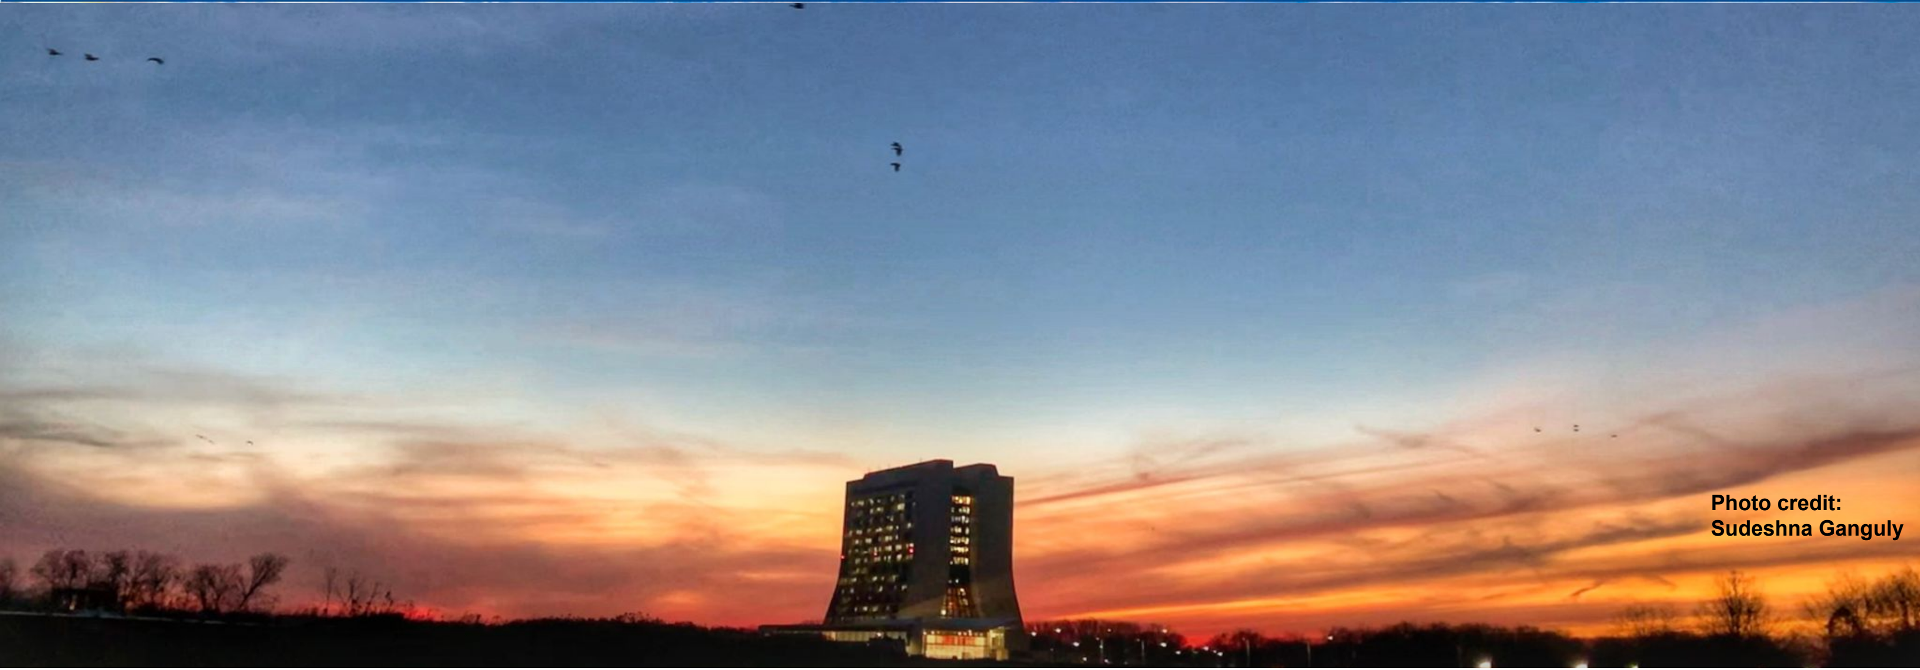

Fermilab Computing Status

Meghna Bhattacharya (FNAL)

Path to Dark Sector Discoveries at Neutrino Experiments Workshop

June 5th, 2023

The Physics Landscape at Fermilab

Fermilab hosts and participates in a wide range of scientific physics programs

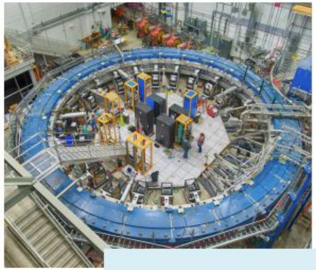
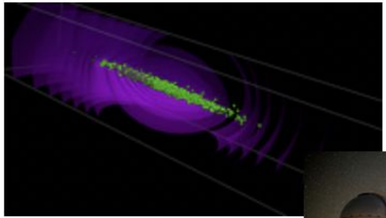


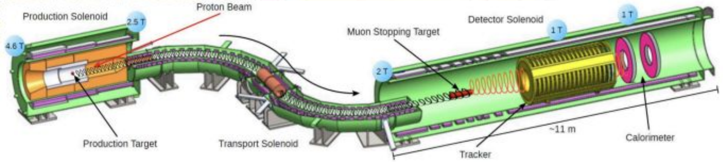


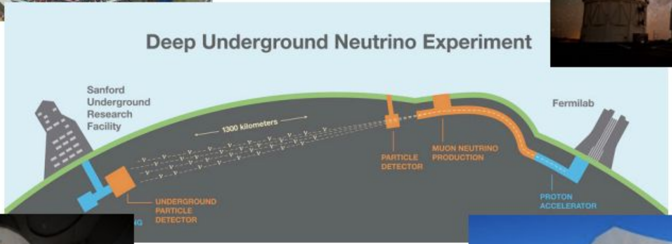



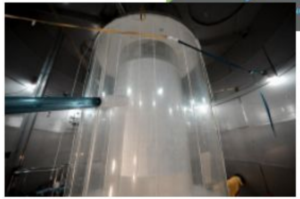
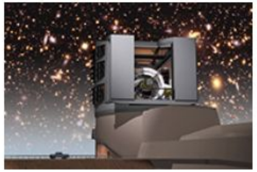
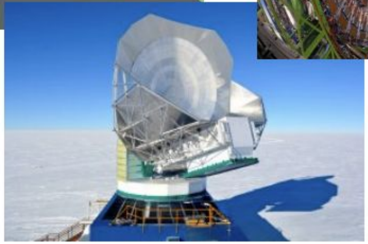


The Physics Landscape at Fermilab

Fermilab hosts and participates in a wide range of scientific physics programs

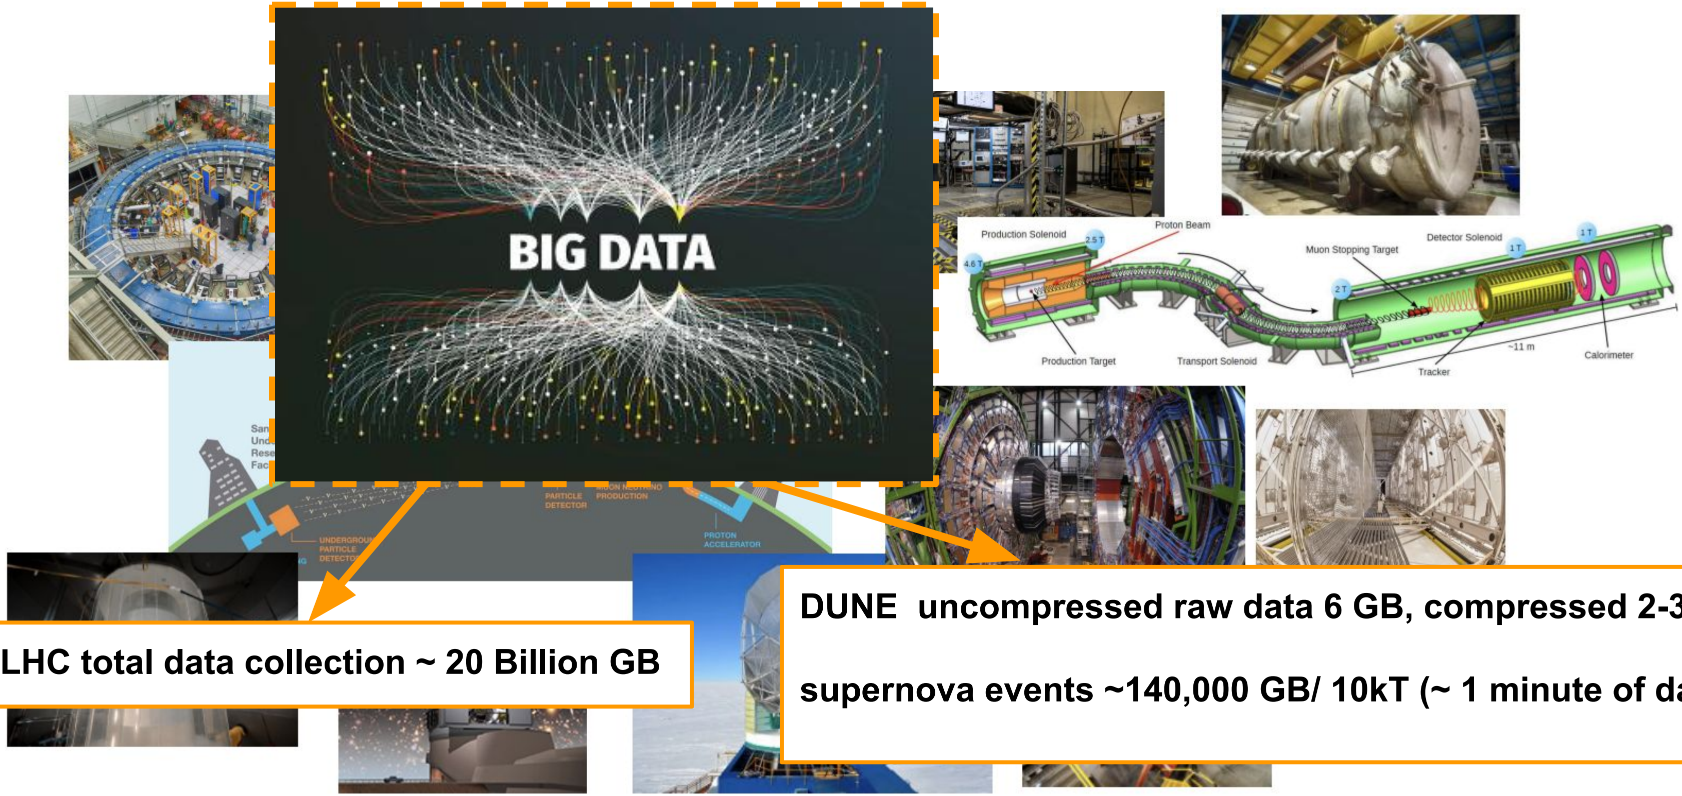

HL-LHC total data collection ~ 20 Billion GB

DUNE uncompressed raw data 6 GB, compressed 2-3 GB

supernova events ~140,000 GB/ 10kT (~ 1 minute of data)

Next Gen LArTPCs

Bigger Detectors → more neutrino interactions expected

Challenge - efficient, fast turnaround data processing to meet physics goals

LAr Mass: 1X → ~ 5X → ~ 500X

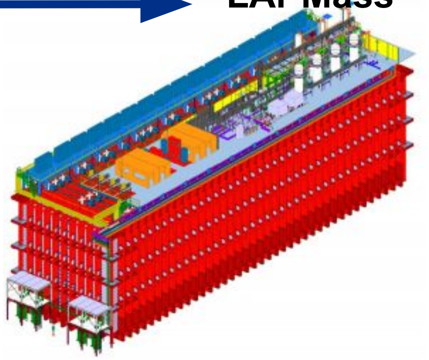

µBooNE

SBN Program

DUNE

#channels: 1X → 5X → ~100X

Data rates: 33 GB/s → 45 GB/s → 5000 GB/s

ICARUS + SBND dCache Usage and Predictions

ICARUS has 1 PB allocated at CNAF (currently 60% used)

Disk Space Used by Pool Group

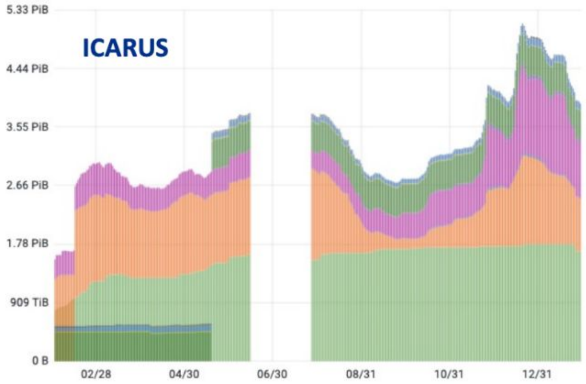

	min	max	avg	current
Fermilab Public dCache:StorageGroup:sbn_SbnDataPools	760 TiB	1.77 PiB	1.52 PiB	1.67 PiB
Fermilab Public dCache:StorageGroup:sbn_SbnAnalysisPools	1.84 TiB	1.85 TiB	1.84 TiB	1.85 TiB
Fermilab Public dCache:StorageGroup:sbn_PublicScratchPools	3.65 GiB	3.65 GiB	3.65 GiB	
Fermilab Public dCache:StorageGroup:icarus_readWritePools	147 TiB	1.37 PiB	907 TiB	806 TiB
Fermilab Public dCache:StorageGroup:icarus_SlowReadWritePools	1.02 TiB	20.2 TiB	8.24 TiB	3.52 TiB
Fermilab Public dCache:StorageGroup:icarus_SlowPublicScratchPools	3.00 TiB	29.6 TiB	7.96 TiB	6.23 TiB
Fermilab Public dCache:StorageGroup:icarus_PublicScratchPools	273 TiB	1.34 PiB	526 TiB	862 TiB
Fermilab Public dCache:StorageGroup:icarus_NovaPrestagePools	801 GiB	60.5 TiB	21.1 TiB	29.1 TiB
Fermilab Public dCache:StorageGroup:icarus_IcarusReadWritePools	427 TiB	454 TiB	450 TiB	453 TiB
Fermilab Public dCache:StorageGroup:icarus_IcarusCalibPools	8.40 TiB	15.1 TiB	12.4 TiB	8.40 TiB
Fermilab Public dCache:StorageGroup:icarus_IcarusAnalysisPools	59.6 TiB	114 TiB	81.0 TiB	112 TiB
Fermilab Public dCache:StorageGroup:icarus_HsmWritePools	810 GiB	8.60 TiB	4.51 TiB	2.60 TiB
Fermilab Public dCache:StorageGroup:icarus_HsmReadPools	194 GiB	5.47 TiB	2.84 TiB	1.24 TiB
Fermilab Public dCache:StorageGroup:icarus_Geant4ReadWritePools	1.02 TiB	20.2 TiB	8.24 TiB	3.52 TiB

