

Estimate of BNL and JLab requirements for the ePIC/EIC computing in 2027-2032

October 28, 2025

Alexei Klimentov, BNL

Amber Boehnlein, JLab

Markus Diefenthaler, JLab

Motivation: The host laboratories for the EIC, BNL and JLab are primarily responsible for developing and implementing EIC computing services and infrastructure, from online delivery of data to its processing and archiving. AI will play a central role in this development, positioning EIC computing as a leader in AI for the online system and computing in general. This development is being conducted in close collaboration with the EIC project and the ePIC collaboration. The ePIC collaboration is co-designing the detector and the computing of the primary experiment at the upcoming Electron-Ion Collider. These systems are being developed in parallel and in close coordination to establish a distributed computing model that enables seamless data processing from the detector to the physics analysis, incorporating AI and streaming readout. Current efforts focus on data transfer and processing from the detector to Echelo 1 sites at BNL and JLab, as well as collaborations with international partners on distributed computing.

We must be fully prepared for the early tests and data challenges planned with the ePIC collaboration. Figure 1 summarizes the deliverables and milestones for the development, deployment, and integration of the ePIC DAQ system, the distributed computing framework, and the autonomous experimentation and control enabled by AI.

It is essential to have tested and commissioned infrastructure and services prior to EIC start-up. The computing infrastructure includes efficient computing through modeling, management, and control; prototyping of computing infrastructure for the early EIC phase (including LAN at both labs, WAN, and a distributed computing model evaluation); and evaluation of new computing architectures. Computing services include databases, data and workflow management, new tools for computing infrastructure control and monitoring, as well as infrastructure for the deployment and integration of AI approaches and methods required for autonomous experimentation and control of the EIC experiments. Early prototyping and integration of services, infrastructure, and AI/ML algorithms are planned to ensure the timely delivery of the Echelon 1 sites.

In fall 2023, host labs established a collaborative entity, the EIC Computing and Software Joint Institute (ECSJI), expected to support the computing and software needs and activities of the EIC. The ECSJI leverages complementary expertise at the two labs and provides visibility to their respective management and stakeholders. This structure also increases the reliability and availability of resources for the ePIC collaboration and other future collaborations. Since its

foundation, the ECSJI has been actively involved in ePIC/EIC computing, particularly in organizing the EIC International Computing Organization (EICO) and the EIC Computing and Software Advisory Committee (ECSAC). The ECSJI works closely with the ePIC DAQ, SW&Computing, and the EIC project in general. The ECSJI is heavily involved in EIC computing, software development, and modeling.

Currently, ePIC/EIC computing opportunistically uses Lab and DOE/NSF resources, which are partly supported by the Labs' program development funds. Both Laboratories support the EIC program in all three core categories: compute, software infrastructure, and storage.

We successfully completed the initial phase of the ECSJI and EIC International Computing Organization (EICO) in 2024/25, and we are ready to start developing the EIC computing infrastructure prototype. This prototype will be integrated with the ePIC international computing resources in CY 2027 and will be ready for the planned DAQ challenges in 2027/2028.

Inputs and assumptions of the estimate

Assumption 1: The estimates presented here are based on the most recent version of the *ePIC Streaming Computing Model Report* ([Zenodo, Record 14675920](#)). In this estimate, it is assumed that in 2034 the EIC will operate at 10% of its expected maximum luminosity of approximately $10^{34} \text{ cm}^{-2}\text{s}^{-1}$.

Assumption 2: To estimate host labs contributions, we assume their share is based on the following allocation fractions.

Assumption	Computing	Storage (disk)	Storage (tape)
Host Labs fraction	100%	100%	100%

It is further assumed that Echelon 1 computing facilities will be sufficient to perform prompt and full processing of ePIC data at the EIC startup. This aligns with the expectation that two data replicas will be archived at BNL and JLab and Echelon 2 sites will not be fully operational by 2032.

Assumption 3 : We assume a typical 99% of uptime for the Echelon1 computing resources.

Assumption 4: A common 5-year replacement cycle is assumed for computing and storage servers. Equipment can be purchased with longer warranty (at a cost) and can also be operated out-of-warranty (at a risk).

Final assumption: The final assumption concerns the EIC schedule. The execution of the current request will cover the period between 01/01/2027 until 12/31/2032. Computing procurement and deployment are aligned with the overall EIC Project schedule and the

associated timelines for the development and deployment of the ePIC DAQ, computing, and AI components, as shown in the figure above.

The funds for purchasing in a CY are mostly allocated in the year prior. Therefore, we use the ePIC estimates for the period between CYs 2028 and 2032.


Risk 1: Given the large increase in computing and storage needs just prior to the start of the EIC program, additional schedule delays may modify the amount of funds needed in this 6-year request.

