International Computing

Wouter Deconinck (Manitoba)
for the ePIC Software and Computing Team

ePIC Computing & Software Review
September 26-27 2024

Overview

- **Global distributed streaming computing model**
  - Strategy: globally distributed computing with distributed storage
    - Expected standardization of platforms (containerization over operating system reqts)
    - Expected standardization of protocols (XRootD, local write-through caches, Rucio)
    - Taking advantage of the WLCG experiences, but mindful of the differences
  - Availability considerations for streaming computing operations
- **Current global computing partners** as an on-ramp to scaled up resources
  - Current platforms: OSG (opportunistic OSPool and ~25% dedicated resources)
  - Integrated international partners: Digital Research Alliance Canada, INFN CNAF (Italy)
  - Expected integration of 10 int'l partners in EICO: STFC UKRI & GridPP (UK), Japan,...
- **Ongoing transition from opportunistic to dedicated infrastructure**
  - Gradual transition centered around streaming simulation and calibration workflows
  - At least in the short term, maintain OSG as the interface to tie sites together

Globally Distributed Computing Model

- "Global" means US non-host-lab plus international contributions (will include US)
- Role of Echelon 2 and 3 centers may evolve with time, but aim for collocation of online computing and storage before replication to secondary sites
- Data distribution policies to be determined by EIC/ePIC data management plan

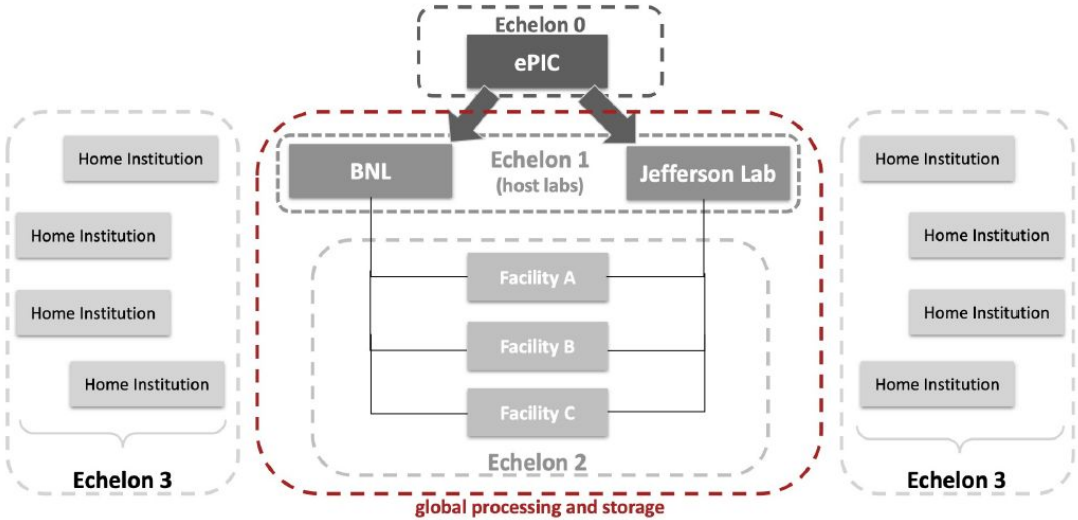

Includes Echelon 2 interconnects like evolved WLCG

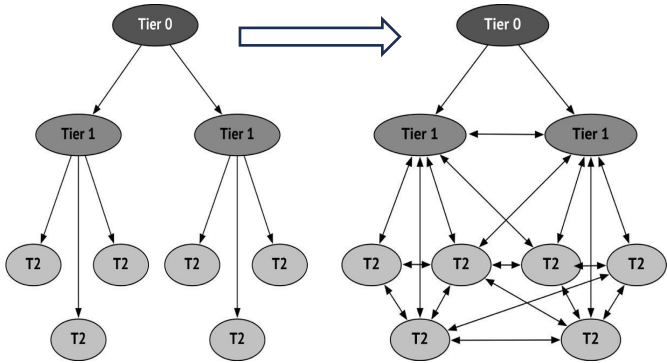

Ramping Up the Globally Distributed Computing Model

Global computing contributions concentrated in:

- **Echelon 2: globally distributed computing facilities** for streaming data analysis
  - Streaming data (in contrast with batch processing) requires potentially higher availability of the sites, and faster adaptability and redundancy overhead for the scheduler
- Echelon 3: home institute computing (largely opportunistic, important for serving users where they are, mainly via services provided at Echelon 1/Echelon 2)

Administrative structures:

- **EIC International Computing Organization** (EICO), led by ECSJI co-directors, administered by the ECSJI
  - Collects resources, requirements, produces accounting reports, and supervises SLA
  - Initial list of international contacts assembled and awaiting first meeting
- **To support EIC/ePIC TDR simulation campaigns**, ePIC has been integrating Canada and Italy into current operations, and aims to extend to other countries