Plots double-count SbnDataPools

Disk Space Used by Pool Group ⌄

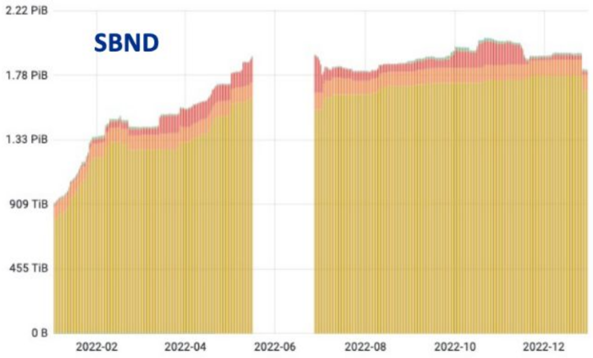

	min	max	avg	current ⌄
Fermilab Public dCache:StorageGroup:sbn_SbnDataPools	760 TiB	1.77 PiB	1.52 PiB	1.67 PiB
Fermilab Public dCache:StorageGroup:sbnd_SbndAnalysisPools	95.3 TiB	118 TiB	105 TiB	110 TiB
Fermilab Public dCache:StorageGroup:sbnd_PublicScratchPools	5.87 TiB	261 TiB	77.5 TiB	19.8 TiB
Fermilab Public dCache:StorageGroup:sbnd_readWritePools	2.46 TiB	22.8 TiB	9.87 TiB	11.2 TiB
Fermilab Public dCache:StorageGroup:sbn_SbnAnalysisPools	1.84 TiB	1.85 TiB	1.84 TiB	1.85 TiB
Fermilab Public dCache:StorageGroup:sbnd_NovaPrestagePools	20.3 GiB	771 GiB	429 GiB	713 GiB
Fermilab Public dCache:StorageGroup:sbnd_SlowPublicScratchPools	518 MiB	4.99 TiB	2.37 TiB	246 GiB
Fermilab Public dCache:StorageGroup:sbnd_Geant4ReadWritePools	23.9 GiB	159 GiB	69.5 GiB	153 GiB
Fermilab Public dCache:StorageGroup:sbnd_SlowReadWritePools	23.9 GiB	159 GiB	69.5 GiB	153 GiB
Fermilab Public dCache:StorageGroup:sbnd_HsmWritePools	264 MiB	36.5 GiB	23.4 GiB	30.9 GiB
Fermilab Public dCache:StorageGroup:sbnd_HsmReadPools	254 MiB	247 GiB	51.2 GiB	9.79 GiB
Fermilab Public dCache:StorageGroup:sbn_PublicScratchPools	3.65 GiB	3.65 GiB	3.65 GiB	

	Analysis (Persistent)	Other Dedicated (Write)
Current	1.9 PB (actual)	0.8 PB (actual)
2023	3.3 PB	1 PB
2024	5.7±1 PB	2 PB
2025	5.9±1 PB	2 PB
2026	6.2±1 PB	2 PB
2027	7.3±1 PB	2 PB

Analysis = IcarusAnalysis + SbndAnalysis + SbnData pool
Other = readWrite pool

Fermilab

DUNE Storage and Disk Requirements

- Combination of disk cache and tape archiving
- Disk ~ 5 - 10 % of LHC experiment by end of 2030
- Tape 10-15% of LHC experiment
- 2 physically separate copies of raw data

Disk currently in use/available
FNAL ~ 11/11 PB
Other ~ 7.6/12.5 PB

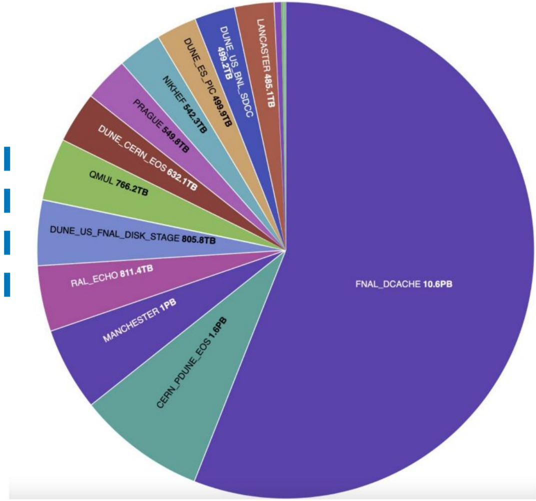

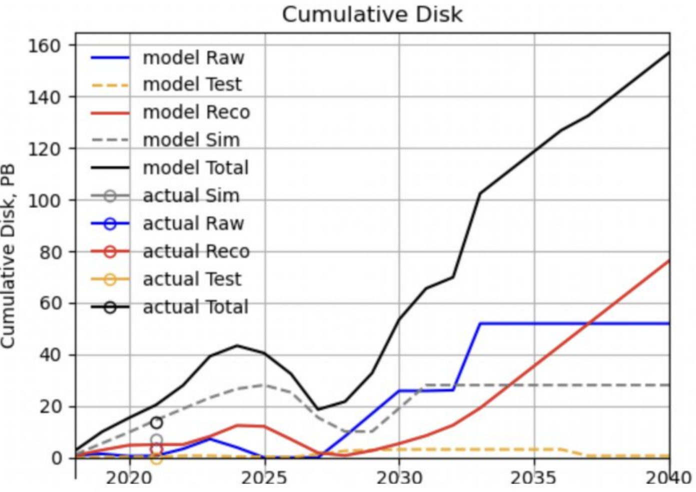

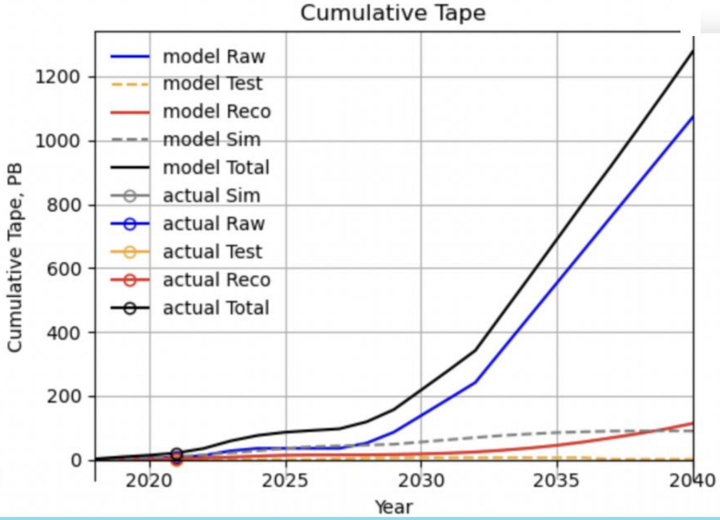

[DUNE Computing CDR]

Fermilab

Fermilab Resources Plans - Storage and Disk

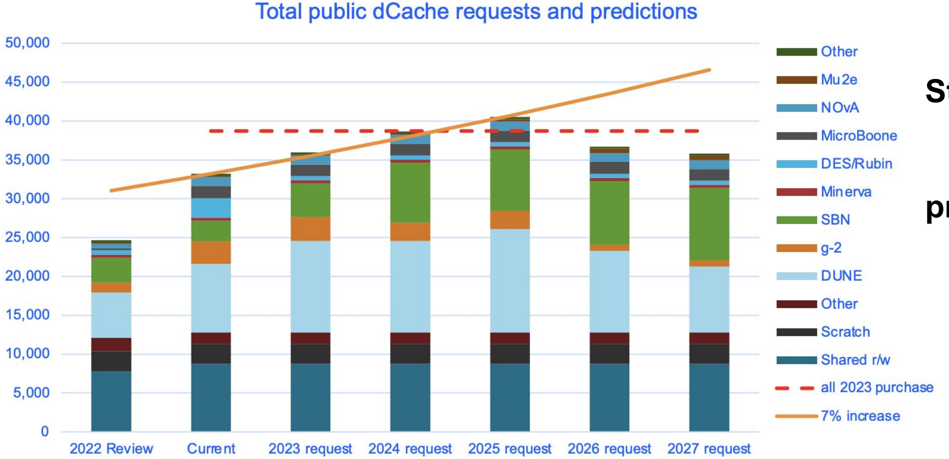

Storage service –

- **Bulk disk**
  dCache, Lustre (Wilson Cluster/LQCD), Ceph (not yet in production)
- **Tape/archival storage**
  Enstore, CTA (not yet in production)

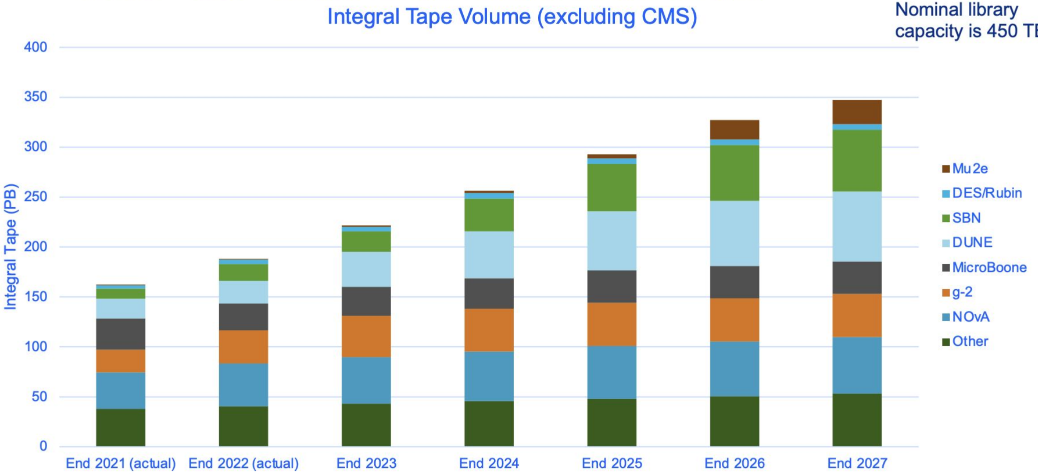

Nominal library capacity is 450 TB

New Tape Library

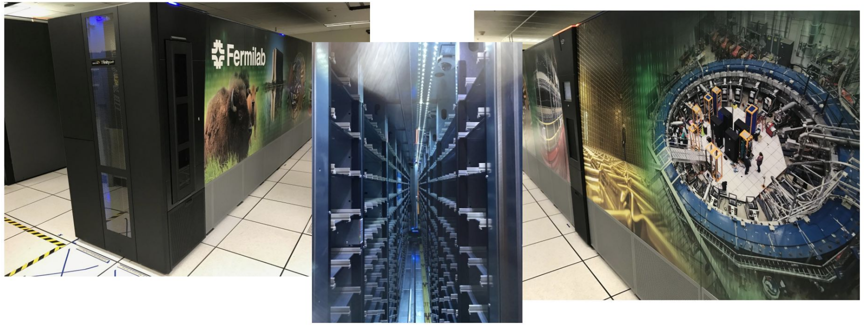

Fermilab

DUNE - Future Flagship Experiment @ Fermilab

- **One of the major physics goals**
  - **study rare (off-beam) events at Far Detector**

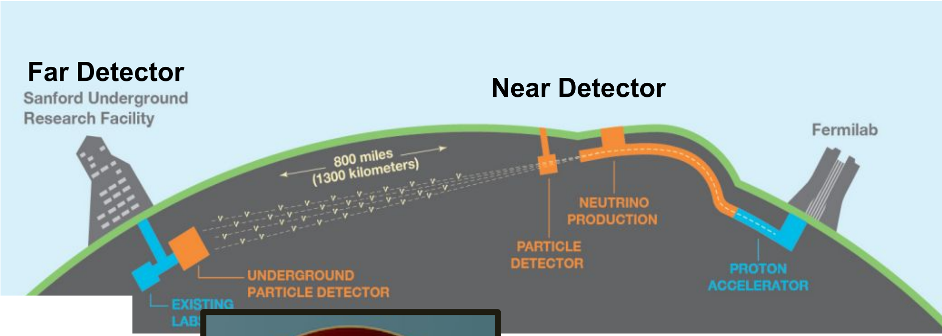

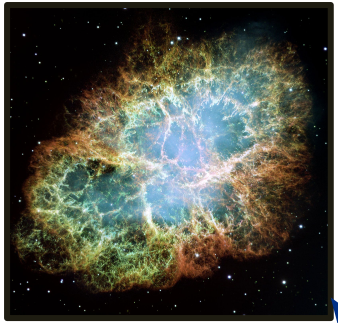

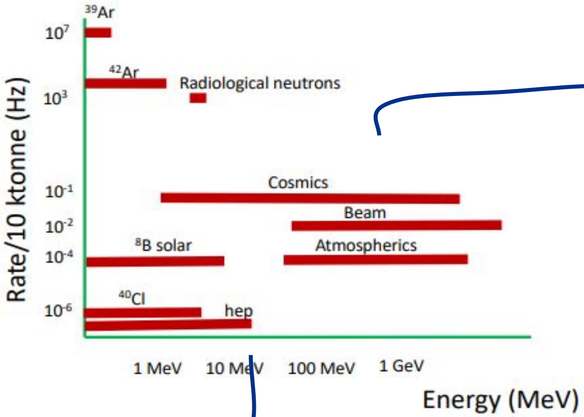

Figure 7.2: Expected physics-related activity rates in a single 10 kt module.

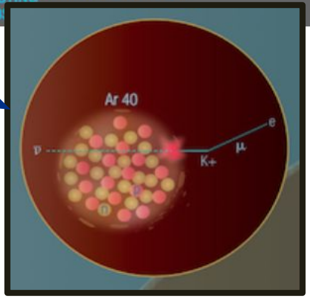

Expected data rate ~1.15 TB/s/ 10kT

Requirement : efficient and continuous data processing

Supernova Candidates

- Sensitive to neutrinos from a (relatively) nearby supernova
- Continuous detector readout for ~ 100s
- ~ 140 TB in 100s at current compression levels/ FD module

Prelim. event reconstruction

↓

source location accurate to < 5°

↓

optical/near IR telescopes follow-up

Challenge → 4+ hours to transfer data on a dedicated 100 Gb/s network, 4 hours processing time

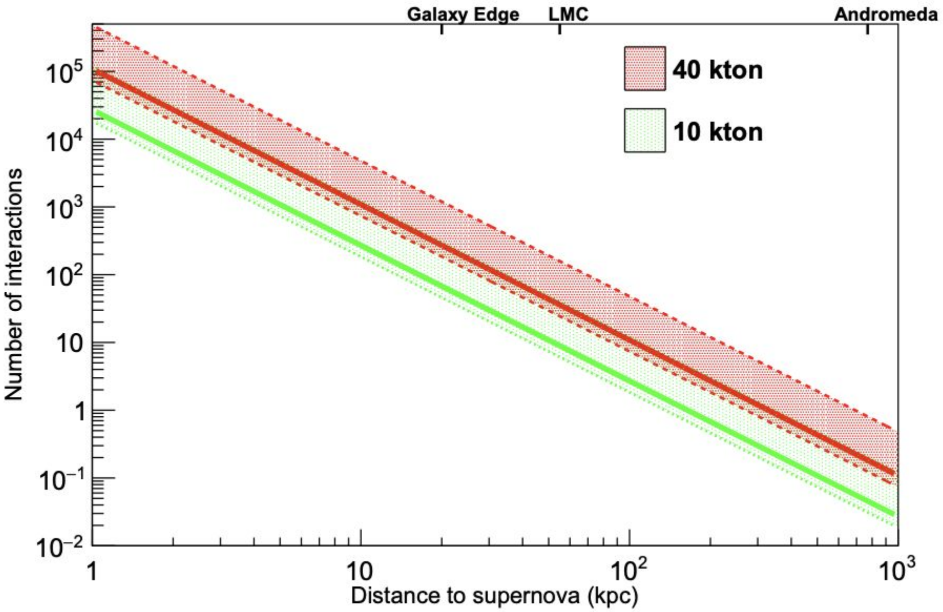

Process	Rate/module	size/instance	size/module/year
Beam event	41/day	3.8 GB	30 TB/year
Cosmic rays	4,500/day	3.8 GB	6.2 PB/year
Supernova trigger	1/month	140 TB	1.7 PB/year
Solar neutrinos	10,000/year	≤3.8 GB	35 TB/year
Calibrations	2/year	750 TB	1.5 PB/year
Total			9.4 PB/year

