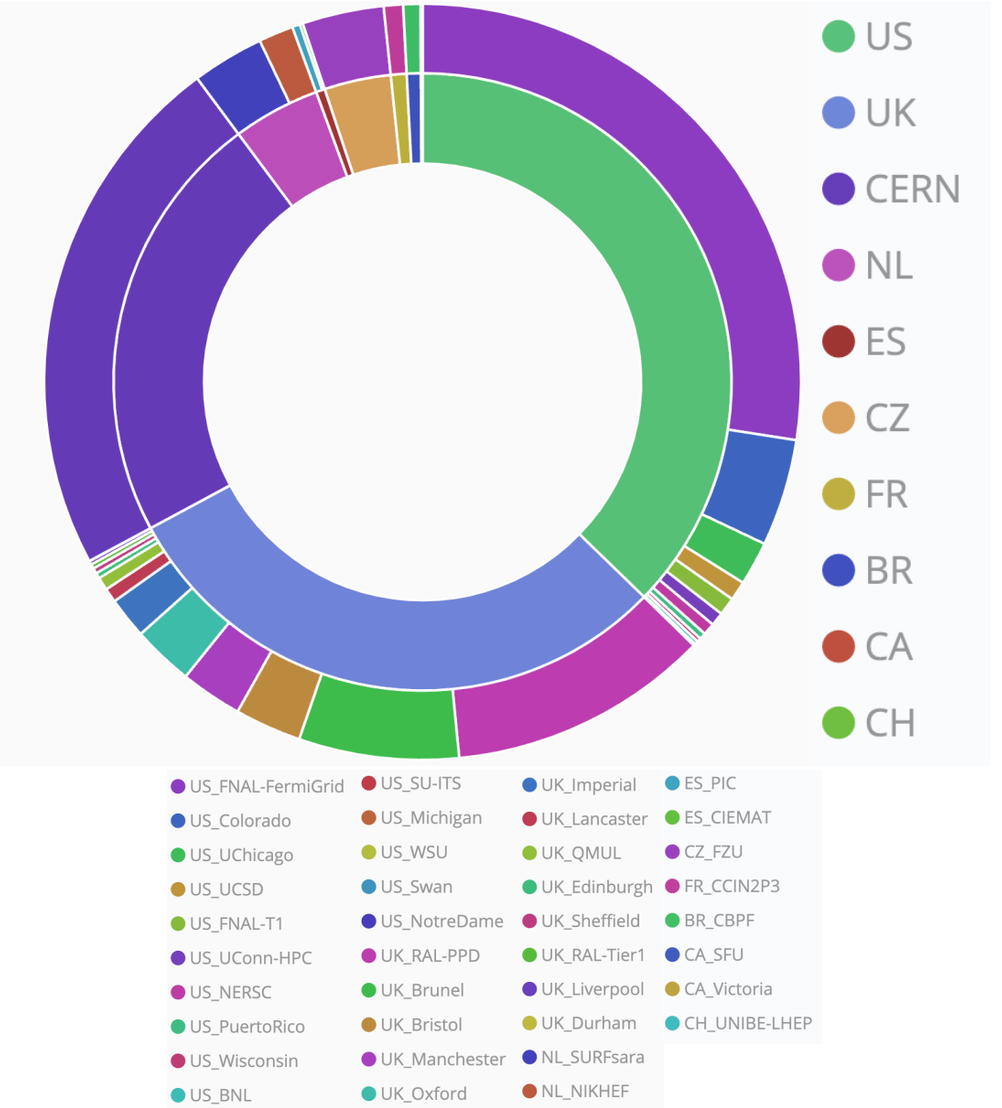

justIN Workflow Management System

- justIN ties together MetaCat, Rucio, and GlideInWMS
  - runs jobscript on dataset specified with an MQL metadata query
  - directs jobs to the optimal sites and handles all the Rucio storage operations
  - justIN successfully tested during WLCG Data Challenge 2024
  - v1.0.1 is now the basis of official DUNE Productions
- Andrew McNab's poster 402 Wed poster session

