THE (EXPERIMENTAL) SCIENTIFIC COMPUTING AT INFN

G. Carlino (INFN NA), P. Lubrano (INFN PG)

ICSC Spoke 3 pre-kick off meeting
Roma, 16/06/2022

Le 5 linee di ricerca
e le commissioni scientifiche nazionali

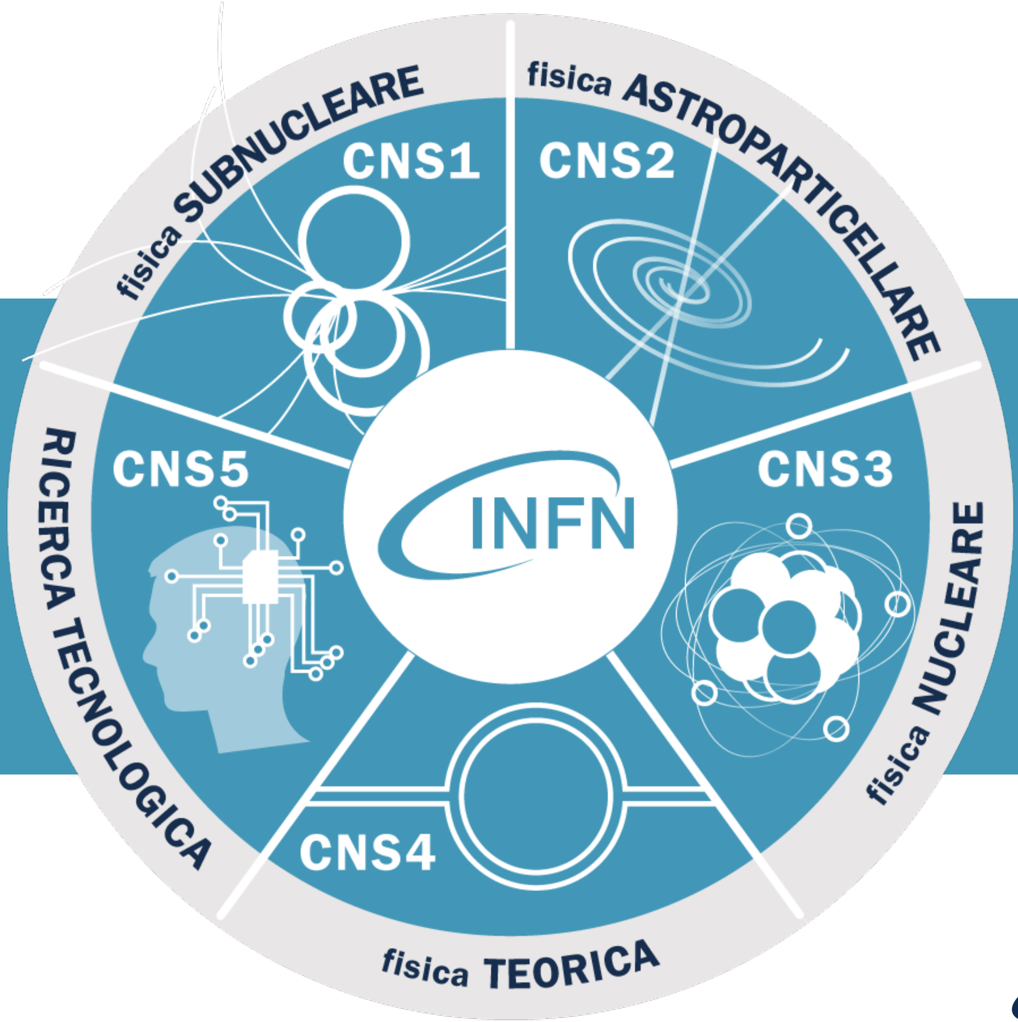

INFN

Strutture
dell'INFN

- 4 Laboratori Nazionali
- 20 Sezioni
- 6 Gruppi associati
- 3 Centri Nazionali e Scuole
- 1 Consorzio internazionale

INFN

La nostra **missione**

Spingersi oltre
le frontiere della conoscenza
I segreti del Big Bang

Sviluppare
nuove tecnologie di frontiera

Formare
gli scienziati e gli ingegneri di domani

Lavorare insieme
con i giovani e i ricercatori di tutto il mondo

INFN

Scientific computing @ INFN : where we are today

INFN computing infrastructures. INFN is a pioneer in the design and implementation of large-scale computing infrastructures and applications

