# 1st Draft of the ePIC Computing Model

A. Bressan,
(University of Trieste and INFN)

# EIC Accelerator Design Overview

- Design based on **existing RHIC Complex;** RHIC accelerator chain will provide EIC Hadrons.
- EIC constructed in Collaboration with **JLab**

- **Hadron storage ring (HSR): 41-275 GeV** (based on RHIC)
    - up to 1160 bunches, 1A beam current (3x RHIC)
    - bright vertical beam emittance (1.5 nm)
    - strong cooling (coherent electron cooling, ERL)

- **Electron storage ring (ESR): 2.5–18 GeV** (new)
    - up to 1160 polarized bunches
        - high polarization by continual reinjection from RCS
    - large beam current (2.5 A) ➔ 9 MW SR power
    - superconducting RF cavities

- **Rapid cycling synchrotron (RCS): 0.4-18 GeV** (new)
    - 2 bunches at 1 Hz; spin transparent due to high periodicity

- **High luminosity interaction region(s)** (new)
    - superconducting magnets
    - 25 mrad crossing angle with crab cavities
    - spin rotators (produce longitudinal spin at IP)

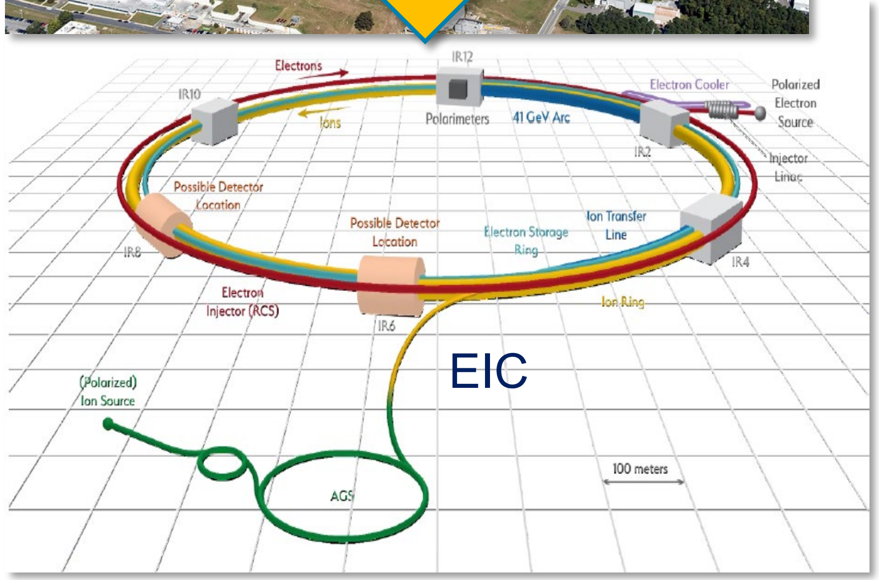

# ePIC at the EIC

**Integrated interaction and detector region (90 m)**
Get ~100% acceptance for all final state particles, and measure them with good resolution. All particles count!

