Kickstarting the ePIC Computing Plan

David Lawrence - JLab

July 18, 2023

DPAP era Computing Plans

Page | Discussion

Software and Computing

Computing Strategy

We envision a distributed approach for computing at the EIC where possible. This was already the case during the proposal period, where we heavily relied on S3 storage at BNL to enable the simulation effort to take place at various computing sites. Scaling to many sites with differing computing infrastructures and practices benefits from using prevalent standardizations where possible over custom solutions and frontends (e.g., no custom schedulers, no custom RDM). Our containers were deployed successfully to the OSG, JLab, BNL, Compute Canada, ALCF/Theta, LCRC/Bebop, NERSC/Cori, and INFN-CNAF.

The components that integrate with the DAQ, by necessity, have to occur close to the detector (on-site at BNL). Here we will benefit from Gaudi's proven track record to integrate well into a DAQ environment (see, e.g., the LHCb High-Level Trigger system). We will immediately distribute the experimental data, after the necessary data reduction, to various tier-2 sites worldwide. Here we can build on the considerable investments made into the ESNet infrastructure. We intend to leverage the future exascale machines at the DOE leadership computing facilities, as well as other HPC and HTC worldwide, for data processing.

We currently anticipate deploying Rucio to enable the distributed data storage, and other solutions involving GridFTP or Globus solutions are also under investigation. The proposed map-reduce-style workflow could either be implemented in a fully distributed fashion, or with primary storage residing at the host labs BNL and JLab. For example, in order to mitigate the risk of saturation at BNL, we imagine a distributed model where analysis products analyzed at different sites could be streamed back with a higher priority to the data center at JLab (before duplication back to BNL during periods of lower network load). This would avoid some of the load issues experienced during heavy read & write periods when relying on a single central hub at BNL.

https://wiki.bnl.gov/athena/index.php?title=Software_and_Computing

View PDF | Download full issue

ELSEVIER

Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment

Volume 1047, February 2023, 167859

Scientific computing plan for the ECCE detector at the Electron Ion Collider

J.C. Bernauer [52 53 54], C.T. Dean [30 34], C. Fanelli [34], J. Huang [4], K. Kauder [4], D. Lawrence [27], J.D. Osborn [46 4], C. Paus [34], J.K. Adkins [33], Y. Akiba [50], A. Albataineh [66], M. Amaryan [44], I.C. Arsene [70], C. Ayerbe Gayoso [35], J. Bae [58], X. Bai [75], M.D. Baker [4 27], M. Bashkanov [83], R. Bellwied [64], F. Benmokhtar [15]…P. Zhuang [60]

Show more

+ Add to Mendeley | Share | Cite

https://doi.org/10.1016/j.nima.2022.167859 ↗

Get rights and content ↗

Abstract

The Electron Ion Collider (EIC) is the next generation of precision QCD facility to be built at Brookhaven National Laboratory in conjunction with Thomas Jefferson National Laboratory. There are a significant number of software and computing challenges that need to be overcome at the EIC. During the EIC detector proposal development period, the ECCE consortium began identifying and addressing these challenges in the process of producing a complete detector proposal based upon detailed detector and physics simulations. In this document, the software and computing efforts to produce this proposal are discussed; furthermore, the computing and software model and resources required for the future of ECCE are described.

https://doi.org/10.1016/j.nima.2022.167859

Worldwide LHC Computing Grid (WLCG)

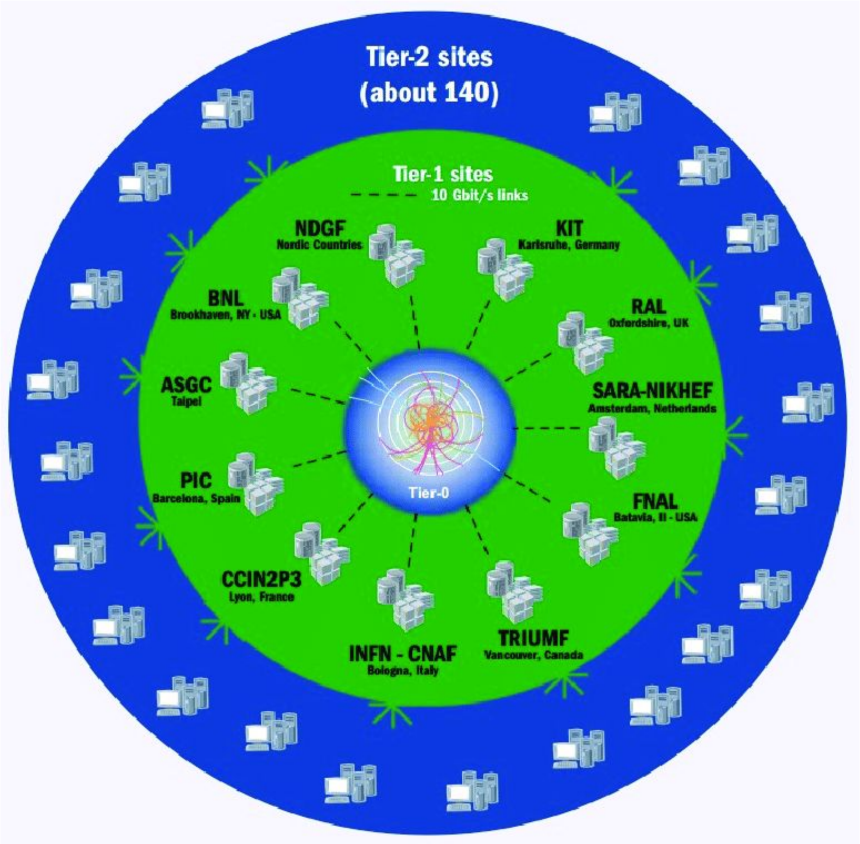

- **Tier 0 - LHC**
  - Store all raw data
  - 20% of LHC computing
  - Initial reconstruction
  - Distribute raw data and initial recon to all Tier 1 sites

- **Tier 1 - multiple sites(13)**
  - Store subsets of raw data
  - Compute for large scale reprocessing
  - Distribute data to Tier 2 sites
  - Store simulated data from Tier 2 sites

