EIC Computing at BNL

Torre Wenaus

PO Nuclear & Particle Physics Software Group Leader

Alexei Klimentov

NPP SDCC Interim Director

EIC Experimental Project Dependencies Review

August 22-23, 2024

Electron-Ion Collider

BNL EIC Host Lab Responsibility in Computing and Software

As the host lab for the EIC, BNL has major responsibility for EIC computing services, computing infrastructure (from online and till data is delivered and to be archived) and prompt reconstruction, making it well positioned to lead in AI for the online system and EIC complex in general.

In the next years BNL should evaluate Computing Infrastructure, including efficient computing through modeling, management, and control, and prototype computing infrastructure for the early EIC phase (including LAN and distributed storage evaluation), as well as evaluation of new computing architectures

1. Develop (together with JLab and ePIC collaboration) Computing Services, including databases, data and workflow management, new tools for computing infrastructure control and monitoring, new AI/ML algorithms for anomaly detection
2. Develop funding profile for computing at BNL
3. Develop AI/ML algorithms and services for EIC.
    1. The 2023 NSAC Long Range Plan stated that "EIC could be one of the first large-scale collider-based programs in which AI/ML is integrated from the start."
4. Evaluate WAN performance (together with JLab and ESnet)
5. Do early prototyping and integration of services, infrastructure, AI/ML algorithms and use RHIC experiments for demonstration analysis chain

The Primary Technical Responsibilities of Host Labs

Oversight for ePIC software and computing designs and execution to provide assurance functions for the host labs and DOE.

Provisioning and operating standard infrastructure solutions consistent with supported lab infrastructures and community best practices.

Support for the EIC International Computing Organization.

Interface for local resources and policies at the respective labs.

On-going computing operations in support of the accelerator and detector design and construction.

Operational support functions for:

- ✓ Experimental data curation.
- ✓ First-pass processing.
- ✓ Data analysis.
- ✓ Support of collaboration(s) and users.
- ✓ Accelerator and detector simulations.

Presented to EIC RRB in 2023 and 2024

EIC Computing and Software Joint Institute

Brookhaven National Laboratory (**BNL**) and Thomas Jefferson National Accelerator Facility (**JLab**), as Electron-Ion Collider host labs, have established a collaborative entity, the **EIC Computing and Software Joint Institute** (**ECSJI**).

- This joint institute is designed to support the computing and software needs and activities of the EIC.
- The Institute was created fall of 2023

ECSJI will leverage **complementary expertise at the two labs** and provide needed visibility to the respective lab management and stakeholders.

- *The advantages of such a structure also include increased reliability and availability of resources for the ePIC collaboration and other future collaborations.*

The success of the EIC, an international scientific endeavor, will benefit from contributions from international partners towards its computing effort.

To facilitate efficient coordination, the Institute will administer the EIC International Computing Organization (EICO), which will include all the contributors to the EIC computing effort.

EIC Computing and Software Joint Institute Scope

The Institute will provide for EIC computing and software matters:

1. A single entity to interface with the EIC project, the ePIC collaboration, theoreticians, and future collaborators.
2. Execution of host lab technical computing responsibilities (slide 3).
3. Maintenance of service level agreements and statements of work outlining the host labs' contribution to the ePIC collaboration concerning computing resources, services, and personnel assigned to work on ePIC computing and software deliverables.
4. A coordinating body for interacting with international partners, providing computing resources as in-kind contributions, including:
   - Assessing resources.
   - Managing the agreements with the sites delivering resources (including service levels).
   - Facilitating and assessing the delivery against the agreements.

The scope may evolve over time and the organization of the institute as well

ePIC Streaming Computing Model

- ePIC streaming computing model is a primary document to define roles for BNL (Echelon0 and Echelon1) and JLab (Echelon1) computing centers and for preparing MoUs with International partners
- Role of Echelon 2 and 3 centers may evolve with time
- Data placement and number of data replicas (data distribution between sites is TBD)
- ECSJI will organize ePIC Computing Model review September 26-27 of 2024
  - Review charge is finalized (D.Dean and A.Deshpande)
    - ePIC S&C plans, and ECSJI itself, will be reviewed
  - Xiaofeng Guo, Program Manager for Nuclear Physics Computing, will attend the review

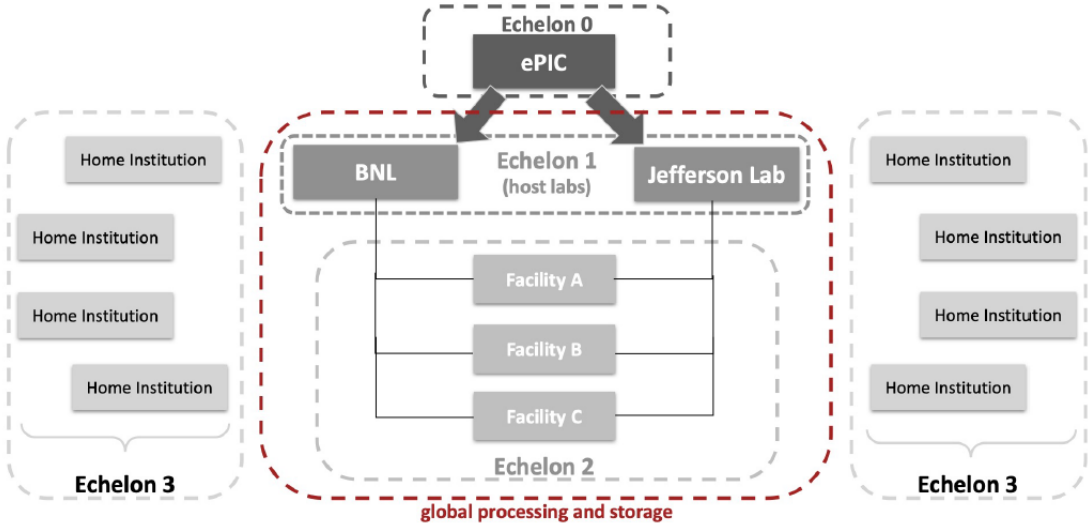

ePIC Streaming Computing Model

This sketch was presented at EIC RRB in May 2024
Paul Mantica : "DOE will support SW and Computing needs of Echelons 0 and 1"