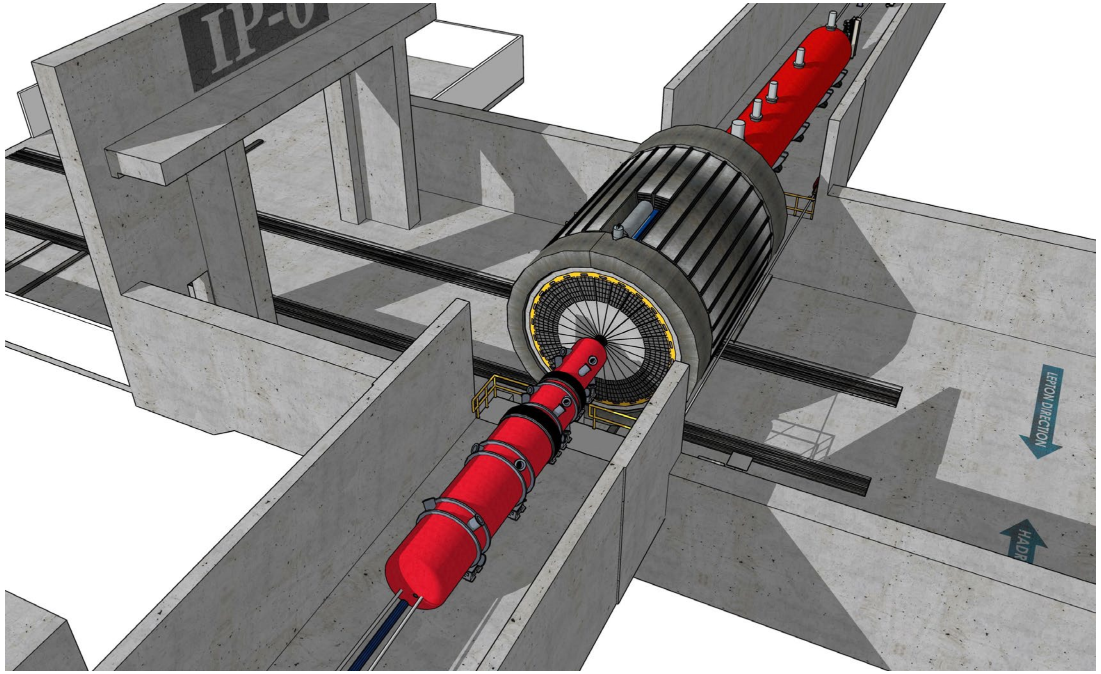

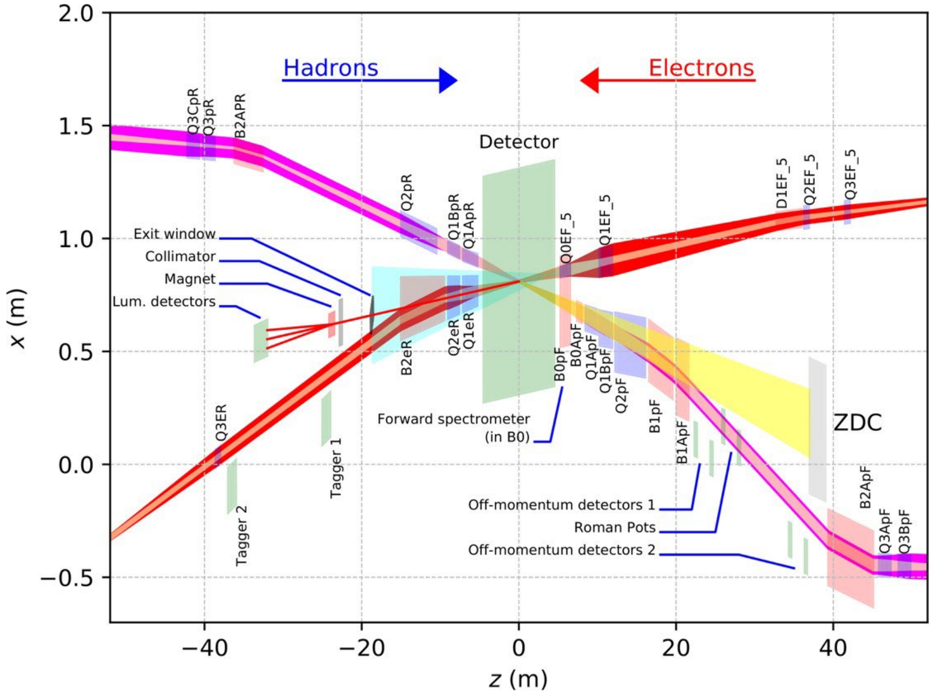


**Compute-Detector Integration**
Extend integrated interaction and detector region into detector readout (electronics), data acquisition, data processing and reconstruction, and physics analysis.

# A Brief Timeline

- EICUG Yellow Report (2020-21)
- Call for proposals issued jointly by BNL and JLab in March 2021
  - Proposals due Dec. 1, 2021
  - ATHENA, CORE and ECCE proposals submitted
- Public DPAP meetings Dec. 13-15, 2021
  - Presentations from proto-collaborations
  - Panel-assigned homework questions
- Second DPAP session Jan. 19-21, 2022
- DPAP closeout March 8th, 2022
  - Final report available March 21st, 2022
  - ECCE proposal chosen as basis for first EIC detector reference design
- Spring/Summer 2022 – ATHENA and ECCE form joint leadership team
  - Joint WG's formed and consolidation process undertaken
  - Coordination with EIC project on development of technical design
- Collaboration formation process started July, 2022
  - First IB Meeting July 18th
  - Charter writing committee formed and active
- First ePIC Collaboration meeting July 26-29, 2022
- ePIC Charter approved Dec. 14th, 2022

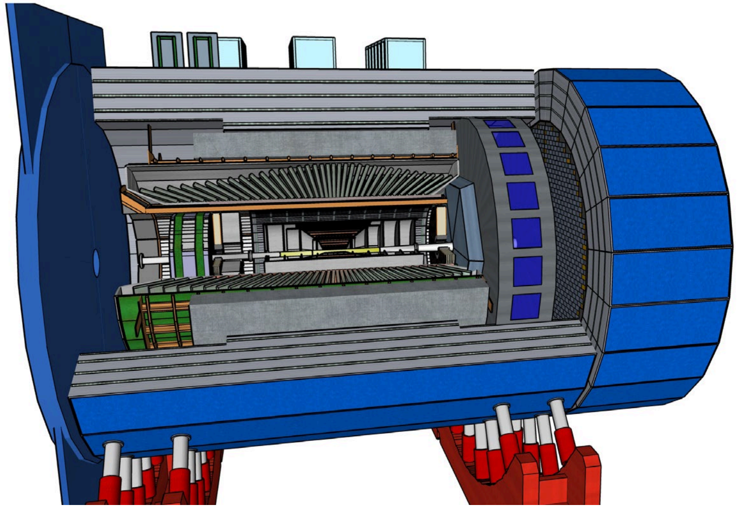

**ePIC Detector**

- To be sited at IP6 (25mr crossing angle)
- Addresses EIC science program as outlined in the EIC white paper and NAS report
- Must be ready for Day-1 EIC operations
- Working towards pre-TDR and CD-2/3A