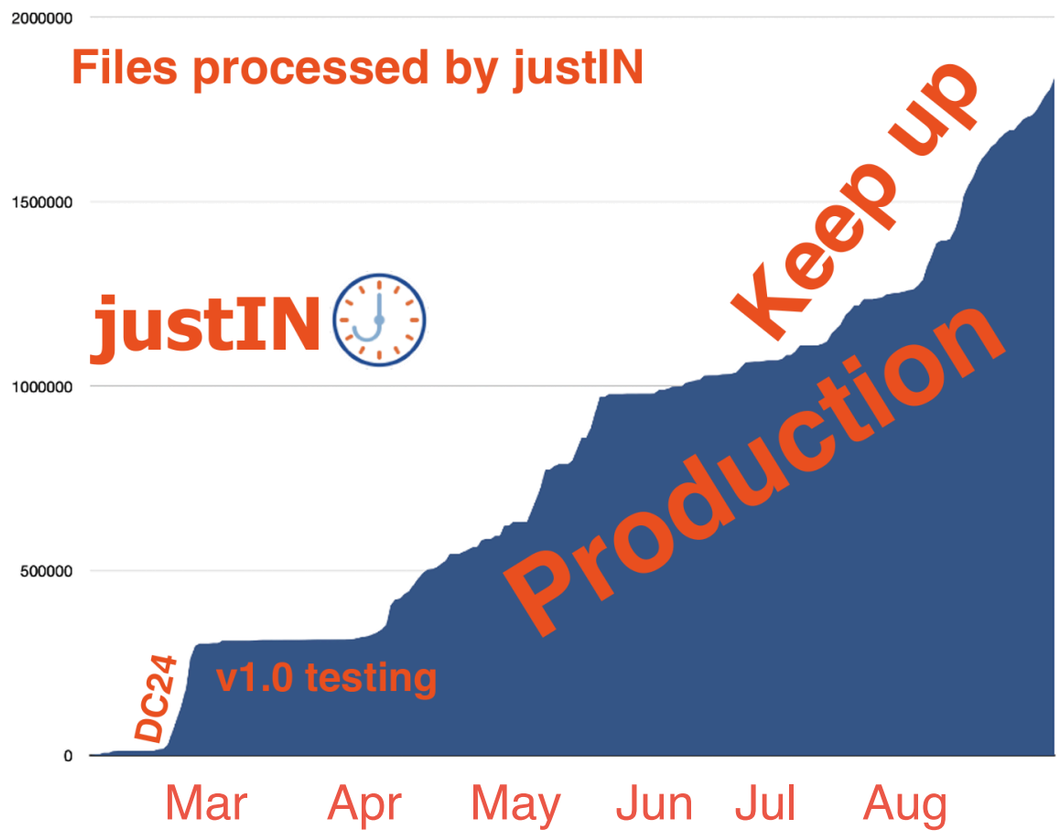

HD ProtoDUNE keep up processing

- beam ran on-and-off 06/19 – 09/16
- automated submission (twice a day)
- 650+ TB of reco files have been produced (reco2 stage)
- need to resubmit for times where site issue occurred
- experience will inform design of production systems to more easily automate (i.e. recovery jobs) & logging
- using new justIN workflow system developed and improving feature set
- details Jake Calcutt talk Thu Development and Testing of DUNE's Production System

Number of files processed daily

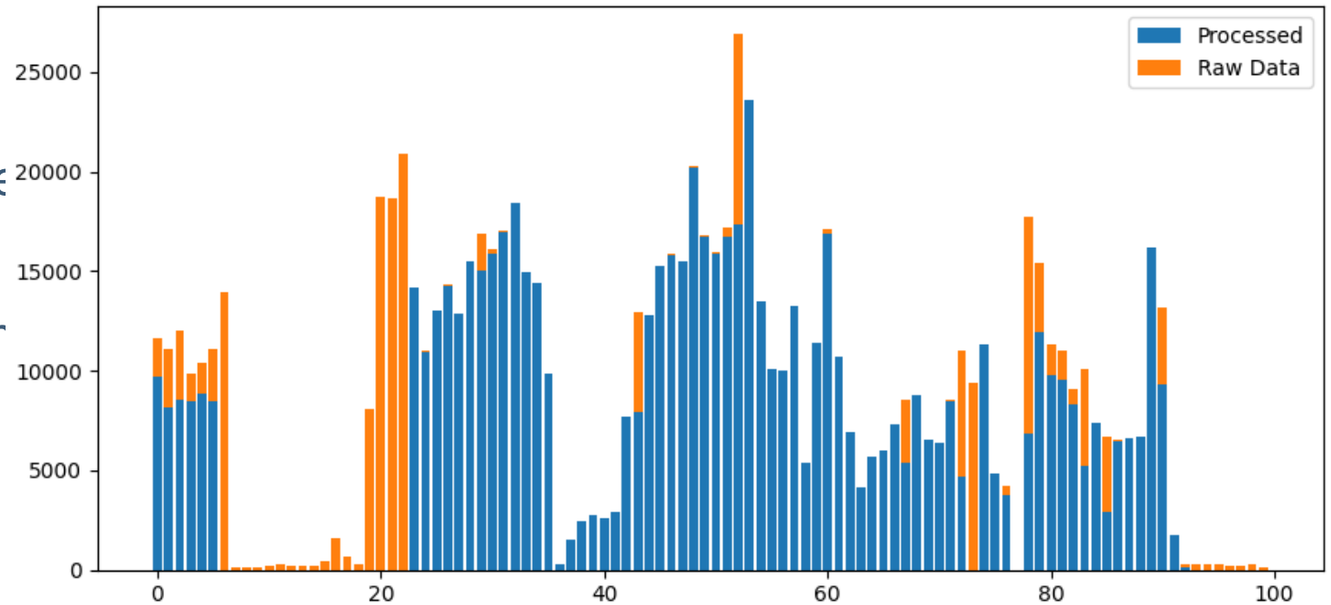

Days since start of beam

Near Detector 2x2 Demonstrator

- prototype using NuMI neutrino beam
- perform neutrino physics measurements at DUNE neutrino energy and on an Argon target
- data pipeline developed for ProtoDUNE was replicated at Fermilab for 2x2-MINERvA LAr prototype
- approx. 1 week of good data taken with (anti)neutrino beam just before July shutdown
  - 1.2 TB of Minerva chambers data
  - 13.2 TB of LAr Light readout
  - 0.6 TB of LAr charge readout.
- working with the 2x2 team to more fully integrate Data Pipeline/Rucio/Metacat into operations at NERSC