Supernova Candidate Processing

Cartoon version of data movement

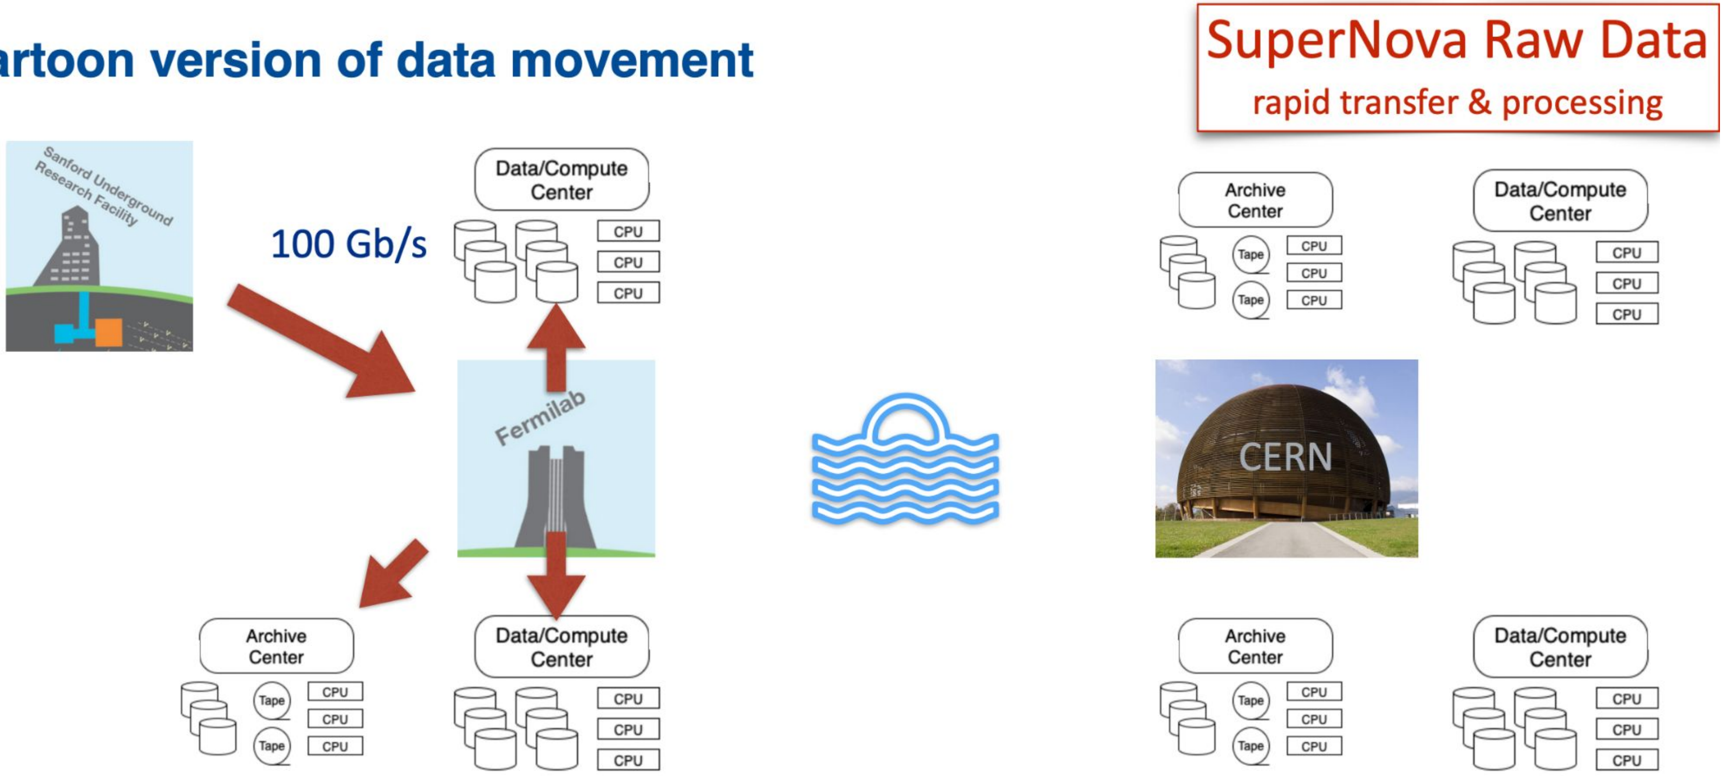

not to scale, not a technical design
it's just a cartoon

Not anticipated to be part of DC24

Supernova Candidate Processing

Cartoon version of data movement

SuperNova Raw Data
rapid transfer & processing

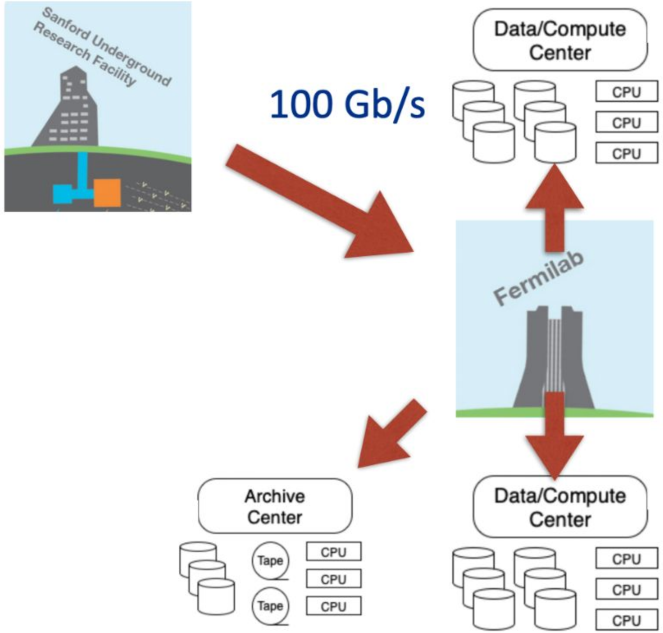

- Limited space and infrastructure (i.e. cooling) at far site → no bulk processing on local farm
- 10,000 – 40,000 present-day CPUs needed for reconstruction → 4-8 hours
  - HPC centers
    - Concern → data transfer in and out
    - Entire workflow - stitching data, output of reco - failure modes - efficiency vs. accuracy trade off
- Must be able to handle large input stream as well as output at similar rate

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

Not anticipated to be part of DC24

Supernova Candidate Processing

Supernova Neutrino Events in LArTPC

Triggers

Supernova Neutrino Interaction Simulations

Supernova Candidate Processing

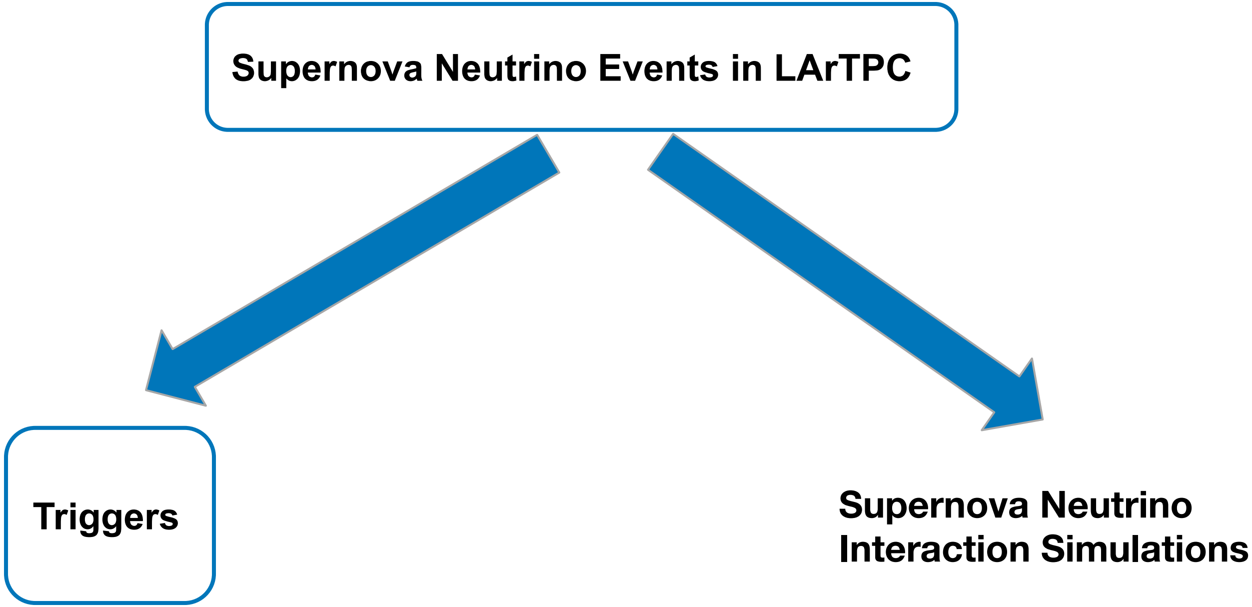

Realtime AI/ML trigger and tagging @ FNAL

- FPGA trigger implementation
  - ~$3M project to explore physics-inspired neural nets (PINNs)
  - Designing efficient edge AI with physics phenomena
  - AI applications on the “Edge” in CMS, DUNE, and accelerator physics

Event Reconstruction

Online/Real-time Triggering and Tagging

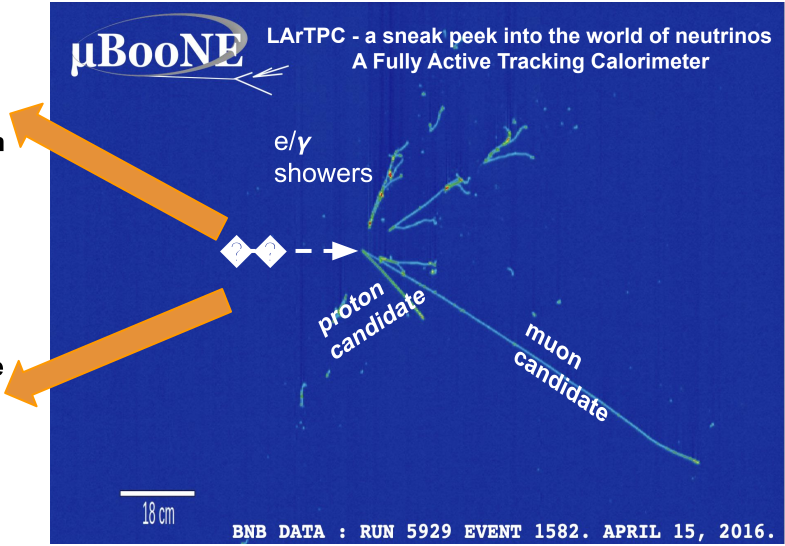

https://www.energy.gov/science/articles/department-energy-announces-64-million-artificial-intelligence-research-high

Realtime AI/ML trigger and tagging @ FNAL

- DUNE component focuses on efficient AI for identifying low-energy LArTPC interactions

Calibration

Ar39 and Michel electrons

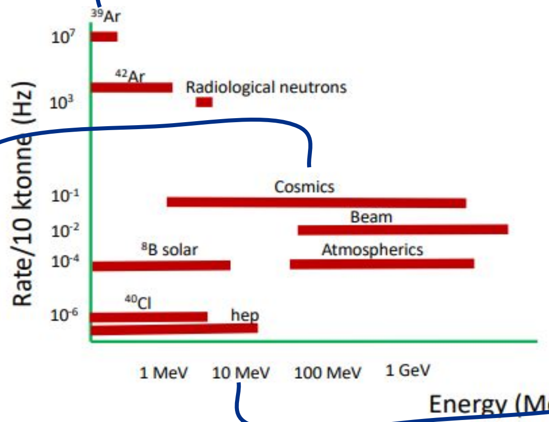

Figure 7.2: Expected physics-related activity rates in a single 10 kt module.

supernova neutrino detection

other low-energy neutrino physics

- Low power and low latencies → FPGA implementation for inference
- Involves collaborators from Columbia, UChicago, and Duke
- Goal - demonstrate proof-of-concept in ICEBERG test facility at Fermilab

Fermilab

Supernova Candidate Processing

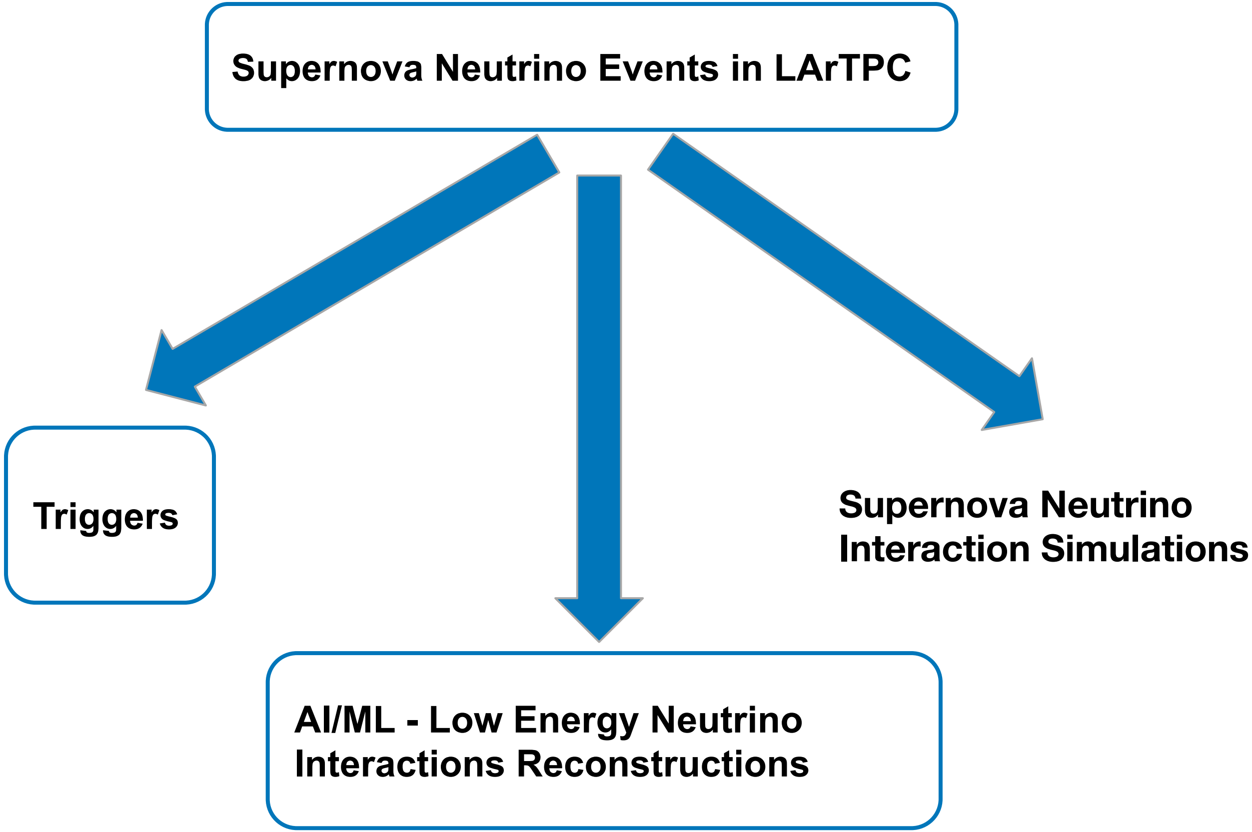

Computing Facilities - Wilson Cluster

- **Wilson Cluster available to all of FNAL either through experiments/departments or specific projects**

- **Available GPUs range from K80s (100) to A100s (4); base OS is SL7 (Alma8/9 in future)**

- **More information https://computing.fnal.gov/wilsoncluster/**

cpu_gce | cpu_gce_test | gpu_gce | gpu_gce_ppc | gpu_ose | knl_gce | Project Usage

Refresh Rate=5 mins. Last updated on: 05-22-2023 9:20 PM

SLURM job status for partition cpu_gce

JobID	Job Name	User	Account	# of Nodes	Partition	QoS	List of Nodes -or- (Reason for job queued)	Priority	GPU Resources	Start Time	Requested Walltime (DAY-HH:MM:SS)	Time Remaining (DAY-HH:MM:SS)	Job State
430111	form	garora	hpt_rad	2	cpu_gce	walltime7d	wcgwn[010-011]	64112	None	2023-05-22T09:46:30	2-02:00:00	1-14:26:25	RUNNING
430119	nick_qv	nbornman	qis_algo	1	cpu_gce	regular	wcgwn009	63465	None	2023-05-22T10:22:07	1-00:00:00	13:02:02	RUNNING