# Compute-Detector Integration to Maximize Science

**Broad ePIC Science Program:**

- Plethora of observables, with less distinct topologies where every event is significant.
- High-precision measurements: Reducing systematic uncertainties of paramount importance.

**Streaming Readout Capability**

- **Capture every collision signal**, including background.
- Event selection using all available detector data for **holistic reconstruction**:
    - **Eliminate trigger bias** and provide accurate estimation of uncertainties during event selection.
- Streaming background estimates ideal to **reduce background** and related systematic uncertainties.

| | EIC | RHIC | LHC → HL-LHC |
|---|---|---|---|
| Collision species | $\vec{e}+\vec{p}, \vec{e}+A$ | $\vec{p}+\vec{p}/A, A+A$ | $p+p/A, A+A$ |
| Top x-N C.M. energy | 140 GeV | 510 GeV | 13 TeV |
| Peak x-N luminosity | $10^{34}$ cm$^{-2}$ s$^{-1}$ | $10^{32}$ cm$^{-2}$ s$^{-1}$ | $10^{34} \rightarrow 10^{35}$ cm$^{-2}$ s$^{-1}$ |
| **x-N cross section** | **50 µb** | **40 mb** | **80 mb** |
| Top collision rate | 500 kHz | 10 MHz | 1-6 GHz |
| $dN_{ch}/d\eta$ | 0.1-Few | ~3 | ~6 |
| **Charged particle rate** | **4M $N_{ch}$/s** | **60M $N_{ch}$/s** | **30G+ $N_{ch}$/s** |