BNL and JLab will need to work on the requirements and submit them to DOE NP

Echelon 0 Computing Options

Plans for DAQ/Electronics Integration/Testing/Installation and needs for Off-Project Support
D.Abbot, J.Landgraf, F.Barbosa

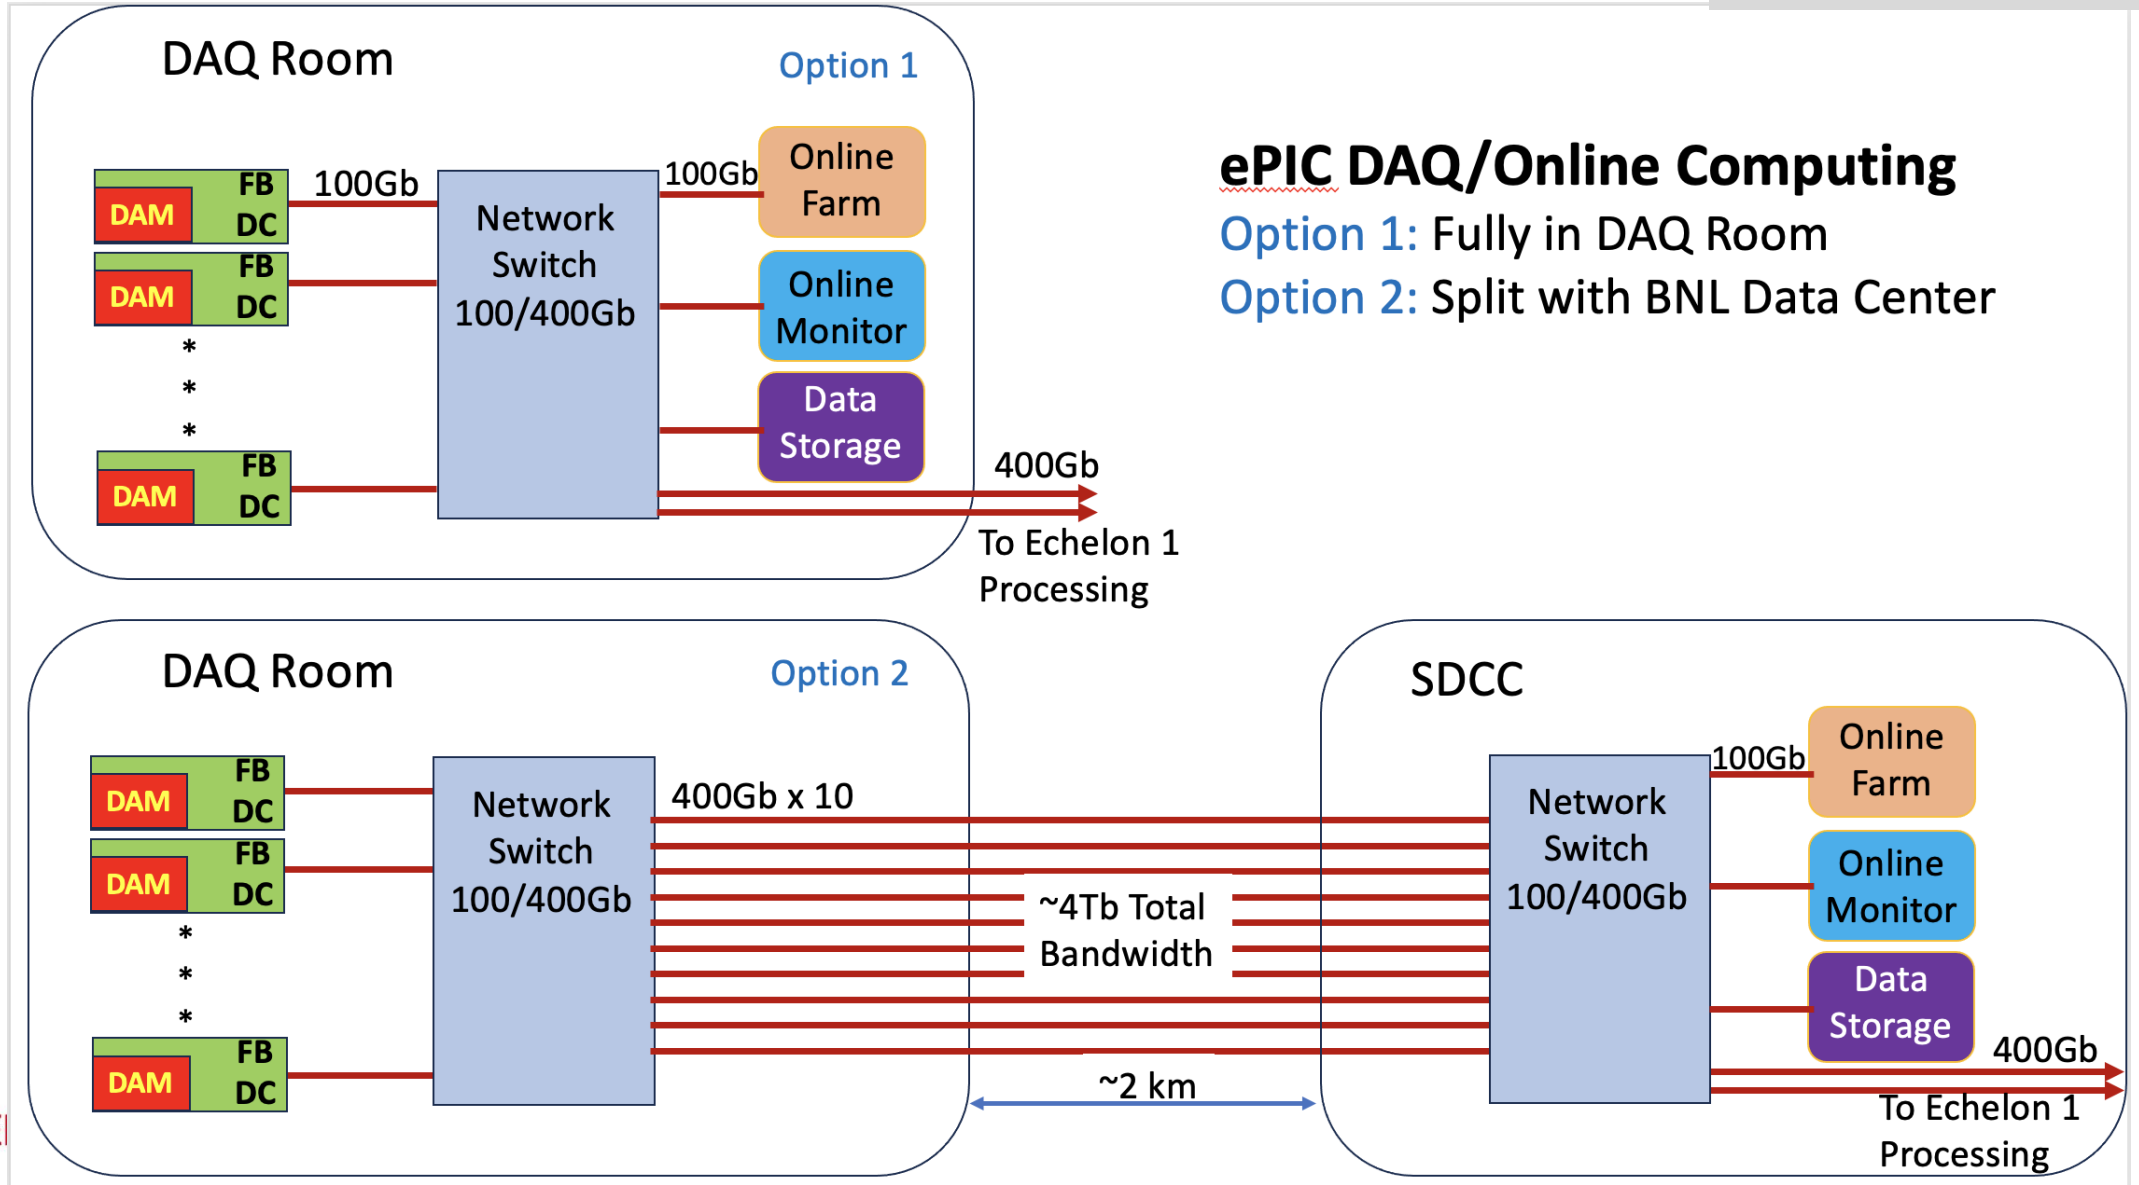

Network Infrastructure

- ePIC DAQ and Computing require network connectivity at BNL between resources in IP6, SDCC and development labs that is bandwidth guaranteed and minimizes routing.

- Early DAQ & Electronics development will be in Labs (Physics).
- Initial computing needs for development would best be placed in a central location (SDCC, 5-6 racks total).
- As construction and installation ramps up at IP-6 additional Echelon 0 computing resources can be placed in both sites.
- For production running it is important to have equal and transparent network access from SDCC to both Echelon 1 sites (BNL and JLAB)