Computing Facilities - Wilson Cluster

- Wilson Cluster available to all of FNAL either through experiments/departments or specific projects

- Available GPUs range from K80s (100) to A100s (4); base OS is SL7

- More information https://computing.fnal.gov/wilsoncluster/

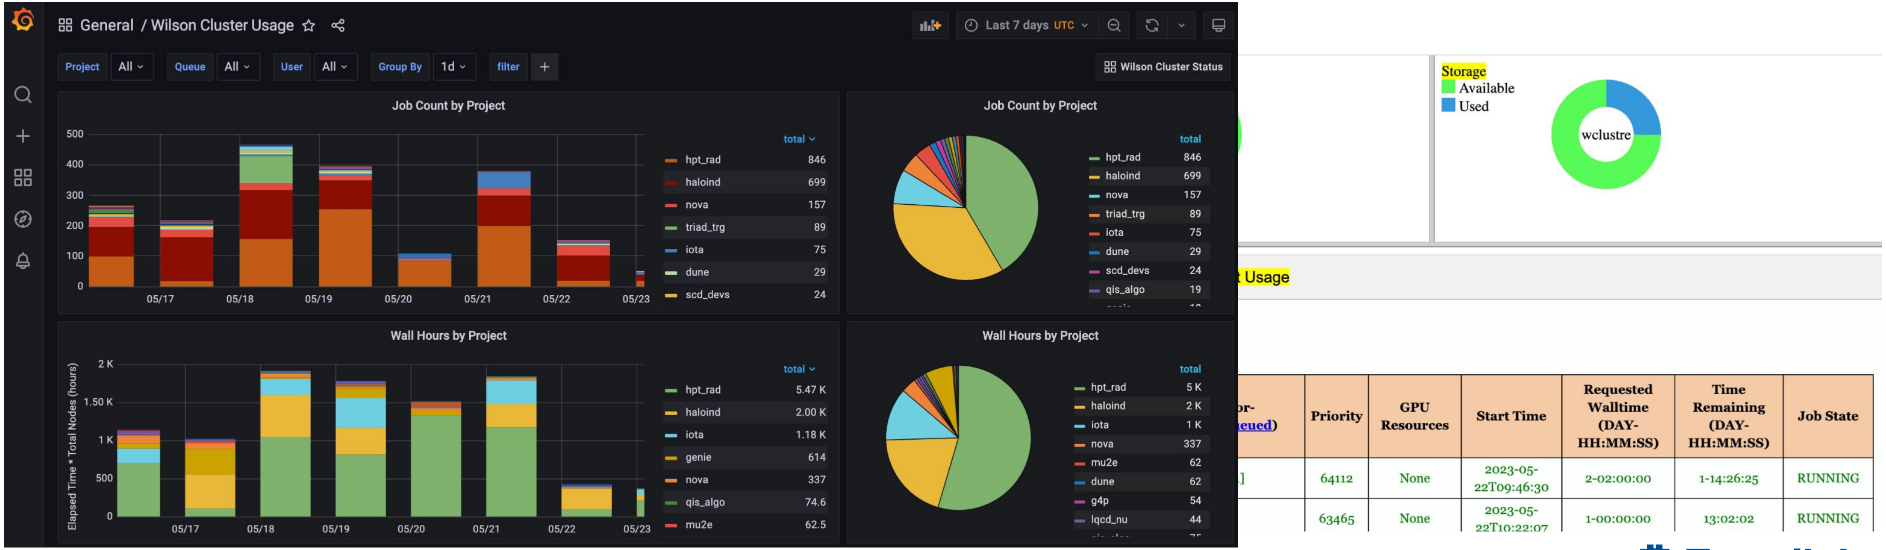

Fermilab

Computing Facilities - Elastic Analysis Facility

- **Jupyter Hub deployment with general CPU and GPU-enabled notebooks available**
- **Highly scalable, customizable, and replicable elsewhere**
- **GPUs are available through analytics-hub.fnal.gov (on-site or on VPN)**

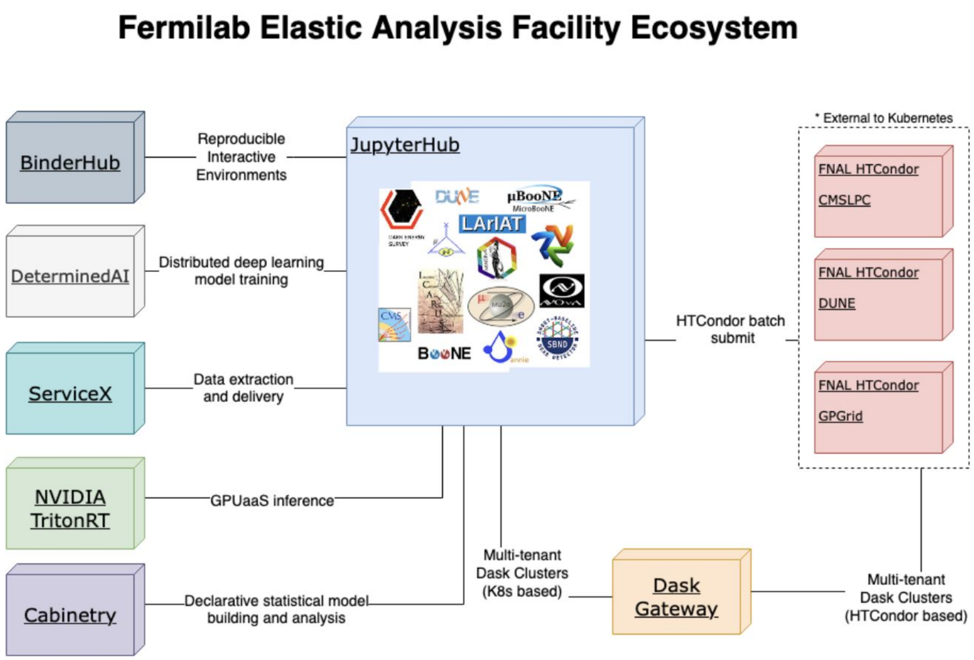

Computing Facilities - Elastic Analysis Facility

- **Jupyter Hub deployment with general CPU and GPU-enabled notebooks available**
- **Highly scalable, customizable, and replicable elsewhere**
- **GPUs are available through analytics-hub.fnal.gov (on-site or on VPN)**
- **Latest documentation is eafjupyter.readthedocs.io CVMFS also available**
- **Anyone with a services account can log in, but follow your experiment's usual instructions**
- **Streaming with xrootd also works to access larger storage elements**

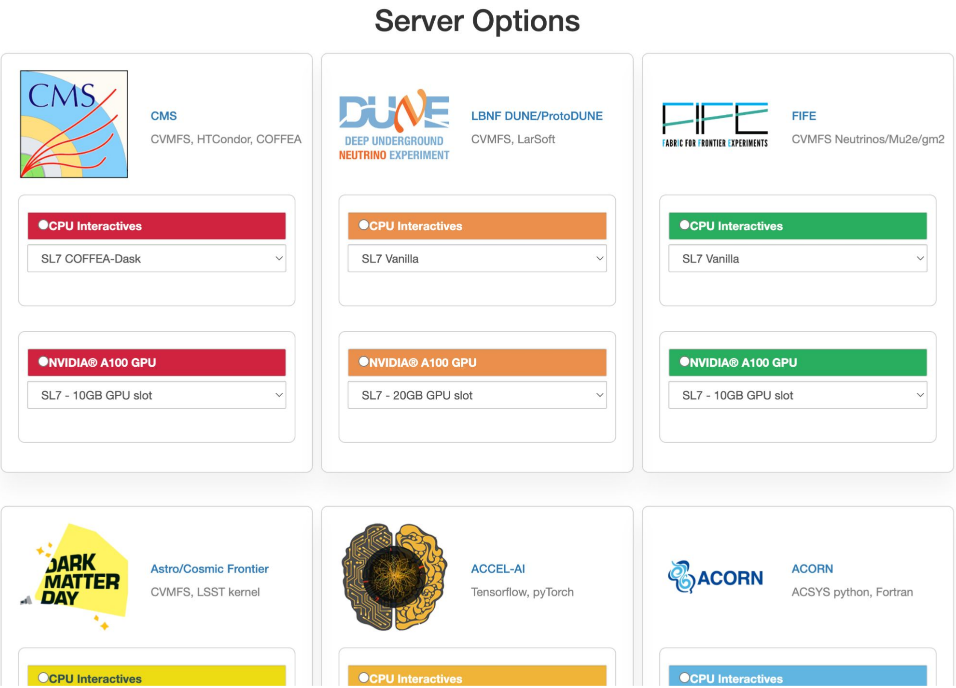

Supernova Candidate Processing

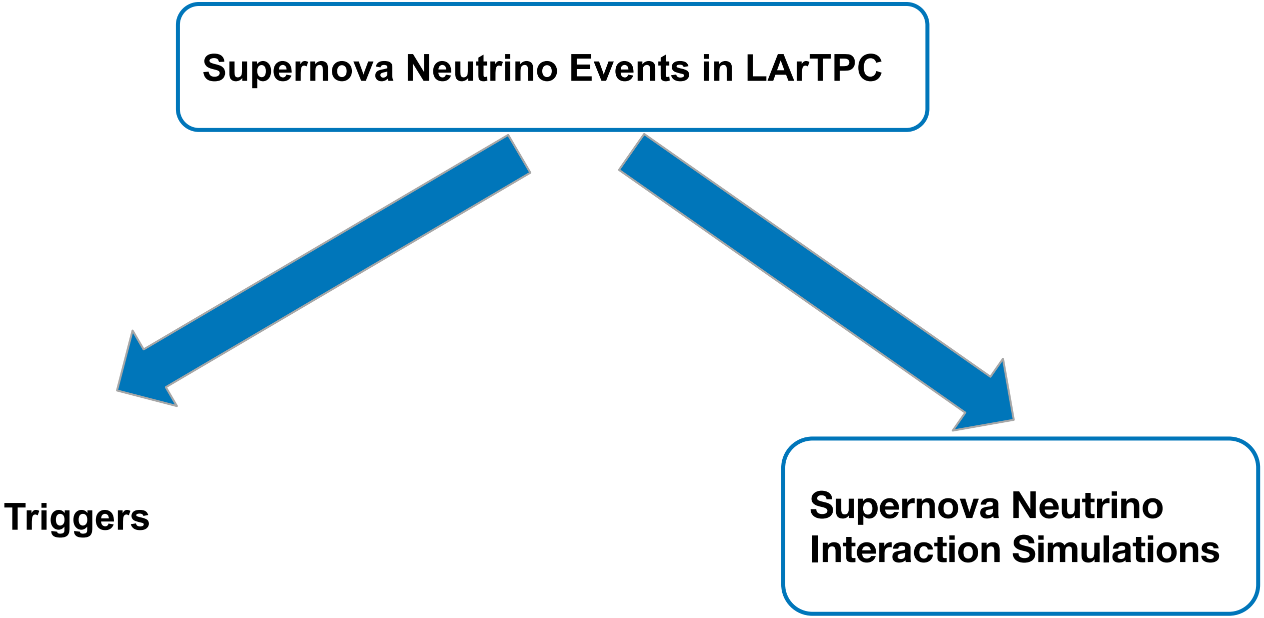

Supernova Neutrino Interaction Simulations

- **Galactic supernovae are rare (~1.6 / century)**

- **DUNE needs high-efficiency trigger → 1/month false positive rate**

- **Trigger design requires detailed interaction modeling**
  - **critical for interpreting future supernova observation [arXiv:2303.17007]**

- **Fermilab-maintained event generator: MARLEY**
  - **Physics models: [Phys. Rev. C 103, 044604 (2021)]**
  - **Implementation: [Comput. Phys. Commun. 269, 108123 (2021)]**

- **Event reconstruction techniques under development by DUNE and other experiments (e.g., MicroBooNE)**

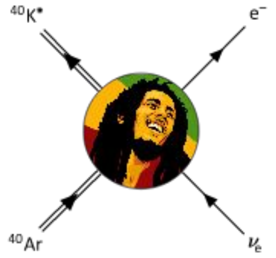

Model of **A**rgon **R**eaction
Low **E**nergy **Y**ields

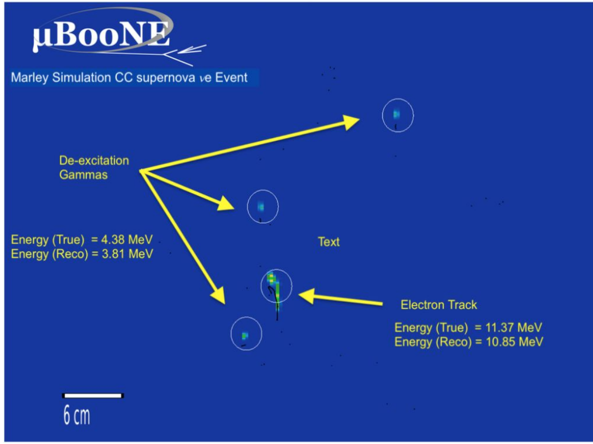

Event Generators and BSM Physics

Simulations for Neutrino Beam Experiments

Beam production

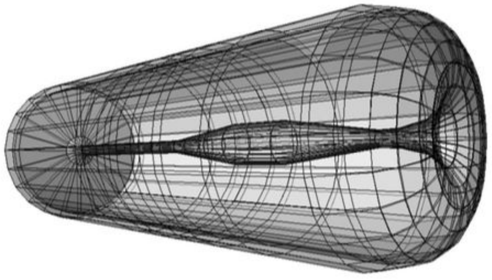

BNB horn geometry from
Phys. Rev. D 79, 072002 (2009)

Flux propagation

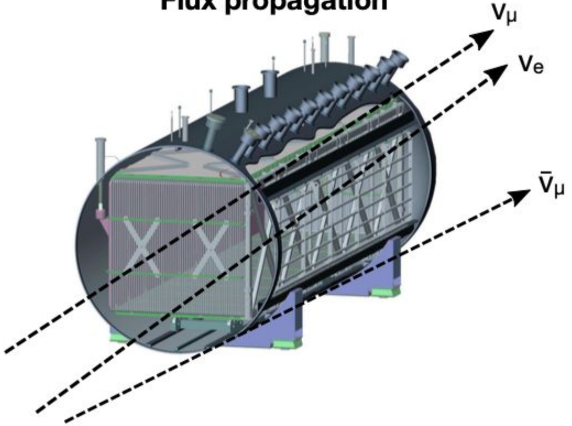

Neutrino interactions

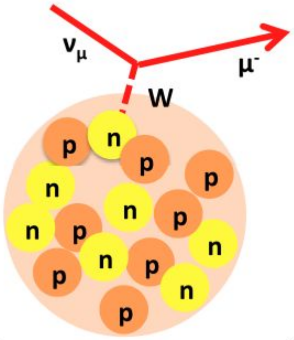

Particle transport

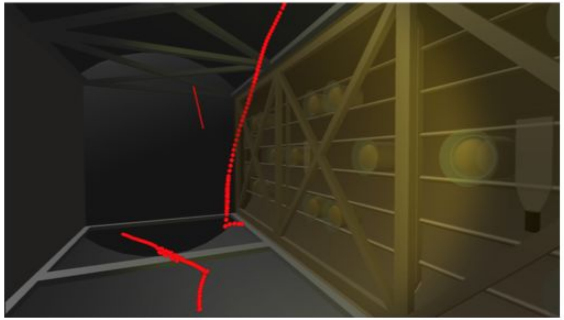

M. Del Tutto, JETP seminar May 2019

Detector response

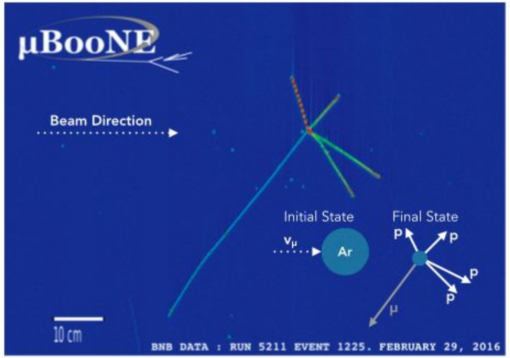

Fermilab plays a key role in developing the software tools that make this all possible

