SCAI
Super Computing Application & Innovation
CINECA

Sanzio Bassini – January 2020

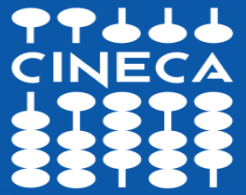

CINECA at a glance

CINECA is a **Consortium of Universities**, founded in 1969 by (MIUR) to support scientific research
CINECA is a not-for-profit organization

Members:

- MPI & MUR
- 67 Italian Public Universities
- 10 Institutions (Anton Dohrn, CNR, CREA, INAF, INDIRE, **INFN**, INRIM, INVALSI, OGS, Policlinic Umberto I°)

In-house provider:

- Supervised by MUR
- No private capital
- 80% of activities towards Consortium members
- 20% Technology Transfer
- Main premises in Bologna
- Annual Budget: 100M€
- Employees: 700

Main Cineca's activities

- High Performance Computing & Technology Transfer
- ICT in house providing for the MPI & MUR
- ICT in house providing for Universities

The HPC – BIG data Exascale race

Projected Exascale Investment Levels

U.S.

- $1 to $2 billion a year in R&D (around $10 billion over 7 years)
- Investments by both the government & vendors
- Plans are to purchases multiple exascale systems each year

EU

- About 5-6 billion euros in total (around $1 billion a year)
- EU: 486M euros, Member States: 486M euros, Private sector: 422M euros
- Investments in multiple exascale and pre-exascale systems
- Large EU CPU funding

China

- Over $1billion a year in R&D (at least $10 billion over 7 years)
- Investments by both governments & vendors
- Plans are to purchases multiple exascale systems each year
- Investing in 3 pre-exascale systems starting in late 2018

Japan

- Planned investment of over $1billion* (over 5 years) for both the R&D and purchase of 1 exascale system
- To be followed by a number of smaller systems ~$100M to $150M each
- Creating a new processor and a new software environment

- **Exascale systems are being designed for HPC, AI, HPDA, etc.**
- **Competitive forces are driving companies to aim more complex questions at their data structures and push business operations closer to real time**
- **Iterative methods will expand the size of data volumes needing to be stored**
- **Physically distributed, globally shared memory will become more important**
- **<u>Artificial Intelligence will grow faster than everything else</u>**

www.HyperionResearch.com

- **HPC is at the forefront of R&D for economically important AI use cases.**
  - HPC shows where the mainstream AI market is headed.

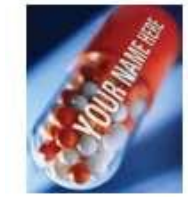

Precision Medicine

Automated Driving Systems

Fraud and anomaly detection

Affinity Marketing

Business Intelligence

Cyber Security

IoT/Smart Cities

Hyperion study:
70% of the most innovative blue chip world wide industries will invest in AI and HPC big data processing as a key technology to make persistent theirs development trends.