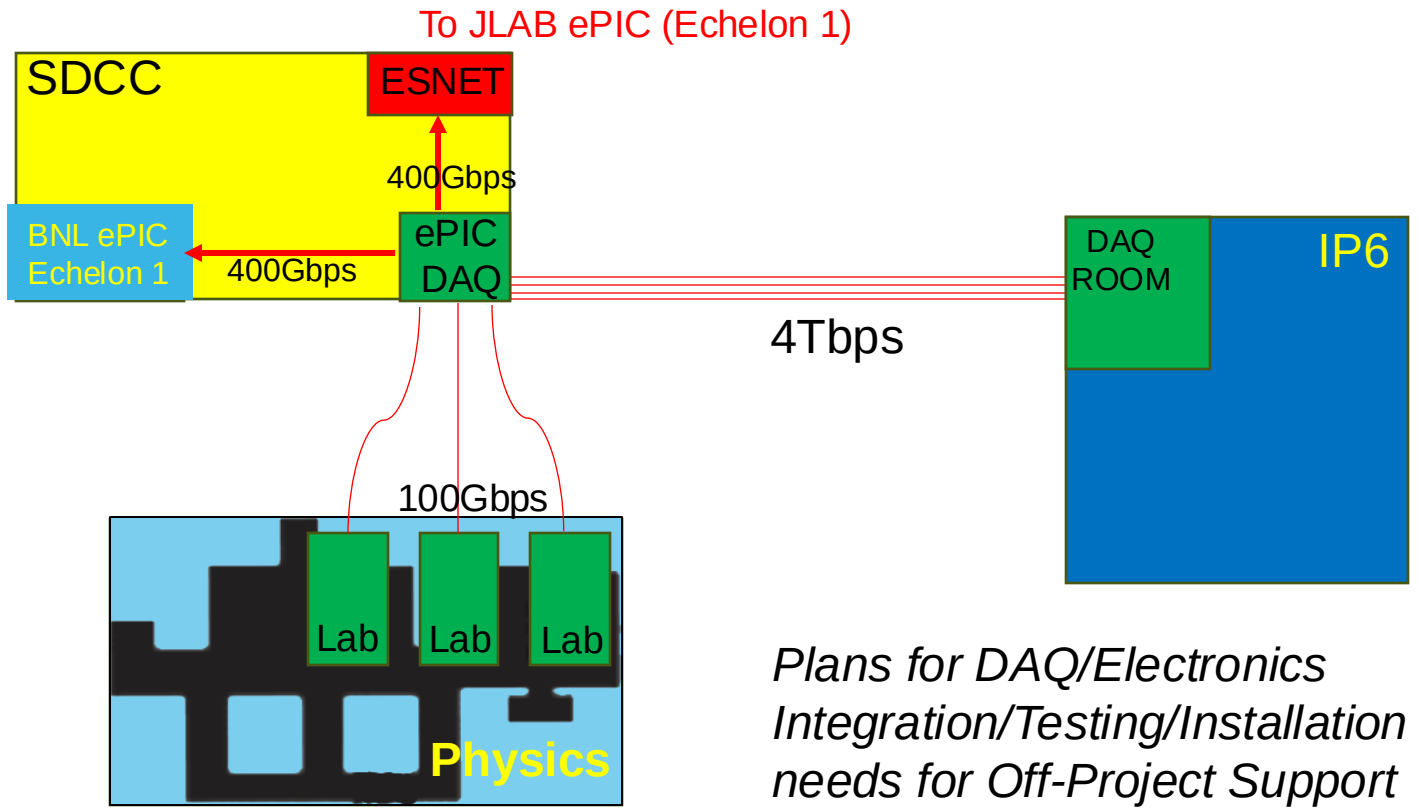

Plans for DAQ/Electronics Integration/Testing/Installation and needs for Off-Project Support
D.Abbot, J.Landgraf, F.Barbosa

ePIC Computing Requirements for Echelon1. Networking

Networking	egress [Gbps]	ingress [Gbps]
RAW data (immediate ~17% of total)	17	
RAW data (contingency)	50	
Monitoring, H&S data, E0/E1/E2 status	1	1

All numbers from ePIC "The streaming computing model" note and do not include Monte-Carlo data [at least 30% should be added]

ePIC network requirements do not include:

- Networking infrastructure at BNL
  - Data transfer between counting house and computing center
  - Data transfer between counting house and "buffers"
- ingress doesn't include data transfer (derived data) from Echelon2 to Echelon1
- egress doesn't include raw data transfer to Echelon2 and derived data transfer from Echelon1 to Echelon2
- Users traffic is not included
- BNL < – > JLab WAN connectivity performance is not addressed

Open questions :

1. What is the required uptime for connectivity from the counting house to the 725 data center and to JLAB

   The current fibre infrastructure within the EIC/RHIC ring doesn't include fibre redundancy to/from the counting house. Also, these current projects don't address the fibre infrastructure and cabling from the counting houses to Building 725.

2. Bldg.515 and Bldg.725 network infrastructure
   1. Bldg 515 (old data center) network infrastructure will be EoL in FY26 and it doesn't support high data rate
      1. The 515 Spine can connect compute top of rack switches at ~40Gbps where in 725 we can support 400Gbps+ to each compute rack.

ePIC Computing Requirements for Echelon1. Storage

Storage	Disk transient [PB]	Disk permanent [PB/year]	Tape [PB/year]
RAW and derived data	11	20	220

All numbers from ePIC "The streaming computing model" note and do not include Monte-Carlo data [at least 30% should be added]. All numbers for nominal luminosity.

ePIC storage requirements do not include :

- storage for intermediate data buffering
- storage for end user analysis and for users
- storage for tape buffer

- *The scale of compute resources is not estimated yet, but 220PB/year is comparable with ATLAS annual data volume for LHC Run3 and require O(3000kHS23) for the first year and 15% increase per year*

US ATLAS approach to have 10% margin for storage and 20% for Compute

Services : data management, workflow management, databases are not included

Tape (assuming 'data carousel' model – active data exchange between disk and tape)

- 1 LTO10 tape library (8,000 slots) can hold 240PB. LTO10 is GA in 1Q25.
- To meet 50Gb/sec (6.5GB/sec) , we should prepare 2 disk arrarys ($105K each) as HPSS internal cache and 28 LTO10 drives (400MB/sec per drive).
- The tape library with 28 LTO drives should cost about $450K. Tape media cost is about $3 per TB.

