2ND AUGUST, 2022

Argonne NATIONAL LABORATORY 75 1946–2021

Data Storage for HEP Experiments in the HPC Era

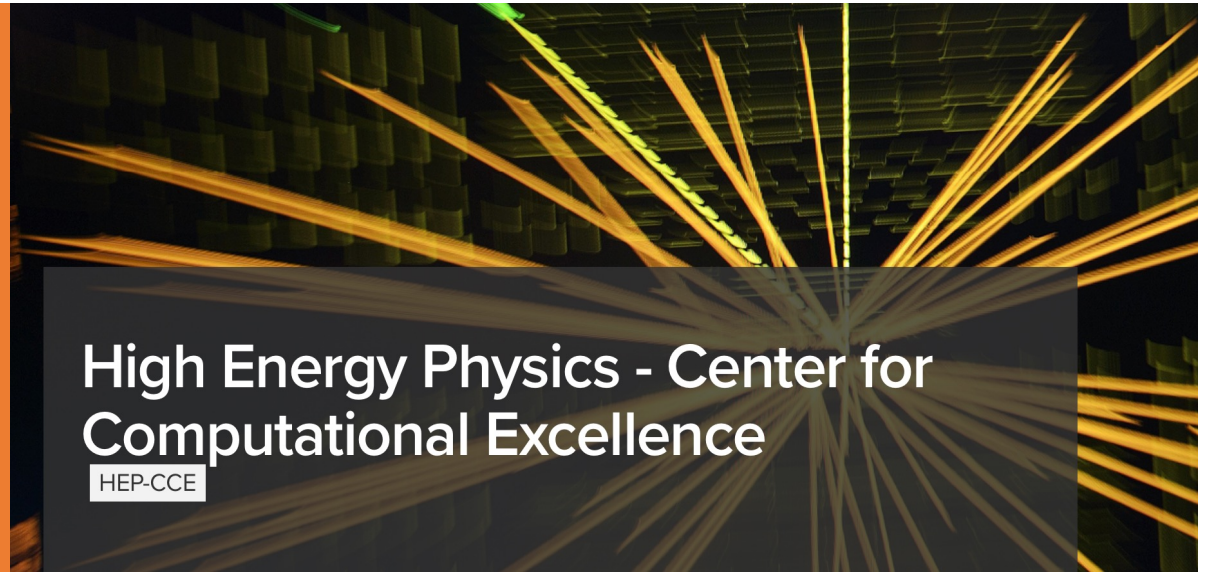

AMIT BASHYAL
Postdoctoral Research Fellow
HEP-CCE
Argonne National Lab

Works shown here are the collective work of HEPCCE IO members.
Introductory slides have been pulled from several good talks given over the years.

U.S. DEPARTMENT OF ENERGY Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

HEP Experiments Past

- HEP Experiments and Data
  - Analysis of particle interactions with matter in atomic or sub atomic level
  - Data = Recording of these interactions with the detector material

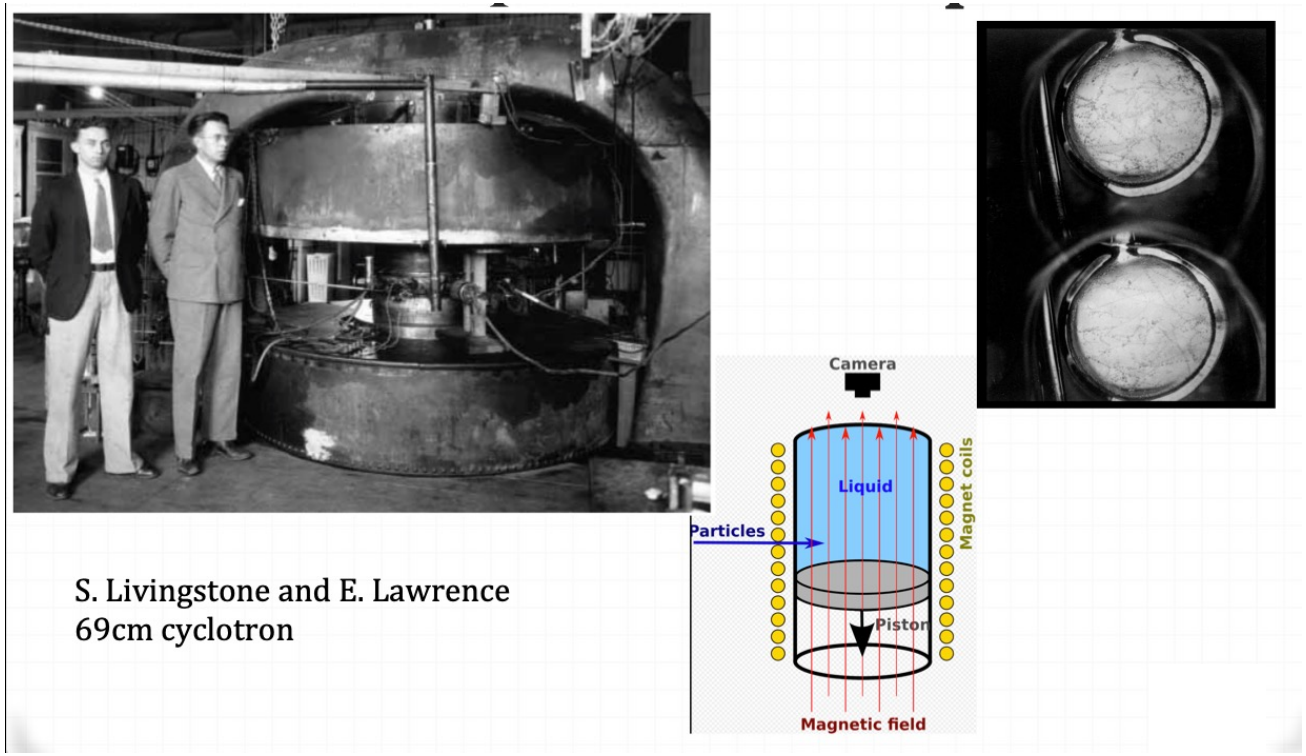

Experimental setup of the 69 cm (circa 1930s) cyclotron.