- Primarily to meet the needs of the latest generations of high energy physics (HEP) experiments: **delivers the LHC experiments O(7-20)% of their computing budgets**
- **now rapidly extending to other communities.**

Ten centers were selected to host WLCG Computing:

- A Tier-1 at CNAF (red in the picture)
- 9 Tier-2s at Bari, Catania, LNF, LNL/Padova, Milano, Napoli, Pisa, Roma, Torino, (yellow in the picture)
- Then came the GRID, the Cloud, …
- They are all still operational, even if their size has increased ~100x since then and their interconnectivity (thanks to GARR-X) reaches multiples of 100 Gpbs

Where we are today

- INFN Distributed Computing federation delivers the LHC experiments O(7-20)% of their computing budgets
- ”non LHC” (VIRGO, Astro, Nuclear, …) is ~ 10-20% of the total
- Sites are of top quality among their peers, and have worked uninterruptedly for the last ~15 years

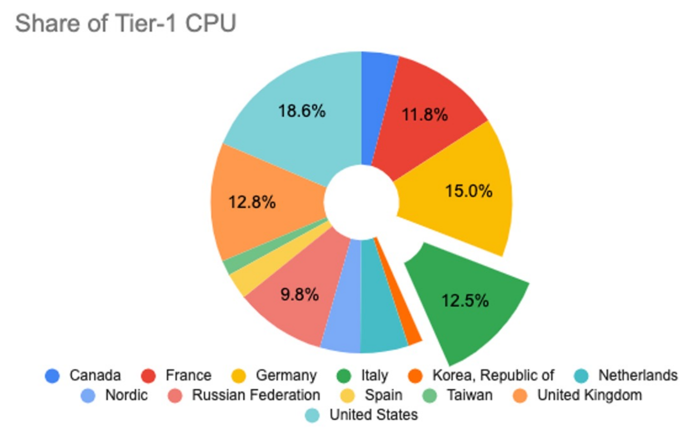

- LHC is the vast majority of resources, but not of manpower / effort: they have become the “standard” computing, which is ~easy to handle. Smaller / peculiar / different use cases are the challenge

The (not-only HEP) challenges

- The next generation of HEP experiments have unprecedented needs for computing, initially close to "impossible"
  - We can do a case study of HL-LHC, FCC-ee, FCC-hh, CEPC, ILC, CLIC to see where / when we expect problems
- Other Scientific domains

ProtoDune 2-3 GB/s (like CMS); Real Dune 80x

SKA up to 2 PB/day

A single genome ~ 100 GB. a 1M survey = 100 PB

CTA projects to 10 PB/y

1 LHC Experiment ~2020: ~ 200.000 CPU Cores; ~ 200 PB disk; ~350 PB tape

Using the same computing model

Need to optimize the computing model

1 HL-LHC Experiment ~2028: ~ 15M CPU Cores; ~ 15 EB disk; ~26 EB tape

The data lake model

- Keep the real value from the experiments safe
  - (RAW) **data** and a solid baseline of **CPU** in owned and stable sites
  - Allow for multiple CPU resources to join, even temporarily
    - Eventually choosing the cheapest at any moment
  - Solid networking: use caches / streaming to access data
- Reduce requirements for Computing resources
  - Commercial Clouds
  - Other sciences' resources
    - SKA, CTA, Dune, Genomics, ...
  - HPC systems