# Expected Data Flow

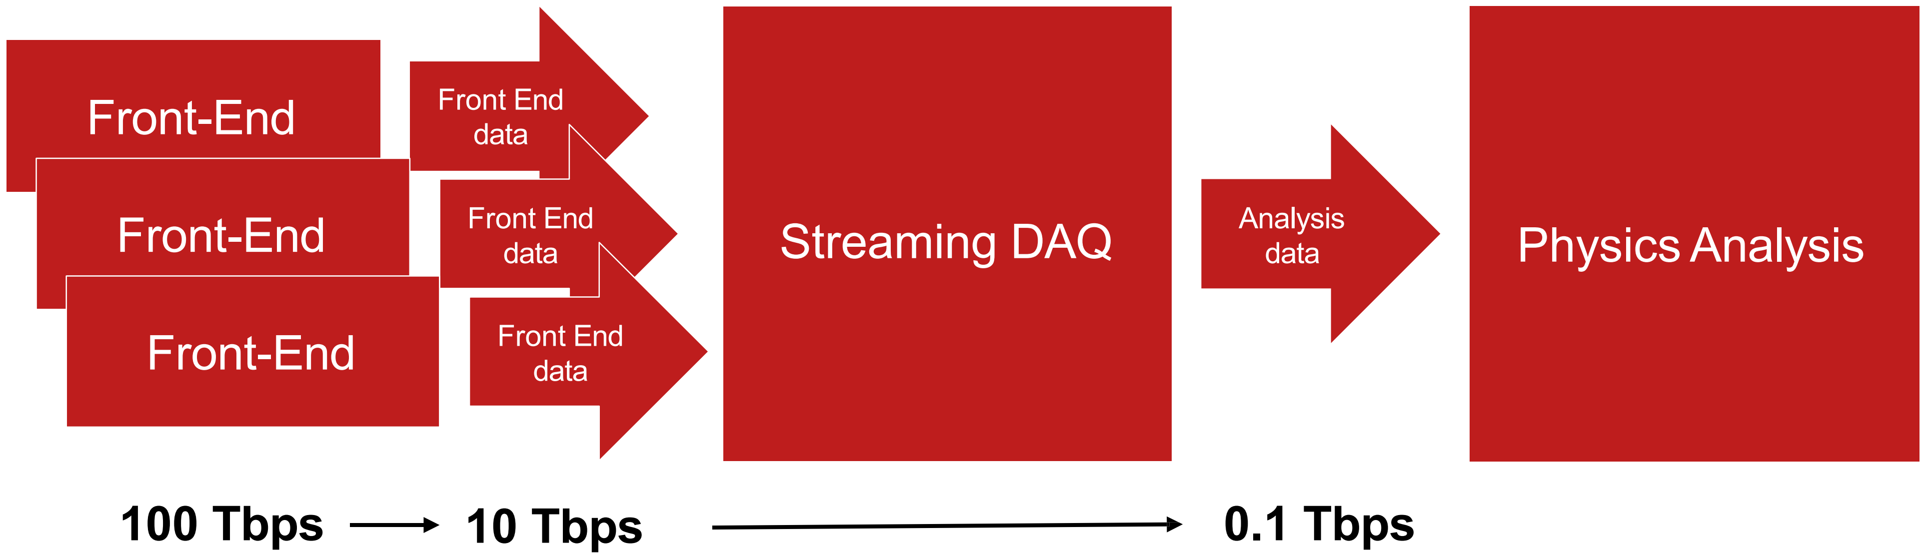

# ePIC Streaming Computing Model

**Four Tiers:**

**Echelon 0**: ePIC Experiment

**Echelon 1**: Host Labs

**Echelon 2**: Global processing and data facilities, includes HPC and HTC resources.

**Echelon 3**: Home institute computing

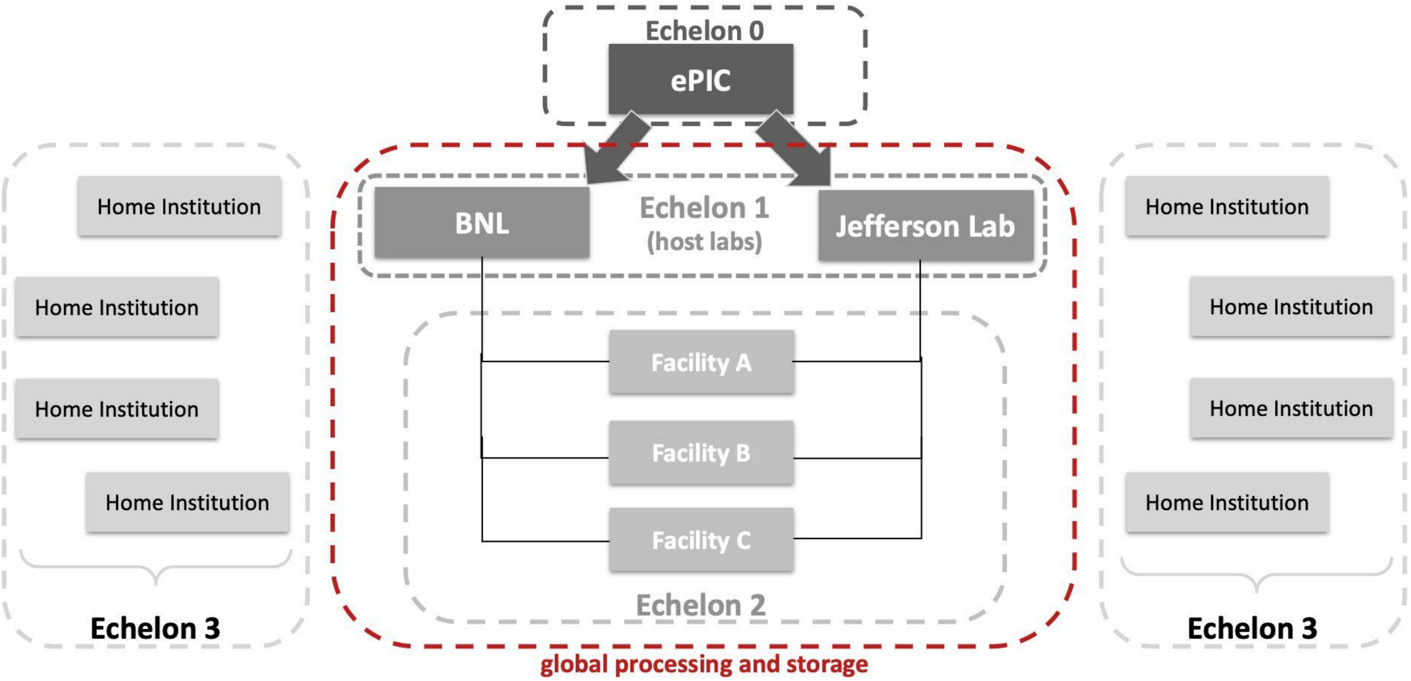


Initial version of a plan set to develop over the next decade.