GEANT4
A SIMULATION TOOLKIT

Genie
UNIVERSAL NEUTRINO GENERATOR & GLOBAL FIT

LArSoft

The GENIE Event Generator

- **Primary beam neutrino interaction simulation for all Fermilab experiments**
  - **Developed by international collaboration**
  - **Primary focus 100 MeV to 10 GeV neutrinos, scope extends further in both directions**
- **Major contributions from Fermilab for version 3 release series**
  - **Eur. Phys. J. Spec. Top. 230, 4449 (2021)**
- **GENIE predictions key for interpreting experimental analyses**
  - **e.g., MicroBooNE's search for anomalous ν_e appearance**
- **March 2023 workshop at Fermilab discussed future development directions for GENIE and similar event generators**
  - **https://indico.fnal.gov/event/57388/**

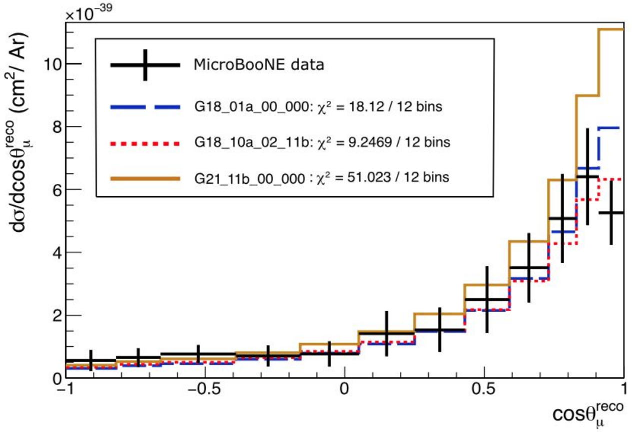

Phys. Rev. Lett. 128, 241801 (2022)

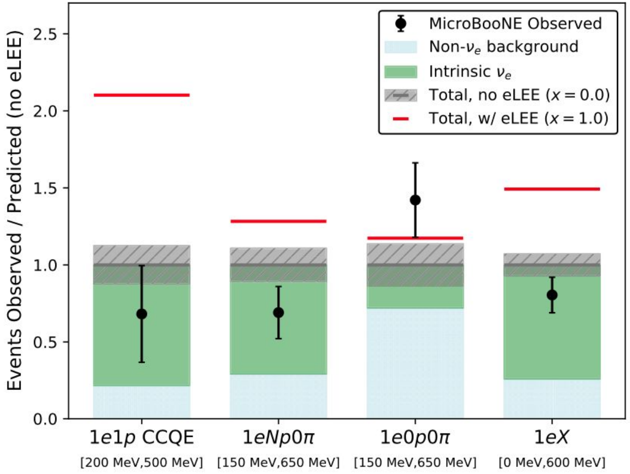

The ACHILLES Event Generator

[Phys. Rev. D 107, 033007 (2023)]

A CHIcagoLand Lepton Event Simulator

- **New theory-driven event generator**
  - **Fermilab-led**
  - **Neutrinos, electrons, BSM**

- **Technical design borrows techniques from collider physics event generators**
  - **Applies these to neutrinos for the first time**

- **Example: Automated leptonic tensor**
  - **[Phys. Rev. D 105, 096006 (2022)]**
  - **Support wide range of BSM models without dedicated development work**

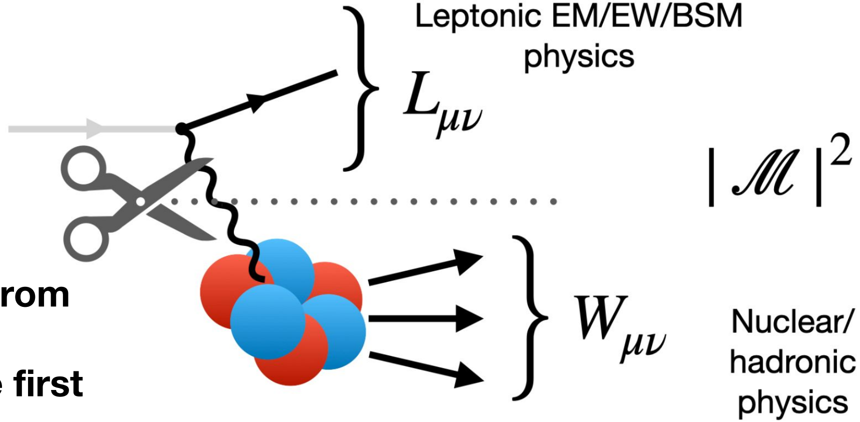

$$|\mathcal{M}|^2 \propto L_{\mu\nu} W^{\mu\nu}$$

Fermilab

ACHILLES approach to automating the leptonic tensor

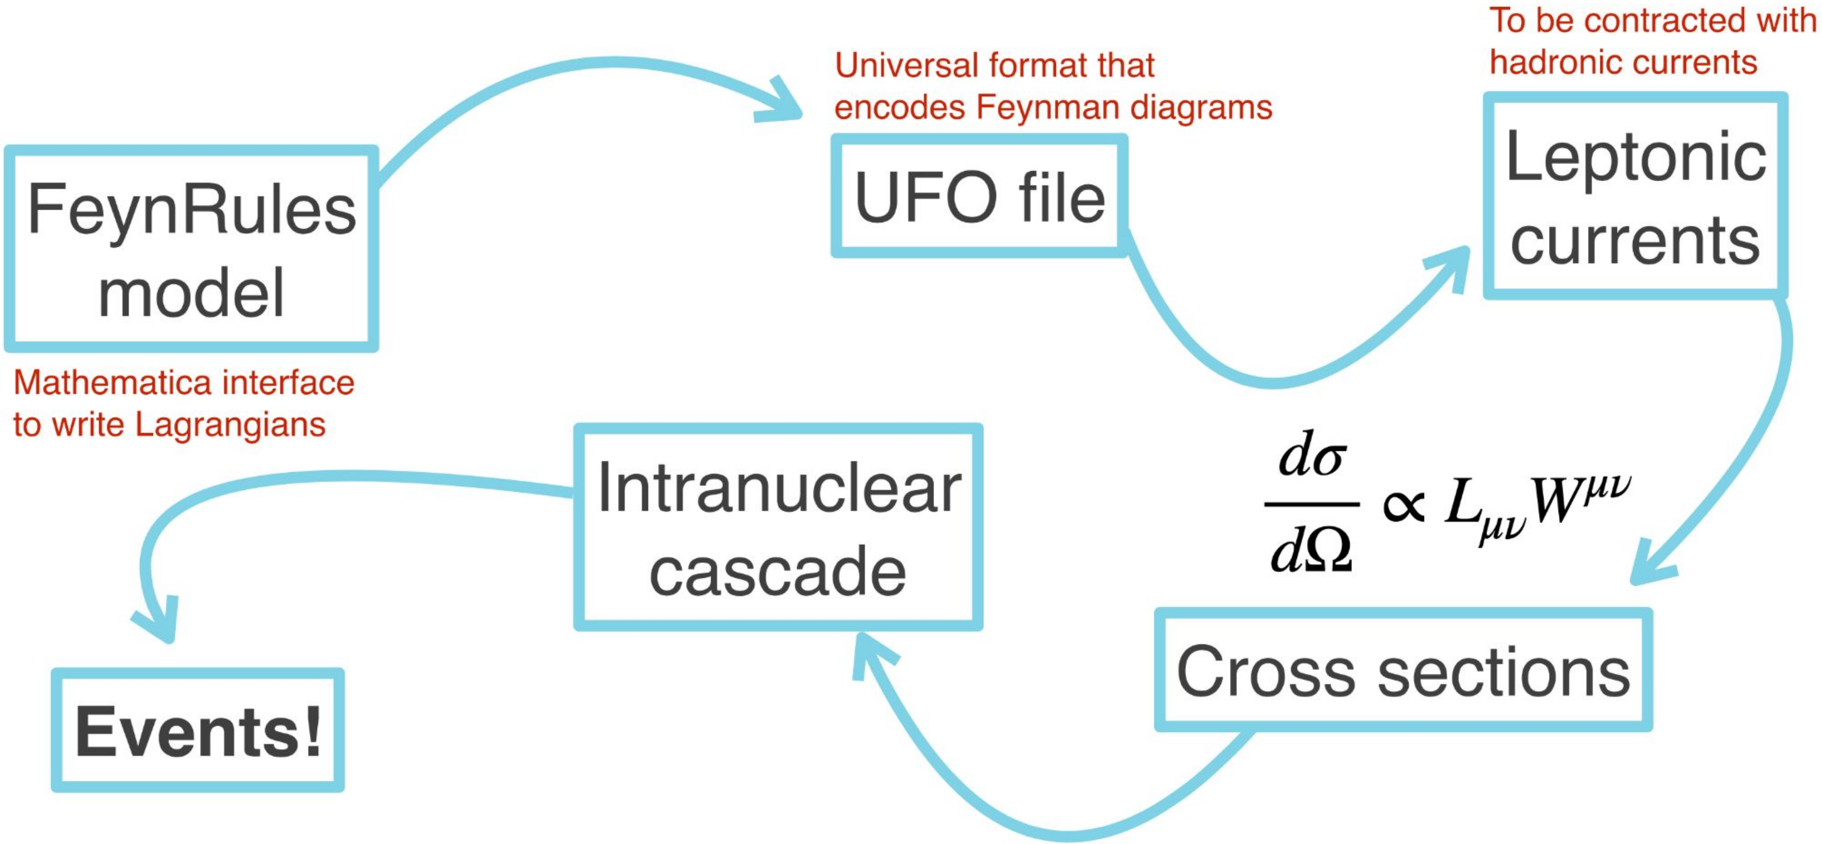

A different take on frameworks than Collider approach

Meld @ FNAL

- **DUNE's current framework (art) originates from a collider-physics experiment, steeped in event-based concepts**

- **Meld - A project for exploring how to meet DUNE's framework needs**

Purpose of Meld - explore more flexible data organizations

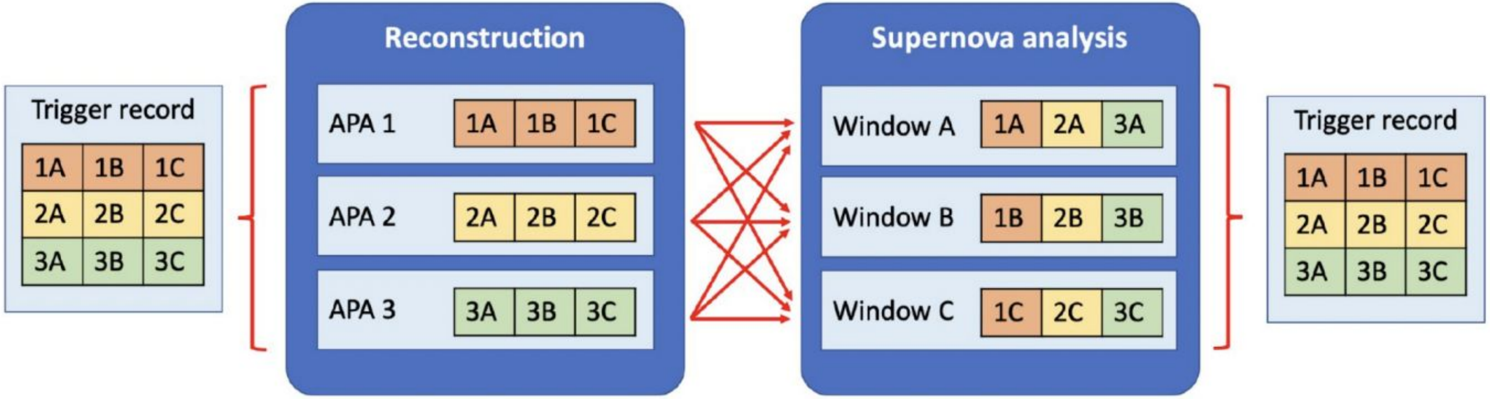

Fermilab

Meld @ FNAL

Physics results are obtained by analyzing the data as a whole

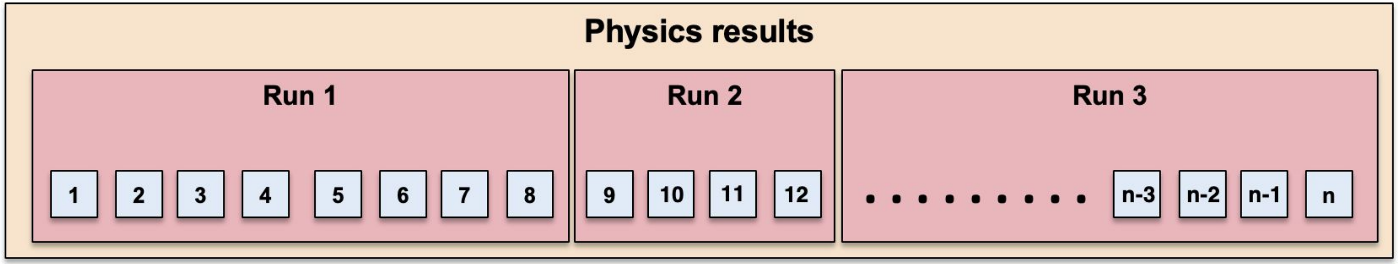

Meld @ FNAL

Physics results are obtained by analyzing the data as a whole

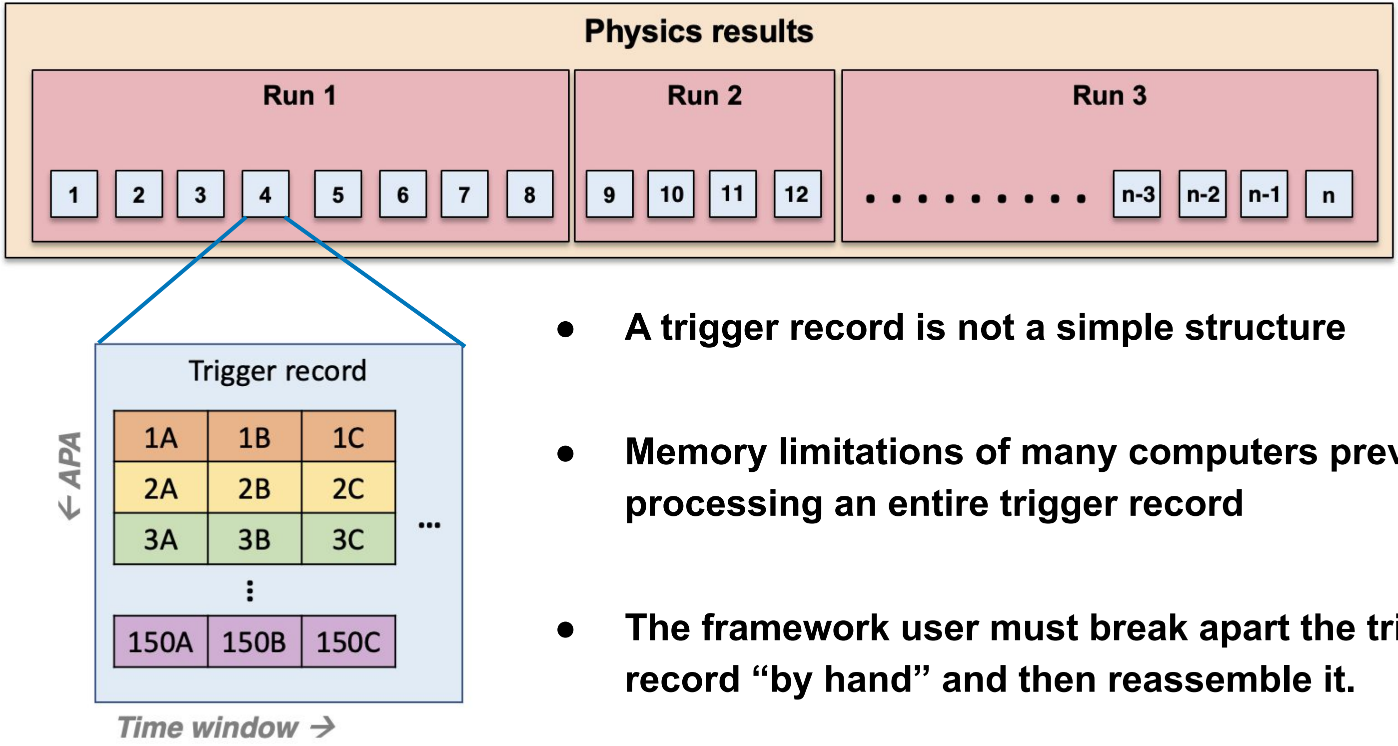

- **A trigger record is not a simple structure**
- **Memory limitations of many computers prevent processing an entire trigger record**
- **The framework user must break apart the trigger record "by hand" and then reassemble it.**