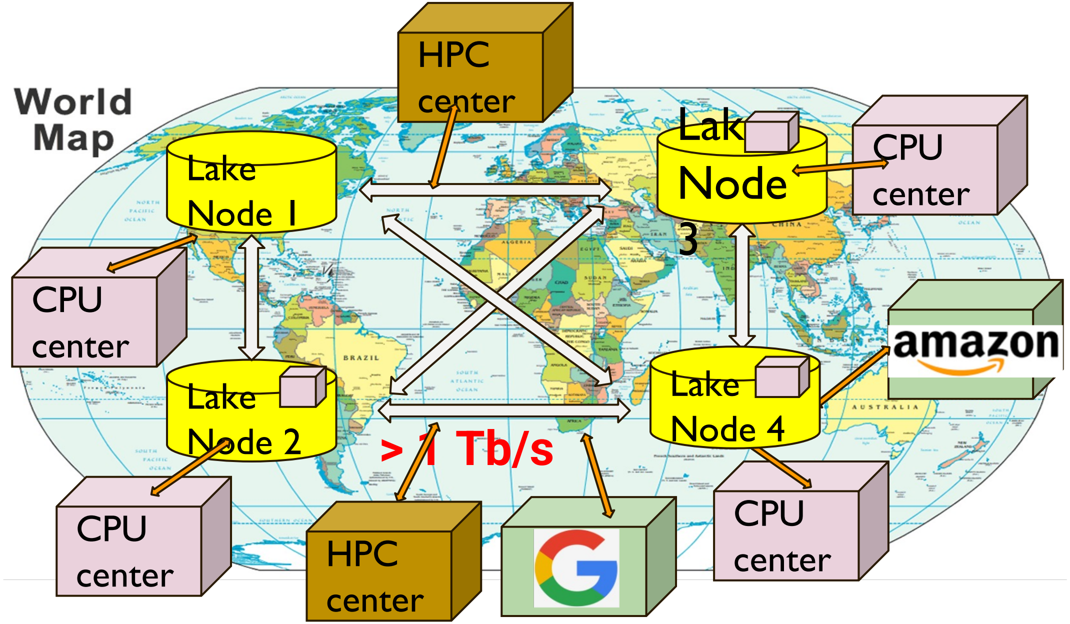

- The biggest gain in storage is the removal of the need to have multiple copies. Remote access (mediated via caches when needed) as efficient as local files
- There is a hidden operational gain: experiments will see just a big storage, and will not have to deal with 100+ computing and storage centers

Supercomputing (HPC)

The world is literally full of Supercomputers. Why ?

- Real scientific use cases
  - Lattice QCD, Meteo, ...
- Industrial showcase ("Country XY is technologically capable")
  - And hence not 100% utilized, opportunities for smart users. Can we be one of them?

Many not trivial problems to solve:

- **Data access** (access, bandwidth, ...)
- **Accelerator** Technology (KNL, GPU, FPGA, TPU, ???, …)
- **Submission of tasks** (MPI vs Batch systems vs proprietary systems)
- **Node configuration** (low RAM/Disk, …)
- **Not-too-open environment** (OS, …)

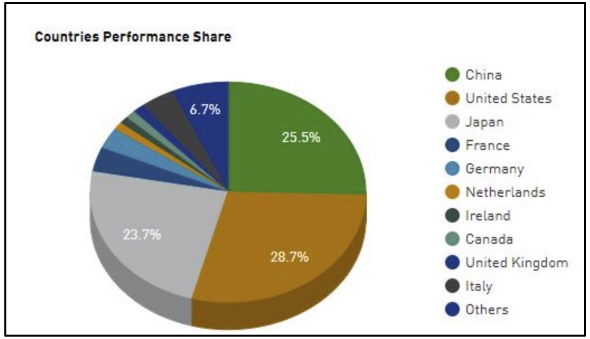

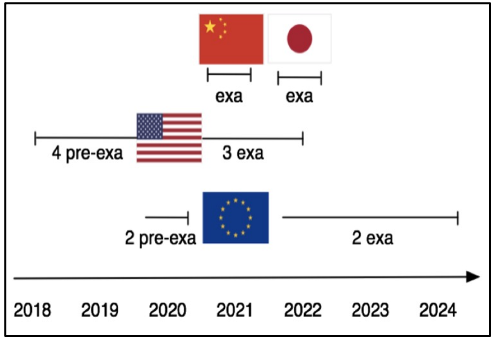

- Demonstration of the capability to execute LHC workflows on HPC systems
- Demonstration of the capability to execute workflows on non standard architectures (Power9)
- Demonstration of the capability to circumvent HPC security via user level tools
- Demonstration of interactive analyses on HPC systems

Enabling CMS Experiment to the utilization of multiple hardware architectures -- a Power9 Testbed at CINECA

T.Boccali[1], A.Malta Rodrigues[2], D.Spiga[3], M.Mascheroni[4] for the CMS Collaboration

[1] INFN Sezione di Pisa, [2] University of Nebraska-Lincoln, [3] INFN Sezione di Perugia, [4]University of California San Diego

Extension of the INFN Tier-1 on a HPC system

Tommaso Boccali[1], *Stefano* Dal Pra[2], *Daniele* Spiga[3], *Diego* Ciangottini[3], *Stefano* Zani[2], *Concezio* Bozzi[4], *Alessandro* De Salvo[5], *Andrea* Valassi[6], *Francesco* Noferini[7], *Luca* dell'Agnello[2], *Federico* Stagni[6], *Alessandra* Doria[8], *Daniele* Bonacorsi[9]

Cloud Computing (INFN-Cloud)

INFN is offering to its users a comprehensive and integrated set of Cloud services through its dedicated **INFN Cloud infrastructure**.