We don't see problems on Data Carousel with HPSS for ePIC at nominal luminosity.

Computing and Networking Schedule. Construction and Commissioning Phase

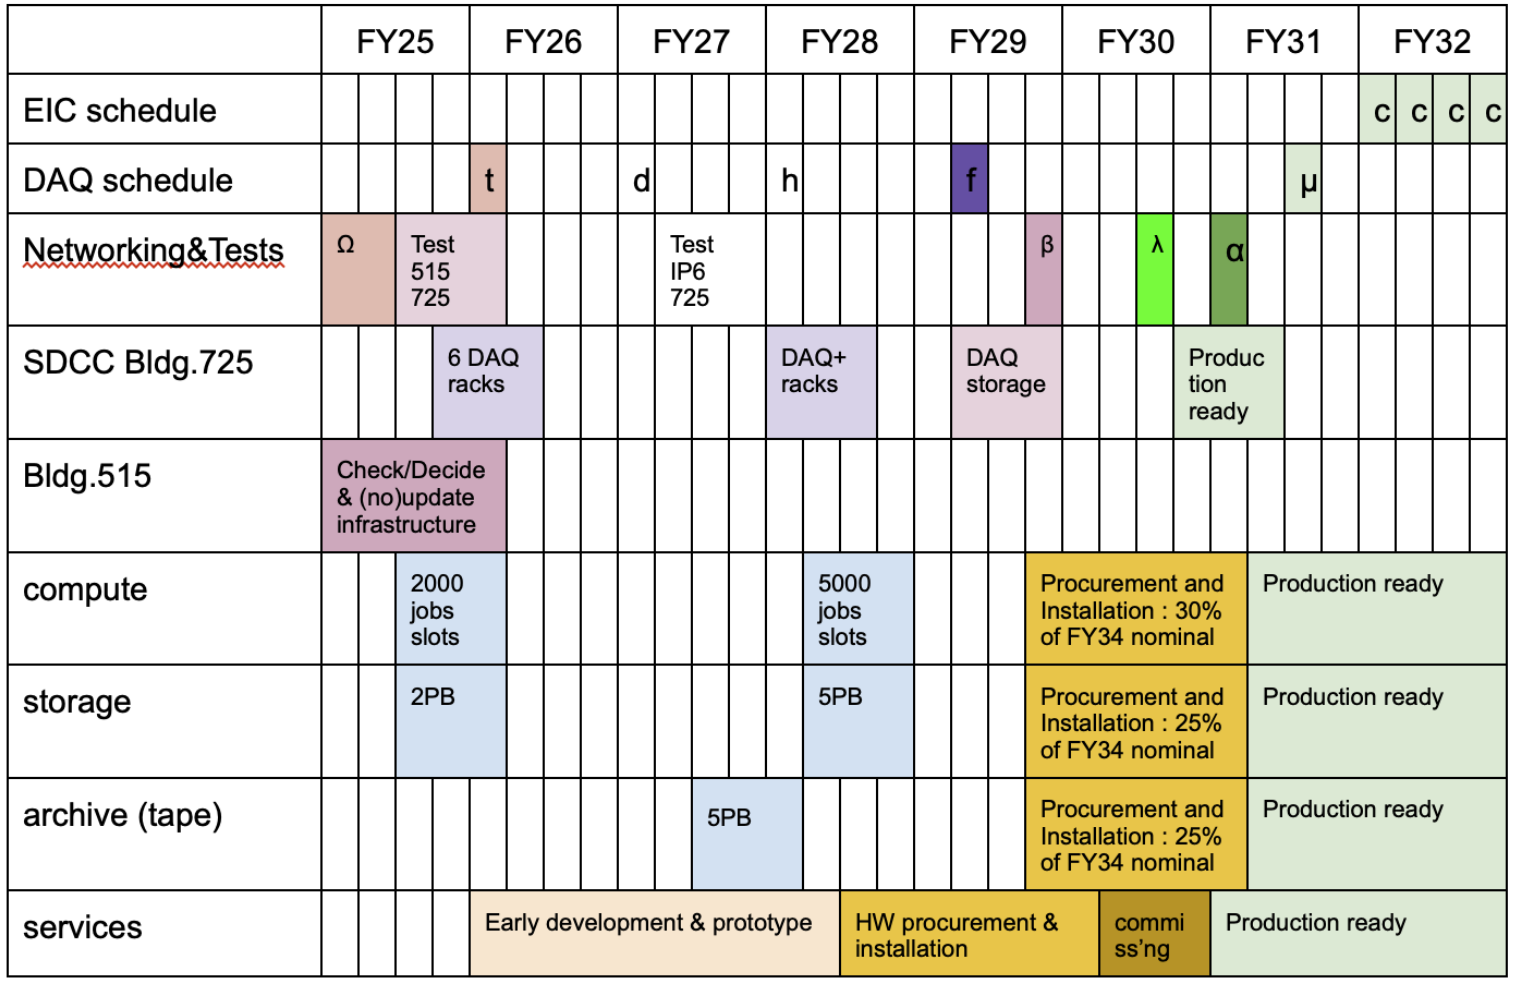

EIC timeline:

- c : commissioning run

DAQ timeline :

- t : DAQ test set up
- h:DAQ full hardware and timing chain
- d : DAQ readout detector in test stand
- f : full DAQ chain
- μ : DAQ ready for cosmics

Networking & tests timeline :

- Ώ : collect and evaluate requirements for EIC/RHIC LAN
- β : BNL LAN Data Challenge
- λ : Echelon0/Echelon1 Data Challenge
- α:Full Dress Rehearsal
  - DAQ → <Echelon 0> → Echelon1/..