HEP Experiments Present

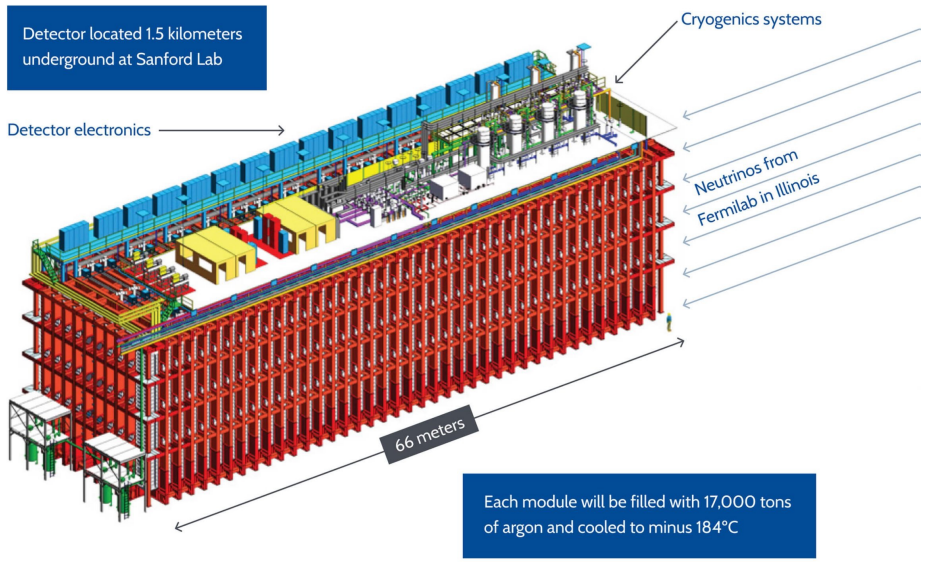

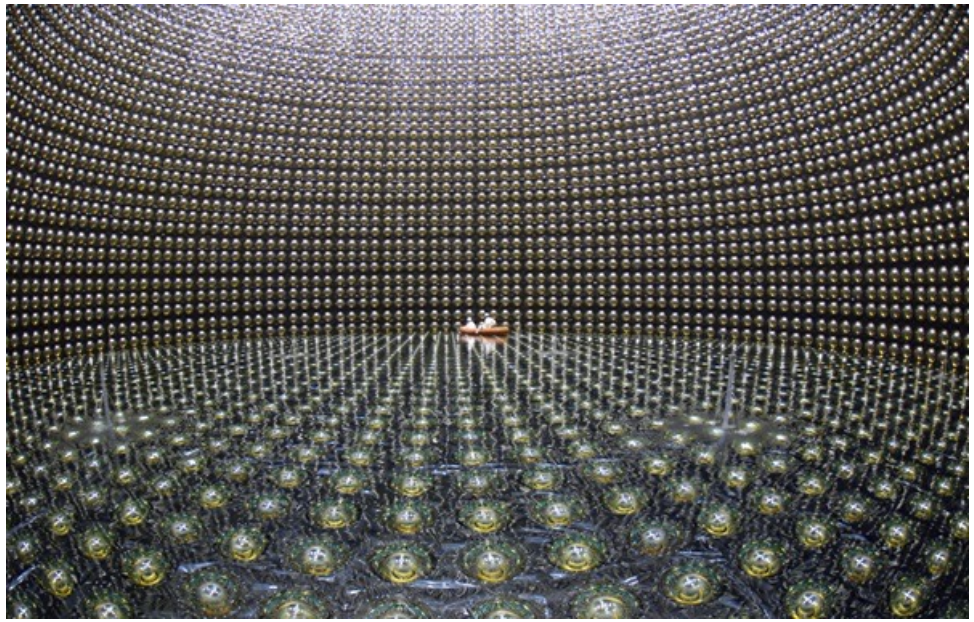

Current HEP experiments try to test the theories that require reaching high energy or great sensitivity.

Complex Experiment Setup = Large Data + Complex Data Management

Most HEP experiments have similar setup to record and store the events.

When something interesting happens in the detector:
If Interesting
Save events: Different parts of detector save different information.

Use different information to create full event:
Particle ID, charge, momentum, location of interaction etc

Store the event for further analysis.

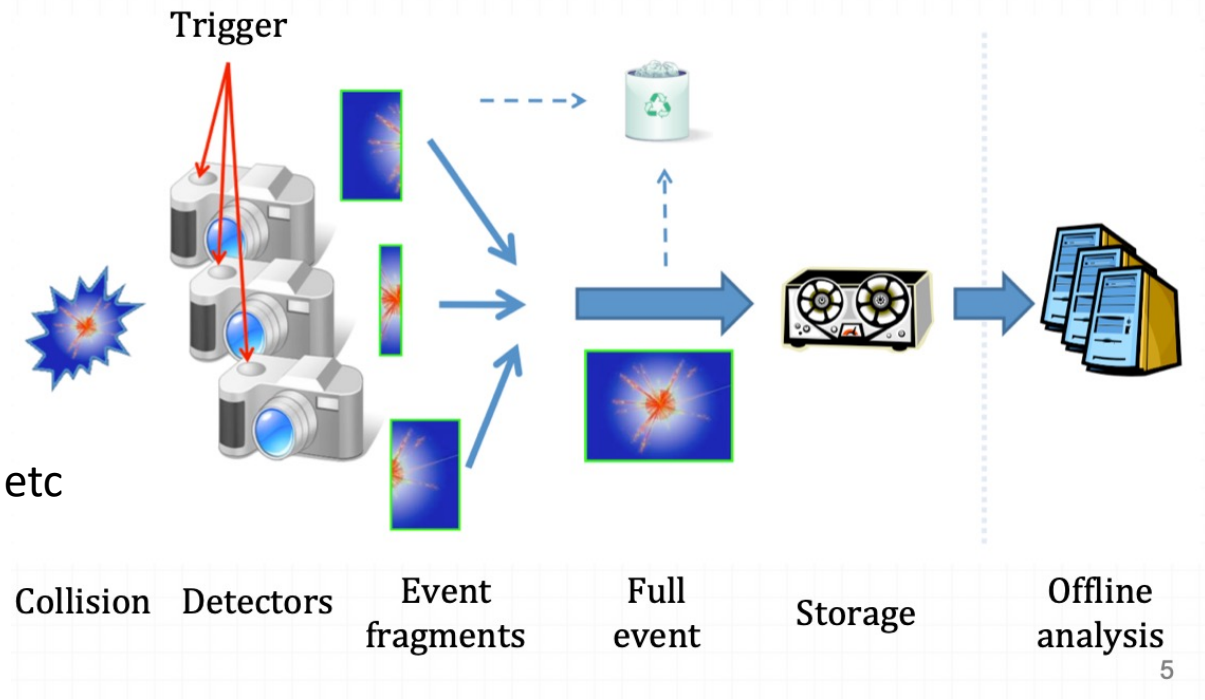

5

Data Production And Storage

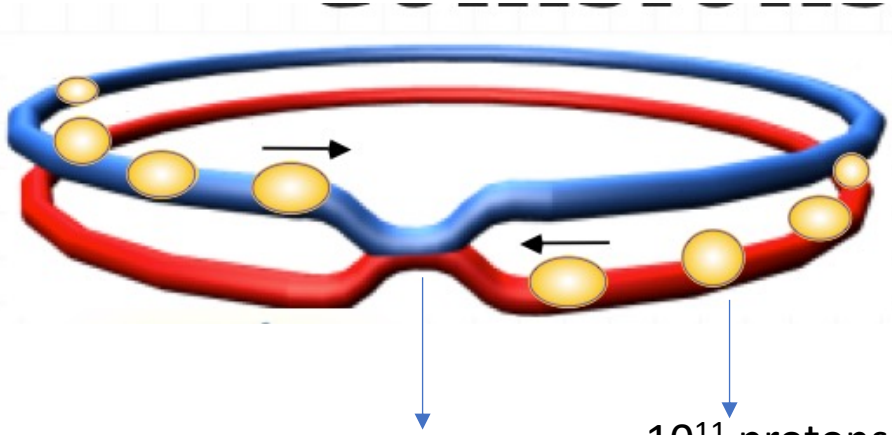

Crossing Rate = 40 MHz

10^{11} protons in 1 bunch

LHC

10^{15} events/yr → 10^{8} events/yr

Perfect world→Store Everything
Real World →Data Storage = $$
Computing capability = $$