The starting point for a National Datalake for research and beyond, building on (existing|renewed|new) e-Infrastructures.

- Built on a thin middleware layer running on top of federated clouds, decoupling physical and logical views via a service composition mechanism

The driving force for the development of all Cloud initiatives at INFN

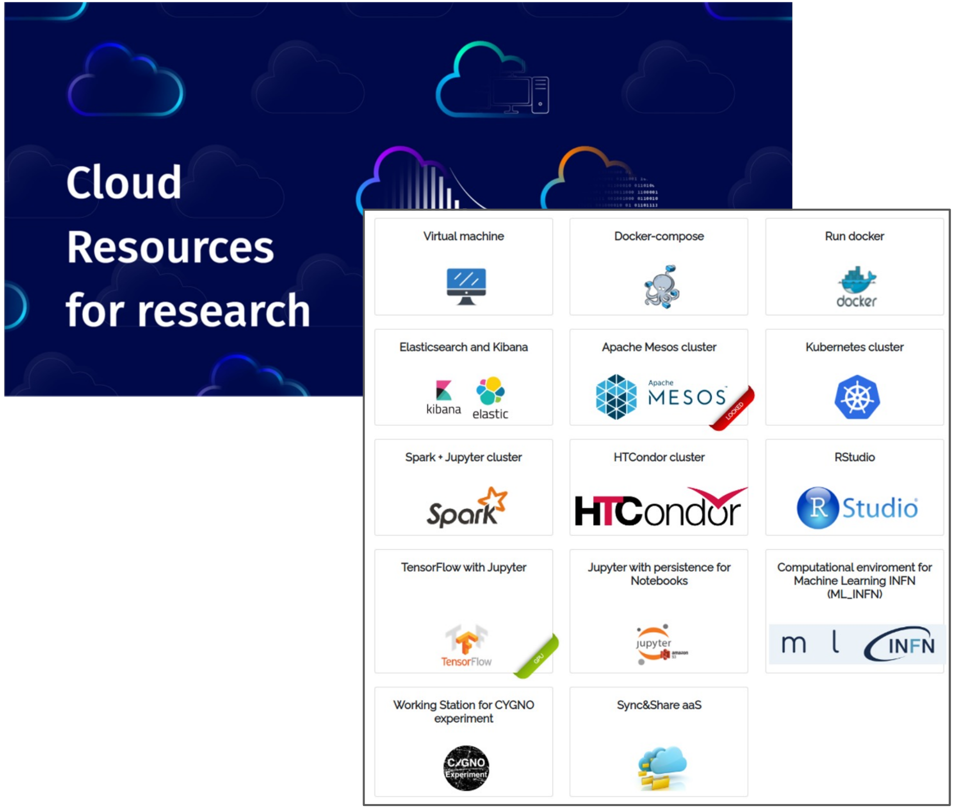

THE OPPORTUNITIES ….

- In **2017**, Bologna won a bid to host the “*European Centre for Medium-Range Weather Forecasts*”
- The Emilia Romagna region decided to repurpose the “*Manifattura Tabacchi **” area to host a technology district, hosting ECMWF <u>and more</u>

Roughly 250x250 m^2

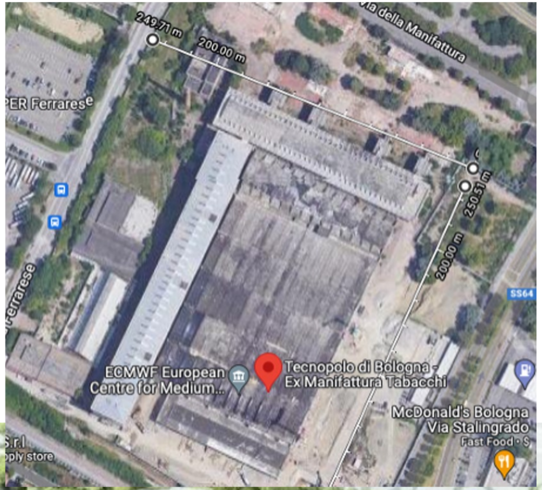

How it will be

(*a former factory of cigarettes)

WHAT CAN THE TECNOPOLO HOST?