https://home.cern/science/computing/grid-system-tiers

Butterfly Model for Distributed Computing

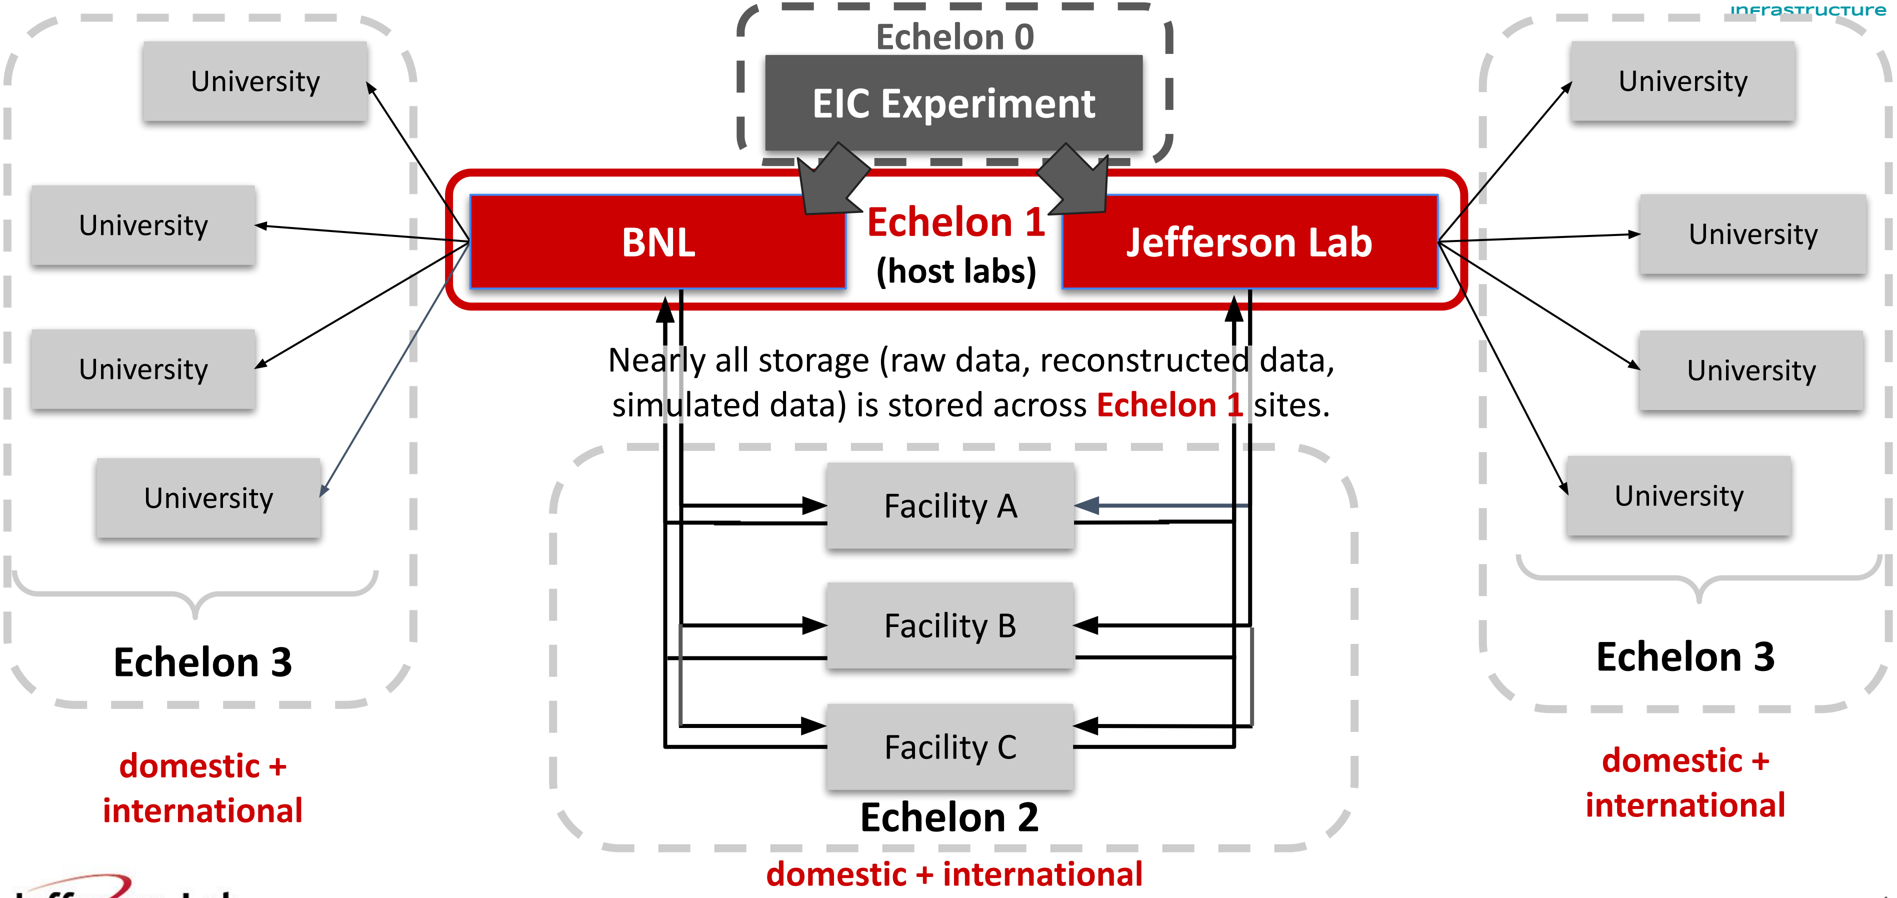

https://science.osti.gov/Funding-Opportunities/Digital-Data-Management

DOE Requirements

To integrate data management planning into the overall research plan, the following requirements will apply to all Office of Science research solicitations and invitations for new, renewal, and some supplemental funding issued on or after October 1, 2014. **These requirements apply to** proposals from all organizations including academic institutions, **DOE National Laboratories,** and others. These requirements do *not* apply to applications to use Office of Science user facilities.

All proposals submitted to the Office of Science for research funding must include a Data Management Plan (DMP) that addresses the following requirements:

1. DMPs should **describe whether and how data** generated in the course of the proposed research **will be shared and preserved**. If the plan is not to share and/or preserve certain data, then the plan must explain the basis of the decision (for example, cost/benefit considerations, other parameters of feasibility, scientific appropriateness, or limitations discussed in #4). At a minimum, DMPs must describe how data sharing and preservation will enable validation of results, or how results could be validated if data are not shared or preserved.
2. DMPs should provide a **plan for making all research data displayed in publications** resulting from the proposed research **open, machine-readable, and digitally accessible to the public** at the time of publication. This includes data that are displayed in charts, figures, images, etc. In addition, the underlying digital research data used to generate the displayed data should be made as accessible as possible to the public in accordance with the principles stated above. This requirement could be met by including the data as supplementary information to the published article, or through other means. The published article should indicate how these data can be accessed.
3. DMPs should **consult and reference** available information about **data management resources** to be used in the course of the proposed research. In particular, DMPs that explicitly or implicitly commit data management resources at a facility beyond what is conventionally made available to approved users should be accompanied by **written approval from that facility**. In determining the resources available for data management at Office of Science User Facilities, researchers should consult the published description of data management resources and practices at that facility and reference it in the DMP. Information about other Office of Science facilities can be found in the additional guidance from the sponsoring program.

4. DMPs must **protect confidentiality, personal privacy**, Personally Identifiable Information, and U.S. national, homeland, and economic security; recognize proprietary interests, business confidential information, and intellectual property rights; avoid significant negative impact on innovation, and U.S. competitiveness; and otherwise be consistent with all applicable laws, regulations, and DOE orders and policies. There is no requirement to share proprietary data.

DMPs will be reviewed as part of the overall Office of Science research proposal merit review process. Additional requirements and review criteria for the DMP may be identified by the sponsoring program or sub-program, or in the solicitation.