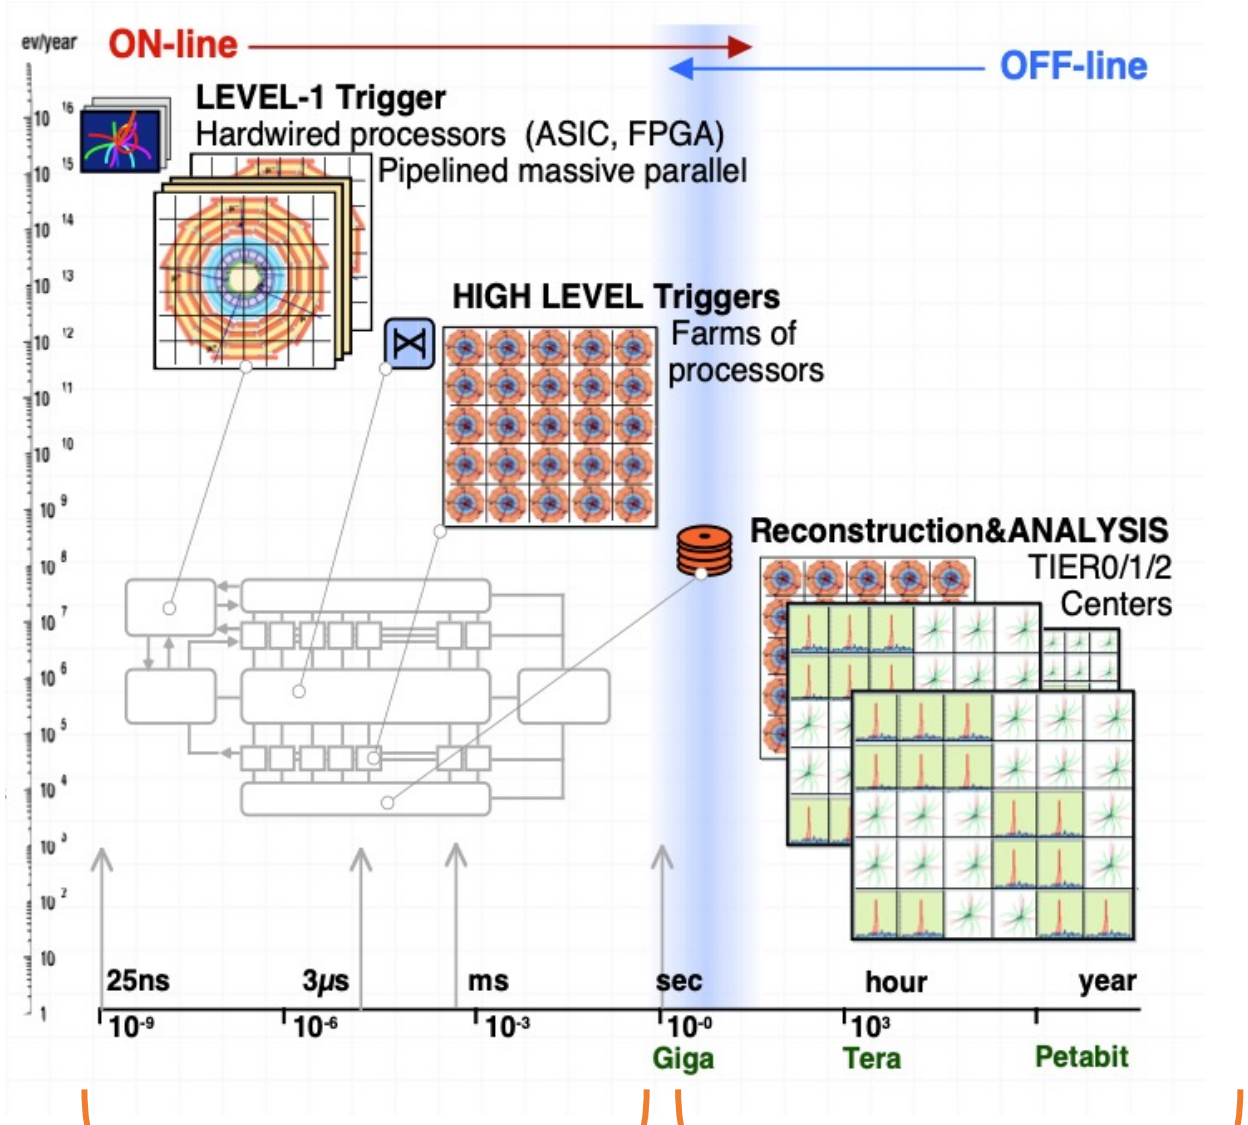

Computing Resources to get Useful physics from Raw Data

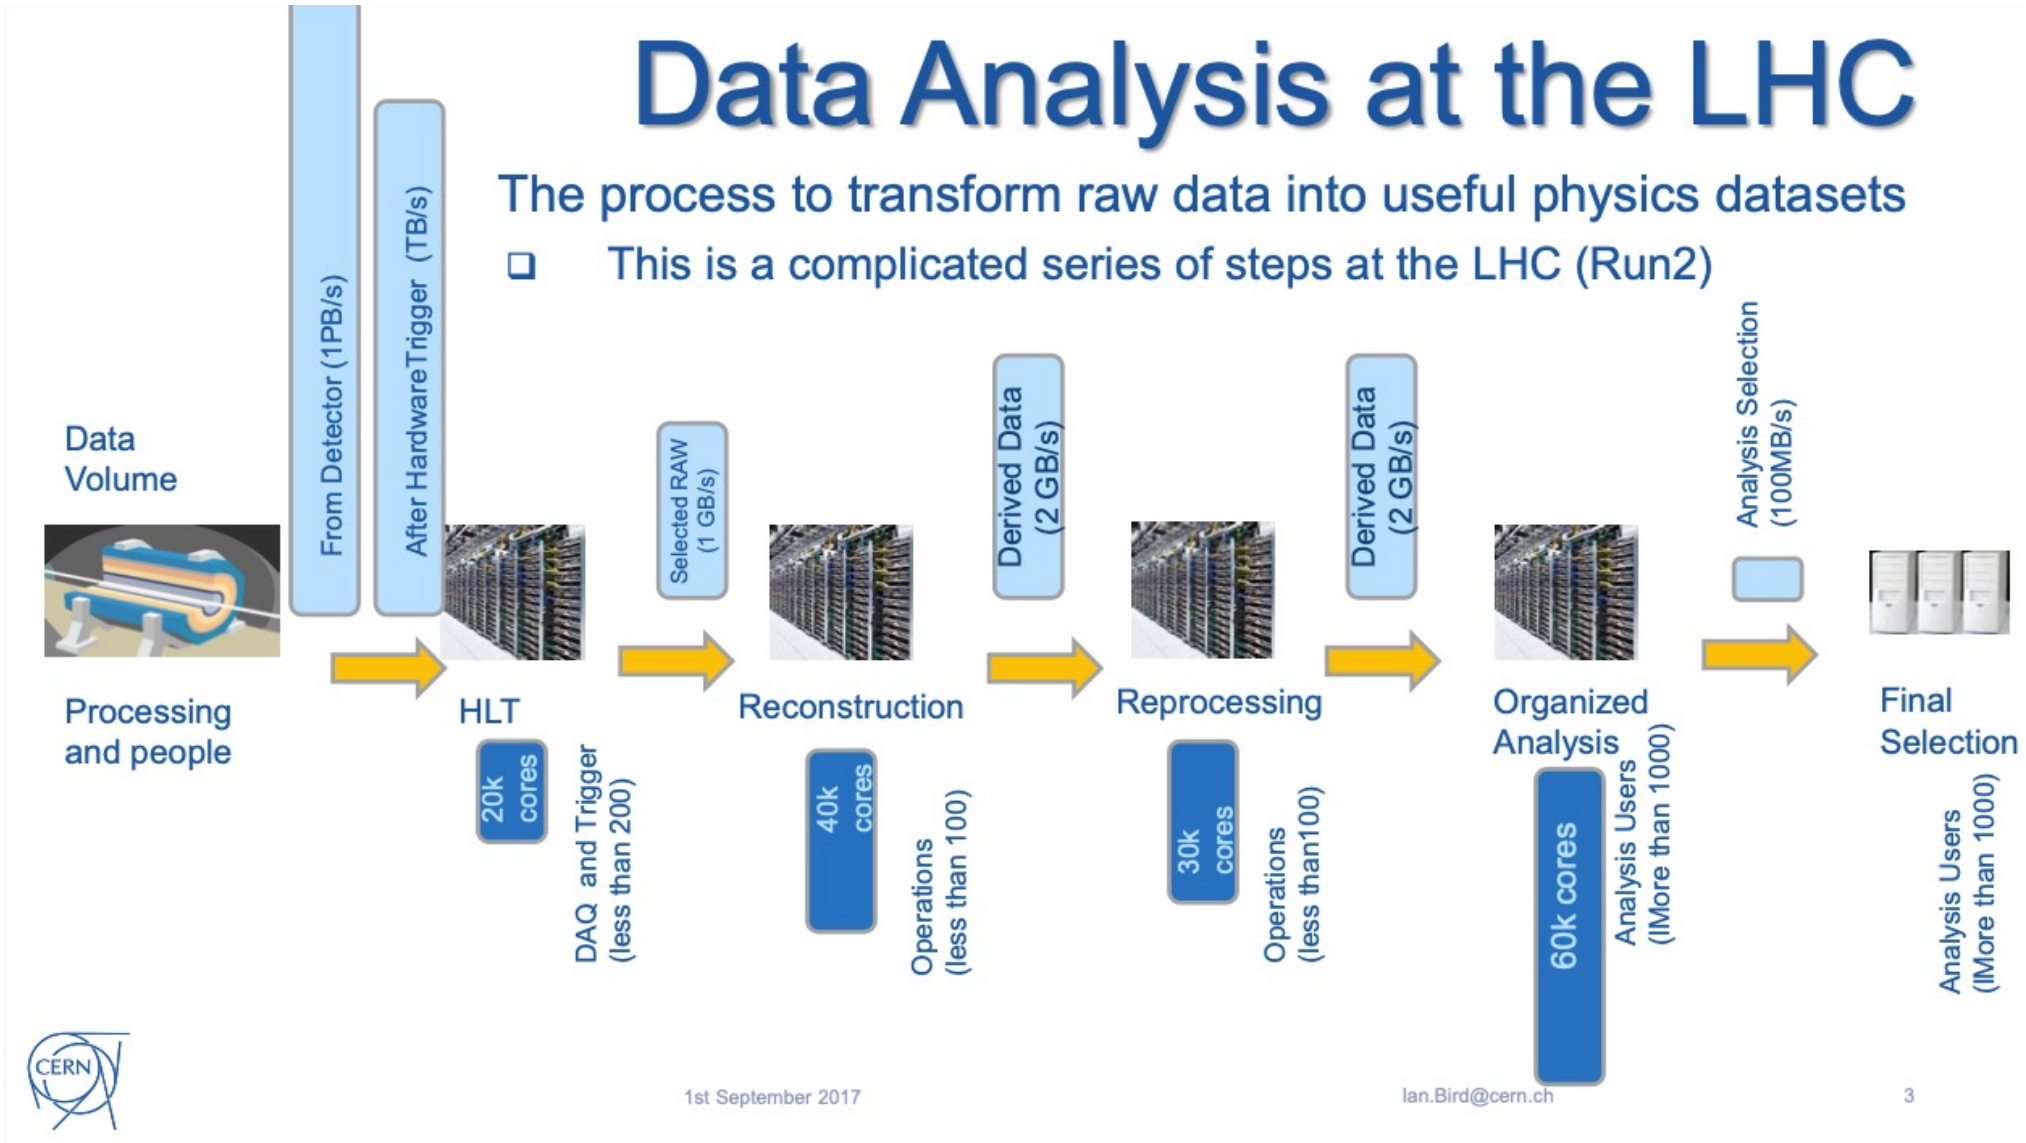

Taken from: https://indico.cern.ch/event/660772/contributions/2696454/attachments/1517232/2368316/WLCG-LigoVirgo-010917.pdf

Data Storage/Computation Future

- Computational Requirements of the HEP experiments
  - Future (Large experiments →Big requirements)
  - ATLAS and DUNE as examples

LHC will go through upgrades

Storage and compute challenges cannot be met by scaling current technological resources.