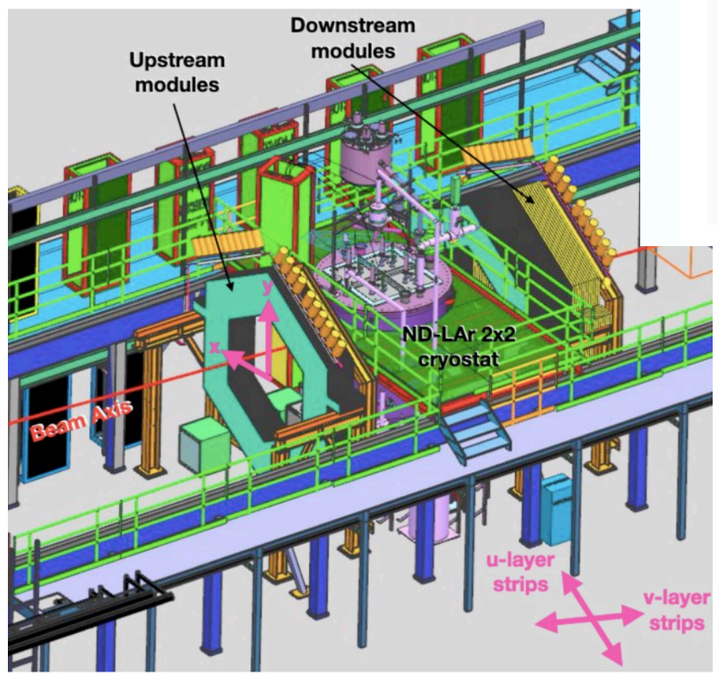

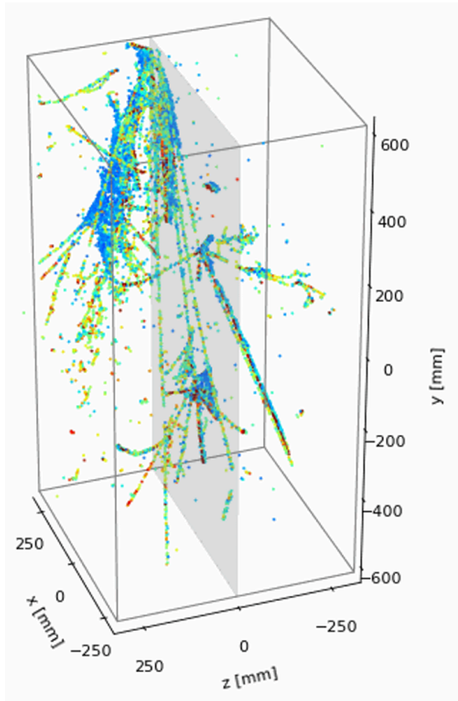

DUNE Database Development

- Conditions DB
  - primary database stores all the metadata
  - Conditions DB at FNAL
    - interfaces
      - Python API can upload and retrieve data
      - Art service can be used to read data
      - C++ interface can be used to read data
      - Metacat filter allows filter on conditions table parameters within Metacat query
    - contains metadata coming from the DAQ run configurations, IFBeam database, and soon the slow control database
  - studies to incorporate the Frontier caching system are under way
  - Nilay Bostan & Ana Paula Vizcaya Hernandez Poster Session Wed
- Hardware Database
  - critical component of detector prototyping and construction
  - iPad interface app developed and deployed for several detector components
  - additional development for iPhone and other interfaces
  - ease of use for people performing construction, QA/QC
  - traditional database backend allows for long term support and stability

Promoting Inclusive and Equitable Research for DUNE-US Computing

- establishing steering committee to develop the procedures for applications and mentor recruitment
- involving professors from many institutions in the program for a diverse group of participants
- provide training and a well-defined career path in computing for the next generation of scientist
- DUNE tracking training, mentorship, career progression, and feedback from collaboration members as DUNE progresses through the decades

- funded by the US Dept of Energy Office of Science and hosted at Rice University

DUNE Core Software & Computing Consortium

- DUNE Computing successfully providing computing resources and processing for the experiment
- successfully participated in WLCG Data Challenge 2024
- operations of data pipelines and production processing for multiple prototypes this year
  - extremely stable during entire beam run
  - throughput capability comparable to ATLAS
- production has been able to utilize extensive resources with justIN for MC campaigns and provide keep-up processing for PDHD
- development has begun on an updated production offline framework
- compute resources from around the world fully integrated into computing model and meeting our current needs

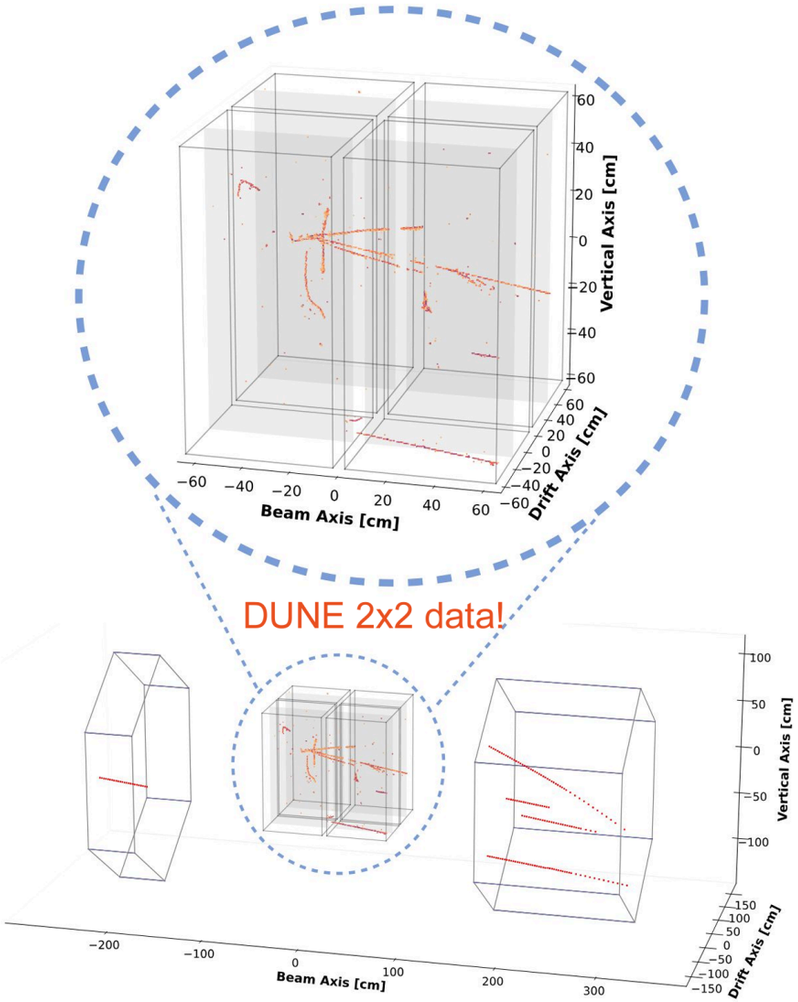

Backup Slides

Computing Model Projections

- Knowledge gained from prototype operations has helped refine the computing model projections
- Additional monitoring development is seen as a key component to matching projections with actual deployment
- working to incorporate accurate model for GPU and CPU usage, especially for ND
- additional data from ProtoDUNE-HD forced re-evaluation of data distribution model short term, but no long-term change yet

