International Data Facilities Perspective on the Computing Model

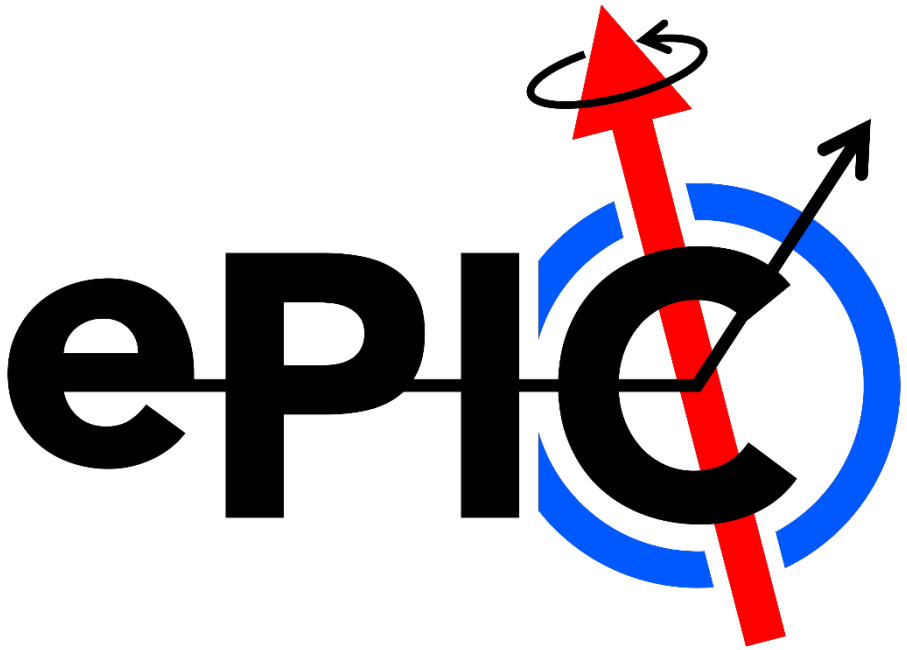

A. Bressan,
(University of Trieste and INFN)

Introduction

- This presentation will address the aspiration of Canada, Italy and the UK in participate to the computing of ePIC and its perspective
- Computing Centers of these countries were already included in the simulation efforts for the Yellow Report and, more recently for large scale simulations of the ePIC detector
- These countries can be considered as a test bench for the development of a Computing Model including international partners

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

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

Digital Research Alliance of Canada

- National computing resources, supporting both HPC and HTC workflows
  - Clusters in past 4 years amount to an overall of 250 kcores: Beluga 50k, Cedar 100k, Graham 40k, Narval 80k, Niagara 80k
  - **Archive** storage on 25 PB scale file systems; additional cloud deployment resources
  - Operated through regional consortia with staff, e.g. Prairies for EIC Canada institutions
- Annual resource allocation competitions (with fast-track continuance option):
  - Access to 80% of total computing resources, with 20% for rapid access requests
  - Peer review process of access
- 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 EIC detector simulations until Open Science Grid (5k-10k core-years)
  - Primary use for EIC Canada researchers and for international contribution proof of concepts

Possible Alliance Resources for ePIC

- Disclaimer: Competitive resource allocation prevents definitive statements
- Likely achievable asks for the EIC Canada collaboration:
  - 2k (2020s) to 5k (2030+) core-years/year
  - Continuing through existing OSG integration into the single-stream simulation campaigns
  - Integrated with services into the streaming computing DAQ (aided by location of developers)
  - Local storage resource associated with site of this scale
- Alliance Prairies consortium local technical services support
  - Personnel for the deployment tools and interface with EIC Host Labs
  - Local 5 PB storage with subatomic physics use cases
- Other possibilities into the future:
  - CFI has allowed the deployment of the ATLAS Canada Tier-1 Centre at TRIUMF
  - This infrastructure would have to follow experiment subsystem infrastructure proposal