EIC Streaming Readout Architecture

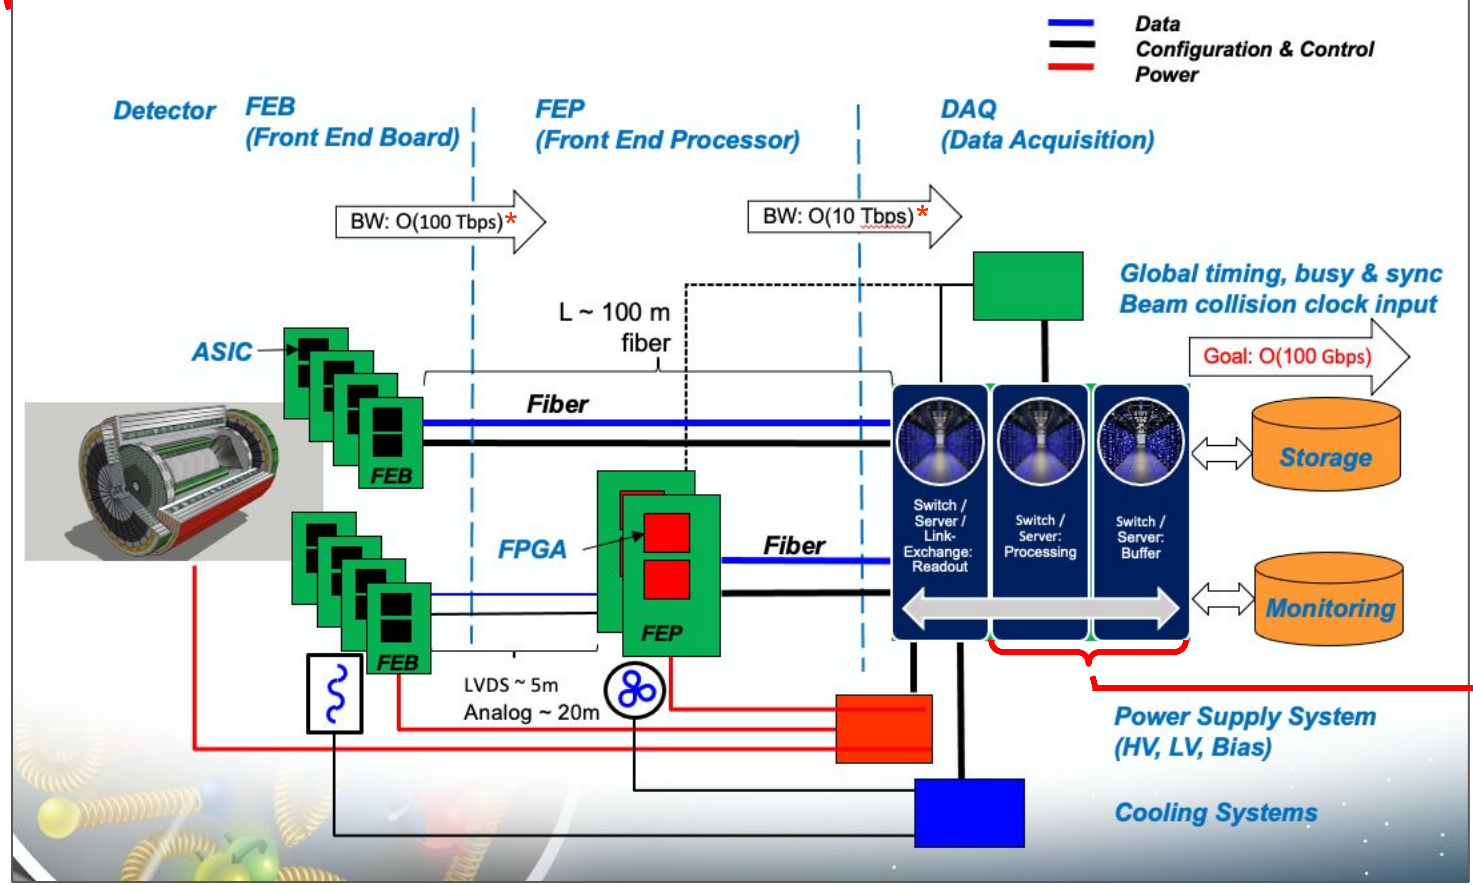

factor of 100 in data reduction

*ATHENA estimates assumed much more suppression at early stages, but still 100Gbps output. (See J. Landgraf talk at SRO X)

EIC Streaming Readout (From Fernando Barbosa's talk at AI4EIC Sep. 9, 2021)

DAM Boards

Data Aggregation Modules

OLD NUMBERS

Detector system	DAM boards	Channel/Fiber Count
Barrel		
Si Tracker	4	100 fibers
uRWell	12	278,000 channels, 576 fibers
AC-LGAD TOF	30	1400 fibers
hpDIRC	5	200 fibers
BECAL	2	9,088 channels , 72 fibers
iHCAL + oHCAL	1	3,264 channels, 26 fibers
Forward		
AC-LGAD TOF	6	300 fibers
dRICH	5	[illegible]20 fibers
FEMC	8	47,85[illegible] channels, 375 fibers
LFHCAL	10	58,590 channels, 460 fibers
Backwards		
mRICH	7	288 fibers
AC-LGAD TOF	3	150 fibers
EEMC	1	2878 channels, 24 fibers
Far-Forward		
B0 Detector, Roman Pots, Off-Momentum Detectors, ZDC	26	7.4M
Far-Backward		
Luminosity Monitor & Low-Q^2 Tagger	18	4.9M
Sum	**138**	

- Transition data to COTS Computing
- Built-in FPGA provides processing/Data aggregation
  - Questions:
    Where does code live? *DAQ/"Offline" boundary*
    Do we need single code base for simulation and real data?
    Emulation?