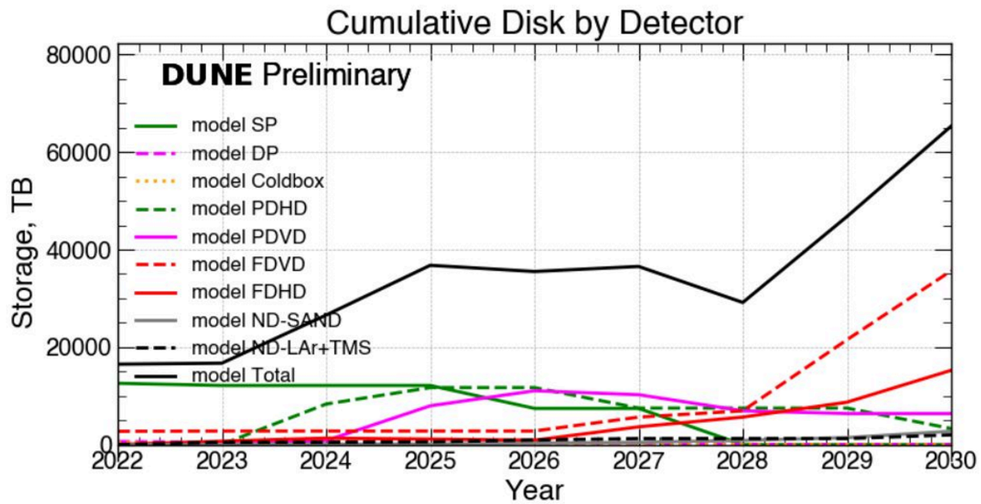

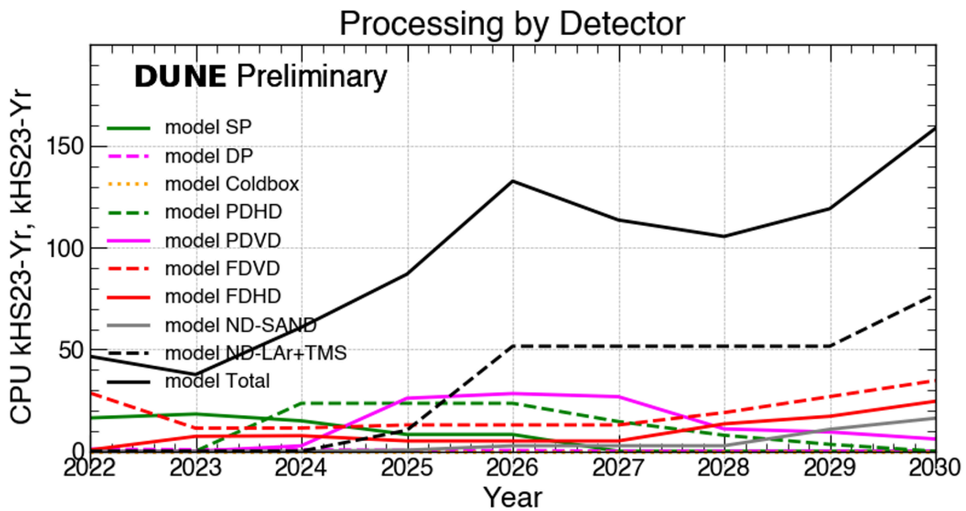

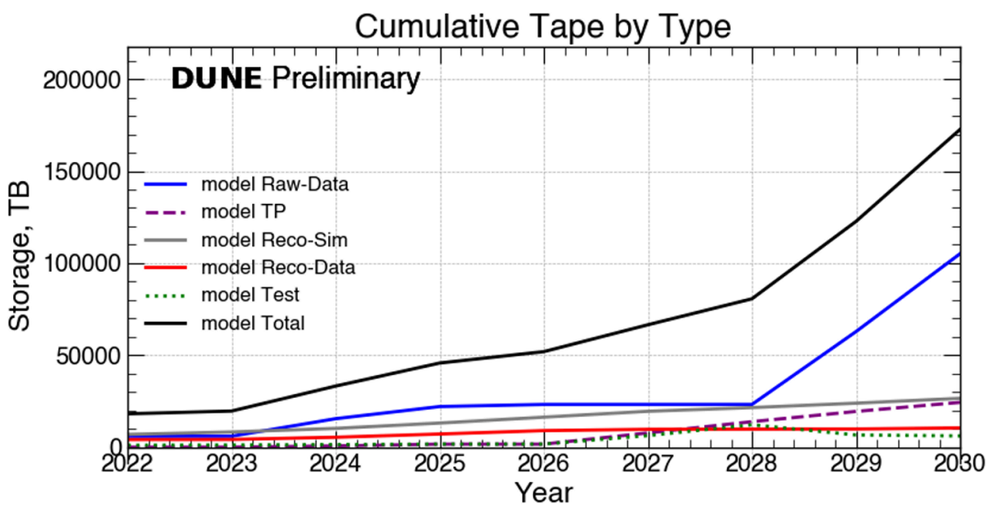

DUNE Core Software & Computing Consortium

- DUNE Computing successfully providing computing resources and processing for the experiment
- successfully participated in WLCG Data Challenge 2024
- operations of data pipelines and production processing for multiple prototypes this year
- development has begun on an updated production offline framework
- compute resources from around the world fully integrated into computing model and meeting our current needs