EIC computing requirements and host labs capacity


With the assumptions outlined above, estimates can be provided of the computing, tape and disk storage that need to be provided by the Echelon1 complex in the next few years. The table and plots below summarizes the Echelon1 sites needs, with the period associated with the equipment purchases :

	Total ePIC/EIC need	Requested to US Echelon1	Total ePIC/EIC need	Requested to US Echelon1	Total ePIC/EIC need	Requested to US Echelon1
Year	computing [cores]	US Echelon1 computing [cores]	disk storage [PB]	US Echelon1 disk storage [PB]	tape storage [PB]	US Echelon1 tape storage [PB]
2027	16,000	8,000	10	2	-	-
2028	16,000	8,000	10	4	-	-
2029	40,000	20,000	15	7	1	1
2030	80,000	30,000	40	20	4	4
2031	120,000	40,000	100	50	40	40
2032	160,000	61,000	400	220	200	200

CPU [kCores] vs Year



Storage [PB] vs. Year



Equipment purchase cost model and purchase cost

The current (2025) WLCG estimated cost of 1 kHS23 is approximately \$6k while the estimated cost of 1PB is approximately \$60k. Due to technology advances, the cost per kHS06 and the cost per PB tend to decrease. Recently, the WLCG has updated the projections of computing and disk storage cost [REF], and has estimated a reduction of 10%/year in the computing per kHS06 and 5%/year disk storage cost per PB.

Risk 2: We assume a 10% and 5% reduction per year in the computing cost per kHS and in the disk storage per PB. External factors like supply-chain interruptions, exchange rates, and additional costs due to importing tariffs may have a large impact on the assumption of cost reduction per year. Changes in architecture (eg, GPU and FPGA computing, or NVMe storage) can also have an impact on the cost.

Purchase year	Computing cost [k\$]	Storage cost [k\$]	Tape cost [k\$]
2027	390	120	-
2028	-	112	-
2029	540	185	10
2030	540	653	20
2031	500	1,460	170
2032	1,650	6,350	775

Including labor, travel, and operational cost

The labor cost for the US Echelon1 complex can be extrapolated from current costs using a 3% yearly cost. The cost of travel and supplies can only be estimated at this time based on the historical average and similar extrapolation. The table below summarizes all the costs per year.


calendar year	Labor [FTE : k\$]		Travel + supplies + Operational Cost [k\$]	Equipment [k\$]	Total [k\$]
2027	3	1100	330	510	1940
2028	4	1450	340	112	1902
2029	5	1900	550	735	3185
2030	10	3750	900	1213	5867
2031	20	6750	2200	2130	11080
2032	28	10,300	3280	8775	22355

Labor Profile 2027-2032

The ramp for the labor is outlined below. The staffing plan initially supports system admin for procurement and deployment of the initial hardware systems. The staffing ramp adds skills each year to support the deliverables for managing the activities, designing and operating the systems including the user support.

This estimate does not include ePIC specific software activities, but Echelon1's personnel will work on distributed computing support, including data and workload management instances at JLab and BNL (and make extraordinary Echelon2 resources available for the ePIC collaboration), CVMFS services, data and workflows orchestration, and collaborative tools and user support.

Labor Profile 2027 - 2032



After 2032, the team is expected to remain at 28 FTEs through the EIC startup phase and the first two years of operation. Subsequently, the team size may be reduced to 22-24 FTEs, based on prior experience with LHC and RHIC. No increase in labor is anticipated beyond 2034, even with the expected ramp-up in luminosity.

Operational costs are projected to remain at approximately \$3.3M/year. Compute, disk storage, and tape capacity will be gradually increased by 55-60 percent in anticipation of the EIC reaching nominal luminosity.

During the first years of data-taking, the Echelon 2 sites will be fully commissioned and integrated into the workflows, providing substantial extraordinary CPU and storage capacity. By that time, algorithms, including AI-based data filtering, are expected to be optimized and tuned, and experience will have been gained regarding end-user analysis needs. A gradual increase in CPU and storage capacity is preferred to avoid the simultaneous retirement of a large portion of resources at the end of their five-year lifecycle.

Timeline and Milestones

Timelines and milestones are aligned with those of the EIC project (DAQ) and the ePIC collaboration. We will review the proposed timeline on a year-by-year basis. The decision to purchase bulk computing and storage in FY29 and FY31 will depend on the EIC project schedule and the readiness of the Echelon 2 sites for ePIC, which provide extraordinary computing in support of the core computing performed at the Echelon 1 sites.

Under the current plan, we retain the flexibility to adjust our pace in response to the EIC schedule. If the EIC project experiences delays, we can slow our activities accordingly; conversely, if it accelerates, we have the capacity to expedite our efforts. The purchase and deployment schedule will be reviewed and discussed annually.

	FY25	FY26	FY27	FY28	FY29	FY30	FY31	FY32
EIC schedule								c c c c
DAQ schedule		t	d	h	f		μ	
Networking & Tests	Ω			Test IP6 BNL JLab		β	λ	α
Echelon0 enclave at BNL		6 DAQ racks		DAQ+ racks	DAQ storage		Production ready	
compute			10% cores		30% cores	50% cores	75% cores	Production ready 61k cores
storage			2PB	4PB	7PB	20PB	50PB	Production ready 220PB
archive (tape)					1PB	4PB	40PB	Production ready 200 PB
services		Early development & prototype				Bulk procurement	Commissioning & Installation	Production Ready

EIC timeline:

c : no beam and cosmics commissioning run

DAQ timeline :

t : DAQ test setup

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 :

Ω : collect and evaluate requirements for EIC/RHIC fibers

β : BNL LAN Data Challenge

λ : Echelon0/Echelon1 Data Challenge

α: Full Dress Rehearsal (DAQ → <Echelon0> → Echelon1/2..)