Desires for CM Requirements

- ePIC Canada will benefit from using general purpose computing providers (with, admittedly, some experience with HEP/NP) where WLCG protocols are not standards. Support for different protocols will be beneficial
- Standards in HEP/NP require extra resource at general purpose providers:
  - Rucio can be a challenge to set up
  - File transfer services: FTS3 (subatomic) vs Globus (general purpose)
  - Federated authentication options: InCommon vs CILogon (general purpose)
- Interfaces and protocols expected by Echelon 1 sites will benefit from breadth over a focus on HEP/NP-specific solutions


INFN

INFN and Computing / Big Data

- A long tradition in **state-of-the-art distributed IT technologies and solutions**, from the first small clusters to Grid and Cloud-based computing.
- INFN operates Grid and Cloud services based on its own:
  - 9 medium size centers
  - 1 large national center, at CNAF (Bologna) – **with an area certified ISO/IEC 27001, 27017, 27018**
- INFN centers are connected through 10-100 Gbit/s dedicated links via the GARR network.
- Collectively the INFN infrastructure currently offers about **140,000 CPU cores, 120PB of disk space, 100PB of tape storage**.
- This Computing Infrastructure will be beneficial for ePIC computing.

The Cloud National Infrastructure for Supercomputing

400 M€ Total Budget

188 M€ Cloud Infrastructure

40 M€ Open Call

40 M€ Innovation & TT

42% Investment South Regions

34 MUR Universities and Research institutions
15 Private Companies
1575 Researchers and Engineers
250 New Temporary positions
250 New PhD
40 % Female