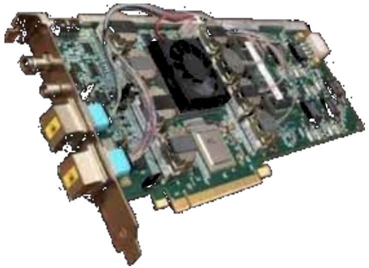

ATLAS FELIX board is an example of a DAM board

COTS = Commercial Off The Shelf
DAM = Data Aggregation Module

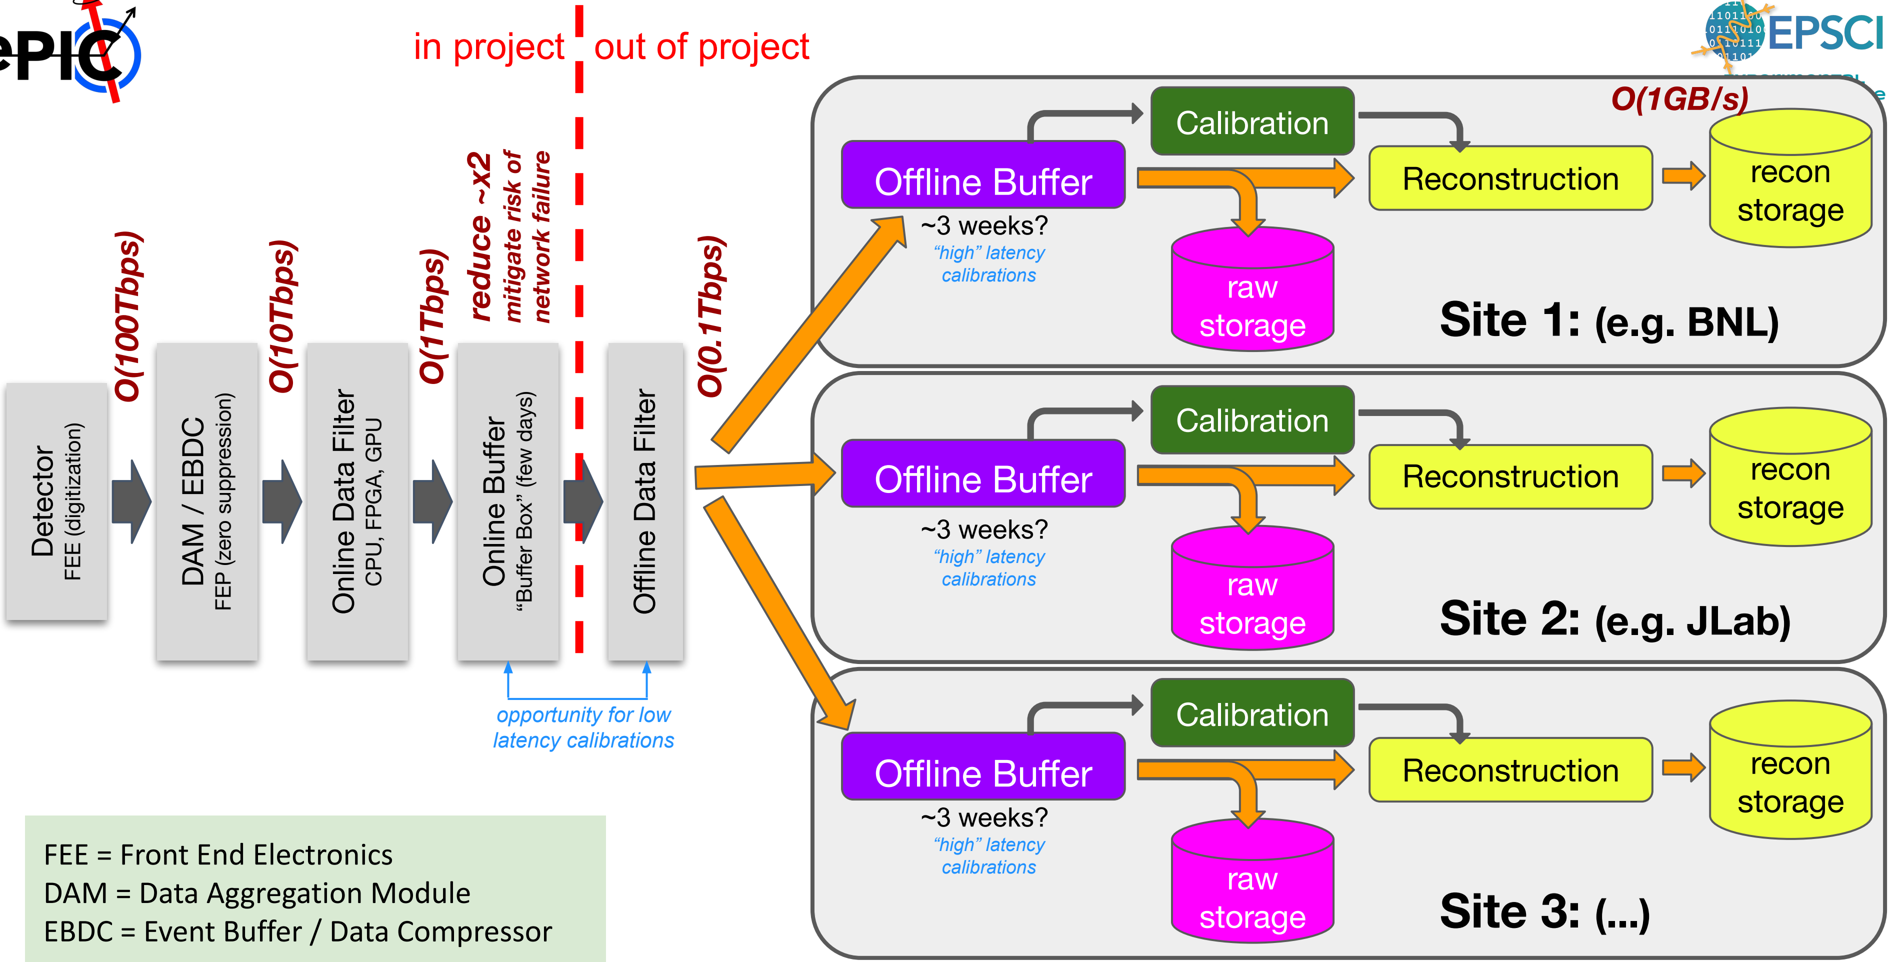

FEE = Front End Electronics
DAM = Data Aggregation Module
EBDC = Event Buffer / Data Compressor

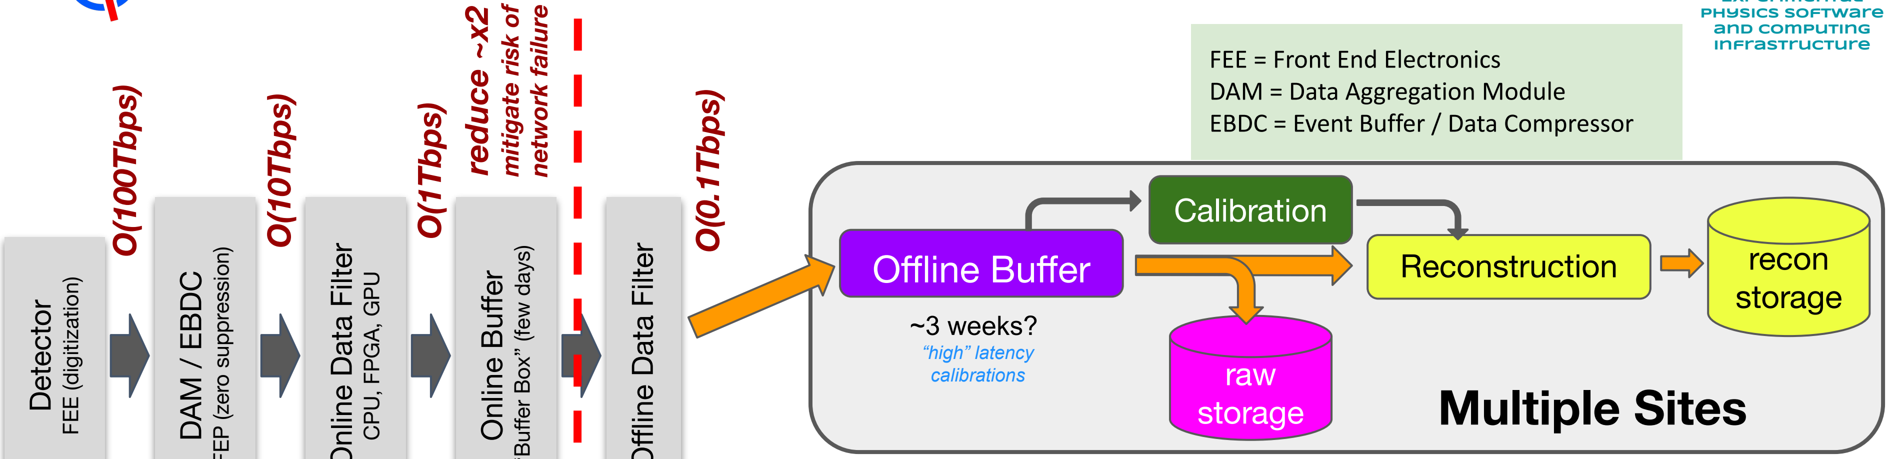

in project

out of project

Raw data will be written to tape as a safety mechanism. We hope to never need to read it!

=

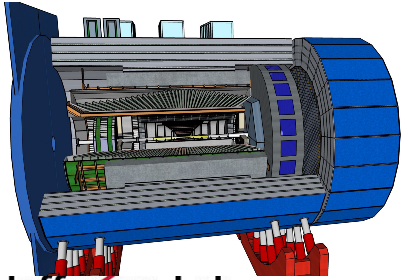

DAQ, BUFFERING, FILTERING, CALIBRATION, RECONSTRUCTION, ANALYSIS ?

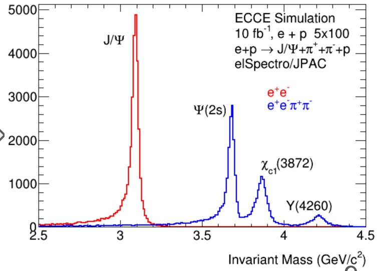

Data Reduction Stages (proposed)

Table 3
Data rates and reduction factors for proposed near real time data flow. Estimated data rate from ECCE detector is $\mathcal{O}$ (100 Tbps). Raw storage will be $\mathcal{O}$ (100 Gbps). Reconstructed object storage will be $\mathcal{O}$ (10 Gbps). Parentheses indicate technologies that could be used, but seem less likely choices.

Stage	Input/Output	Reduction factor	Technology options
Compute Interface (e.g. FELIX)	100 Tbps/10 Tbps	$\times 10^{-1}$	FPGA
Online event filter	10 Tbps/1 Tbps	$\times 10^{-1}$	FPGA, (GPU), CPU
Online buffer	1 Tbps/0.5 Tbps	$\times 5 \times 10^{-1}$	$< disk >$
Offline event filter	0.5 Tbps/100 Gbps	$\times 2 \times 10^{-1}$	FPGA, GPU, CPU
Reconstruction	100 Gbps/10 Gbps	$\times 10^{-1}$	(FPGA), GPU,CPU
Total	**100 Tbps/10 Gbps**	$\times 10^{-4}$	

Network bandwidth

https://www.nngroup.com/articles/law-of-bandwidth/#:~:text=Nielsen's%20law%20of%20internet%20bandwidth%20states%3A,grows%20by%2050%25%20per%20year

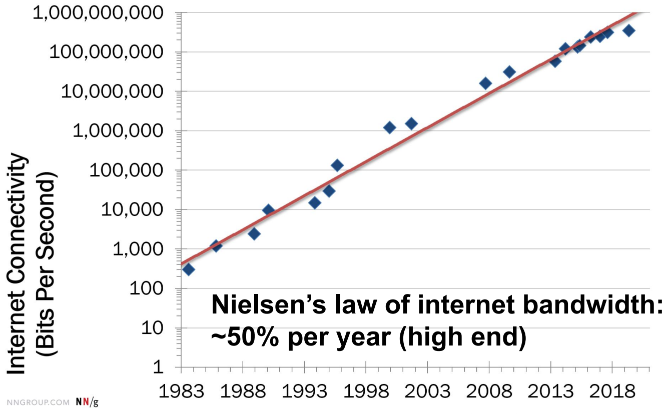

Current BNL ESnet bandwidth is **2x200Gbps**. If **EIC** produces **100Gbps**, it would take **25%** of the total while the experiment is running.