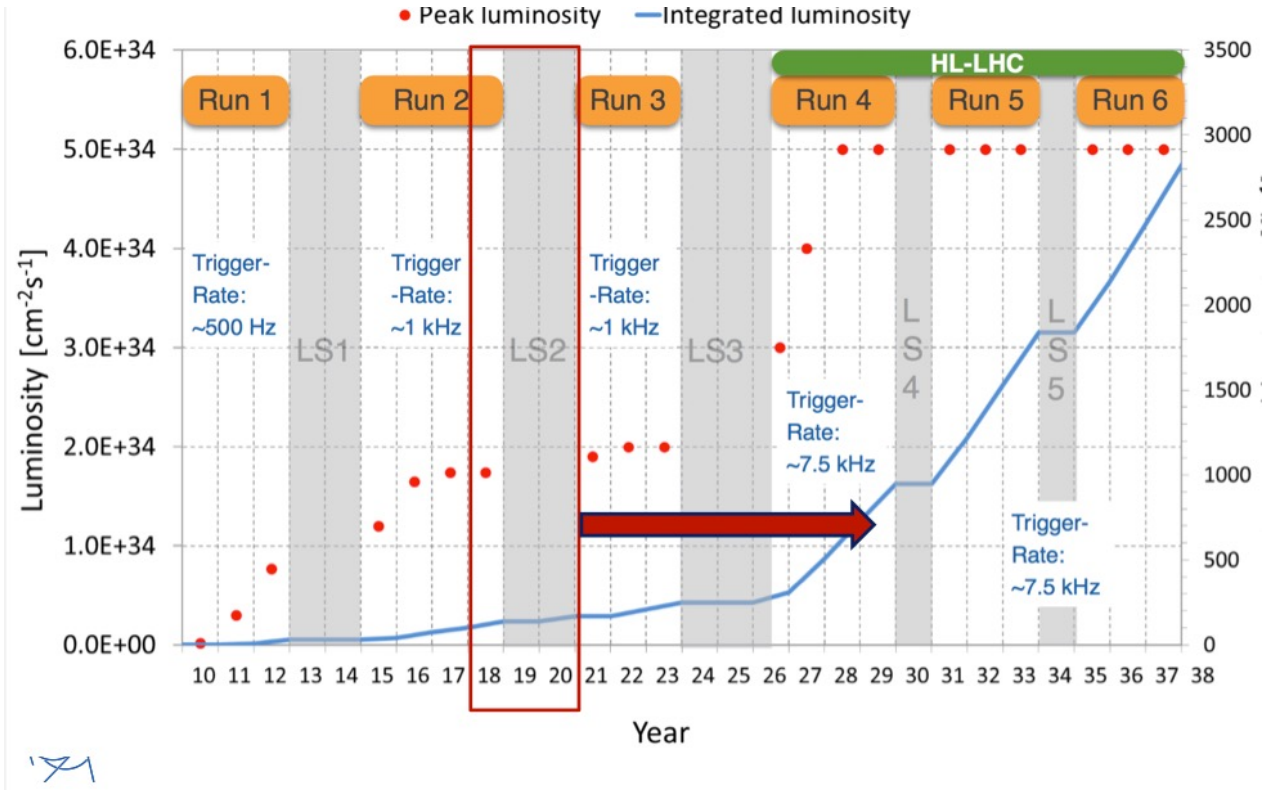

ATLAS: Computational needs in the Future

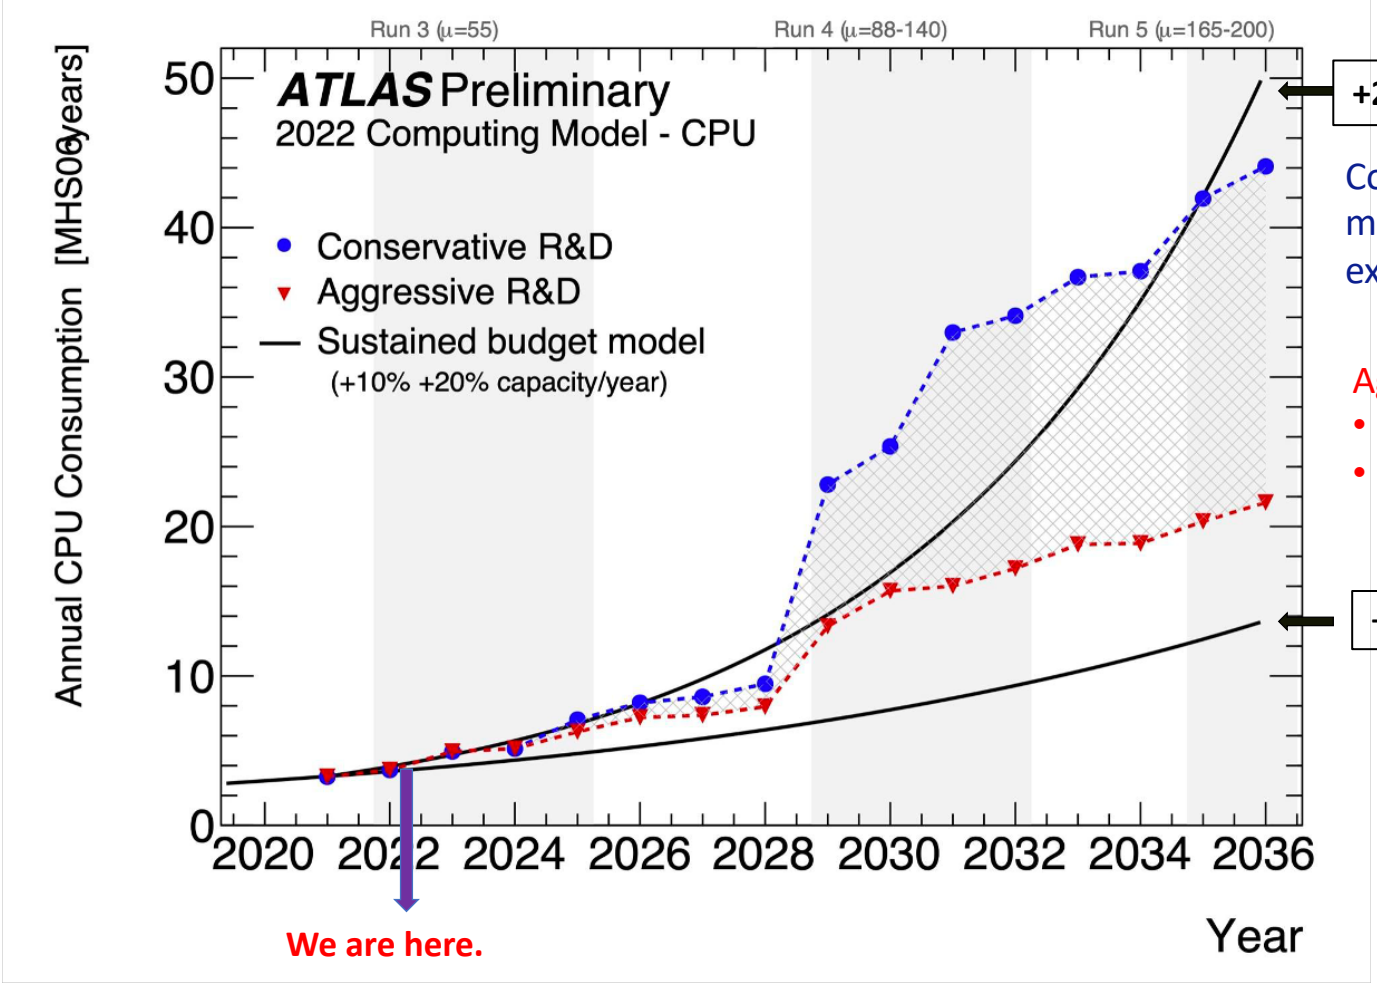

+20% improvement in capacity annually

Conservative: Existing person power level maintained at current level with current level of expertise.

Aggressive: Increase in expertise

- Increase the person power
- Increase the expertise of existing person power

+10% improvement in capacity annually

DUNE: Unique Challenges & New Opportunities

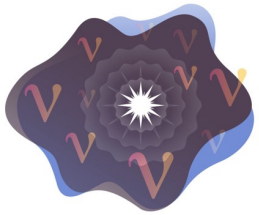

Super Nova induced event $\approx$ $100\ TBs$ in ~100 seconds

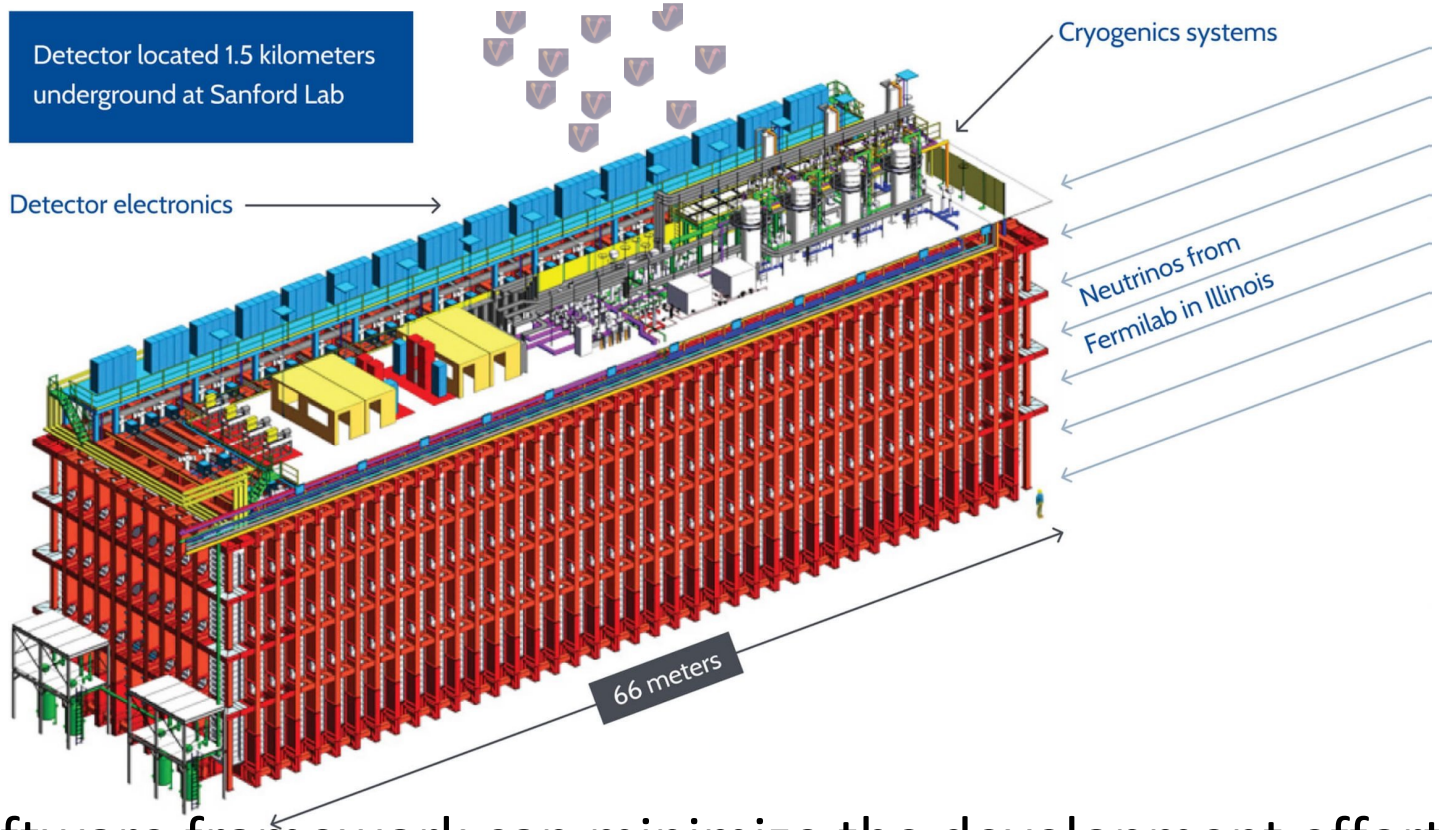

Beam induced event readout size $\approx 6\ GB$

DUNE Software framework can minimize the development effort by leveraging the advances made in HPC ecosystem (parallelism, GPU accelerators etc)

Computing Resources Used in HEP Experiments

- **High Throughput Computing ('Grid')**
  - Distributed systems with a relatively slow network (loosely-coupled jobs)
  - Batch processing with a large number of relatively independent jobs
- **High Performance Computing ('Supercomputing')**
  - Parallel systems with nodes designed for compute-intensive tasks and a fast network (tightly-coupled jobs)
  - Batch processing with a small number of large individual jobs