HPC Systems @ Cineca

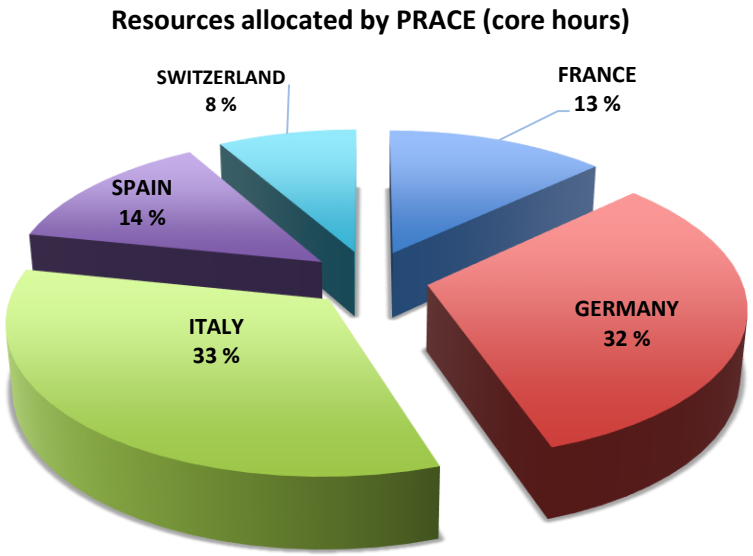

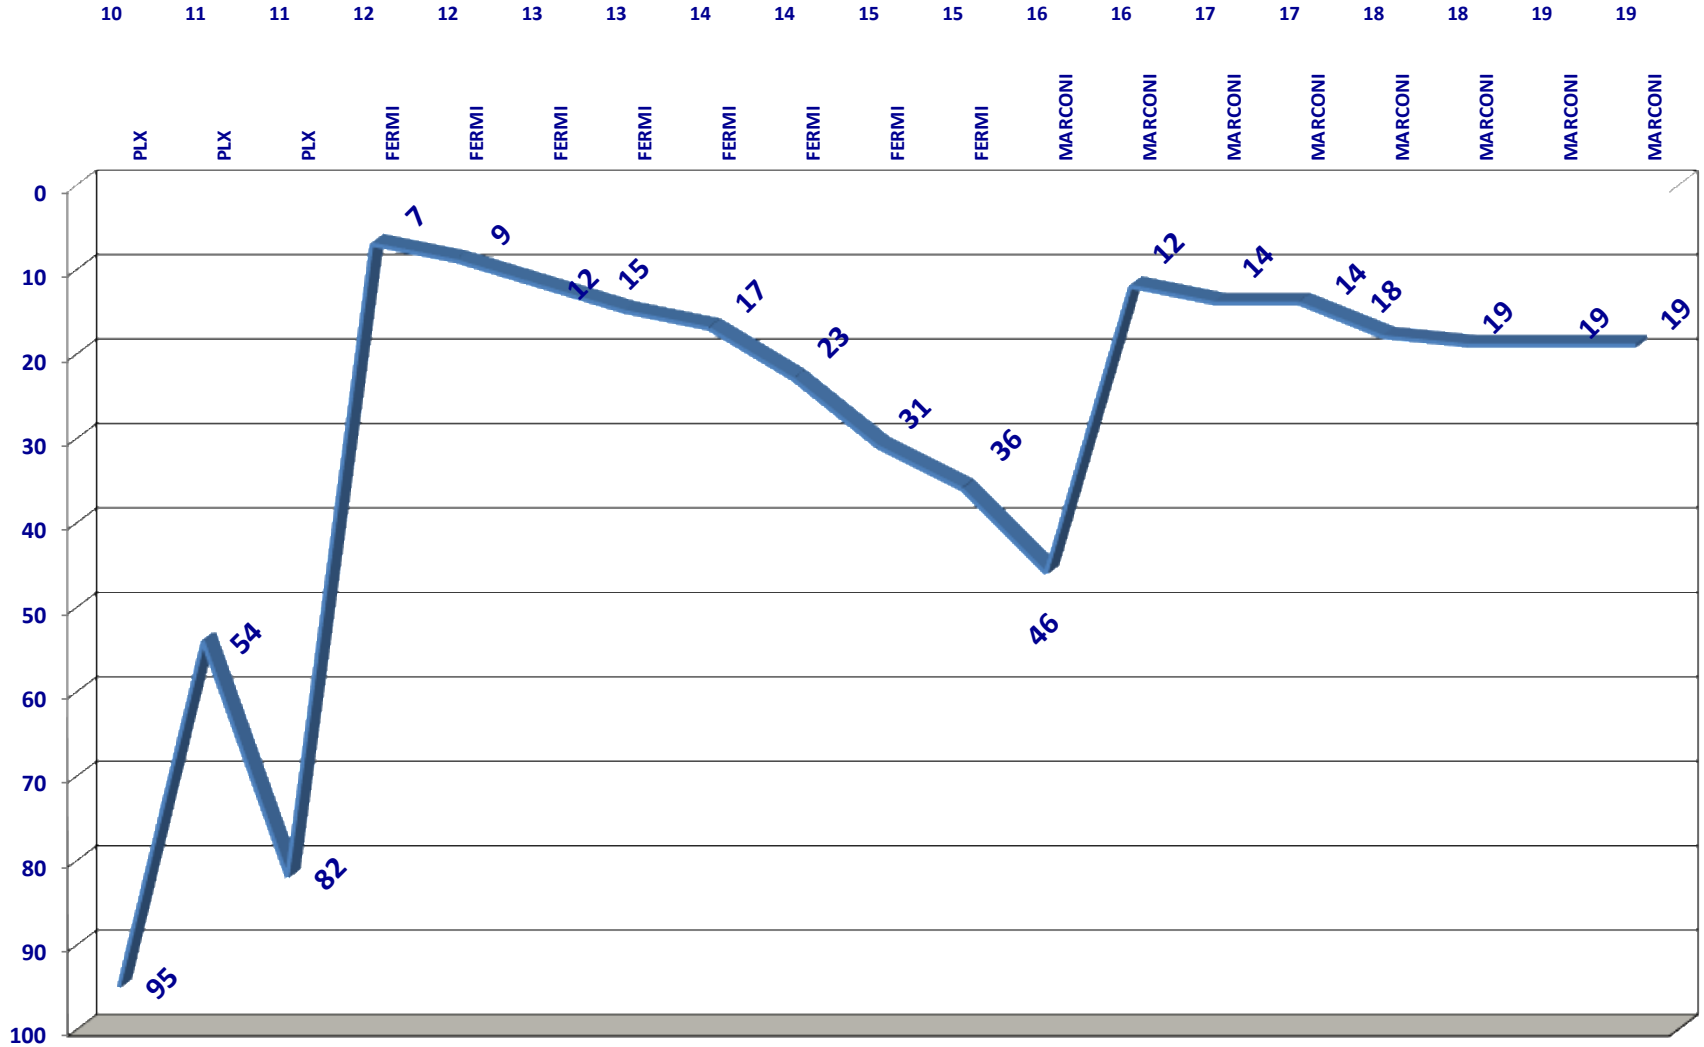

CINECA HPC architecture infrastructure

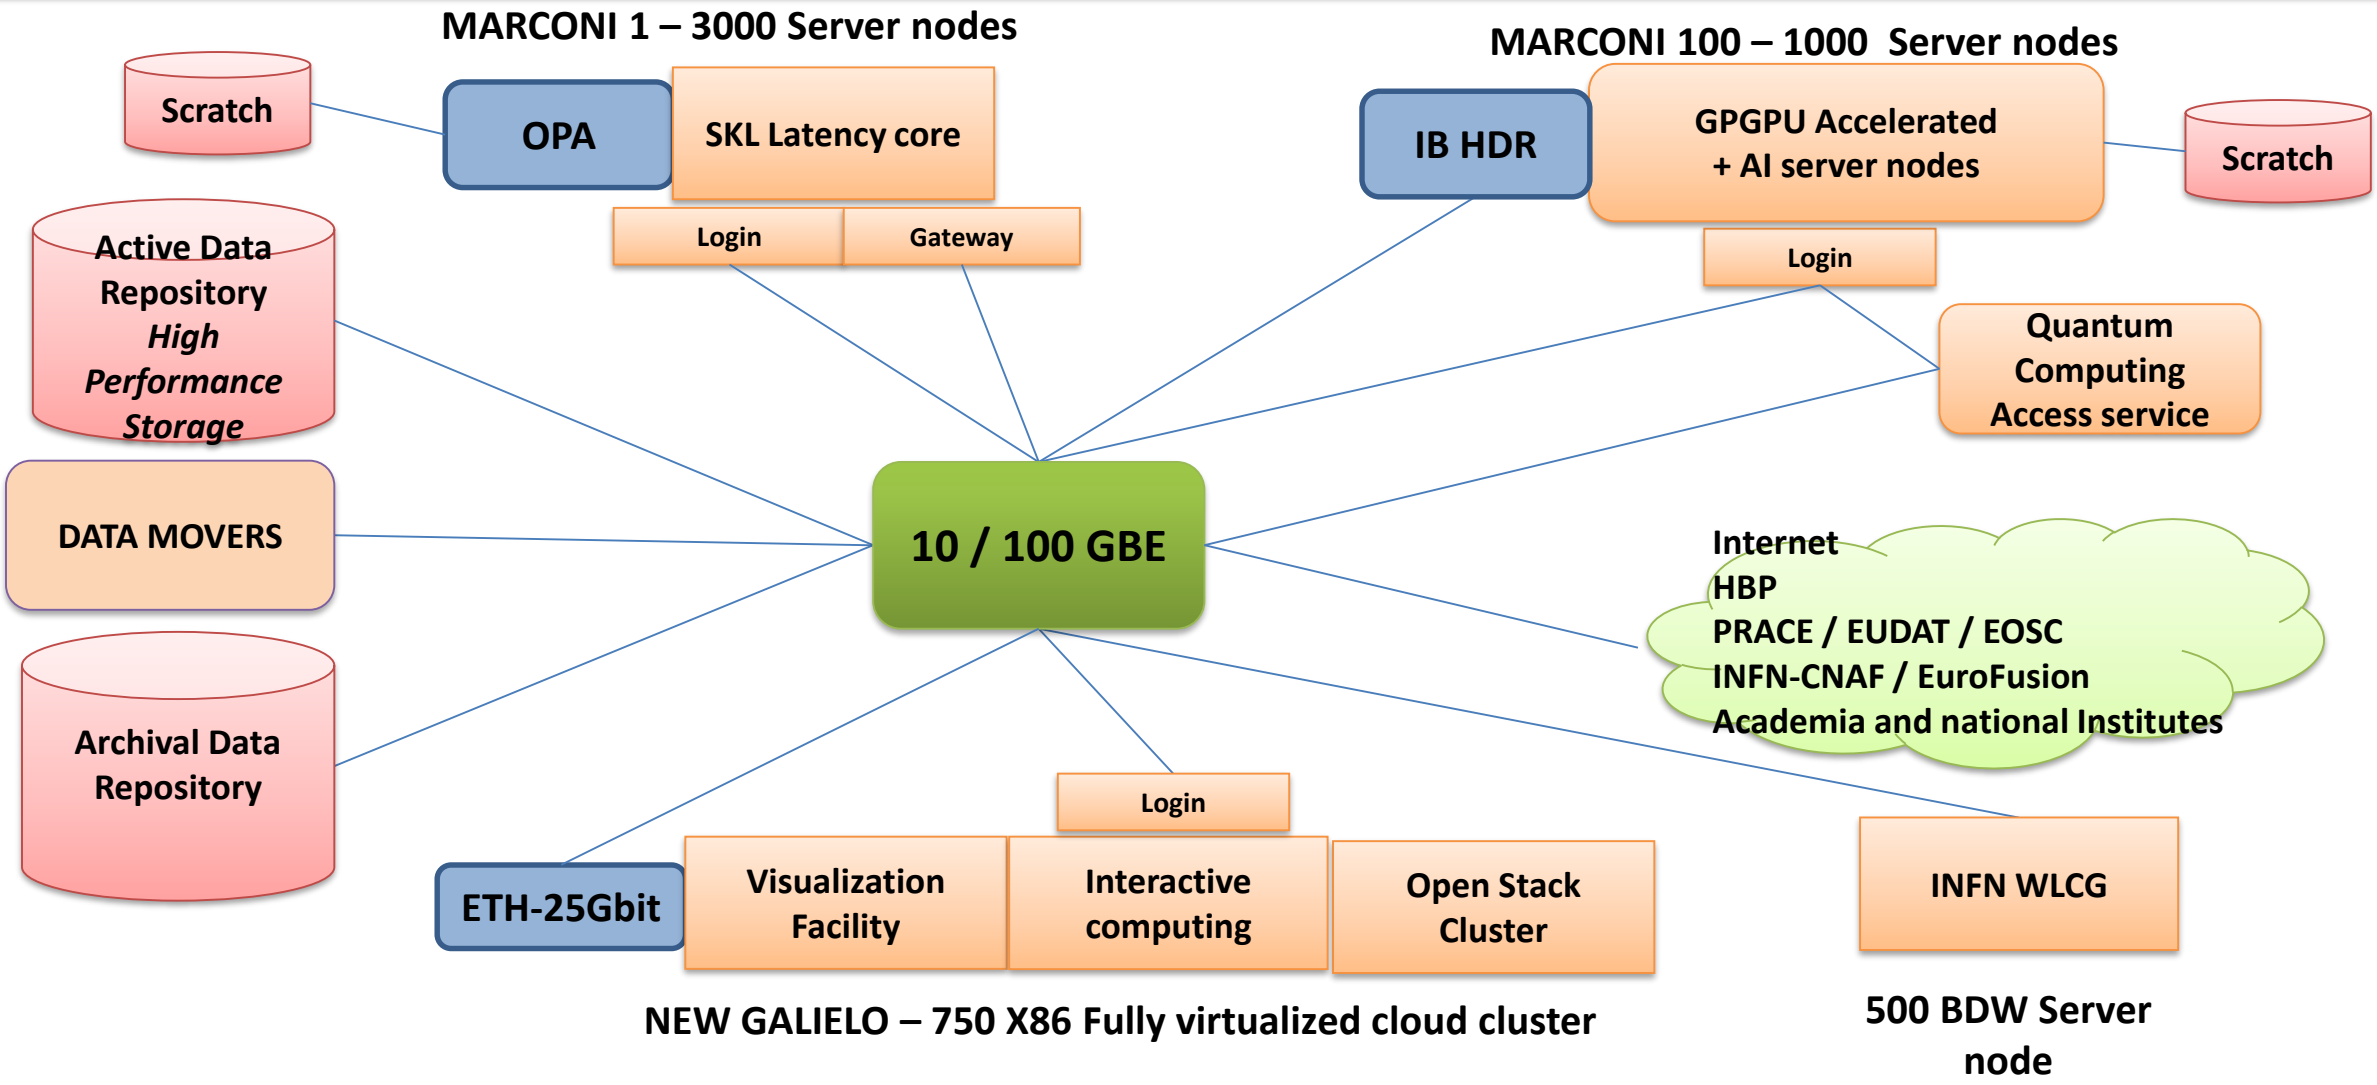

Supporting Innovation

PRACE SHAPE SME HPC Adoption Programme in Europe

FORTISSIMO

bi-rex Big Data Innovation & Research Excellence

CINECA

Industrial Area
Chemistry
Life Science
Engineering