# Use Case to Echelon Mapping

**Echelon 0**: ePIC Experiment
**Echelon 1**: Host Labs
**Echelon 2**: Global processing and data facilities
**Echelon 3**: Home institute computing

| Use Case | Echelon 0 | Echelon 1 | Echelon 2 | Echelon 3 |
|---|---|---|---|---|
| **Stored Data Streaming and Monitoring** | ✓ | ✓ | | |
| **Alignment and Calibration** | | ✓ | ✓ | |
| **Prompt Reconstruction** | | ✓ | | |
| **First Full Reconstruction** | | ✓ | ✓ | |
| **Reprocessing** | | ✓ | ✓ | |
| **Simulation** | | ✓ | ✓ | |
| **Analysis** | | ✓ | ✓ | ✓ |
| **Modeling and Digital Twin** | | ✓ | ✓ | |

# Echelon 0: Worst-Case Data Rate

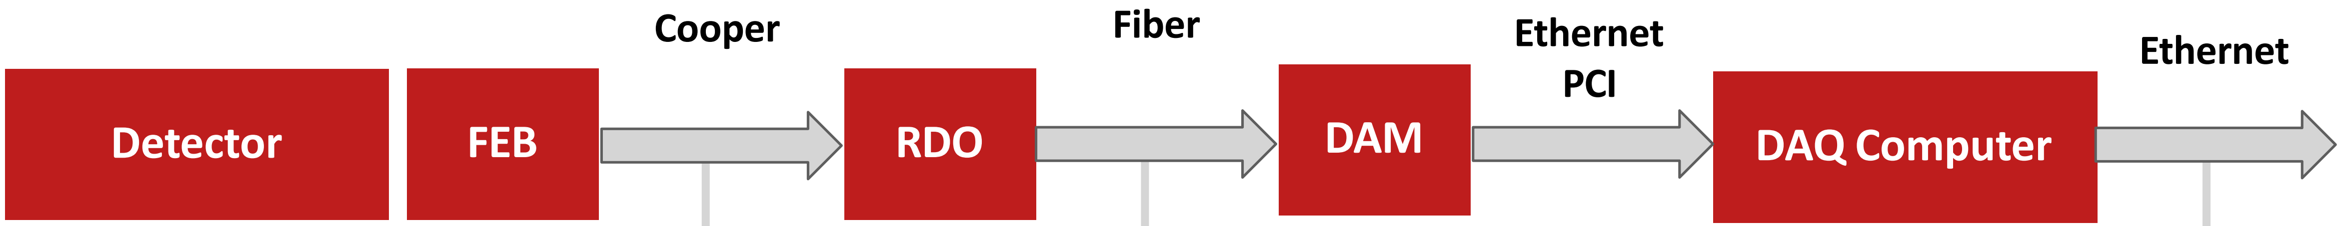


| Aggregate | 2.0 Tbps |
|---|---|
| Signal + Background | 400 Gbps |
| Noise | 1.6 Tbps |

| Aggregate | 2.0 Tbps |
|---|---|
| Per RDO (Avg.) | 0.7 Gbps |

| Aggregate | 96 Gbps |
|---|---|
| Signal | 38 Gbps |
| Synchrotron Radiation | 0.1 Gbps |
| Electron Beam | 22 Gbps |
| Hadron Beam | 4 Gbps |
| Noise | 32 Gbps |

# Echelon 0: Networking and Storage Requirements

- **Max Luminosity and Bandwidth**:
  - Expected maximum luminosity ≈ $10^{34}$ $cm^{-2}s^{-1}$.
  - Corresponding bandwidth ≈ 100 Gbps.
- **Design Considerations**: 100 Gbps between Echelon 0 and Echelon 1 at the host labs.
- **Outgoing Bandwidth**: Required: 200 Gbps

| Resource | Type | Amount |
|---|---|---|
| **Outgoing bandwidth** | Raw data | 200 Gbps |
| | Monitoring, slow controls, misc. meta data | ≤ 1 Gbps |
| | **TOTAL** | **≤ 201 Gbps** |
| **Incoming bandwidth** | Monitoring, calibrations | ≤ 1 Gbps |
| **Storage** | Disk (outgoing data buffer w/ 24hr) | 1PB |

# Echelon 1: Networking and Storage Requirements

| Resource | Type | Amount |
|---|---|---|
| **Outgoing bandwidth** | Raw data – *immediate* (1/6 of total) | 17 Gbps |
| | Raw data – *replay* (contingency) | 50 Gbps |
| | Monitoring, slow controls, misc. meta data | 1 Gbps |
| | **TOTAL** | **≤ 68 Gbps** |
| **Incoming bandwidth** | Monitoring, calibrations, slow controls (from Echelons 1–2) | 1 Gbps |
| **Storage** | Disk (temporary) | 1PB |
| Raw+recon. only, no sim. | Disk (permanent) | 20PB / yr |
| | Tape | 220PB/yr |

Values shown are for a single Echelon 1. There will be two.

# The Partnership

- Brookhaven National Laboratory (**BNL**) and Thomas Jefferson National Accelerator Facility (**JLab**), as Electron-Ion Collider 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.