This is an early assessment and will be discussed and updated after the ePIC Computing model review in September 2024
The procurement schedule, WAN testing and data challenges are also subject to discussion and agreement with JLab.

Computing and Networking Schedule. Science Phase

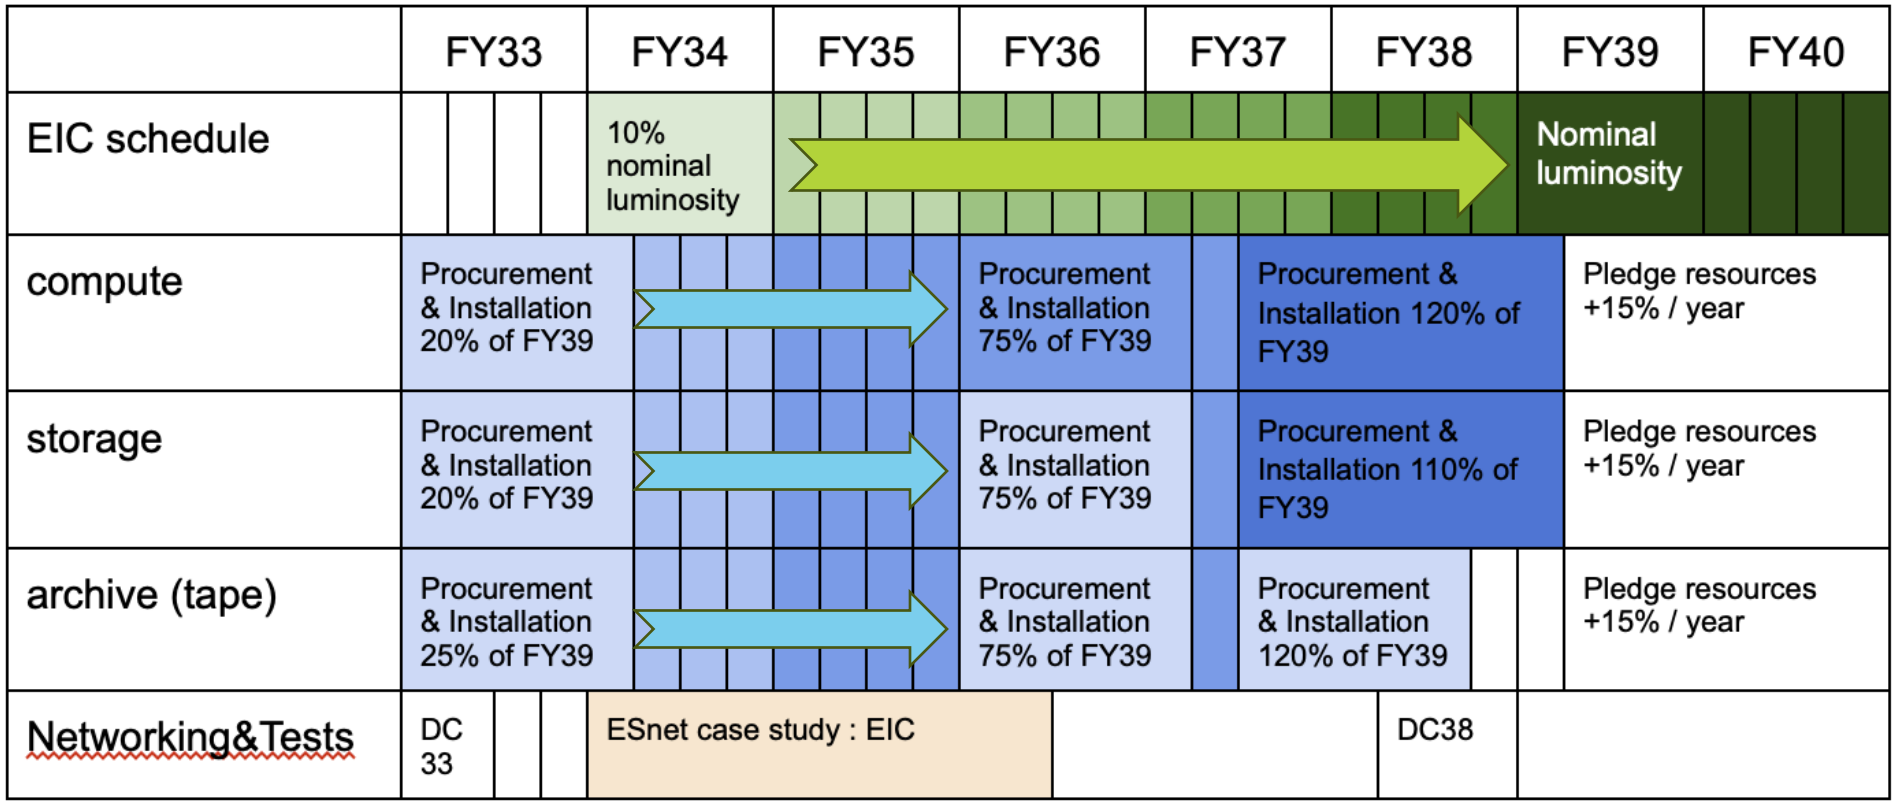

This is an early assessment and will be discussed and updated after the ePIC Computing model review in September 2024
The procurement schedule, WAN testing and data challenges (DC) are also subject to discussion and agreement with JLab.