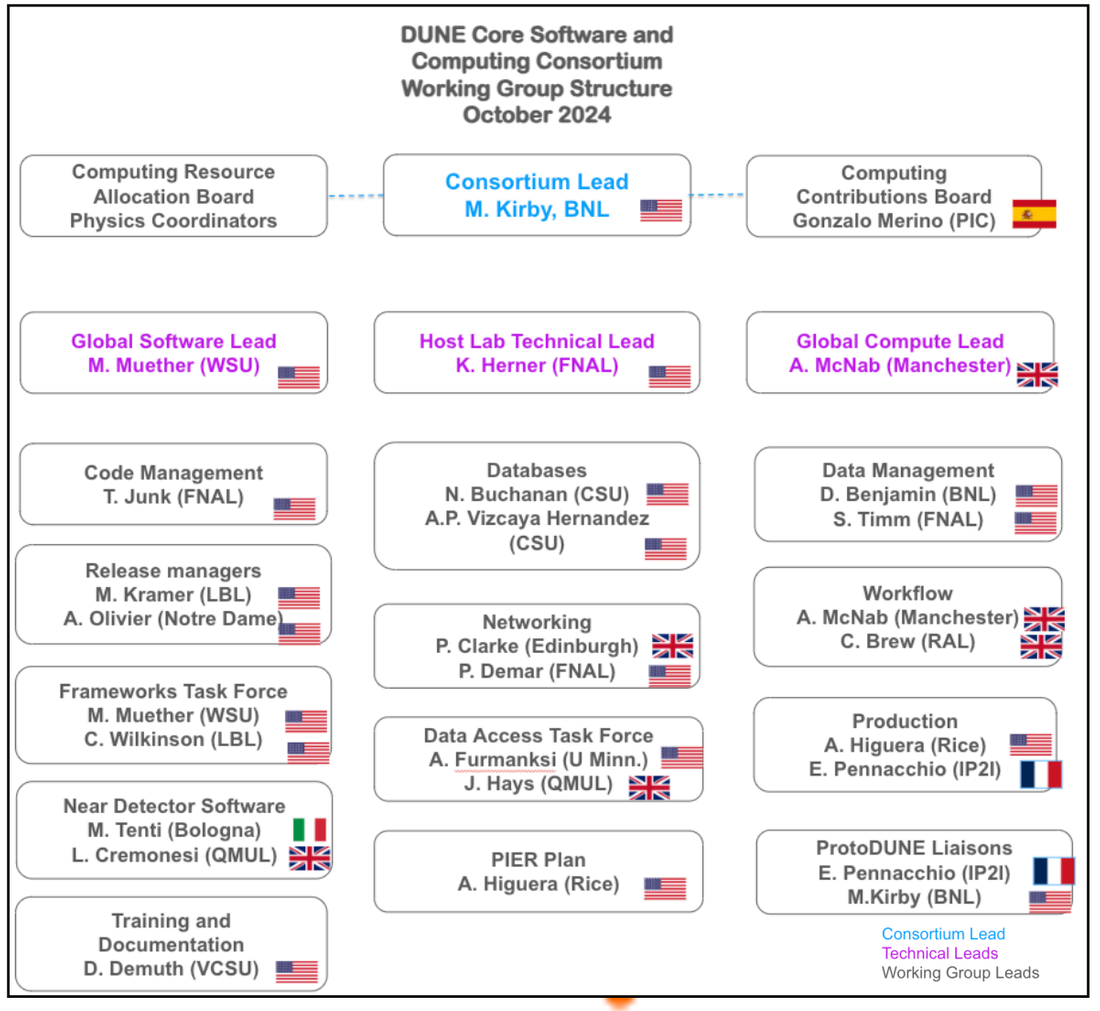

DUNE WLCG Data Challenge 2024

"FD" Raw Data
archival storage

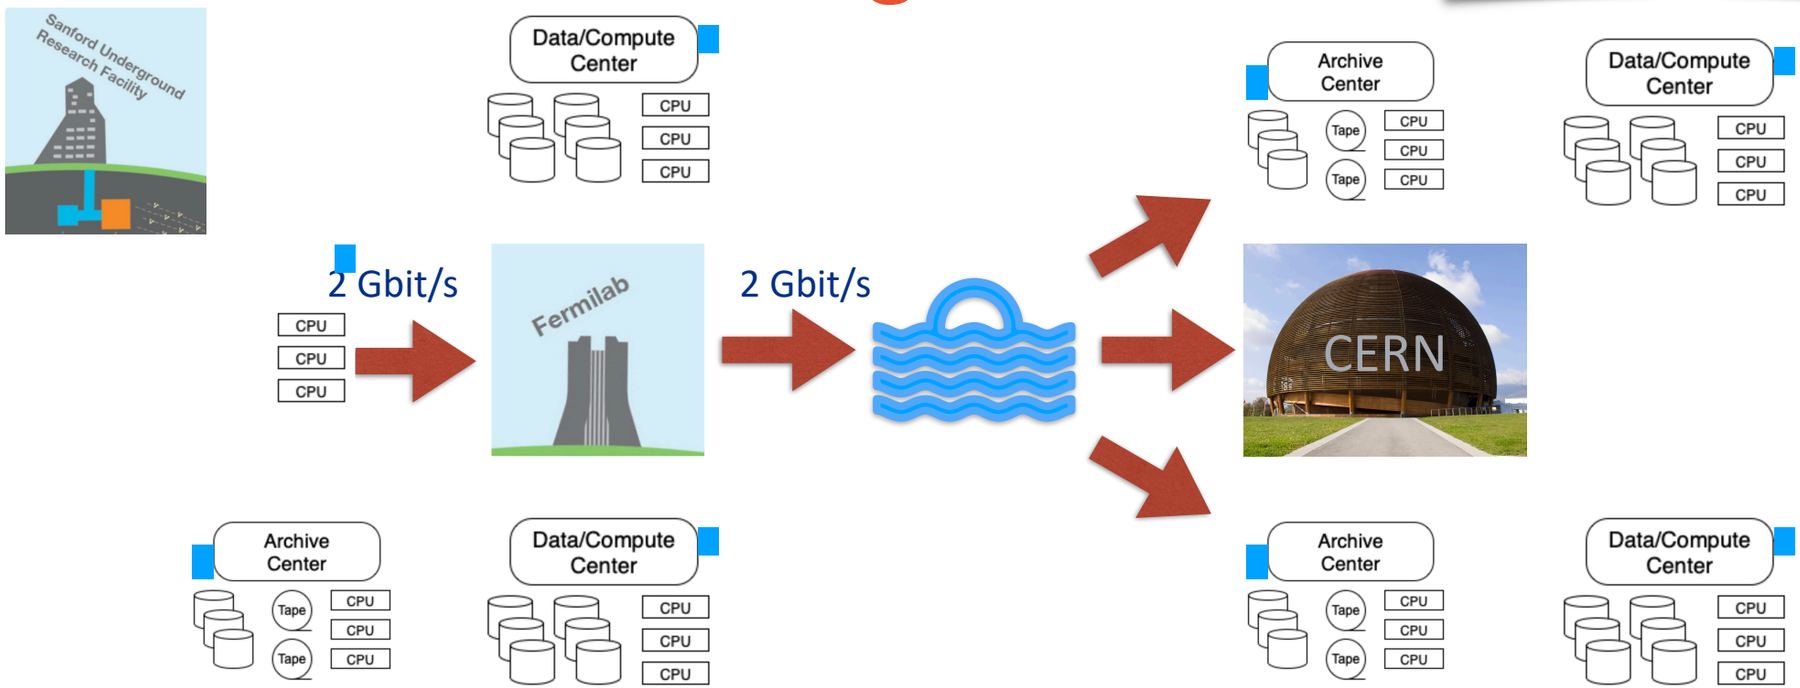

not to scale, not a technical design, it's just a cartoon

- Simulate the archival of 25% of the raw data rate from the Far Detector