PIANO NAZIONALE di RIPRESA e RESILIENZA

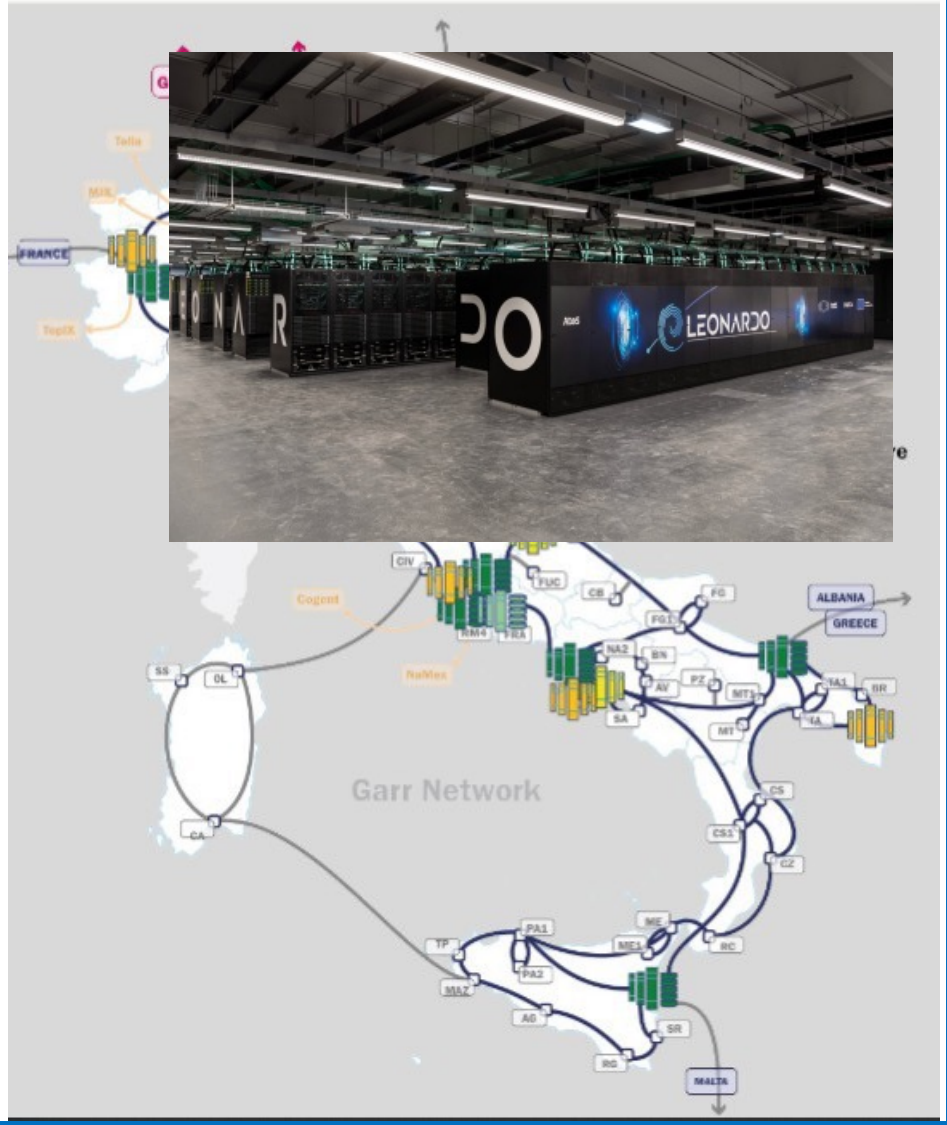

Today LHC Experiments at CNAF

- 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

Science and Technology Facilities Council

UK Digital Research Infrastructure

- 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

- GridPP
  - Federates the Rutherford Appleton Laboratory (RAL) Tier-1 site and 16 Tier-2 sites

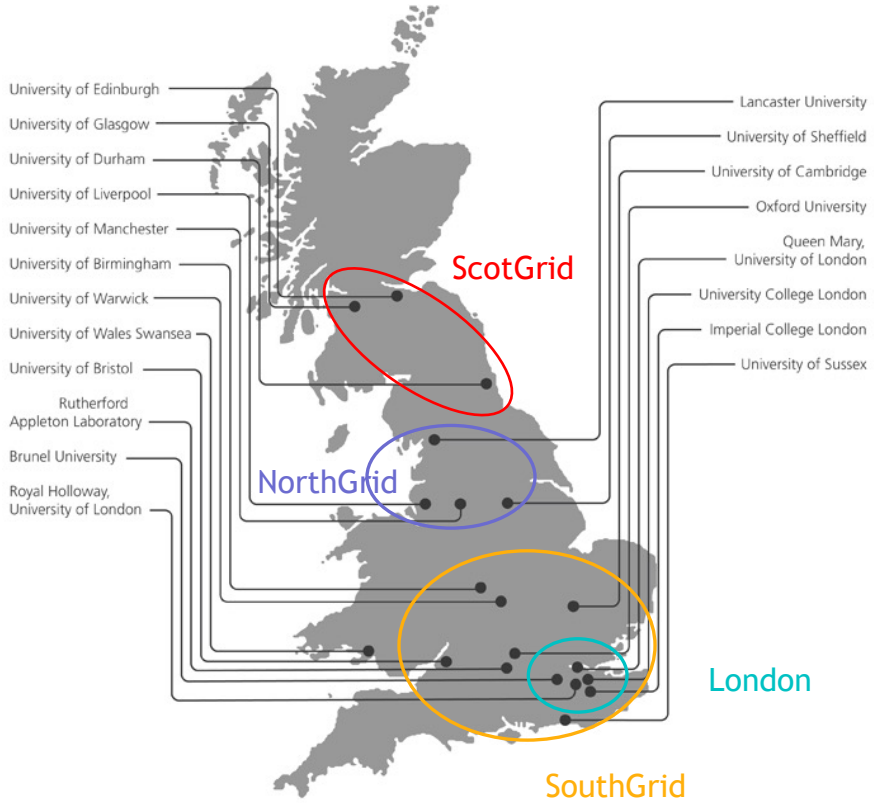

- RAL Tier 1 (2021)
  - 70 PB raw disk storage
  - 42,000 logical cores
  - 240 PB tape storage (2023)
  - Approx. 10% of existing CPU and disk storage is available for non-LHC experiments
  - Resource is allocated through IRIS
  - Future planning assumes EIC will apply for funding for its own allocation, managed by IRIS and implemented by GridPP
  - Exploiting synergies with WLCG data management is highly desirable

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.