Status and Computing Resources for EIC Today

- Presently, support and computing resources for the EIC at BNL and JLab are provided mostly opportunistically
- at Jefferson Lab, 2000 job slots and 1PB of storage allocated for the EIC
- at BNL, HPC cluster is set up to support EIC project
  - There are immediate requests to setup 1 PB of storage and <at least> guarantee the same number of job slots at BNL as at JLab

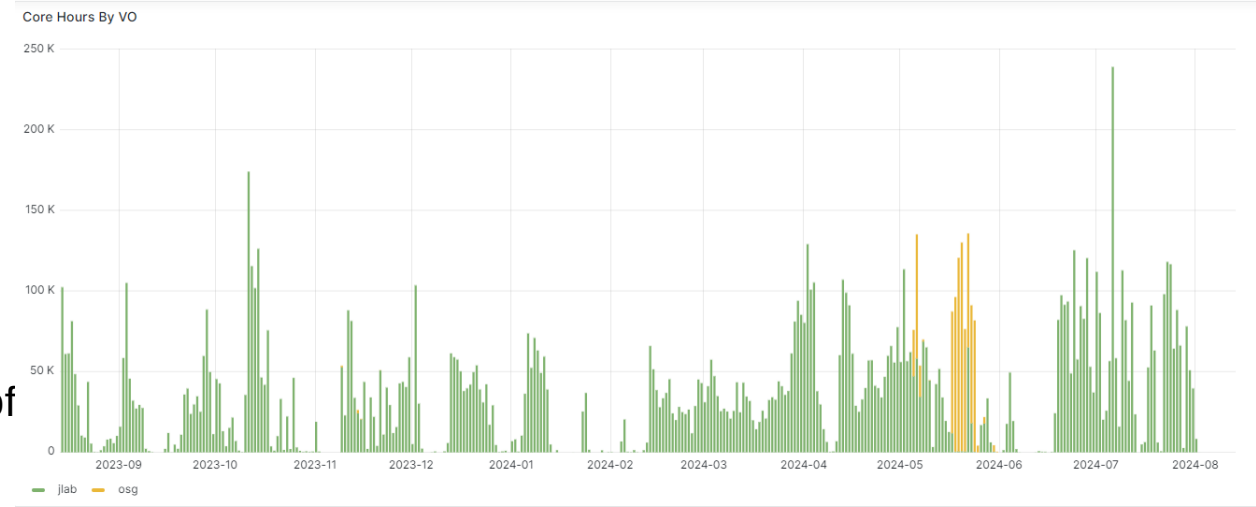

- There is local funding (*at limited scale*) available at BNL for computing development
  - We made computing services, modelling and infrastructure prototype the compute focus of the FY25 PD proposal, approval on that would bootstrap us strongly into an important activity
- The ePIC/EIC is using the Open Science Grid (OSG) and anticipating the inclusion of partners from WLCG-affiliated institutions (INFN IT, GridPP UK, Taiwan,…) and DOE ASCR Facilities

- EIC will use WLCG networking infrastructure : LHCOPN/LHCONE - one of areas where we need a discussion with WLCG

WLCG – World LHC Computing Grid

EIC Software at BNL

- Small effort level, 1 full time person (Kauder), fractions of others (Potekhin - infrastructure, Osborn – reconstruction/tracking)
- Good representation in ePIC S&C (Simu convener - Kauder, S&C deputy coordinator - Wenaus, Infrastructure/collab tools - Potekhin, Reconstruction - Osborn once he has time)
- Simulation is the focus, received tremendous boosts from funded proposals in FY23
  - EIC simulation infrastructure (ESI) (LDRD, 3yrs, 2 postdocs, fractional staff)
  - DOE NP project on AI based detector design (AID2E), now leveraging also HEP-CCE (2yrs, 1 postdoc, fractional staff)
  - We now have effort commensurate with expertise, and better balance (& strong collaboration) with JLab
- We would benefit tremendously if we could do the same on reconstruction!
  - ePIC track reconstruction is ACTS based, from ATLAS, and used by sPHENIX. We have busy experts (Osborn) we could leverage if they could direct postdoc(s)
- We made AI/ML the software focus of the FY25 PD proposal, approval on that would bootstrap us strongly into an important activity
  - Again drawing on existing activity/expertise
    - Jin Huang & Brett Viren on LS4GAN (trustable AI with quantified systematics) and projects spawned from it
    - Jin et al on streaming DAQ, data reduction, compression
    - AI/ML in simu: ESI and AID2E
    - Large scale AI/ML workflows & applications: PanDA/iDDS, REDWOOD (ASCR project on resilient workflows, Alexei PI)
    - And in long term R&D, LLM tech applied to ‘Large Particle Model’
    - All in fruitful collaboration with CSI

List of Required Actions

1. Set up compute and storage resources at BNL for accelerator (EIC) and ePIC (to be sufficient to address current needs)
2. Evaluate networking infrastructure and LAN throughput performance at BNL
3. Do early prototyping and integration of services, infrastructure, AI/ML algorithms and use RHIC experiments for demonstration analysis chain
4. Build on very strong collaboration with JLab and others on software
    1. and building on BNL strengths - achieved in simulation, potential exists in reconstruction, AI/ML growing
5. Work with JLab on compute requirements including WAN requirements (with ESnet)
    1. Develop funding profile for computing at BNL and present it to EIC/NPP and then to DOE
6. Work with ePIC collaboration and JLab on ePIC computing model
    1. Intermediate data buffering at BNL
    2. Role of international partners in computing
7. Develop (together with JLab and ePIC collaboration) Computing Services, including databases, data and workflow management, new tools for computing infrastructure control and monitoring, new AI/ML algorithms for anomaly detection

BNL early (FY25) funding is needed to satisfy requirements 1,2,3 & 7

Plenty of uncertainties in EIC/ePIC S&C, but a lot of collaborative goodwill.