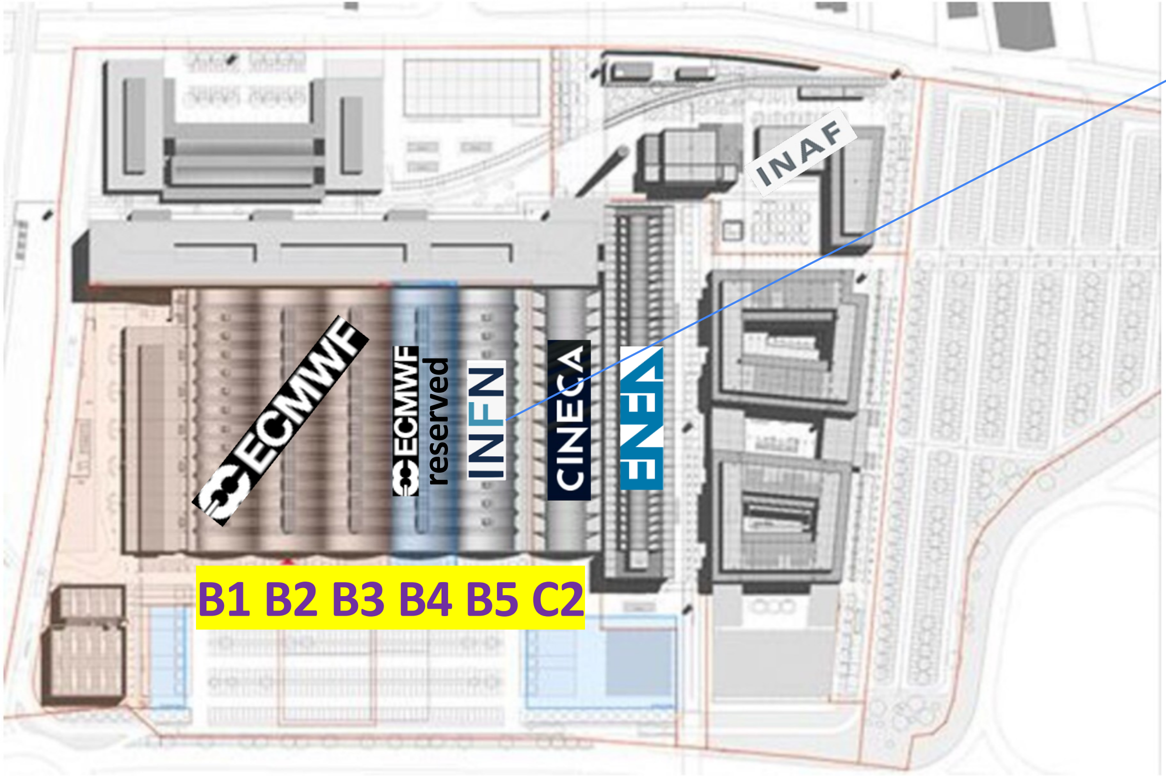

Each of the 6 “botti” (barrels) is ~5000m^2 of usable IT space

Same architect and design of the ”Sala Nervi” in the Vatican

The INFN Unit in Spoke 3

- **Gianpaolo Carlino (Napoli, Atlas) WP1**
- **Sara Cutini (Perugia, Fermi/Virgo) WP4/WP5/WP6**
- **Stefano Della Torre (Milano Bicocca, AMS2/Litebird) WP2/WP3/WP6**
- **Matteo Duranti (Perugia, AMS2/HERD) WP2/WP4/WP6**
- **Valerio Formato (Roma2, AMS2/HERD) WP2/WP4/WP6**
- **Fabio Gargano (Bari, Fermi/HERD) WP2/WP4/WP5/WP6**
- **Dario Gasparrini (Roma2, CTA/Fermi/HERD) WP4/WP5/WP6**
- **Martina Gerbino (Ferrara, Euclid/InDark) WP2/WP4/WP5**
- **Massimiliano Lattanzi (Ferrara, Euclid/InDark) WP2/WP4/WP5**
- **Pasquale Lubrano (Perugia, Fermi/NA62) WP1**
- **Mario Nicola Mazziotta (Bari, Fermi/HERD) WP2/WP4/WP5/WP6**
- **Melissa Pesce Rollins (Pisa, Fermi/XRO) WP4/WP5**
- **Giovanni Signorelli (Pisa, Litebird/LSPE/MEG) WP2/WP3**