Opportunities from EIC Globally Distributed Computing Models

DOE sees Echelon 1 is its operational responsibility to ensure the EIC science program

Global computing in Echelon 2 will **enable faster scientific turnaround**

- Echelon 2 can leverage Echelon 1, and increase the speed and turnaround time to scientific results with significant compute resources during data taking
- Streaming data analysis is only one component of producing results: total computing effort will include Monte Carlo, which requires extensive resources

Global computing in Echelon 2 will **engage the global community**

- After detector subsystem construction, Echelon 2 will present opportunities for significant participation in EIC running, for impact on responsibilities of international calibration teams, and for international infrastructure expansion
- Most importantly, significant compute participation will engage global community

WLCG Resources for EIC Globally Distributed Computing Model

WLCG already implements a globally distributed computing model. We would be remiss if we ignored this, and we would be glad to take advantage of this.

Some ePIC member institutions push us towards WLCG standards since well known.

Evolution of WLGC to the benefit of and **with contributions by ePIC** could center on:

- **Streaming data network access**: moving from batch processing of offline jobs to real-time processing of data streams
  - Imposes networking considerations: unable to run network-disconnected, while network security moves towards restricting this functionality. How to balance these constraints?
- **Higher availability requirements**, or increased redundancy at current SLA levels
  - Inability to process streaming data means dropping events from streaming capability.
  - WLCG SLAs in MOUs were based on batch processing, where inavailability does not have an immediate impact.

Current Integration of International Partners in Distributed Computing Model

Strategy:

- Prepare for actual streaming data taking by gradually **moving towards simulation campaigns that mimic streaming data taking**, including automatic calibration workflows

Distributed computing:

- **Open Science Grid** as scheduler of opportunistic and dedicated resources (already including Canada; currently bringing Italy online)

Distributed storage:

- **Commissioned Rucio instance at JLab**, with integration in simulation campaigns and additional RSEs anticipated at BNL & Canada over the next year

Anticipated Expansion of International Partners in Distributed Computing Model

Current computing primarily scheduled through Open Science Grid, which many international WLCG sites already support (e.g. DRAC, INFN, others)

Inclusion of partners from WLCG-affiliated institutions and DOE ASCR facilities:

- Interest from DOE labs to contributed: ORNL
- Interest from UK, Japan, Taiwan to join

Considerations for EIC as we move towards scaled-up integration:

- Need better ability to utilize guaranteed resources in use-it-or-lose-it models
  - Working towards constant base-load simulation campaigns in addition to bursty monthly campaign model; balance of agile strategy and longer term planning
- Need better accounting to demonstrate resource stewardship to providers
  - Working with OSG on some of the job accounting limitations

Canadian infrastructure available for the EIC

- Digital Research Alliance of Canada: National computing resources supporting both HPC and HTC workflows
  - Operated through five regional consortia, with resources allocated based on best fit
  - EIC Canada's computing contribution is operated by staff from Prairies, with current resources usage concentrated at Simon Fraser University in WestGrid

- Annual *resource allocation competitions* (with fast-track continuance option), peer review process of access, coordination of subatomic physics through SPNT group

- EIC Canada has held active allocation for 4 years
  - ~50 core-years/year, 100 TB during proposal and initial ePIC development phase
  - Supported majority of ATHENA simulations until Open Science Grid (5k-10k core-years)
  - Primary use for EIC Canada researchers and international contribution proof of concepts
  - Already integrated in OSG at 180 core-years/year and 140 TB distributed storage, about 10% of current needs

Digital Research Alliance of Canada | Alliance de recherche numérique du Canada

Future - Canadian infrastructure available for the EIC

- Canada aims for approximately 10% of the U.S. contribution to EIC computing.
- Due to direct involvement in ePIC Software & Computing, ability to be a testbed towards development of international strategies before rolling out more broadly
  - DRAC appointment on **streaming readout grand challenges**, joint with Jefferson Lab
- In the future, EIC Canada could have the opportunity to make a contribution through dedicated infrastructure.
  - Precedent exists in ATLAS Canadian Tier-1 located at Simon Fraser University (SFU).
  - SFU leads a consortium of Canadian institutions, including TRIUMF, operating the Tier-1 Data Centre, which receives funding from the Canada Foundation for Innovation (CFI).
- For EIC Canada to expand with dedicated “Tier-1”-like resources, a formal proposal is anticipated after the completion of accelerator and detector subsystem construction (also requested of CFI).