If high-end internet bandwidth grows by roughly 50% per year, then in 10 years, bandwidth should grow by a factor of

$$\sim 1.5^{10} = 58$$

(i.e. marketplace should provide required technology)

Thus, if the BNL bandwidth grows at a rate consistent with Nielsen's Law, a **100Gbps** experiment would use **less than 0.5%** of the total offsite bandwidth.

A **1Tbps** data stream would use **~5%** of the total offsite bandwidth.

nota bene!

- Network upgrades must be motivated by requirements.
- They tend to happen in big steps, not continuously.
- Major network upgrades require planning with long lead times.

EJFAT

Channel Aggregation + Three Tier Horizontal Scaling

Streaming systems may take advantage of ***in-network*** time slice **building** with smart load balancers

Colors → Events
Shapes → Channels (ROCs)

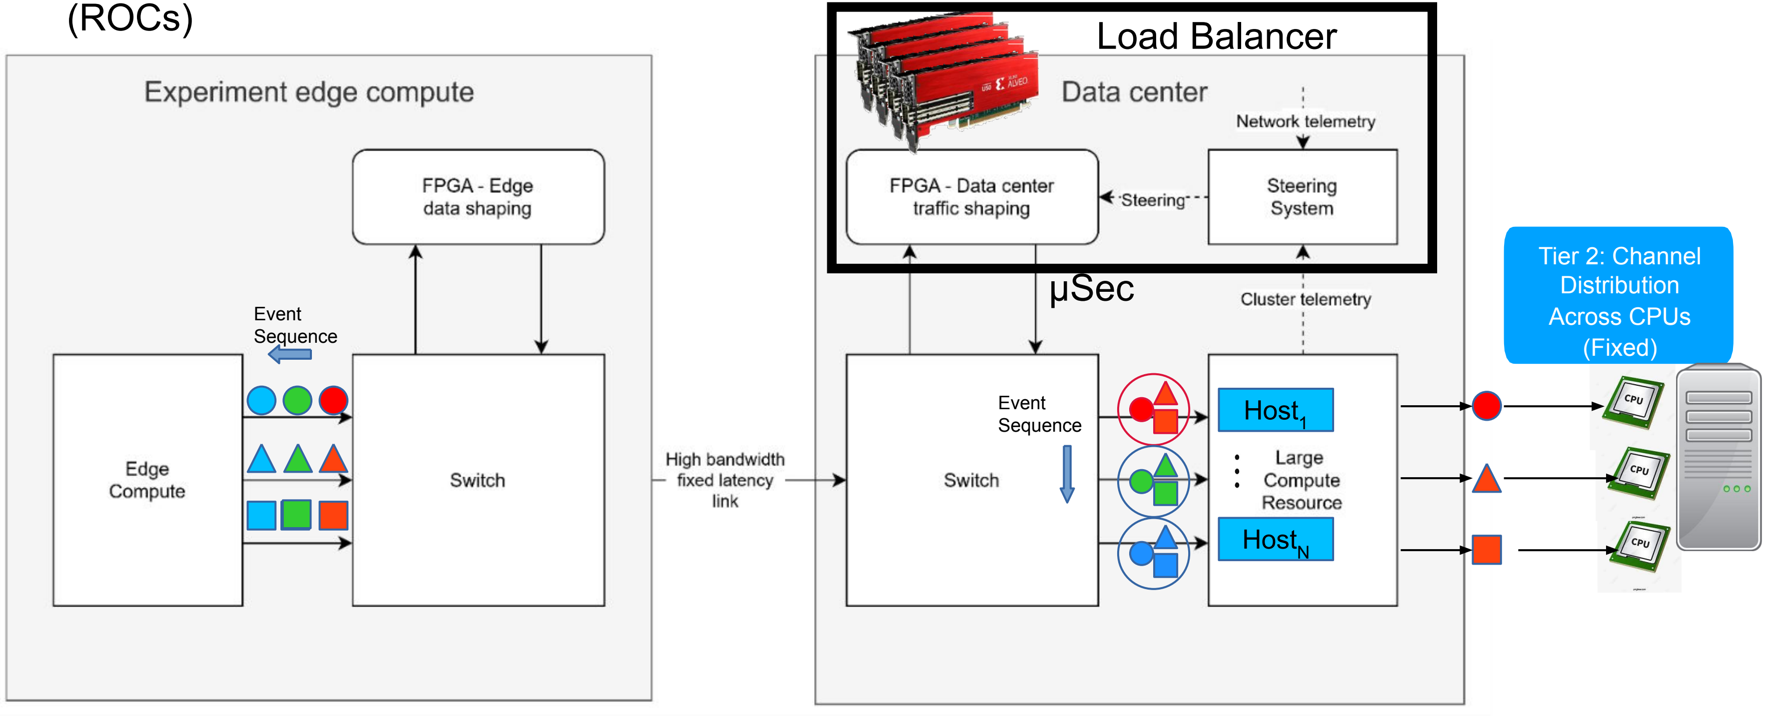

This slide from CHEP2023

See talk by Michael Goodrich *May 11, 2023, 11:15 AM* (Track X)
https://indico.jlab.org/event/459/contributions/11804/

Simulation Storage and Compute

Near term simulation campaigns. Expect drop off as collaboration transitions from design focus to construction.
Ramp back up for data challenges just prior to production.

Year	Number of events [$\times 10^6$]	Storage [TB]	CPU-core hours [Mcore-h]
2022	200	50	45
2023–2024	100	25	22.5
2025–2028	50	12.5	11
2029–2030	500	125	110
Total	1600	400	354

Simulation requirements for first 3 years of experiment production. Assumes 10% of measured cross-section is of interest and that we will need O(10) times more simulation than that. (i.e. factor of "1").

Year	Number of events [$\times 10^9$]	Storage [PB]	CPU-core hours [Mcore-h]
Year-1	120	30	11000
Year-2	600	150	55000
Year-3	5400	1300	490000

Tape Storage for Raw Data

Table 4

Estimate of **raw data tape storage** needed for **first 3 years of EIC running**. Values are estimates assuming ramp up to full luminosity by year 3. Numbers for the first two years are estimated for the purposes of this exercise and do not come from an external source. n.b. each value represents *only* the needs for data produced in that year and *not* a cumulative total.

New storage	Year-1	Year-2	Year-3
Luminosity	10^{33} cm^{-2} s^{-1}	2×10^{33} cm^{-2} s^{-1}	10^{34} cm^{-2} s^{-1}
Weeks of physics running	10	20	30
Operational efficiency	40%	50%	60%
Data rate to storage	6.7 Gbps	16.7 Gbps	100 Gbps
Raw data storage (no duplicates)	4PB	20PB	181PB
Recon storage	0.4PB	2PB	18PB
Total Storage (no duplicates)	7PB	35PB	317PB

Live Disk Space

Table 5

Estimate of the **disk** storage **needed** for the **first 3 years** of EIC running. The **temporary disk** is used to hold the raw data for a **3 week period** while calibrations are derived and reconstruction is done. The **permanent disk** is for holding the **reconstructed data**. This will be cumulative so collaborators will have access to recon data from all years.

Total disk	Year-1	Year-2	Year-3
Disk (temporary)	1.2PB	3.0PB	18.1PB
Disk (permanent)	0.4PB	2.4PB	20.6PB
TOTAL	1.6PB	5.4PB	38.7PB

Raw Data Compute

Table 6

Estimates of **CPU** needed for **reconstruction** of **raw data**. The number of seconds per event is highly dependent on the type of processor being used. Number of events comes from total raw data storage estimate in Table 4. **Calibration** is assumed to be **5% of reconstruction time**.