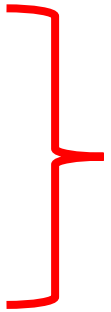

Most HEP Experiments Rely on This

HEP Computational Requirements and HPCs Performance

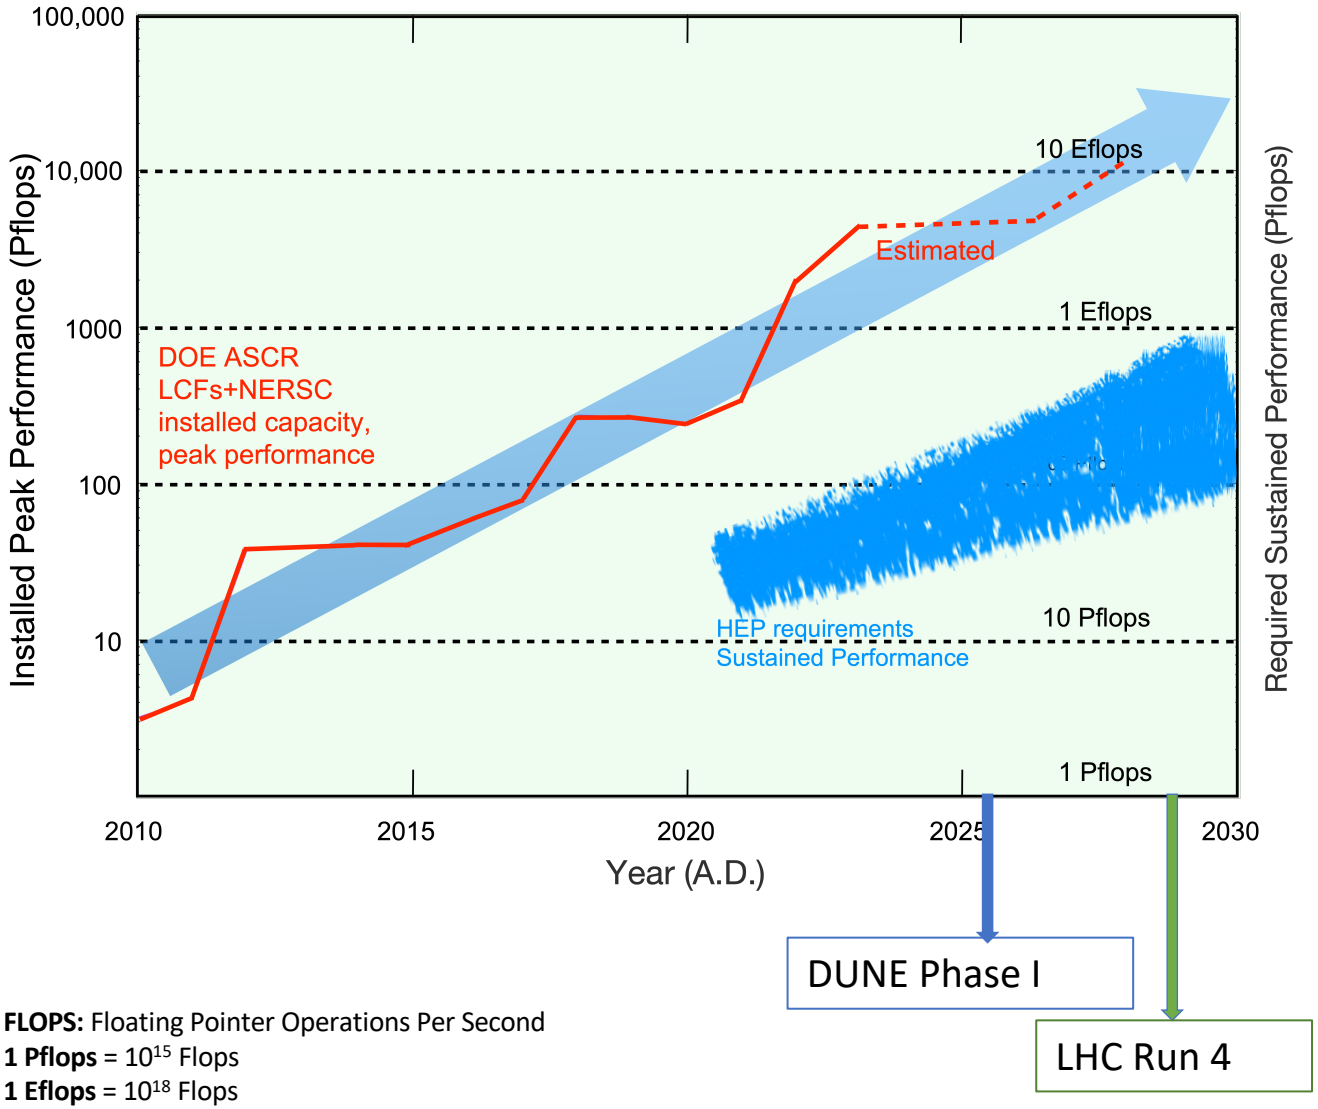

FLOPS: Floating Pointer Operations Per Second
1 Pflops = 10^{15} Flops
1 Eflops = 10^{18} Flops

Overall Outlook:

- Increased computing requirements over coming years.
- HPC's can fulfill the computing needs through the era of DUNE and LHC run 4.

HEP community could benefit by leveraging the HPC ecosystem.

HEP Requirements for HPC friendly Storage systems

ROOT has been the workhorse of the HEP experiments.

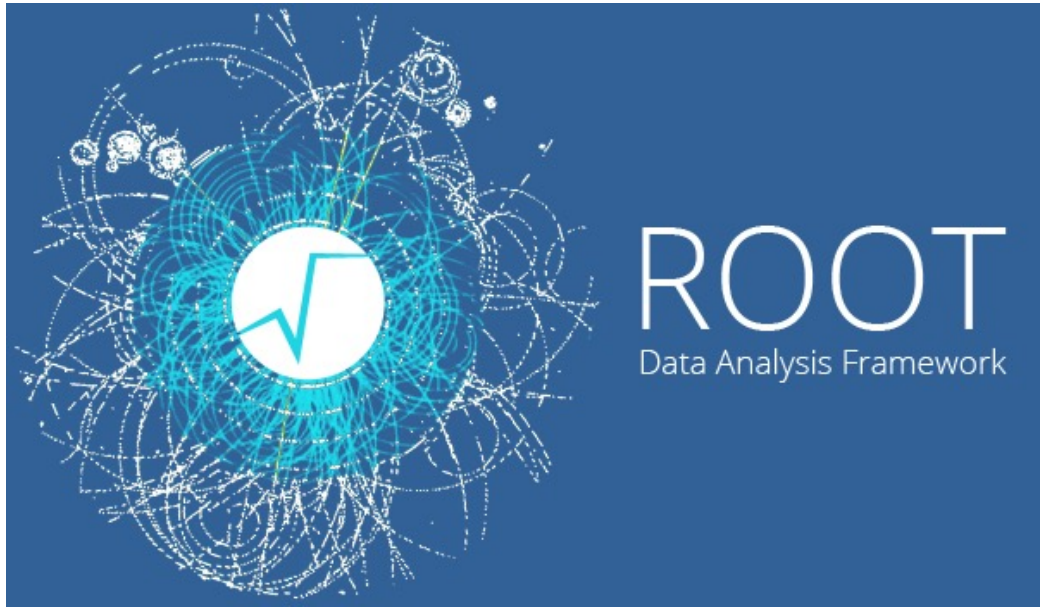

- Data Storage
- Data Analysis

Will continue to play a significant role in HEP field.

HPC has different software/hardware architectures.

- HDF5/PnetCDF optimized for HPC

HPC Friendly file systems

HPCs use parallel file systems for data-storage and access.

- HDF5, PnetCDF etc provide high level I/O libraries
  - Interface between user and parallel file systems

High level libraries to deal with otherwise complex parallel I/O

Exploring HDF5 as a use case for HEP data

- Using ROOT to serialize HEP data products (tested with ATLAS, CMS and DUNE data) and store into the HDF5 Format.
  - Multi threading process
  - Collective I/O currently being pursued.