Meld @ FNAL

- Challenge - Processing data from different runs in parallel
- Framework must store and provide information at coarser level than just trigger record.
- Technologies exist for parallel processing (TBB, MPI, etc.), but they do not support hierarchical data groupings very well

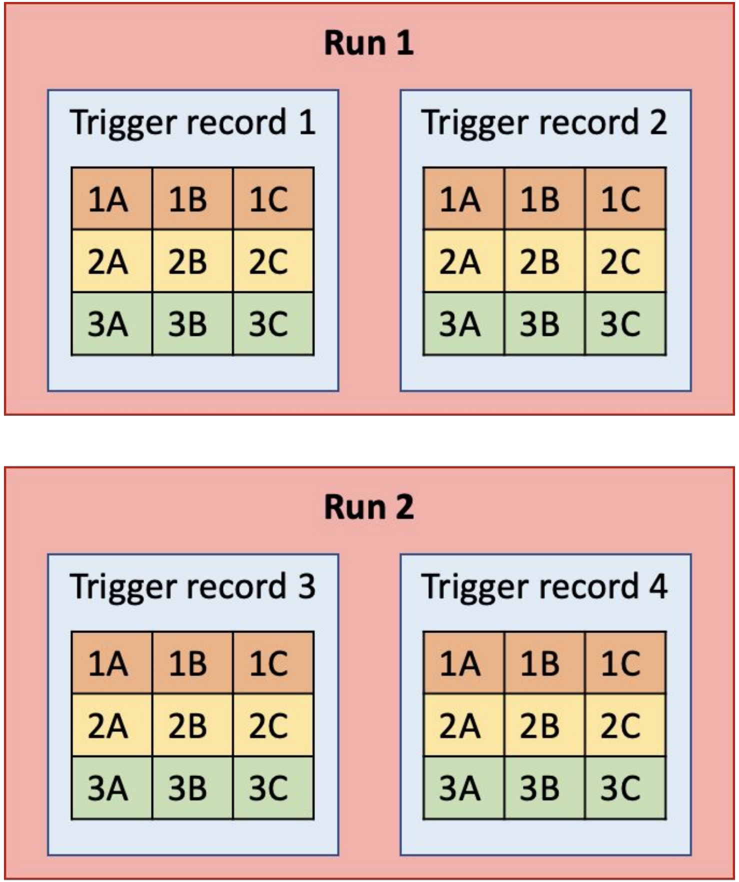

Meld @ FNAL

- **Challenge - Processing data from different runs in parallel**

- **Framework must store and provide information at coarser level than just trigger record.**

- **Technologies exist for parallel processing (TBB, MPI, etc.), but they do not support hierarchical data groupings very well**

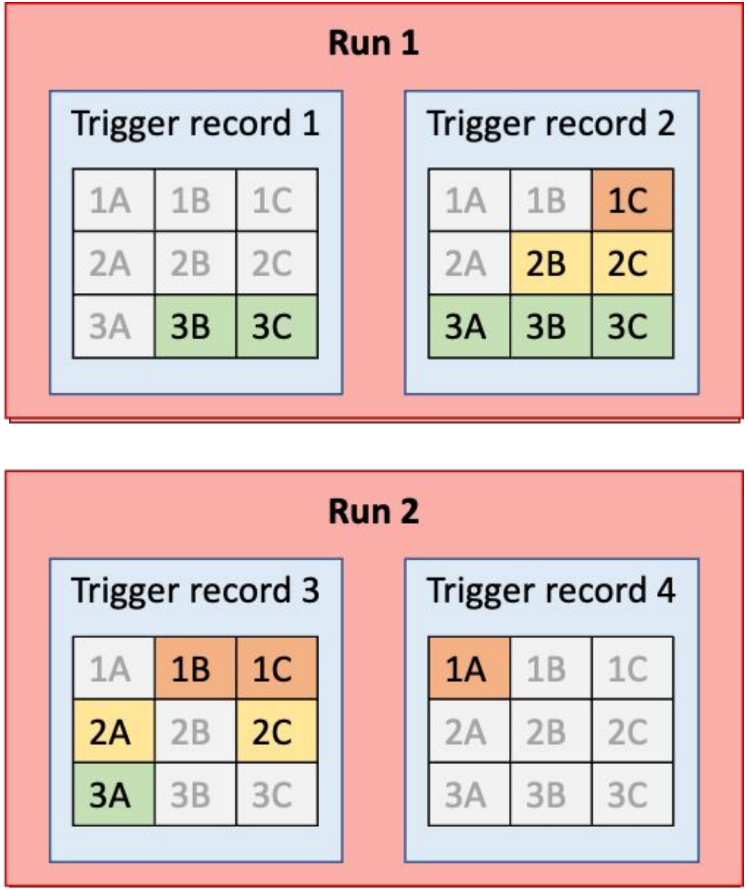

Meld @ FNAL

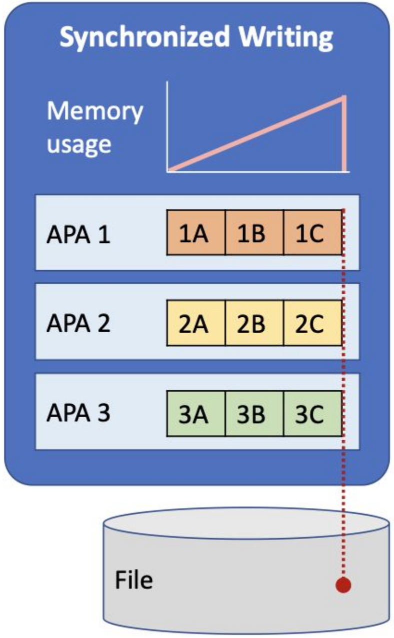

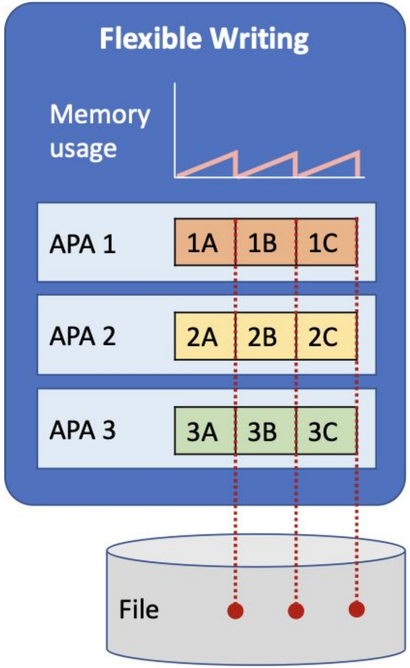

Collider approach - write physics data at the end of processing trigger record

Aim to reduce a program's memory usage, data should be written to a file during the processing of the trigger record

A Paradigm Shift in HEP

- **HEP faces unique high-throughput computing challenges from massive data rates**
- **Advanced computing techniques**
  - **Enable deeper insights and improve performance**
  - **Improve operational efficiency**
  - **Ultimately accelerate time-to-physics and discovery**

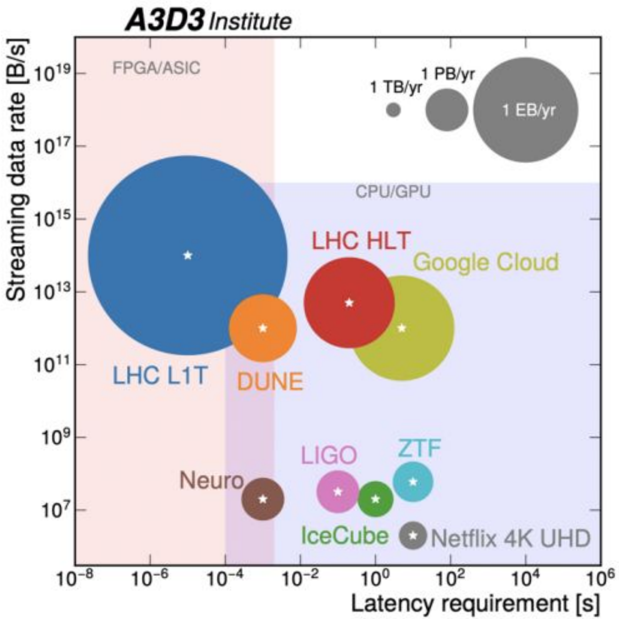

Evolve HEP computing infrastructure
Storage technologies, analysis facilities, heterogeneous computing (e.g. GPUs)

Leverage multidisciplinary computational & domain science expertise
Federal HPC facilities and commercial cloud, specialized services, modern software stacks

Embrace AI/ML for HEP and also HEP for AI/ML
Develop AI capabilities for HEP science, support HEP contributions to broader AI advances

High Performance Computing (HPC) - Gen Z DOE Supercomputers

- Future HEP Experiments -
  - Order of magnitude increase in data rate
    - Data & processing complexity within existing frameworks
    - “Buy more CPUs” - not an option
- Explore parallelism
    - Future HPCs – CPUs + GPUs
      - fast turn around processing and regular raw data processing
      - Code portability from CPU to GPU crucial

High Performance Computing - HEP CCE Efforts

- HEP-CCE (Centre for Computational Excellence) → 3 year pilot project funded by DOE
  - 6 Experiments, 4 National labs across US
  - Intensity, Energy and Cosmic Frontiers
- Goal - Exploit features of HPCs efficiently
  - Develop and test strategies to overcome HEP community wide computational challenges
    - PPS: Portable Parallelization Strategy
    - IOS: I/O and Storage on HPC Platforms
    - EG: Event Generators
    - CW: Complex Workflow on HPC

High Performance Computing - HEP CCE Efforts

PPS: Portable Parallelization Strategy

- Investigate software portability solutions
  - Kokkos, OpenMP
- Evaluate ease of porting - Porting, building, performance, code impact

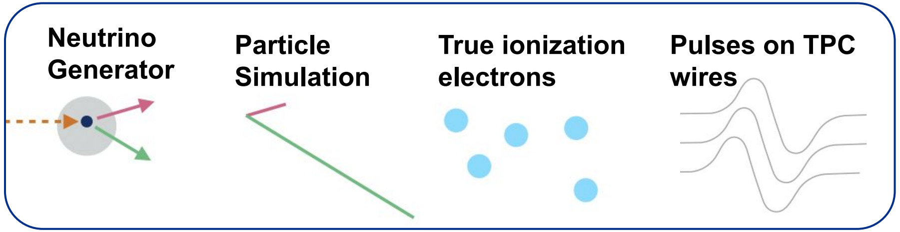

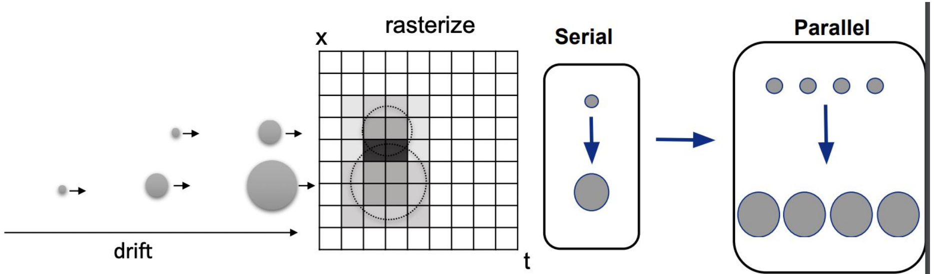

IOS: I/O and Storage

- Investigate HDF5 as intermediate event storage for HPC processing
- HPC friendly Data Model : Experiment agnostic
- Parallel I/O of the HEP data using MPI (Message Parsing Interface) and HDF5 libraries

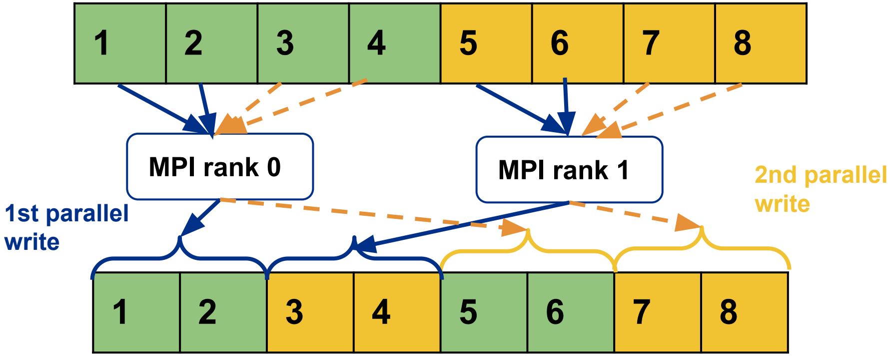

https://www.anl.gov/hep-cce

Computing Facilities - HEPCloud

- **Unified interface to Grid, Cloud, and HPC resources**
- **Currently used mainly to run CMS, NOvA, DUNE workflows on NERSC supercomputers**

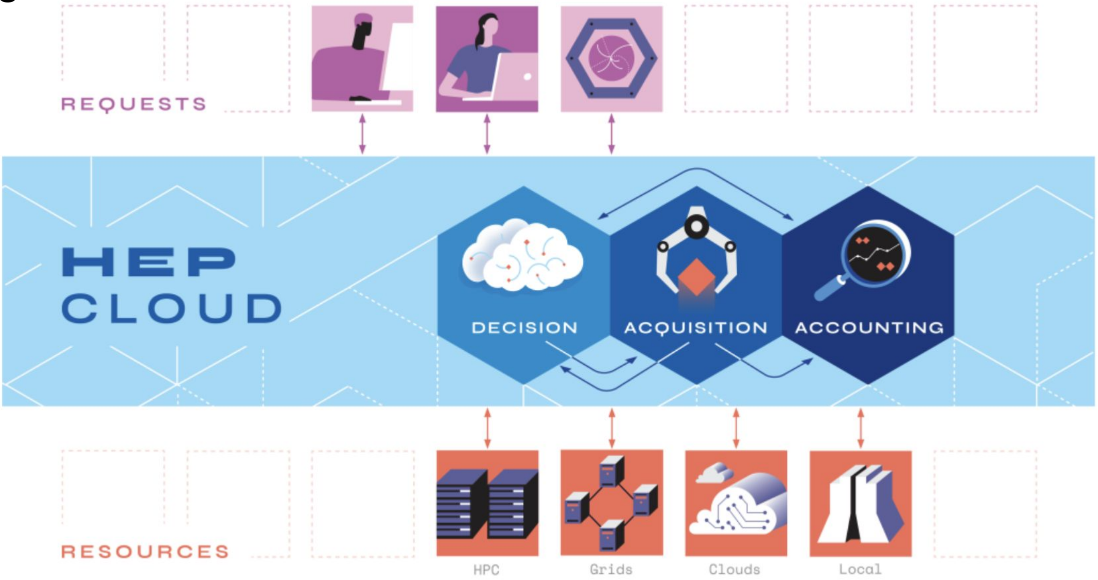

https://hepcloud.fnal.gov/

Fermilab

Computing Services Enabling Physics Discoveries

Computing Services

Supporting local and remote users sharing resources across Fermilab and the Open Science Grid