Geophysics
AI & ML
Cultural Heritage

Main Activities
Molecular Dynamics
Material Science Simulations
Geophysics Simulations
Fluid dynamics Simulations
Applications Developments
Engineering Applications
Code Parallelization
Code Optimizations
Graphics interfaces & VR

Dompé
IDI FARMACEUTICI
Tetra Pak
ALSTOM
SACMI
Chiesi
dallara
BOMBARDIER
AleniaAeronautica
PILATUS
Cremascoli Ortho Group
FERRETTIGROUP
PASSION. INNOVATION. EXCELLENCE.
TEXA
LUNA ROSSA CHALLENGE
ORACLE RACING
FIAT CRF CENTRO RICERCHE FIAT
NOLAN
Ferrari
LAMBORGHINI
UnipolSai ASSICURAZIONI
FONDAZIONE CASSA DI RISPARMIO IN BOLOGNA
PHILIP MORRIS
FERRERO
Prysmian Group
eni

Current active contractual agreement with Industry and SME

- Altran: Engineering and R&D services (Torino)
- Amet: Engineering Company (Torino)
- Chiesi farmaceutici: Innovation in healthcare (Parma)
- Dompè Farma: Innovation in drug design (L'Aquila)
- Eni Oil & Gas (Milano)
- Elica hoods (Ancona)
- Ferretti Yachts (Forlì)
- Nolan Group: motorbike helmet (Bergamo)
- OlsaGroup: Optical Lighting Systems (Torino)
- Unipol Group: insurance (Bologna)
- Philip Morris International (Bologna)
- Ferrero (Cuneo)