CPU compute	Year-1	Year-2	Year-3
Recon process time/core	5.4 s/ev	5.4 s/ev	5.4 s/ev
Streaming-unpacked event size	33 kB	33 kB	33 kB
Number of events produced	121 billion	605 billion	5443 billion
CPU-core hours (recon-only, 1 pass)	181 Mcore-h	907 Mcore-h	8165 Mcore-h
CPU-core hours (calib-only)	9 Mcore-h	45 Mcore-h	408 Mcore-h
2020-cores needed to process in 30 weeks	38k	189k	1701k

old ECCE numbers (*Fun4All*)

simulation+recon: ~9s/event
reconstruction: 5.4s/event

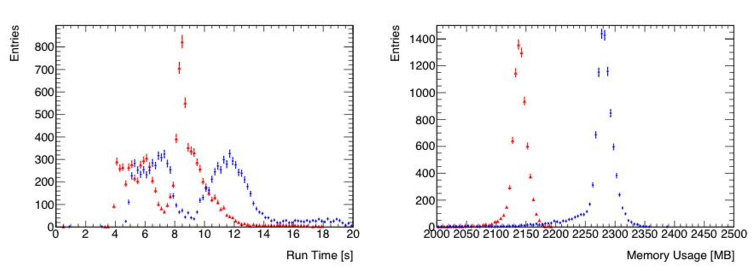

red = ep10x100 pythia8
blue = 18x275 pythia6

recent ePIC Simulation Campaign

(courtesy of Sakib Rahman)

sim. + recon. time

	A	B
1	Dataset	Time Per Event (s) (Excludes startup time)
2	BACKGROUNDS/10x100/merged_lowstat/bgmerged_10000_ep_noradcor_10x100_q2_10_100	8.9
3	DIS/NC/5x41/minQ2=1	2.7
4	DIS/NC/5x41/minQ2=10	3.4
5	DIS/NC/5x41/minQ2=100	4.5
6	DIS/NC/18x275/minQ2=1	13.0
7	DIS/NC/18x275/minQ2=10	13.0
8	DIS/NC/18x275/minQ2=100	16.8
9	DIS/NC/18x275/minQ2=1000	19.6
10	DIS/NC/10x100/minQ2=1	5.7
11	DIS/NC/10x100/minQ2=10	6.6
12	DIS/NC/10x100/minQ2=100	7.8
13	DIS/NC/10x100/minQ2=1000	8.9
14	DIS/CC/5x41/minQ2=100	2.8
15	DIS/CC/18x275/minQ2=100	14.3
16	DIS/CC/18x275/minQ2=1000	14.1
17	DIS/CC/10x100/minQ2=100	5.1
18	DIS/CC/10x100/minQ2=1000	5.2
19	SIDIS/pythia6/ep_5x41/hepmc_ip6/noradcor/ep_noradcor.5x41_q2_0_1	2.1
20	SIDIS/pythia6/ep_5x41/hepmc_ip6/noradcor/ep_noradcor.5x41_q2_1_10	2.6
21	SIDIS/pythia6/ep_5x41/hepmc_ip6/noradcor/ep_noradcor.5x41_q2_10_100	3.1
22	SIDIS/pythia6/ep_5x41/hepmc_ip6/noradcor/ep_noradcor.5x41_q2_100_1000	3.9
23	SIDIS/pythia6/ep_10x100/hepmc_ip6/noradcor/ep_noradcor.10x100_q2_1_10	4.8
24	SIDIS/pythia6/ep_10x100/hepmc_ip6/noradcor/ep_noradcor.10x100_q2_10_100	5.7
25	SIDIS/pythia6/ep_10x100/hepmc_ip6/noradcor/ep_noradcor.10x100_q2_100_1000	6.9
26	SIDIS/pythia6/ep_10x100/hepmc_ip6/noradcor/ep_noradcor.10x100_q2_1000_100000	8.5
27	SIDIS/pythia6/ep_18x275/hepmc_ip6/noradcor/ep_noradcor.18x275_q2_0_1	8.8
28	SIDIS/pythia6/ep_18x275/hepmc_ip6/noradcor/ep_noradcor.18x275_q2_1_10	10.2
29	SIDIS/pythia6/ep_18x275/hepmc_ip6/noradcor/ep_noradcor.18x275_q2_10_100	11.0
30	SIDIS/pythia6/ep_18x275/hepmc_ip6/noradcor/ep_noradcor.18x275_q2_100_1000	13.0
31	SIDIS/pythia6/ep_18x275/hepmc_ip6/noradcor/ep_noradcor.18x275_q2_1000_100000	18.2

slide courtesy of Sakib Rahman Univ. Manitoba

Computing Needs

- 2-3 tagged production campaigns per month
- 1 default detector config but multiple test configs is possible based on demand
- Benchmarked core year estimates for different campaigns for default config:
  - MM.YY.0 ~ 20 coreyears
  - MM.YY.1 ~ 30 coreyears
  - MM.YY.2 ~ 100 coreyears
- Each job requests ~3 GB memory and 2 hours on remote node. , Output may occupy ~2GB diskspace.

Memory Usage over Time

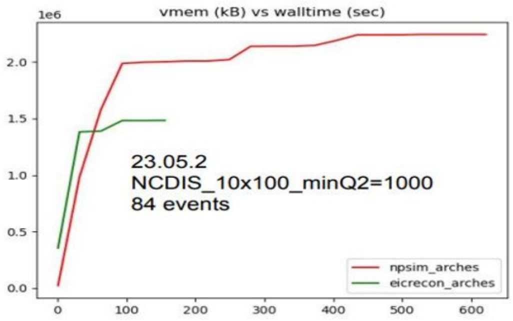

Monthly Campaigns Strategy

MM.YY.0 — Week1: Verification & Validation

MM.YY.1 — Week2: Targeted Development

MM.YY.2 — Week3+4: Development & Special Requests

Compute Usage in June, 2023

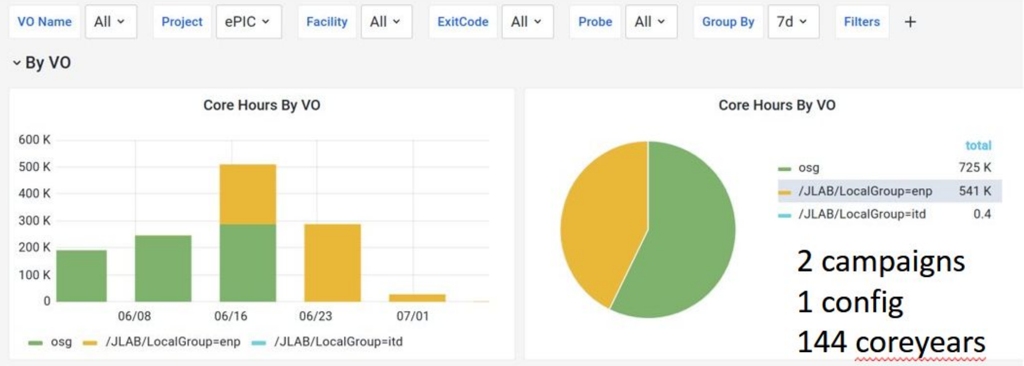

n.b. EECE number of 45M core-hours = **428 core-years/mo**

Challenges

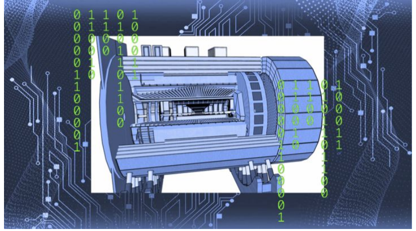

- Network Bandwidth?
  - Technology won't be a particular challenge, but **accurate estimates** will be needed in the next **~3 years** to motivate requirements
- Compute Resources?
  - Overall flops needed is unclear. (Large discrepancy between estimated and current actual core-years.)
  - How much of these are realized using **heterogeneous resources** and how to properly distribute to **compute centers** that have them is the challenge
- Storage?
  - Not technically challenging assuming normal COTS technology scaling
  - **Federated storage and federated read/write access** needs to be implemented (*already well underway*)
- Holistic streaming readout
  - Need detailed **stream-level simulation** to fully test software at all stages of stream processing
    - Realistic data rates at all stages
    - Software Trigger bias
  - What stages can/should alternative architecture hardware be used?
- Schedule **Data Challenges** with multiple sites using distributed model
  - Authentication
  - Monitoring

Other Challenges

- Copies of raw data at two host labs provides duplication at different locations.
- Copies of all reconstructed data at all Echelon 1 sites?
  - Possibly for duplication/archival purposes
  - University access should be channeled through catalogue tool (**e.g. Rucio**) to allow access to all files while enabling smart management of local storage resources
- Which portions of the data are each Echelon 1 site responsible for reconstructing?
  - Need mechanism to orchestrate (e.g. JIRIAF? DIRAC?, PANDA?)