  - translates to 2 Gbit/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

DUNE WLCG Data Challenge 2024

"FD" Raw Data
raw processing

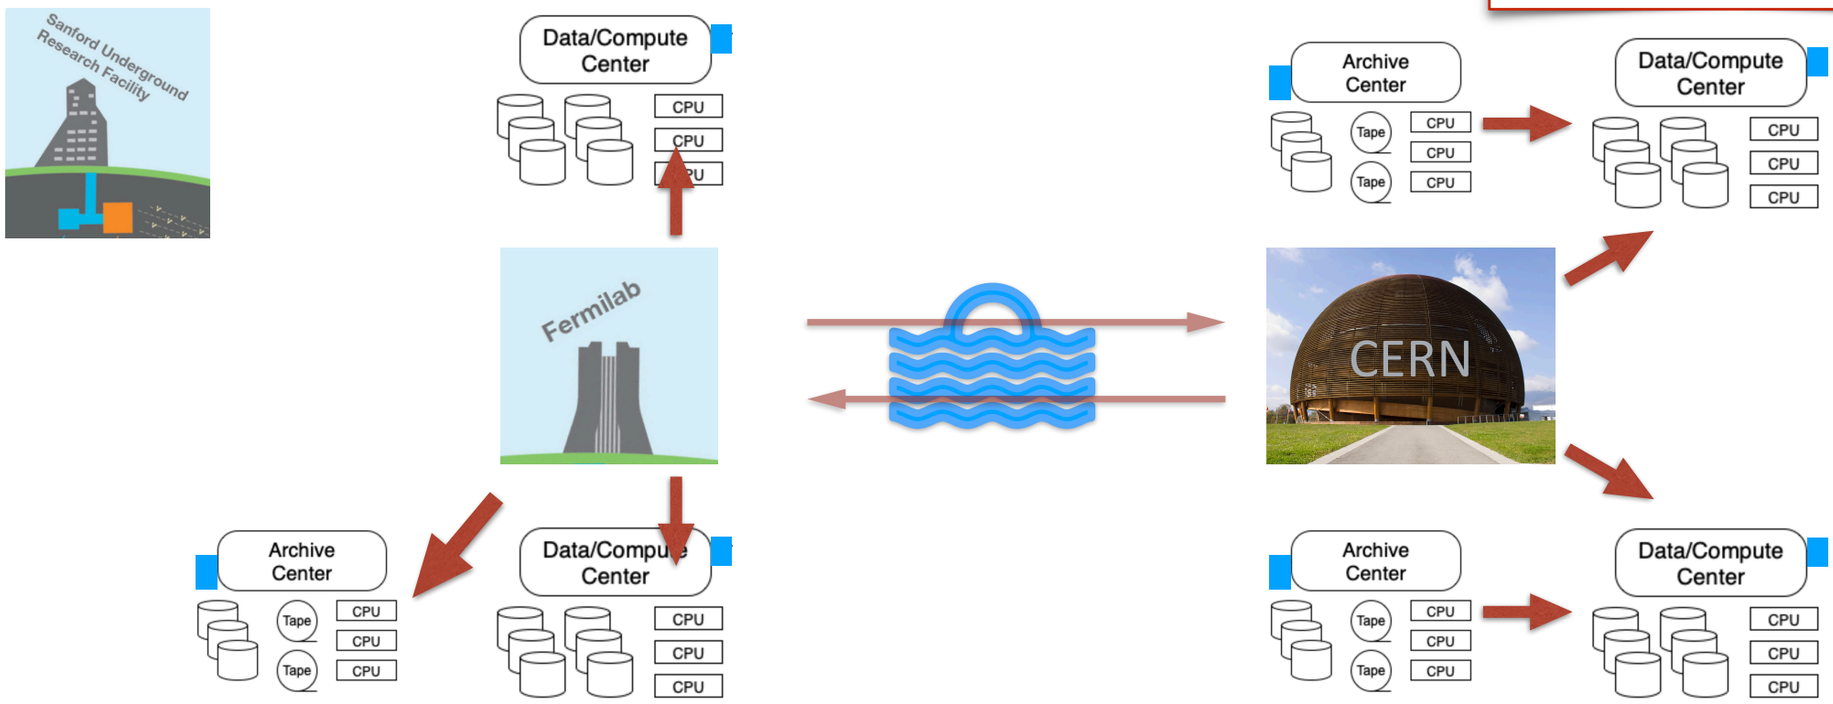

not to scale, not a technical design it's just a cartoon

- 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

Data Management and Data Pipeline from NP04 (PDHD)