Digital Research Alliance of Canada | Alliance de recherche numérique du Canada

INNOVATION
Canada Foundation for Innovation
Fondation canadienne pour l’innovation

UK infrastructure available for the EIC

- Science and Technology Facilities Council (STFC)
  - Part of UK Research and Innovation (UKRI)
  - Supports research in Nuclear Physics, Particle Physics and Astronomy
  - IRIS project set up in 2018 to improve federation of STFC digital infrastructure
  - Infrastructure partners include **GridPP** (High Throughput Computing) and **DiRAC** (High Performance Computing), amongst others
  - **GridPP** supports HTC needs of nuclear and particle physics, in particular the UK's contribution to the WLCG but also non-LHC experiments

UKRI Science and Technology Facilities Council

www.ukri.org/councils/stfc/

iris

www.iris.ac.uk

GridPP UK Computing for Particle Physics

www.gridpp.ac.uk

UK Digital Research Infrastructure

- Scale of available resources

  10% of GridPP is available for use by non-LHC experiments; allocated through IRIS

  IRIS helps support science across the globe

  Examples include:

  Compute power to process and simulate data

  Short- & long-term data storage

  Tools to help move and exploit data

- Integration with ePIC/EIC S&C

  Stephen Kay (York) ePIC S&C User Learning

- Near-term plans (1-3 years)

  Establish use of GridPP non-LHC resources for ePIC; test compatibility of tools for moving data and deployment of ePIC software for simulation and analysis

- Long-term plans (4+ years)

  Develop bid for dedicated resources for ePIC in future expansion of GridPP

  Perhaps a federated European data centre with Italy and France

Italian infrastructure available for the EIC

- INFN operates its own Grid and Cloud services:
  - Nine medium-sized centers
  - One large national center at CNAF (Bologna) collocated with LEONARDO (pre-exascale supercomputer)
  - Successfully contributing to the WLCG for two decades and will support the LHC in the High Luminosity area
- A major national initiative has recently been approved: ICSC
  - National Research Centre for High-Performance Computing, Big Data, and Quantum Computing
  - The EIC could benefit from this global infrastructure
- Modest resources are available for EIC until detector construction is completed
- Italian EIC community has been yearly applying (and getting approved) for a share at CNAF since 2020

Italian infrastructure available for the EIC

Andrea Bressan
Domenico Elia

INFN Computing Centers are WLCG sites, structured by the LHC community.

Entering the WLCG federation by ePIC will strongly facilitate the use of INFN computing resources.

Should ePIC not become a member of the WLCG federation, the use of WLCG tools will stay as a key element for contributing resources from the INFN computing infrastructure

Italian infrastructure available for the EIC

Andrea Bressan
Domenico Elia

From the present 5kHS06, in the next years we will request INFN to further increase our shares at CNAF, in line with the simulation needs. INFN monitors closely the use of the allocated resources and partial use has consequences.

On a longer time scale, the participation of INFN will be defined via MoU that has to be negotiated. Different possible scenarios based on the CM, on the fraction of FTEs and other variables.

As a rule of thumb, INFN participate to the computing with resources proportional to the total resources needed by the experiment scaled by the fraction of the Italian FTEs

A larger amount of IT resources (not just scaling with FTE) might be part of specific in-kind contribution (TBD)

As already stated, the upper bound and assuming a pure FTE scaling, and assuming the ePIC computing needs are equal to today LHC, the figure of the investments that one may foresee at CNAF and other IT centers are:

< 20k logical cores or running job

< 20 PB of storage

Japanese infrastructure available for the EIC

- Scale of available resources and their demonstrated use in similar efforts
  - ~1000 CPU cores and 6 PB storage.
    - These resources are used for low-energy nuclear physics
- Current integration with the ePIC and EIC computing efforts (if any)
  - Not yet but integration will start soon
  - SPADI-Alliance is discussing to build distributed computing systems among different facilities and universities. They will be connected by SINET6 (max 400Gbps). Data from J-PARC, RIBF, and EIC will be distributed among those facilities and universities.
- Anticipated near-term (next 3 years) changes in integration with ePIC and EIC computing efforts,
  - JP requested big grant to the funding agencies. If all go well, we expect $0.5M-1M for computing every year and the resources will grow up every year accordingly.
- Anticipated long-term (4+ years) changes in integration with ePIC and EIC computing.
  - Same as above. Our requested budget profile assumes that $0.5M-1M for computing every year will last until 2030.

Summary

Globally distributed computing model is similar to existing global computing efforts, but also differs in some aspects:

- streaming readout adds requirements for network access and higher availability and redundancy

Current international participation includes Canada, and we are bringing Italy online right now, centered around the Open Science Grid.

Additional international participation is envisioned, and a few selected partners were presented, primarily based on their readiness to contribute within next 3 years.