Two methods of data mapping methods currently being explored

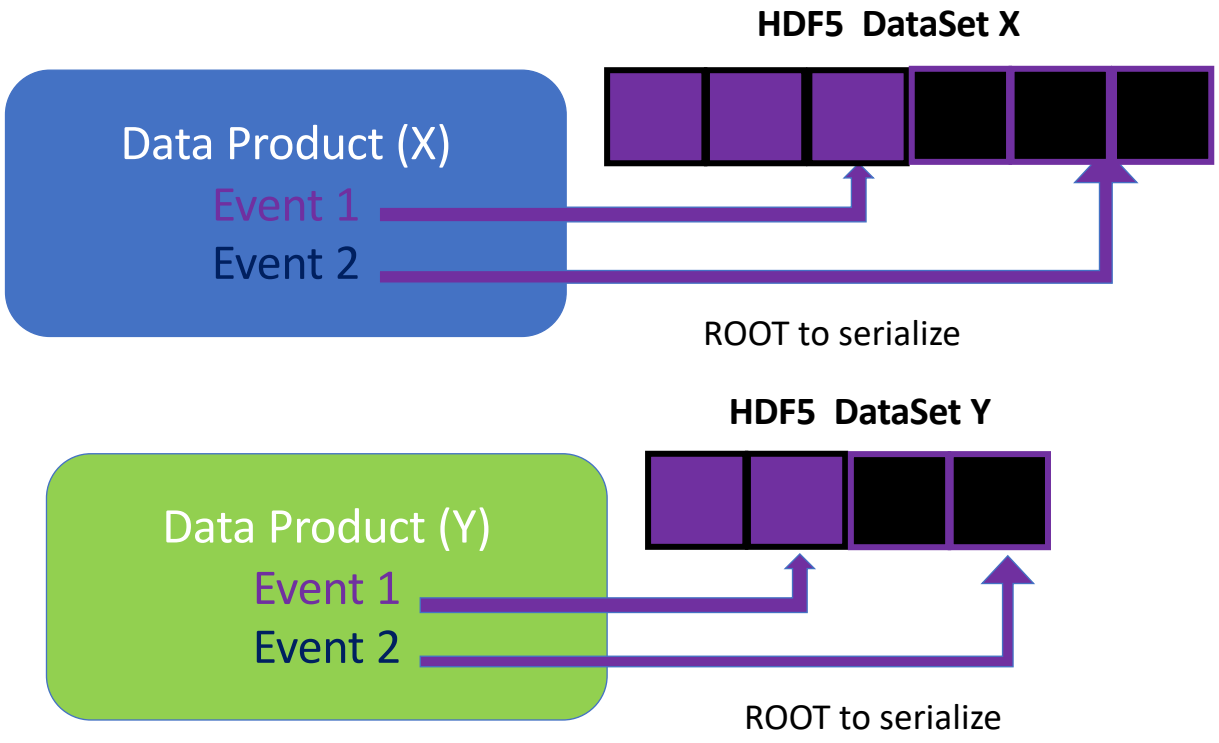

And X and Y navigational metadata dataset

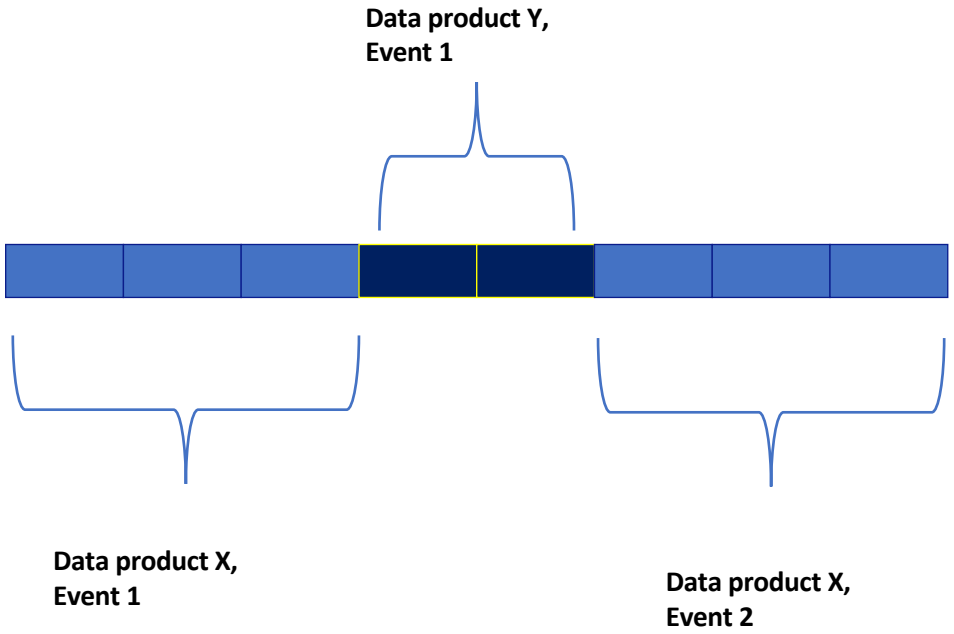

And 1 navigational metadata dataset

On going Performance Studies

The results itself should be taken with grain of salt (or sugar?) but basically it shows that how we map our data affects the I/O performance.

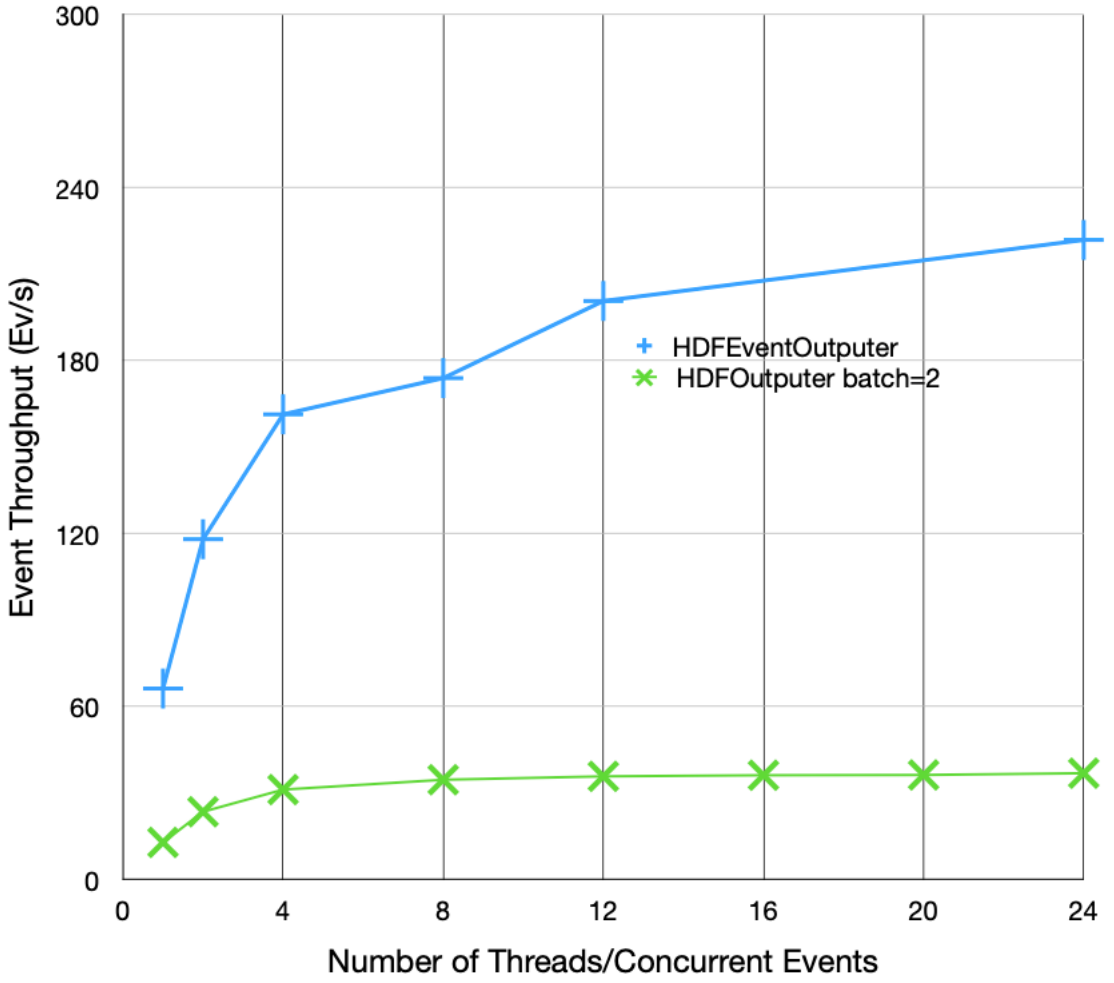

Using Profiling tool to understand I/O

Using Profiling Tool to look at the computing performance.

Using 36 parallel processes to process 1000 events

Shows time taken by each process to write the events into individual file which is later read by the merging process before merging into single file.

File used is ROOT.

Currently working/testing for the HDF5 file formats.