- Pipeline utilizes sequences of tools (FTS3, Metacat, Rucio, and custom ingest/declaration daemons)
- Early estimates anticipated 2-3PB of beam data to be written
- Aug 10-11 consecutive days 100+ TB
- wrote 4.5PB of data
- Moved about 2PB of protodune-hd raw data to other disk sites in one week to accommodate additional data taking
- files moved to BNL, PIC, NIKHEF, FNAL, PRAGUE, RAL-PP, SURFSARA

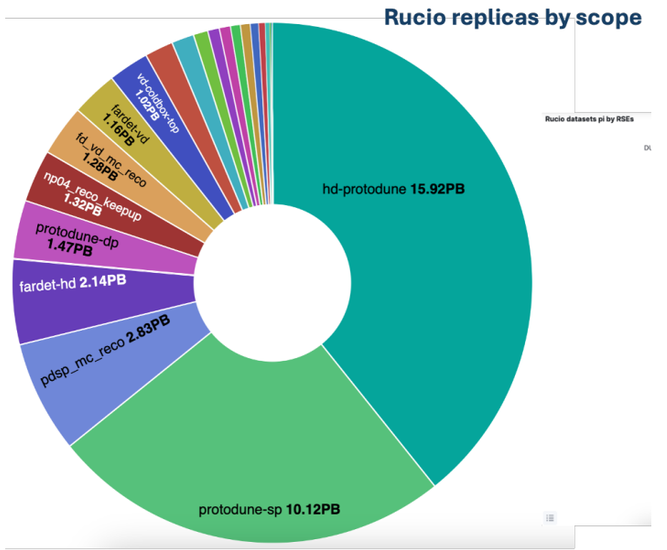

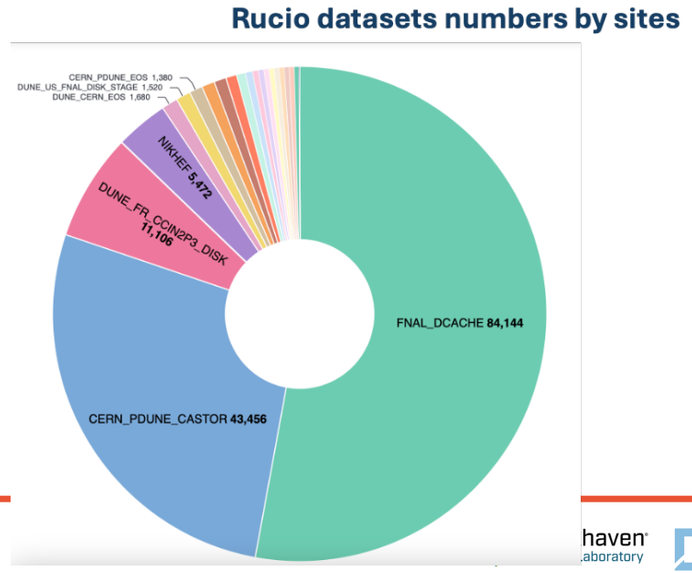

Software management Status

- Nine releases since May 2024
  - Release schedule driven by LArSoft releases
  - DUNE collaborators want new LArSoft features, move to Spack may loosen tight coupling to LArSoft sched
- LArSoft v10 (four release candidates so far)
  - a refactored geometry API and handling of pixels
  - validated for FD1 and FD2 workflow
  - Other experiments also need to sign off
- Working to transition to Spack for build and development environment
  - Status: Two Spack releases of LArSoft available in CVMFS: v09_81_00 (Nov. 2023), and v09_90_01 (May 2024)
  - initial DUNE recipes made by Patrick Gartung for v09_81_00 , upgrade these to v09_90_01
  - Some UPS products provided by SciSoft in each LArSoft release are not present in the LArSoft Spack releases, like genie_phyopt and genie_xsec
  - adapted to StashCache for distro of these products. update appropriate environment variables

Tutorials and Documentation

• Constant improvement to the Computing Basics tutorial using Software Carpentry template
- https://dune.github.io/computing-basics/
• Hardware Database tutorial
• List of tutorials at: https://wiki.dunescience.org/wiki/Computing_tutorials
• Schedule a LArSoft Tutorial Workshop Feb 3 -7, 2025 following CERN CM
- utilize material developed by DUNE-UK groups as basis for one-week intensive tutorial/workshop
- Agenda of 8th UK LArSoft Tutorial to see previous material
- Contact Dominic Brailsford, Andrzej Szelc, or Michael Kirby for information

Monitoring development for Data Management

- Rucio monitoring
  - Monitoring files/replicas transfers
  - Monitoring Storage Element allocation, usage
- Edinburgh developing a monitoring/alert system
  - protoDUNE data pipeline monitoring
    - Monitoring FTS3 transfers: e.g. EHN1 DAQ => CERN EOS disk => distributed disks/tapes
    - Monitoring Ingest daemon and Declaration daemon status
    - Monitoring storage usage of DAQ machines and Storage Elements (quota alert)
- 2x2-Minerva data pipeline monitoring will be available soon
- Computing services monitoring, to give an overview status of each services
  - Now status monitoring of IngestD, DecalD, JustINdb are available
  - Plan to add more services from MetaCat, offline db, etc
- •Alert system
  - Alerts are sending to DUNE-DATA-MGMT-MONITORING-ALERT@fnal.gov and slack #data-management-monitoring-alert
  - Currently using Grafana alerts as a temporary solution, will be fully functional once migrated to OpenSearch alerts.