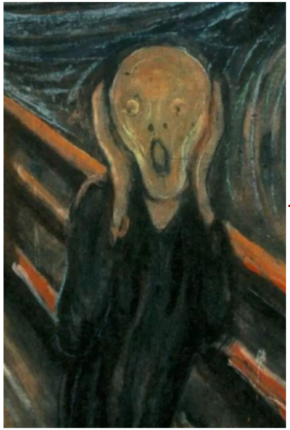

- What mechanism (software) do we use to manage calibration latency?
  - Need re-calibration requirements for all detector systems
    - *Quantifiable criteria that does not require human-in-the-loop!*
  - Trigger re-calibration procedure (e.g. automated job or expert notification)
  - Monitor for calibration completeness and trigger reconstruction jobs
  - “Offline” alarm system for calibration failures
  - How to distribute calibrations to all Echelons and ensure consistency between sites
    - Hypothetical:
      Data processing is assigned in 2-hour chunks to Echelon 1 sites
      Detector system A only needs recalibrating every 24 hours
      Which site generates the calibration?

Data Quality M

Machine Learning
Online Data
CMS Elect

Anomaly Detection S

- AE-based anomaly detection and l

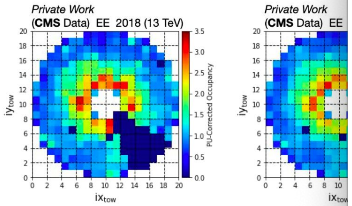

[1] **Input map** — Anomaly present
[2] **AE-recons** — Anomaly not r

- **Anomaly tagging threshold** is ob
  - Based on loss values of anom choose a **loss threshold** that

courtesy Kyungmin Park et al..

see also: http://cds.cern.ch/re

2023.07.18.ePIC_Computing_
AI4DQM workshop (25 August

indico.jlab.org/event/725/

JLab | Publications | EPSCI | GlueX | EIC | Robotics | Data Science | Personal | Other Bookmarks

Public | US/Eastern | D. Lawrence

AI4DQM workshop

25 August 2023
Virtual
US/Eastern timezone

Enter your search term

Overview
Registration

There is a growing desire to use AI/ML to monitor data quality as evident by the increasing number of presentations on the topic coupled with the demands of autonomous systems employed during data taking (e.g. Streaming Read-Out and triggering systems). The field at large would benefit from an exchanging of ideas, tools, and techniques that have been developed for AI/ML based data quality monitoring. The AI4DQM workshop seeks to provide a forum for the discussion of data quality monitoring both in technical solutions as well as sociological challenges with adoption in a fairly informal, virtual, setting. This one day workshop will consist of a series of quick talks with room for discussions.

Starts 25 Aug 2023, 09:00
Ends 25 Aug 2023, 17:00
US/Eastern

Virtual
https://jlab-org.zoomgov.com/j/1616655908?pwd=bThhemI0NC9sbC9kaGV3VHNlblIRdz09
Go to map

Thomas Britton
Torri Jeske

Registration
You are registered for this event.
See details

or Data Quality Monitoring

nt/459/contributions/11393/

monitoring of plots
odels (e.g. Inception-v3)
any number of classifications

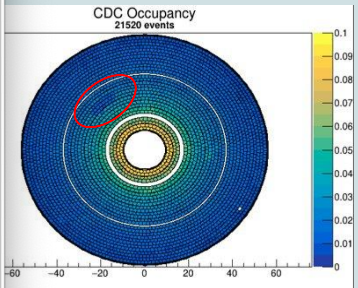

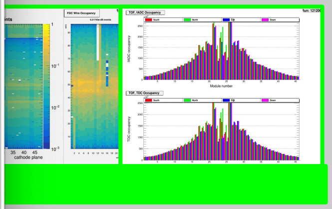

3, 11:45 AM Hampton Roads II
/459/contributions/11393/

Summary

- Distributed Computing **Butterfly Model** can utilize multiple **compute-only-facilities** while maintaining storage and central operations across a few Echelon 1 sites

- Streaming Readout Computing Model will likely require **review of Host Labs' Data Management Plans** with DOE

- Estimates for **raw data sizes** are important to ensure required **network support** will be in place when needed
    - older estimates indicate this will not be an issue

- Computing resource **requirements** may increase **exponentially** over the first few years of running
    - smart purchasing plans will require accurate numbers

- Fully **automated calibration** should be **required** for all **sub-detectors**
    - re-calibration criteria clearly defined
    - possible iterative procedures and inter-detector dependencies
    - monitoring
    - offline alarm system

BACKUPS

DAQ: Overview

- **Current design calls for a Streaming Data Acquisition System (SRO)**
  - Widely recommended by experts: EIC Computing Consortium, EIC Yellow Report
  - No need to wait for all signals from single crossing to read out data
  - Removes nearly all deadtime
  - Less restrictions for filter criteria and potentially less bias

FEE = Front End Electronics
DAM = Data Aggregation Module
EBDC = Event Buffer / Data Compressor

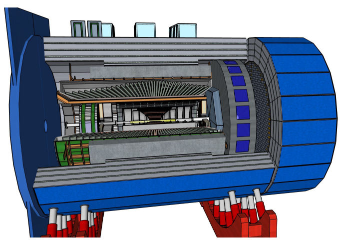

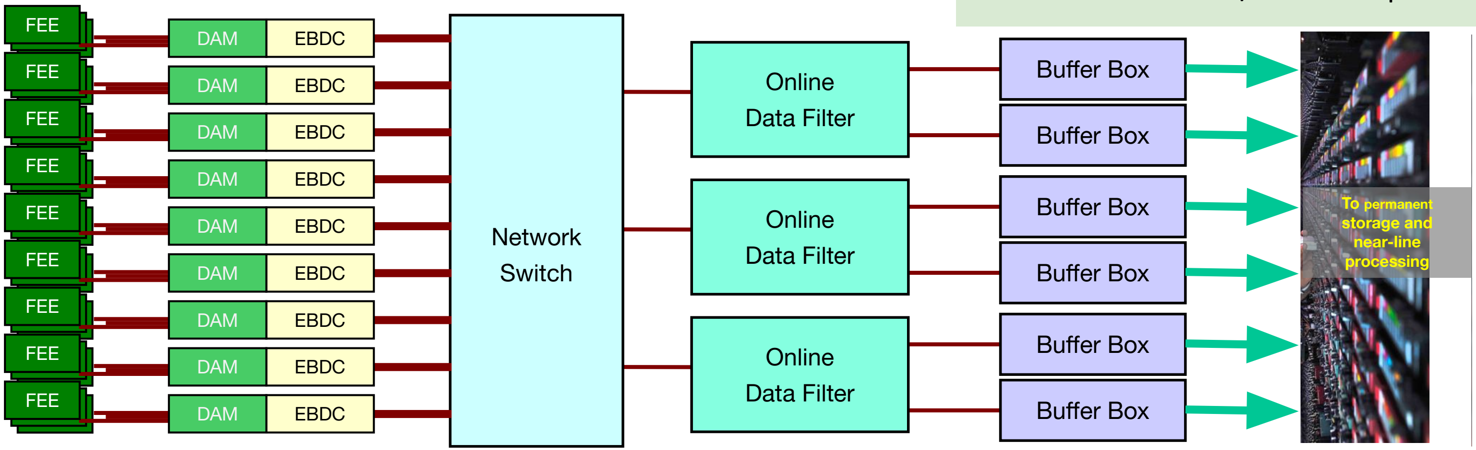

Heterogeneous Hardware

- Anticipate supporting heterogeneous hardware at all stages of pipeline
  - GPU, TPU, CPU, FPGA, …