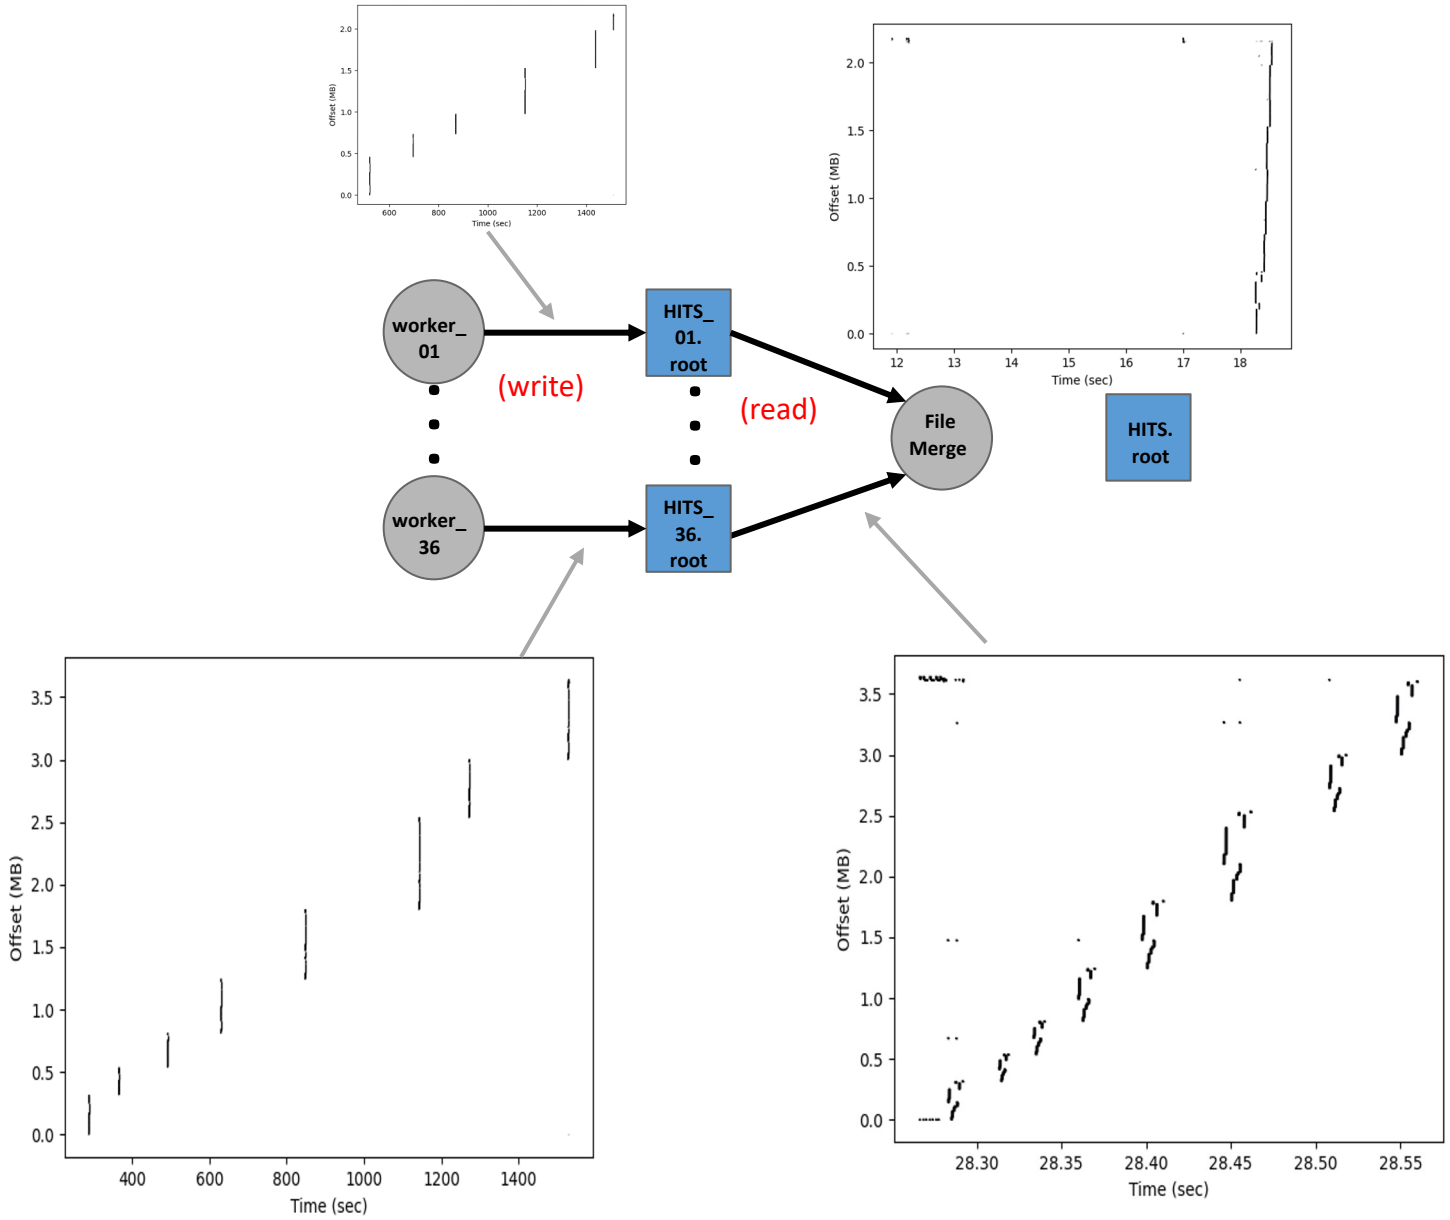

Testing Data Models to offload into the GPUs

HPCs use accelerators like GPUs to increase their performance.

GPUs cannot process complex data models.

HEP data models that can be offloaded into the GPUs.

Figure shows a diagram where computing load is divided between CPU and GPU.

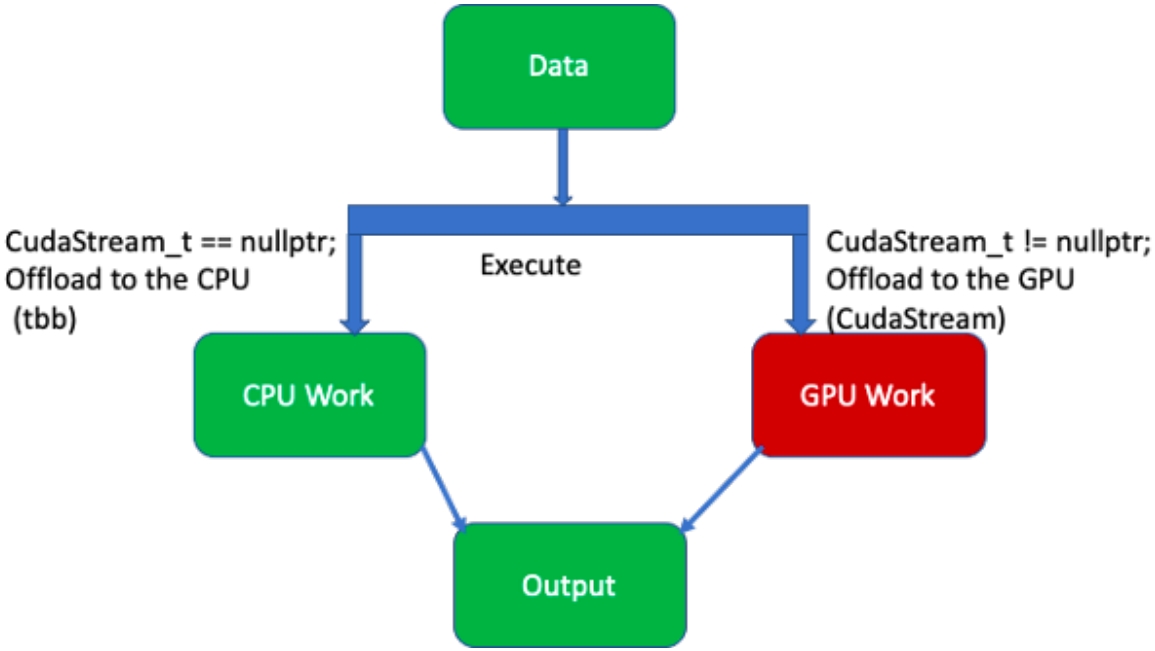

Conclusion

- Most of the works are exploratory
  - Future computing challenges of HEP requires a change in current compute/storage models
  - Finding ways to expand the roles of HPCs in HEP
- ROOT: Major tool for HEP data flow (processing + storage)
- HPC Friendly Storage
  - Study of file systems like HDF5 which are optimized for HPCs
  - Development of data models to optimize I/O
  - Development of data models to utilize GPU resources