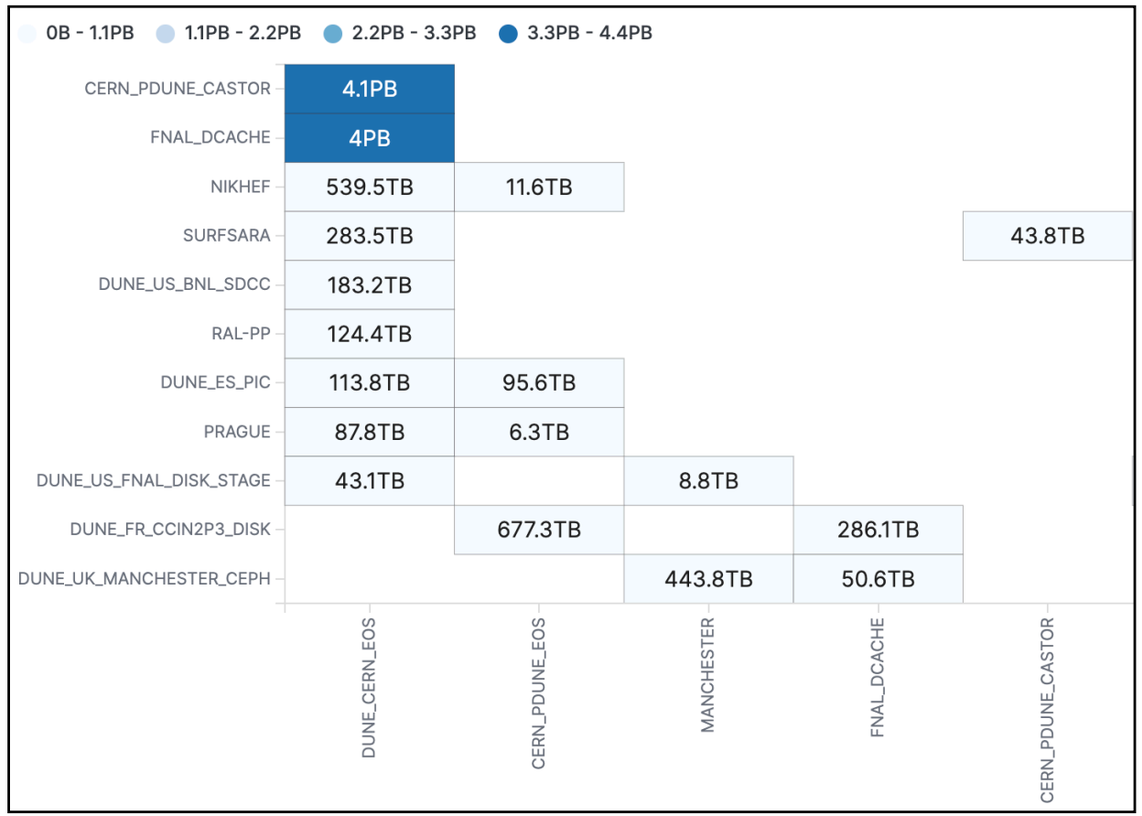

File Transfer Map

Scientific Linux 7 -> Alma Linux 9 Transition

- •Bare SL7 is no longer permitted at Fermilab – EOL Summer 2024: Now AL9
- •Transition to Spack is not finished
  - –Marc Paterno is new CSAID Spack team lead
  - –Kyle Knoepfel is working on MPD, an MRB replacement
  - –Status: Two Spack releases of LArSoft available in CVMFS: v09_81_00 (Nov. 2023), and v09_90_01 (May 2024)
  - –initial DUNE recipes made by Patrick Gartung for v09_81_00 , upgrade these to v09_90_01
  - –Some UPS products provided by SciSoft in each LArSoft release are not present in the LArSoft Spack releases, like genie_phyopt and genie_xsec
  - –Products like these, and dune_pardata, have no source repositories and are not "built". They have tarballs at scisoft.fnal.gov though
  - – adapted to StashCache for distro of these products. update appropriate environment variables

DUNE Software R&D - Framework

- Highest priority DUNE-US S&C deliverable is the development of a new software framework dedicated to DUNE use-case.
  - Deal with disparate data types with varying granularity:
- ND/FD, Beam/HE/LE, Prototype data
  - Support a wide variety of users
  - Support a variety of existing sim+reco algorithms
  - Utilize modern computing facilities efficiently – grid/HPC
  - Include Core services, e.g. Persistency, Accelerators,...
  - Be supported for the lifetime of DUNE

- ●Taskforce set up in October 2023 to explore framework-related needs of DUNE
  - Requirements document v0 completed January 11, 2024
  - Iteration between this team and the TF since early February to digest and clarify requirements

DUNE Software R&D - Framework Interfaces

- **Persistence**
  - The new software framework will enable DUNE to process data with **varying granularity**.
  - To do so efficiently, it needs **persistence infrastructure** that can read, write and store data according to that granularity.
  - New Input/Output infrastructure development led by DUNE-US
- Since the new framework is targeting the full lifetime of DUNE, it's Persistence must include modern ROOT container such as **RNTuple**
  - Some DUNE experts participated in HEP-CCE Review of RNTuple API
  - HEP-CCE is developing RNTuple backend for CAF
- **Accelerators**
  - The new software framework will enable DUNE to process data on **HPC**.
  - Most HPC use Accelerators, such as GPU, to provide majority of processing cycles
    - Since GPUs are more power efficient
  - Therefore the framework needs to provide transparent Accelerator Integration
    - Instructions/portability
    - And data model

Quick Reminder about DUNE