Distributed Disk Management

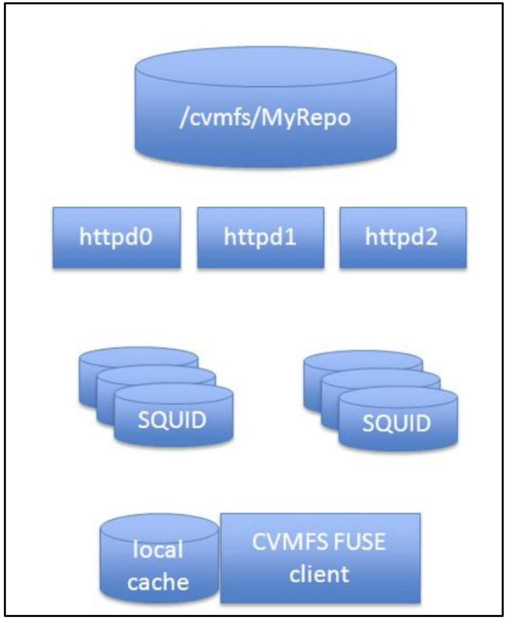

Dan Bradley

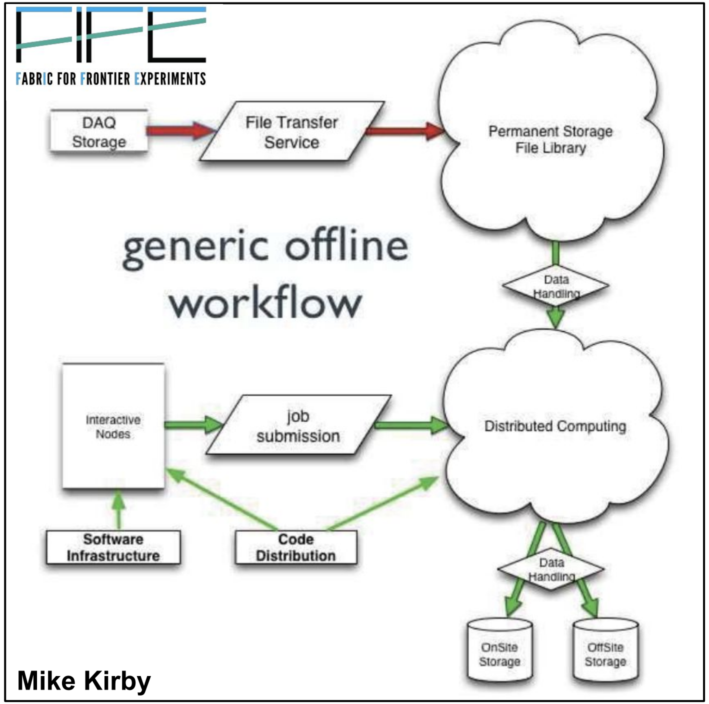

Mike Kirby

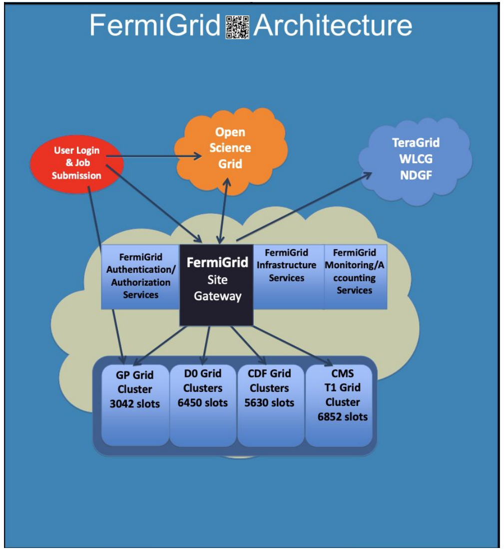

Fermilab

Computing Services

Supporting local and remote users sharing resources across Fermilab and the Open Science Grid

FABRIC FOR FRONTIER EXPERIMENTS

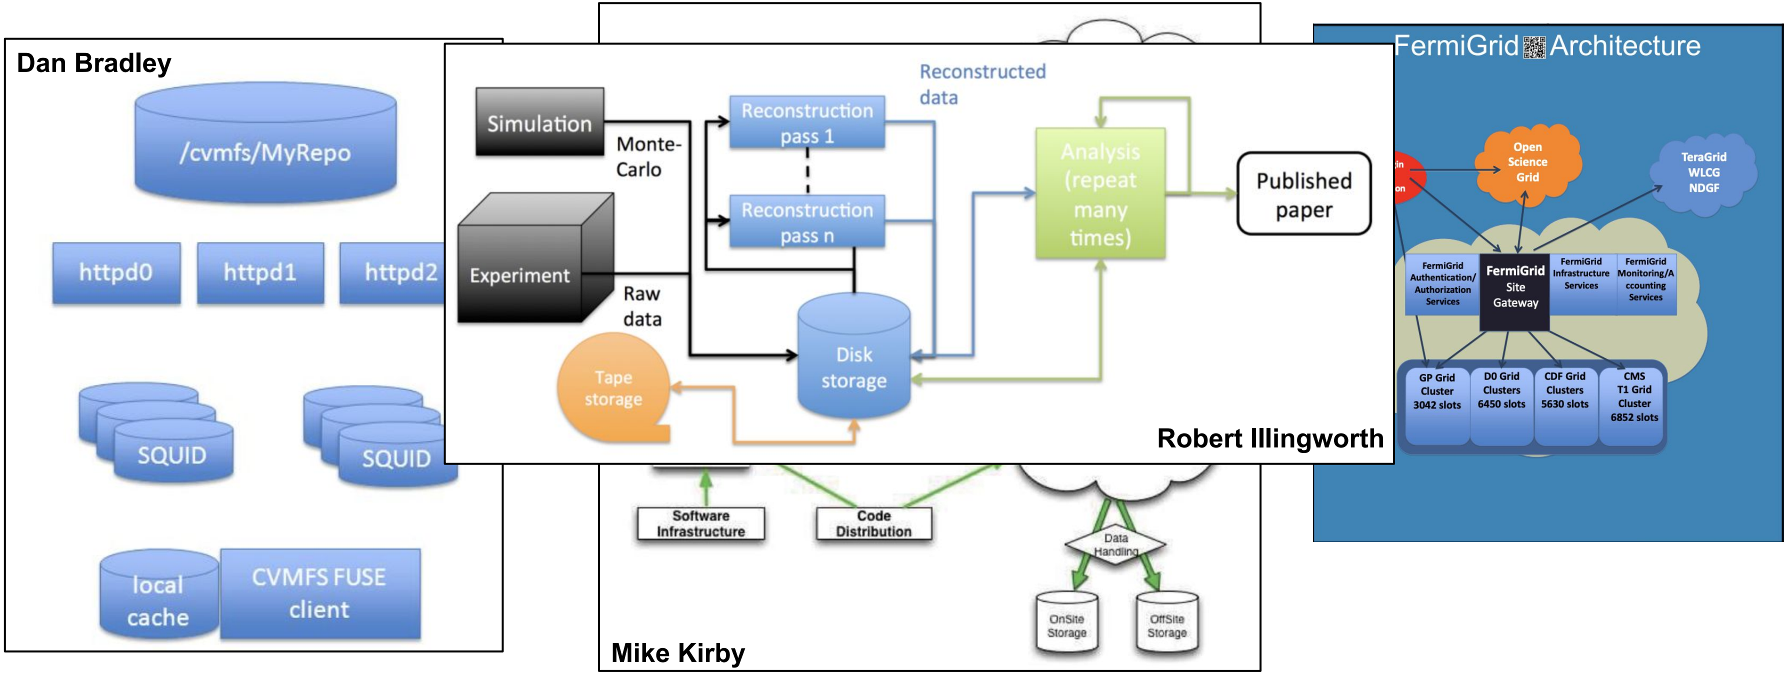

Computing Services

- Support and develop
  - common computing tools
  - data monitoring tools

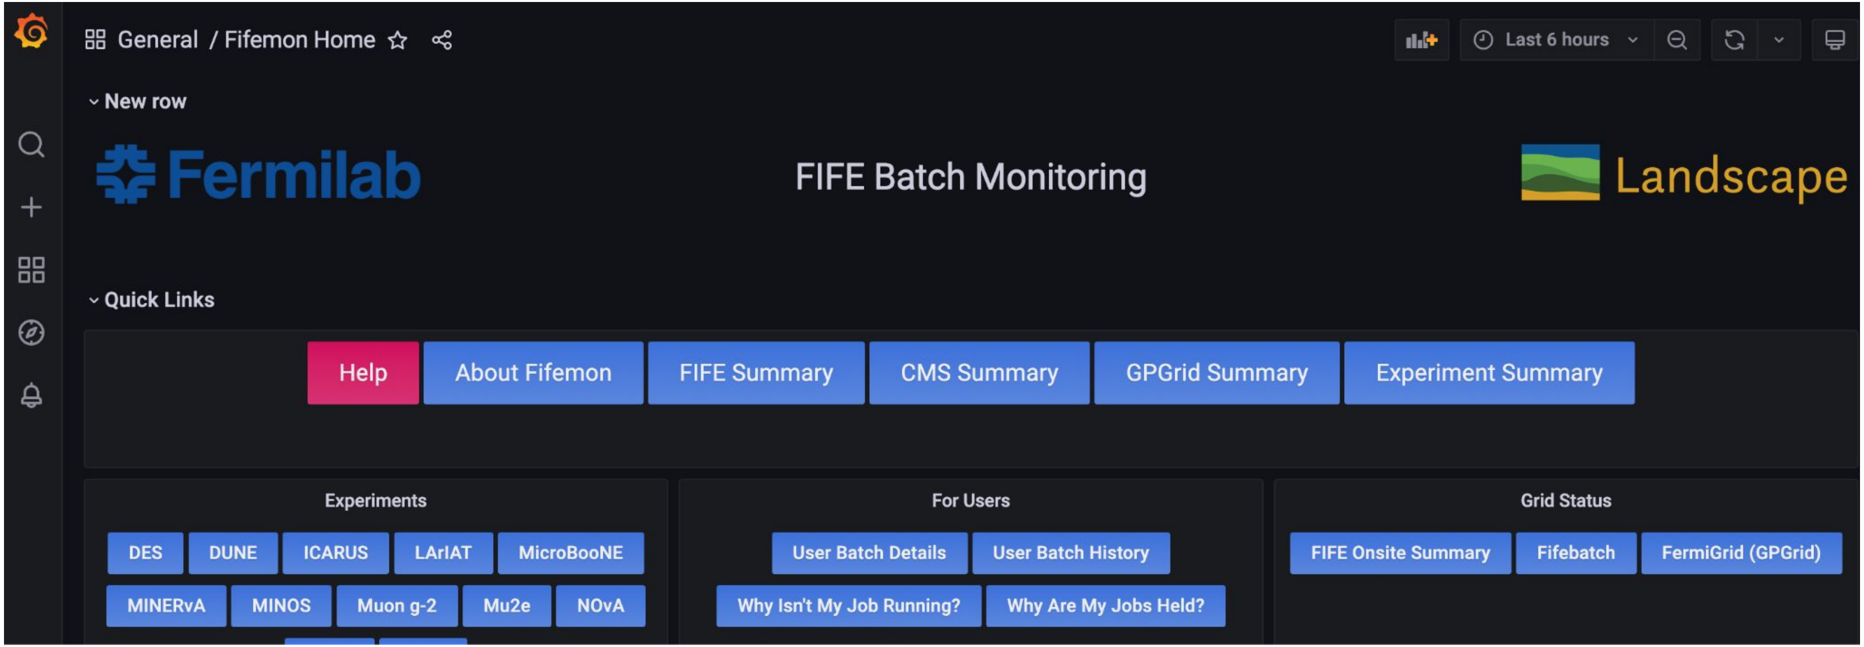

Summary

- Fermilab offers a wide range of services for experiments
- Big Data challenge for Next-gen HEP experiments
  - A paradigm shift
    - Hardware Accelerators such as FPGA, ASIC
    - High-Performance Computing (HPC) resources for data processing
    - Increasing AI/ML applications
- Experience beyond ROOT
  - Analysis facilities with python, Julia
- Quest continues
  - Computing R&D projects to address challenges
  - Fermilab Frameworks Workshop June 5th- 7th

A Plethora of Tools to Support Neutrino Physics and HEP

SciDAC
Scientific Discovery through Advanced Computing

LArSoft

FIFE
FABRIC FOR FRONTIER EXPERIMENTS

High Energy Physics - Center for Computational Excellence
HEP-CCE

Genie
UNIVERSAL NEUTRINO GENERATOR & GLOBAL FIT

dCache.org

Enstore

ROOT
Data Analysis Framework

python™

C++

SCIENTIFIC LINUX

XRootD

art

G4

Thank you!

Many Thanks to Steven Gardiner, Kyle Knoepfel, Ken Herner, Burt Holzman, Mike Kirby, Giuseppe Cerati, Mike Wang, Erica Snider, Andrew Norman, Robert Harris and everyone within CSAID

Back up

dCache

Rob Illingworth

- dCache provides grid accessible bulk storage and tape interface
- Main cache is backed by tape; data staged on access
  - 8.7 PB – aim to maintain 30 day lifetime; +1 PB since last year
- Scratch is another shared resource; LRU file removal, but not tape backed
  - 2.5 PB – similarly aim for 30 day lifetime; no change since last year
- Dedicated tape-backed areas are allocated to specific experiments primarily for raw & production data
  - 9.3 PB; +1.3 PB since last year
- Persistent space is permanently resident on disk (not tape cache) under experiment control
  - 7.9 PB; +4 PB since last year: a move towards more experiment managed tape recall
- Outside FCRSG scope (small experiments, external customers, some unallocated space)
  - 1.4 PB

For More Details

- Childers, Taylor, et al. “Porting CMS Heterogeneous Pixel Reconstruction to Kokkos.” *vCHEP 2021.* arXiv:2104.06573v1. Slides.
- Dong, Zhihua, et al. “Porting HEP Parameterized Calorimeter Simulation Code to GPUs.” *Frontiers in Big Data*. arXiv:2103.14737v2. Slides.
- Kortelainen, Matti J., et al. “Performance of CUDA Unified Memory in CMS Heterogeneous Pixel Reconstruction.” vCHEP 2021. Paper. Slides.
- Pascuzzi, Vincent R., Goli, Mehdi. “Achieving Near Native Runtime Performance and Cross-Platform Performance Portability for Random Number Generation Through SYCL Interoperability.” arXiv:2109.01329
- Yu, Haiwang, et al. “Evaluation of Portable Acceleration Solutions for LArTPC Simulation Using Wire-Cell Toolkit.” *vCHEP 2021.* arXiv:2104.08265v1. Slides.
- HEP-CCE Collaboration, Portability: A Necessary Approach for Future Scientific Software, Snowmass White Paper

https://www.anl.gov/hep-cce

HPC and DUNE Computing

- **Utilizing HPC → Computing R&D**
  - **I/O and Storage**
  - **How do we transfer data to LCFs**
  - **Existing frameworks and data formats in HEP not HPC friendly**