# The Scope

**The Institute will provide for EIC computing and software matters:**

1) A single entity to interface with the EIC project and the ePIC collaboration.
2) 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.
3) A coordinating body for interacting with international partners, providing computing resources as in-kind contributions, including:
   - Assessing resources.
   - Managing the MOUs with the sites delivering resources (including service levels).
   - Facilitating and assessing the delivery against the MOUs.
4) Execution of host lab responsibilities.

# Organization & Governance

- The Institute aims to provide efficient support to the EIC while acknowledging organizational differences at the two labs.
- The proposed governance model ensures that EIC experiments are well supported in matters of computing and software, the Institute's performance is monitored, and reporting is clearly defined.

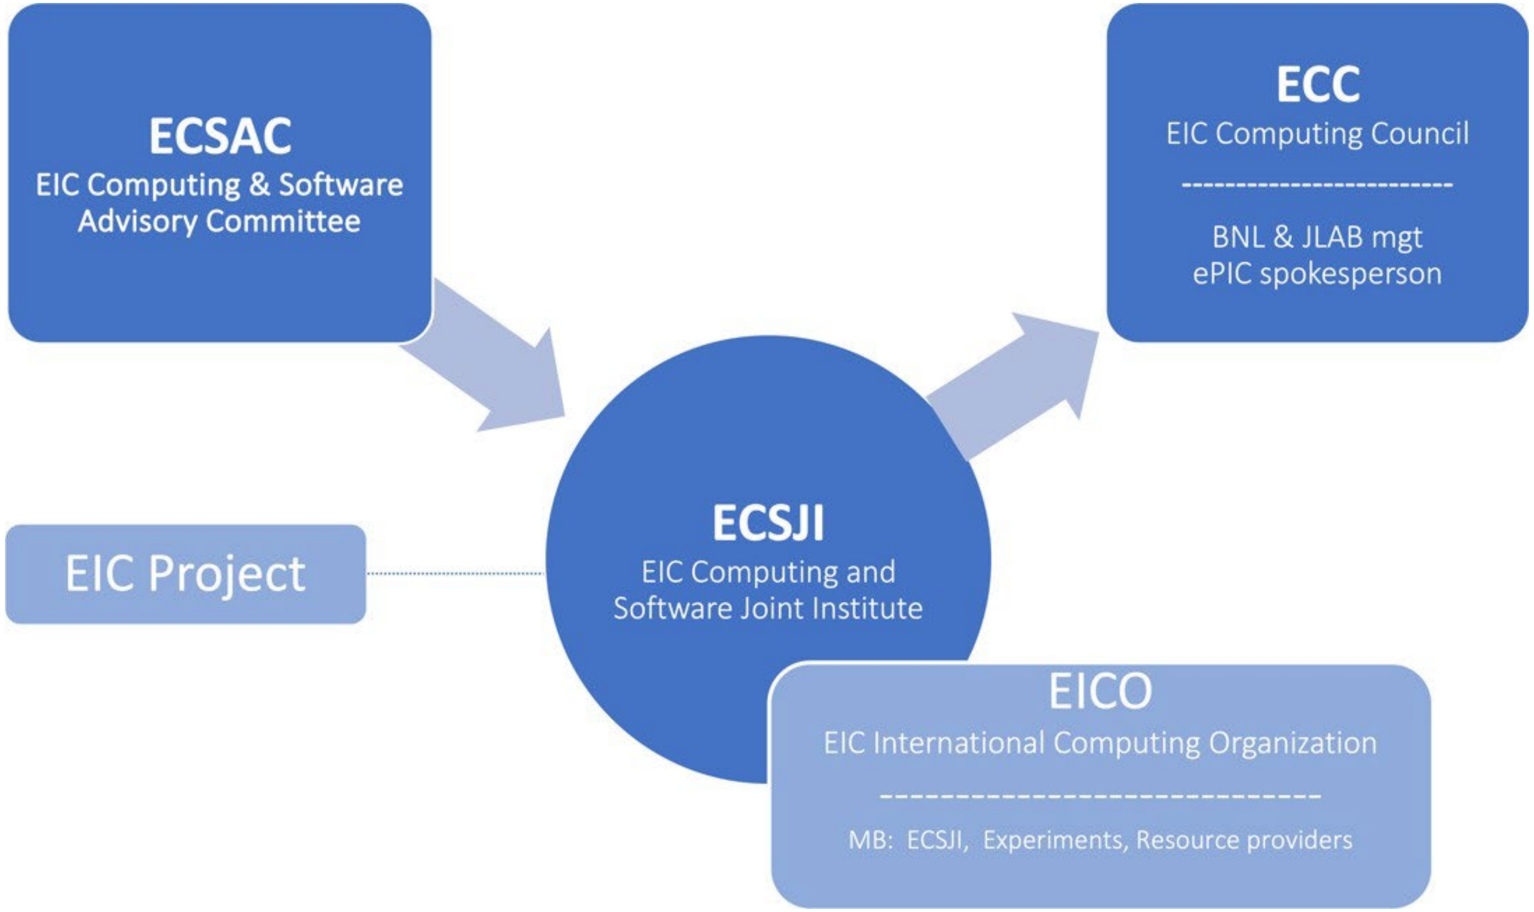


*The proposed governance structure is composed of two central bodies, designed to facilitate communication, coordination, escalation of issues, and conflict resolution.*

# Institute Management

- **Composition**: Management will comprise two co-directors, each 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 host lab management.
- **Duties and accountability**:
  - The management will be responsible for organizing the Institute to deliver on the previously defined responsibilities.
  - The management will maintain a multi-year operation plan for the host labs, providing matrixed staff members to support activities.
  - Institute management will provide a yearly report to the host labs' management.