- *Highly collaborative and free of inter-lab tensions (at least between NPPS and JLab) in software.*
- *It is very important BNL in particular focuses now on tension-free collaboration with ePIC and JLab on the computing side. Tensions have been a substantial problem in the past.*
- *We should craft the EIC Computing and Software Joint Institute to facilitate and support a cohesive collaborative community between labs, US institutes and internationally.*

Complimentary Slides

Topics to be addressed (from D.Lissauer and H.Ma)

- BNL Host Lab Responsibilities in Computing and SW for EIC
- ePIC Computing Model
- DOE NP support for Computing at BNL
  - Basic assumptions of computing needs
    - Compute, storage, tape
    - Services
      - Data Management
      - Workflow Management
    - Computing for end users
- Computing Funding Profile, when we should start and how long will it take
- ECSJI - EIC Computing and Software Joint Institute
  - EICO - EIC International Computing Organization
- ePIC Software and Computing review Sep 26-27

The EIC Computing and Software Joint Institute was created fall of 2023

Boehnlein, Amber | amber@jlab.org
Lancon, Eric | elancon@bnl.gov

The EIC Computing and Software Joint Institute (ECSJI)

Brookhaven National Laboratory (BNL) and Thomas Jefferson National Accelerator Facility (JLab), as EIC host Labs, are creating a joint structure, the EIC Computing and Software Joint Institute (ECSJI), incorporating parts of BNL and JLab facilities to support the EIC and computing and software needs and activities. ECSJI will leverage complementary expertise at the two Labs and provide needed visibility to the respective Lab management and stakeholders. The advantages of such a structure also include increased reliability and availability of resources for the ePIC collaboration.

The success of the EIC, an international scientific endeavor, will benefit from contributions from international partners towards its computing effort. To facilitate efficient coordination, the institute will administer the EIC International Computing Organization (EICO), which will include all the contributors to the computing effort.

Scope of the EIC Computing and Software Joint Institute

This institute will provide for EIC computing and software matters:

1) A single entity to interface with the EIC project and the ePIC collaboration,
2) Maintains Service Level Agreements and statements of work outlining the host labs' contribution to the ePIC collaboration concerning computing resources, services, and personnel assigned to work on ePIC computing and software deliverables,
3) A coordinating body for interacting with international partners providing computing resources as in-kind contributions. This includes assessing resources, managing the MOUs with the sites delivering resources (including service levels), and facilitating and assessing the delivery against the MOUs,
4) Execution of host Lab responsibilities as detailed below.

Organization & Governance

The institute aims at providing efficient support to the EIC while acknowledging the differences in the organization at the two Labs. The proposed governance model ensures that the EIC experiment(s) are well supported in matters of computing and software, the institutes' performance is monitored, and reporting is clearly defined.

The Institute Management

- *Composition:* the management will comprise two co-Directors; each is nominated by one Lab. The co-directors are currently Eric Lancon (BNL) and Amber Boehnlein (JLab).
- *Reporting*: the institute's management will report jointly to the two host Lab management.
- *Duties and accountability*:
  - The management will be responsible for organizing the institute to deliver on the responsibilities defined above.
  - The management will maintain a multi-year operation plan for the host Labs, providing matrixed staff members to support the activities.
  - The management will provide a yearly report to the host Lab's management.

Boehnlein, Amber | amber@jlab.org
Lancon, Eric | elancon@bnl.gov

Responsibilities

The Host Lab's Responsibilities

The primary technical responsibilities of host Labs include and are not limited to the following:

- Oversight for ePIC software and computing designs and execution to provide assurance functions for the host Labs and DOE,
- Provisioning and operating standard infrastructure solutions consistent with supported Lab infrastructures and with community best practices,
- Support for the EICO,
- Interface for local resources and policies at the respective Labs,
- On-going computing operations in support of the accelerator and detectors design and construction,
- Operational Support Functions for:
  - Experimental data curation,
  - First-pass processing,
  - Data analysis,
  - Support of collaboration(s) and users,
  - Accelerator and detector simulations.

ePIC Collaboration Responsibilities

The ePIC collaboration responsibilities include and are not limited to the following:

- Developing and documenting a cost-effective computing model tailored to the experiment's needs, with the concurrence of the host Labs,
- Developing and maintaining multi-year resource plans,
- Report ePIC status in computing and software to the EIC-RRB,
- Identifying with input from the host Labs, a Computing and Software coordinator who serves as Point of Contact to ECSJI,
- Developments of Software Algorithms,
- Production operations.

Haiyan Gao
Digitally signed by Haiyan Gao
Date: 2023.09.29 10:38:27 -04'00'

Haiyan Gao
Associate Laboratory Director
Brookhaven National Laboratory

David J. Dean
Digitally signed by David J. Dean
Date: 2023.09.29 09:58:17 -04'00'

David J. Dean
Deputy Director for Science
Jefferson Lab