- AI/ML support
  - Expected at all stages (design, simulation, filtering, calibration, reconstruction, analysis, …)
  - Type of hardware depends on AI model (some models run faster on CPU)

- Custom hardware-specific algorithms
  - Only if necessary
  - Easier for front end hardware-wise, but limits diagnostics/emulation offline
    - Tools like HLS or hls4ml can help apply single code base to multiple architectures

Raw Data Requirements *(estimated during ECCE DPAP)*

	year-1	year-2	year-3
Luminosity	$10^{33}\text{cm}^{-2}\text{s}^{-1}$	$2 \times 10^{33}\text{cm}^{-2}\text{s}^{-1}$	$10^{34}\text{cm}^{-2}\text{s}^{-1}$
Weeks of Running	10	20	30
Operational efficiency	40%	50%	60%
Disk (temporary) 3 week buffer	1.2PB	3.0PB	18.1PB
Disk (permanent) recon data	0.4PB	2.4PB	20.6PB
Data Rate to Storage	6.7Gbps	16.7Gbps	100Gbps
Raw Data Storage (no duplicates) raw data	4PB	20PB	181PB
Recon process time/core	5.4s/ev	5.4s/ev	5.4s/ev
Streaming-unpacked event size	33kB	33kB	33kB
Number of events produced	121 billion	605 billion	5,443 billion
Recon Storage	0.4PB	2PB	18PB
CPU-core hours (recon+calib)	191Mcore-hrs	953Mcore-hrs	8,573Mcore-hrs
2020-cores needed to process in 30 weeks	38k	189k	1,701k

ECCE Computing Plan https://doi.org/10.48550/arXiv.2205.08607

CPU

Simulation focused on next four years where detector refinement will be key

Year	Number of Events [$\times 10^6$]	Storage [TB]	CPU hours [kHours]
2022	200	50	450
2023	100	25	225
2024	100	25	225
2025	50	12.5	110
Total	450	112.5	1010

OLD NUMBERS

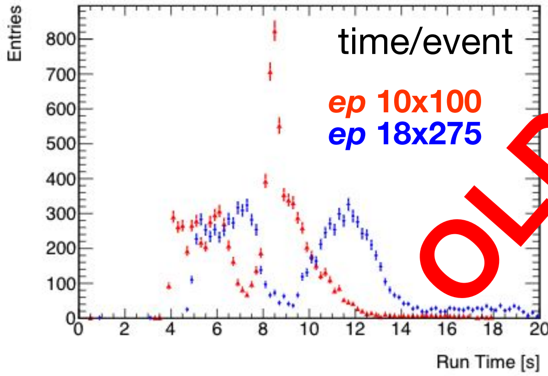

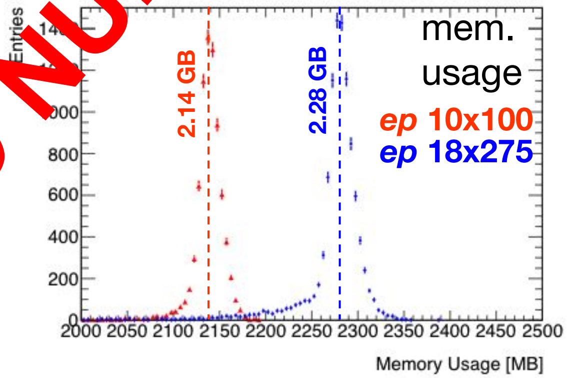

n.b. not all cores are equal. These reflect rough average of cores currently in use (i.e. some modern and some several years old)

Current External bandwidth of BNL and JLab

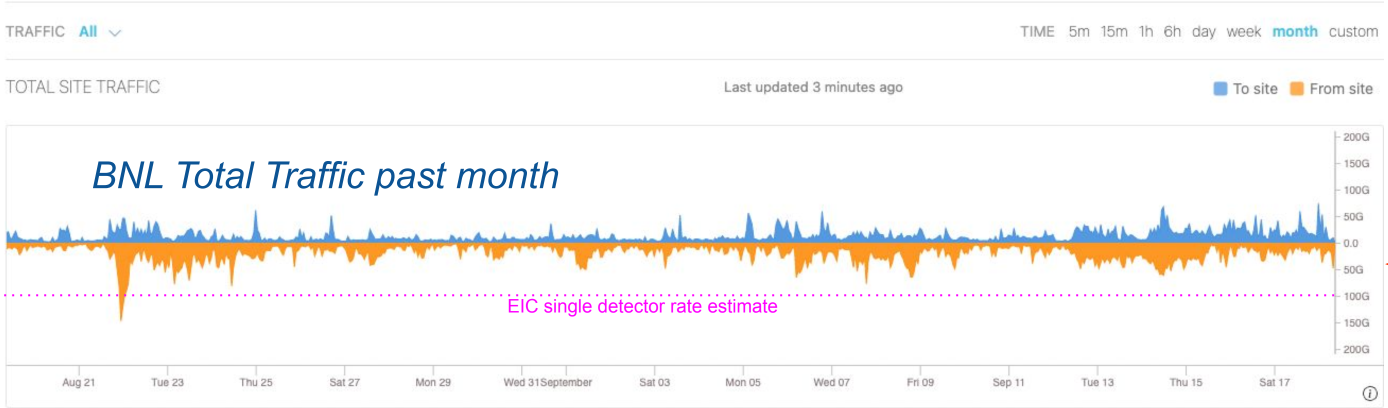

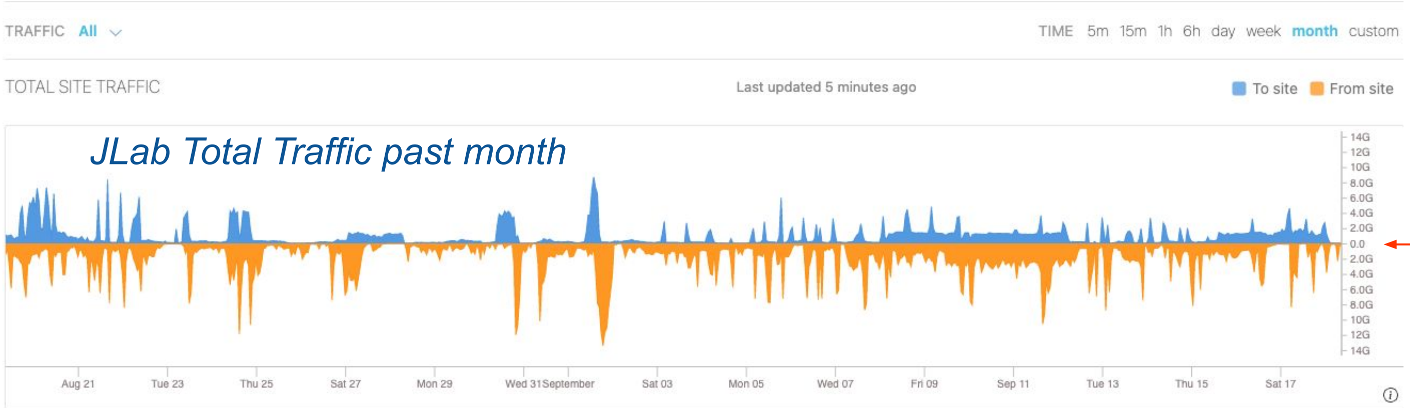

n.b. different scales

BNL currently has **2x200Gbps** bandwidth to ESNet

JLab currently has **2x10Gbps** bandwidth to ESNet (soon to be 2x100 Gbps)