# Echelon 2: Global ePIC Computing

- **ePIC is an international collaboration and so is its computing**:
  - **Echelon 2** includes **global resources contributed by collaborating institutions**.
  - Achieving scientific goals relies on effectively using Echelon 2's resources.
  - Design of computing model aims for **effective integration and management**.
  - Echelon 2 resources must have **MOUs specifying service requirements**.
  - These MOUs **assure technical compatibility with the ePIC computing model**.
  - **ePIC commits to facility integration** and robust testing/validation, includes monitoring and diagnostics.
- **Connectivity**
  - Echelon 2 sites connect equally to both Echelon 1 Host Labs.
  - Connection is ultimately through ESnet network backbone.

# Perspective

- ePIC is progressing in developing its Computing Model
- The level of possible support that will be presented from Canada, Italy and United Kingdom is not a commitment by Funding Agencies, but it is based on "reasonable" expectations, while discussions are ongoing
- The choices made while developing the ePIC Computing Model may help members of the International community in the negotiation with their Funding Agency

# Today LHC Experiments at CNAF

- Today level of support of for LHC experiment
  - 10 MHS days for ALICE, ATLAS and CMS, i.e $30 \div 35\ k$ logical cores
  - About 20 PB increase/year for ALICE, ATLAS and CMS
  - About 6-7 PB increase/year for ALICE, ATLAS and CMS
- HL-LHC will scale previous numbers in 10 years by
  - $\sim \times 10$ for CPUs
  - $\sim \times 5$ for DISKs
  - $\sim \times 10$ for TAPE

- As upper bound and assuming a pure FTE scaling, the amount of investments that one may foresee at CNAF and other IT centers are:
  - $20 \div 30$ k logical cores or running job
  - **About 20 PB** of storage/year
  - **About 5 PB** of disk per year
- A larger amount of IT resources (not just scaling with FTE) might be part of specific in-kind contribution (TBD)

# Desires for CM Requirements

- INFN Computing Centers are sites of WLCG, build on for and driven by the LHC community
- Accounting, virtual organizations, certificates (or future tokens) as well as GRID tools are elements of WLGC
- Present WLCG organization foresees 1/3 of resources for TIER0, 1/3 for TIERs1 and 1/3 for TIERs2. A CM that include replicas/reconstruction outside Echelon1 will be relevant (maybe mandatory) for International Institutions
- The ePIC CM adopting these tools will have a positive relevant impact for the possible INFN contribution to computing

UKRI

Science and Technology Facilities Council

# UK Digital Research Infrastructure

- Aspiration
  - Replication of all the reconstructed EIC/ePIC data (at least) across two or more federated data centres in Europe (Italy, UK, …)
- Assumptions
  - Planning assumes that the UK will comprise 10% of EIC users
  - UK data centre equivalent to EIC Echelon 2 but with storage for reconstructed data
  - Storage estimates based on EIC data rate of 100 Gbps at highest luminosity (2035+) being similar to Run 2 LHC (2018)

- Potential UK EIC Data Centre
  - Disk storage
    - 33% of reconstructed data
    - 6 PB/year (estimate)
  - CPU
    - 10% of overall requirement
    - 20,000 logical cores (estimate)
    - Similar to GridPP now (2021) but a factor of 2 lower to account for simpler event topology and increased performance
    - Event reconstruction, Monte-Carlo and user analyses
- Desires for CM requirements similar to Italy.

# BackUp

# RAW DATA mirroring for HL-LHC

- At HL-LHC, both the ATLAS and CMS experiments will produce $\sim 350$ PB of RAW data per year.
- The traffic from CERN to the T1s for RAW data quasi real time export will be $\sim 400$ Gbps per experiment (7M seconds of LHC data taking per year).
- + $\sim 100$ Gbps needed for other data formats
- The data at the T1 needs to be staged from tape and exported to the T2s for processing (assuming most of that processing happens outside the T1). This fix the need of T1-T2 bandwidth

# Main drivers for the next 10 years of SC

- Infrastructure
  - Renew infrastructures to be ready for the High Luminosity-LHC (HL-LHC) era, up to ~2035 or more
  - Use more compact computing (from today's ~20 kW/rack to 80 or more)
  - Lower the PUE (power usage effectiveness), be greener
  - Extend and expand networking for a future-proof infrastructure