Summary

- We have presented the aspirations of Canada, Italy and UK in participating in the computing of the future ePIC experiment
- These Countries are front runners and proof of principle for the development of a Computing Model including international participation
- A Computing Model that uses distributed resources in as many steps as possible will help the commitment of International Institutions

BackUp

RAW DATA mirroring for HL-LHC

- At HL-LHC, both the ATLAS and CMS experiments will produce ~ 350 PB of RAW data per year.
- The traffic from CERN to the T1s for RAW data quasi real time export will be ~ 400 Gbps per experiment (7M seconds of LHC data taking per year).
- + ~ 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

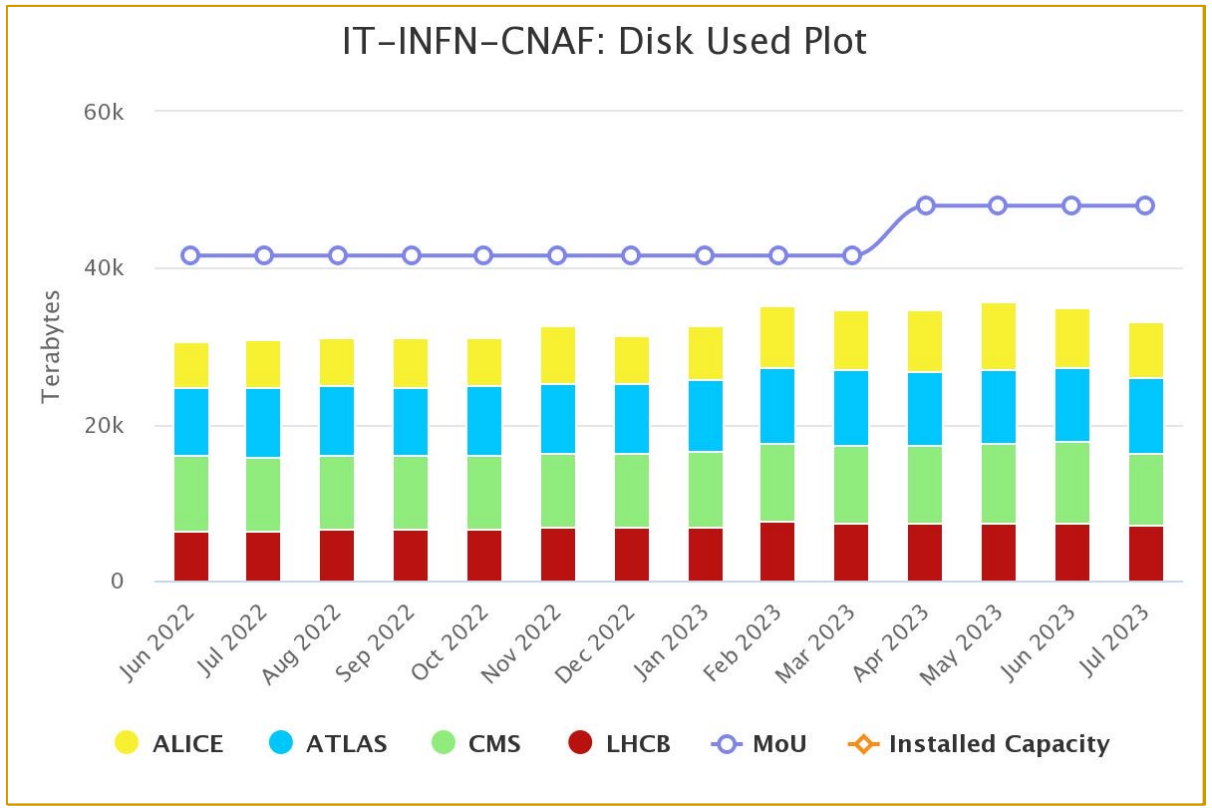

- About 6-7 PB increase/year for ALICE, ATLAS and CMS
- ATLAS/IT about 270 FTEs
- CMS/IT about 200 FTEs
- ALICE/IT about 130 FTE (110 paying members for pledge)
- Or a total of 600 FTEs
- Estimate for ePIC: 50 ± 10 FTEs

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)