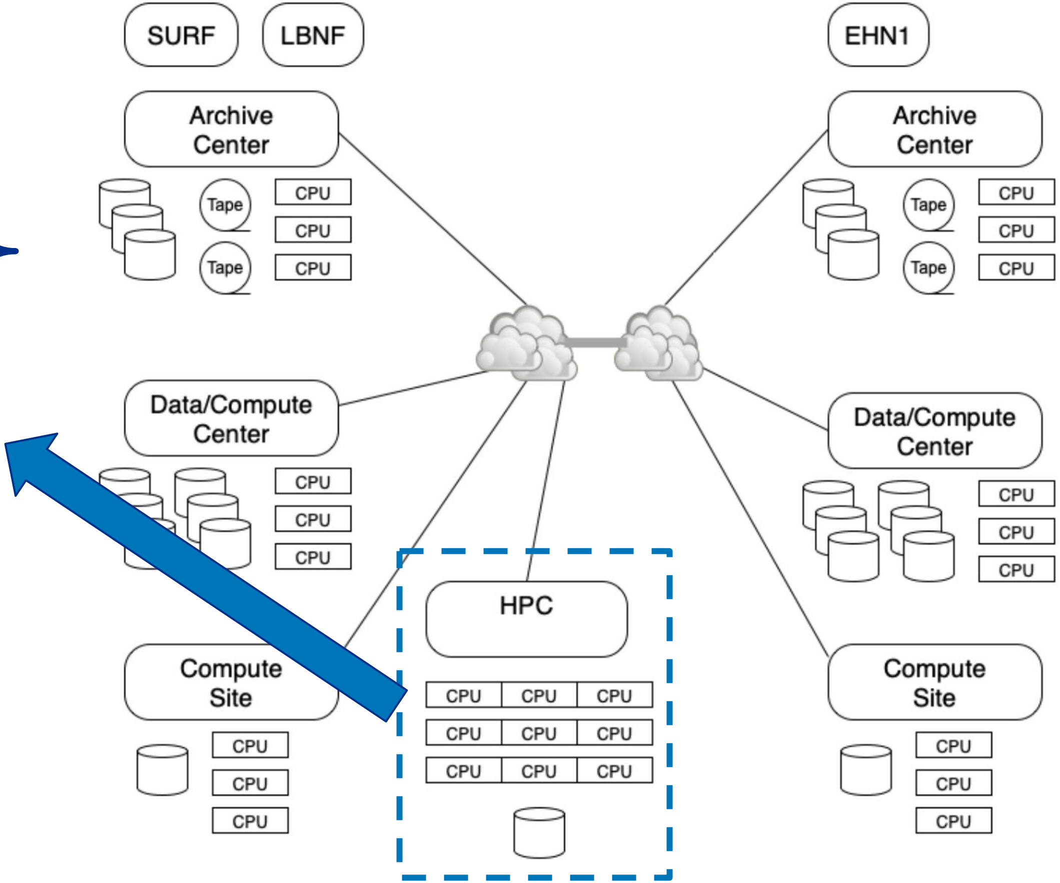

Fermilab

Realtime AI/ML trigger and tagging @ FNAL

- Efforts focused on :
  - Developing a data-driven electronics noise model for ICEBERG
    - crucial for producing realistic simulated samples
  - Testing and developing ML based techniques:
    - to process raw 1D waveforms from individual wire channels to detect regions of interest and denoise waveforms
    - To process raw 2D wire plane data to identify signals such as those from Ar39 decays and separate them from noise background

Supernova Candidate Processing

- Limited space and infrastructure (i.e. cooling) at the far site means → no bulk processing on a local farm

- 10,000 – 40,000 present-day CPUs needed for reconstruction to finish within a few hours of event (goal: preliminary direction before event rises in optical bands)
  - HPC centers
    - Concern → data transfer in and out

- Must be able to handle large input stream as well as output at a similar rate
  - Run standard data reco or make a slimmed-down, faster version? Speed vs. accuracy tradeoff?

- Implications for additional network paths? What are those requirements? Costs?

The ACHILLES Event Generator

[Phys. Rev. D 107, 033007 (2023)](#)

A CHIcagoLand Lepton Event Simulator

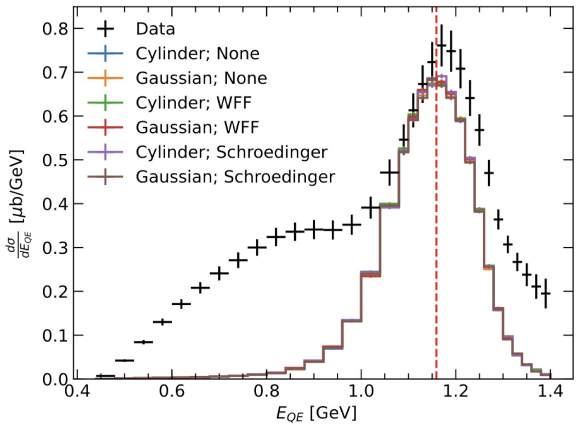

- New theory-driven event generator
  - Fermilab-led
  - Neutrinos, electrons, BSM

- Technical design borrows techniques from collider physics event generators
  - Applies these to neutrinos for the first time

- Example: Automated leptonic tensor
  - [Phys. Rev. D 105, 096006 (2022)](#)
  - Support wide range of BSM models without dedicated development work

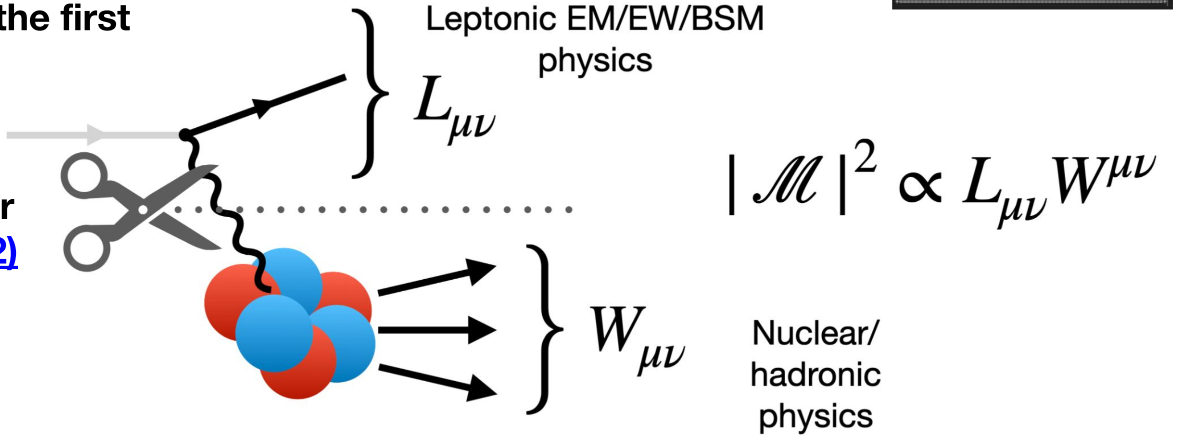

HEPCloud Contd.

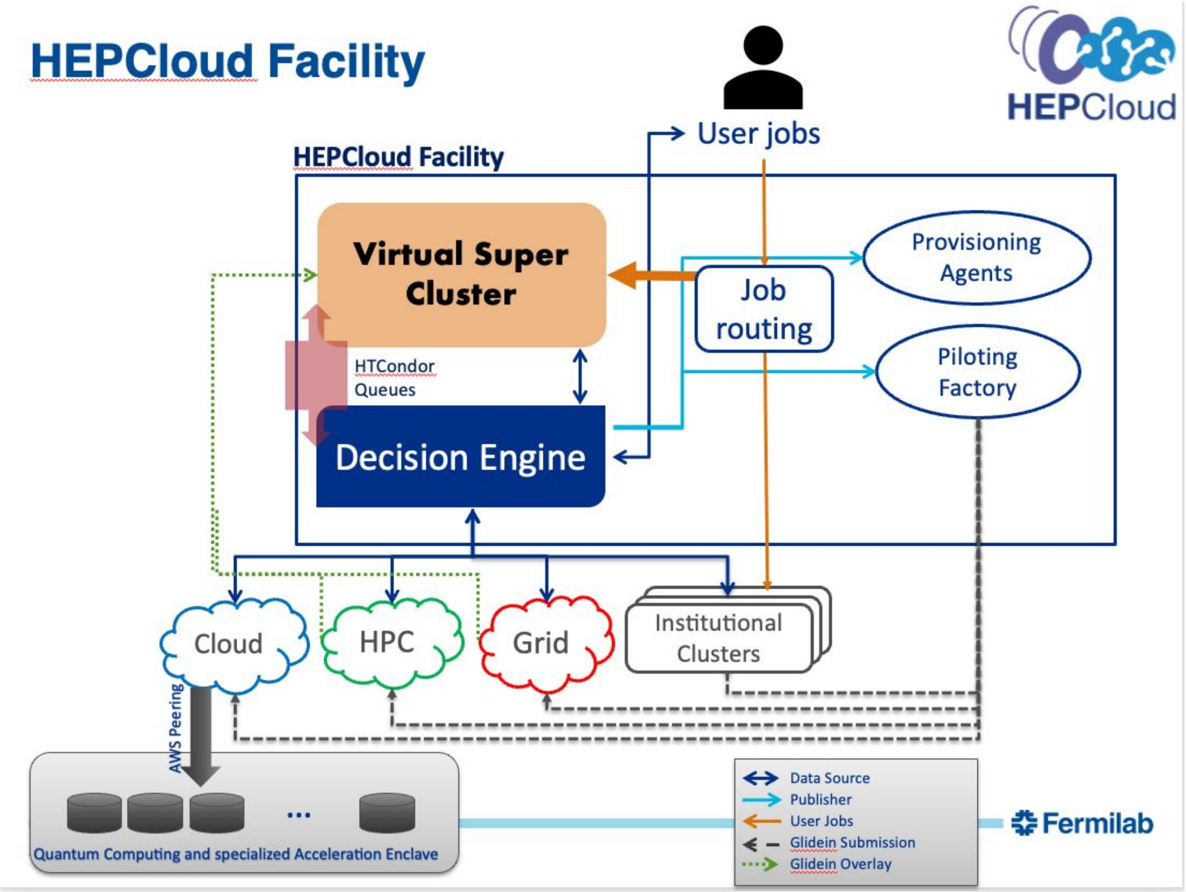

Computing Services

Support and develop common computing tools for all experiments

- **Data management and submission**
- **Software distribution and build systems**
- **Source code version control systems and repositories**
- **Access to Open Science Grid and High-Performance Computing centers**
- **Interactive computing machines**
- **Online and Offline software frameworks**
- **Interactive Analysis Tools**

High Performance Computing (HPC) - Gen Z DOE Supercomputers

- Future HEP Experiments -
  - Order of magnitude increase in data rate
    - Data & processing complexity within existing frameworks
    - “Buy more CPUs” - not an option

- Explore parallelism
    - Future HPCs – CPUs + GPUs
      - fast turn around processing and regular raw data processing
      - Code portability from CPU to GPU crucial

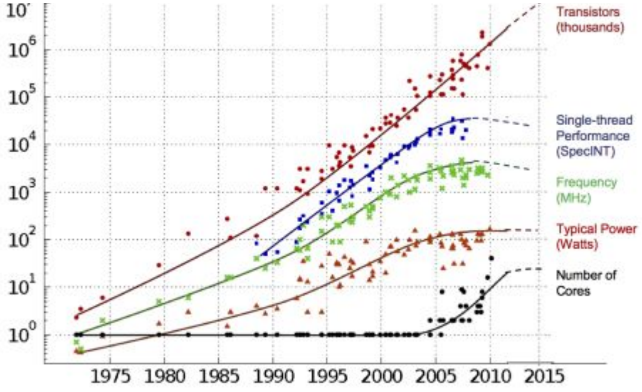

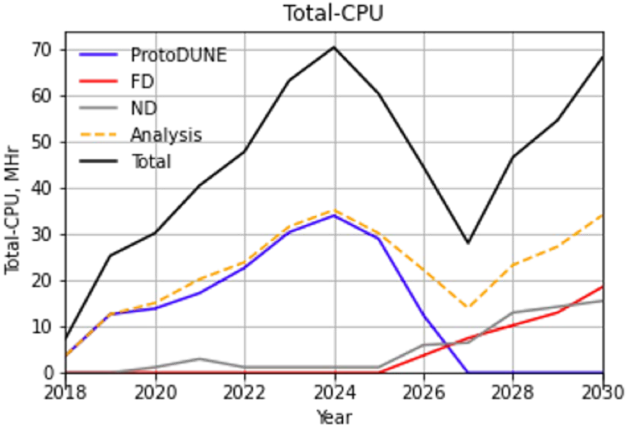

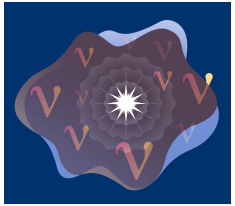

Fermilab

High Performance Computing - HEP CCE Efforts

PPS: Portable Parallelization Strategy

- **Investigate software portability solutions**
    - **Kokkos, OpenMP**

- **Evaluate ease of porting - Porting, building, performance, code impact**

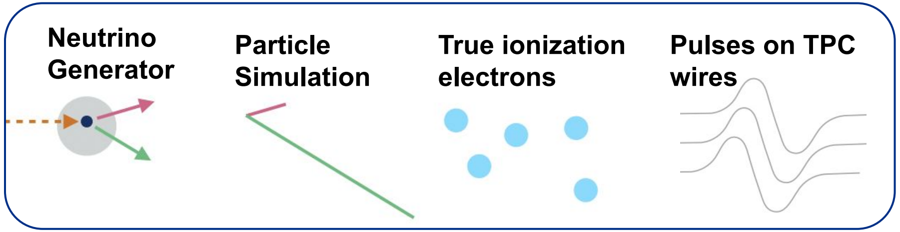

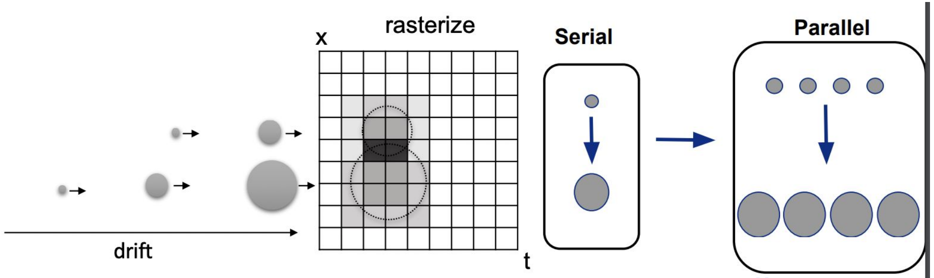

IOS: I/O and Storage

- **Performance of ROOT I/O in HEP workflows on HPC systems**

- **Investigate HDF5 as intermediate event storage for HPC processing**

- **HPC friendly Data Model : Experiment agnostic**

- **Parallel I/O of the HEP data using MPI (Message Parsing Interface) and HDF5 libraries**

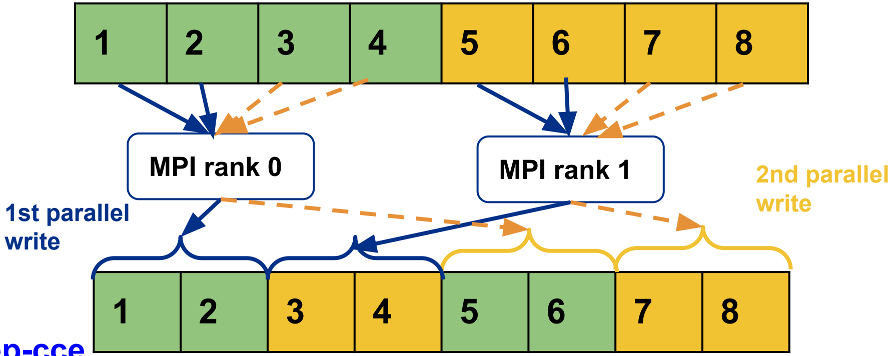

https://www.anl.gov/hep-cce