- Foster and simplify the utilization of more viable technologies (Eur/task or J/task), like GPUs, FPGA, ... down to Quantum when available
- Be more efficient, elastic and resilient
- Pervasive use of geographically distributed storage ("the datalake")
- Abstract from physical machines, and form a national pool of resources and services ("the Cloud")
- Extend elastically to external providers such as HPC@CINECA or other cloud providers (via "dynamic federations")

# Today LHC experiments at CNAF

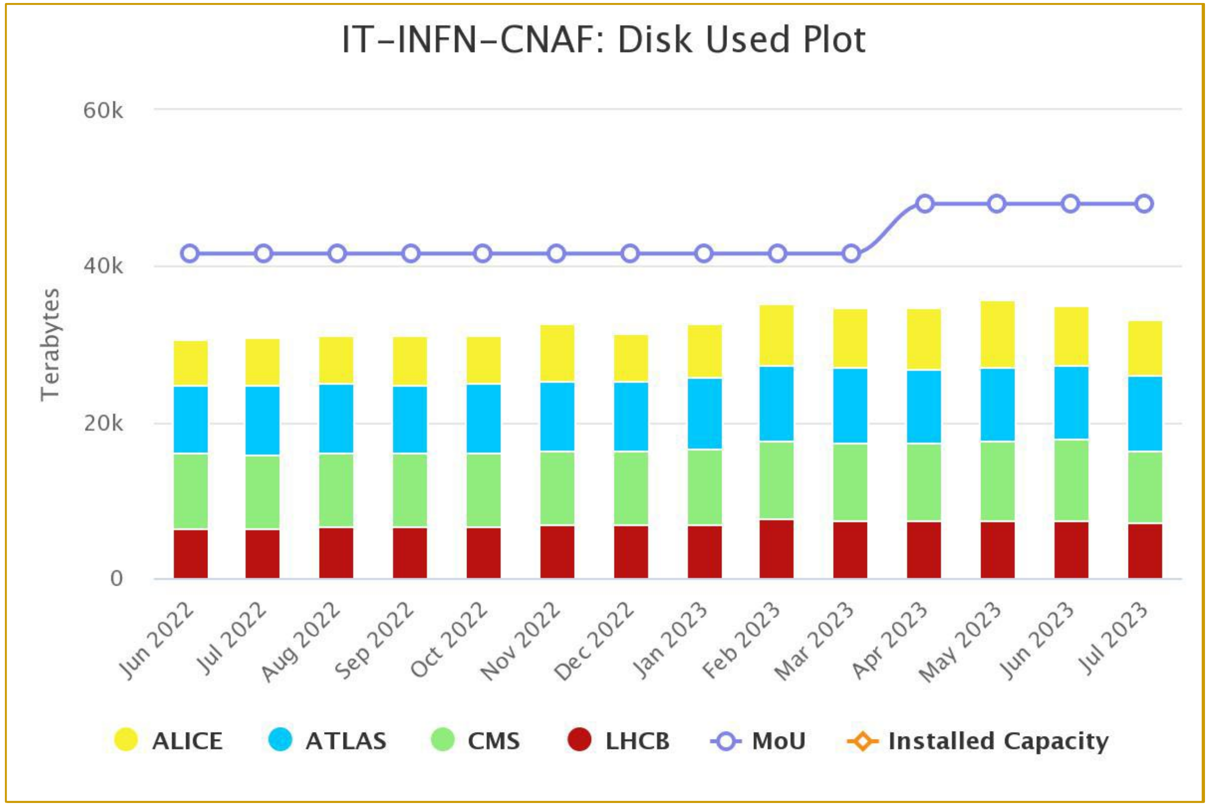


- About 6-7 PB increase/year for ALICE, ATLAS and CMS

- ATLAS/IT about 270 FTEs
- CMS/IT about 200 FTEs
- ALICE/IT about 130 FTE (110 paying members for pledge)
- Or a total of 600 FTEs

- Estimate for ePIC: $50 \pm 10$ FTEs

# ePIC running TODAY with 50 FTEs

- IF ePIC was running today, assuming a level of support comparable to LHC experiments (as upper bound) and assuming a pure FTE scaling, the amount of investments that one may foresee at CNAF and other IT centers are:
  - About 3000 (virtual)COREs or running job
  - About 2 PB of storage/year
  - About 0.5 PB of disk per year
- HL-LHC will scale previous numbers in 10 years by (with some uncertainties)
  - $\sim \times 10$ for CPUs
  - $\sim \times 5$ for DISKs
  - $\sim \times 10$ for TAPE
- A larger amount of IT resources (not just scaling with FTE) might be part of specific in-kind contribution (TBD)