Prospect: FEV Powertrain; Ferrari Automobiles

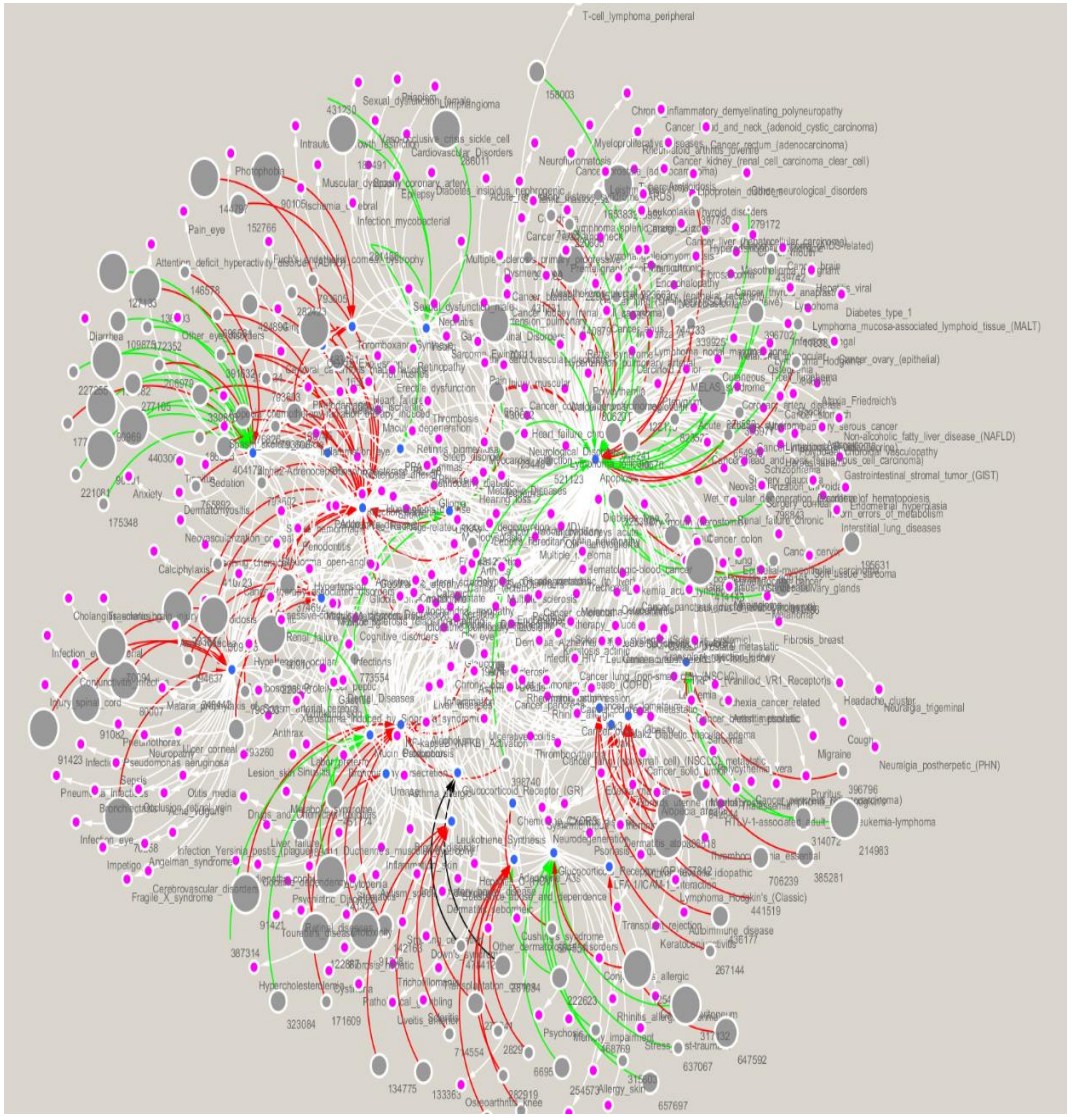

CINECA e Eni: i punti salienti di una partnership ventennale (1999-2019)

GigaFlops

TeraFlops

PetaFlops

ExaFlops

1° Cluster AlphaSC

10.000 CPU x86

Hybrid (CPU+GPU)

HPC3 & HPC4

HPC X (Intersect)

1999 2004 2005 2008 2009 2012 2013 2014 2016 2017 2018 2019

CPU Rental
Commercial Apps

Data Center CINECA

1° Eni Proprietary Application: KTA

eDVA **1.0**

GPU Code Porting

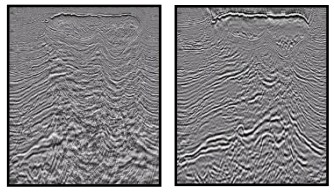

RTM 60 Hz

HPC1 HPC2

GDC ENI

Servizi & supporto 24/7

FWI

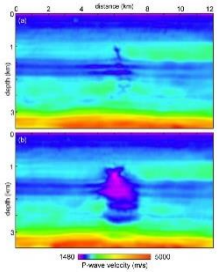

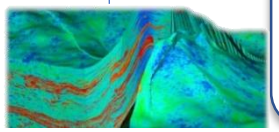

Reservoir Modeling Echelon

eDVA **10.0**

INFN HPC architecture infrastructure

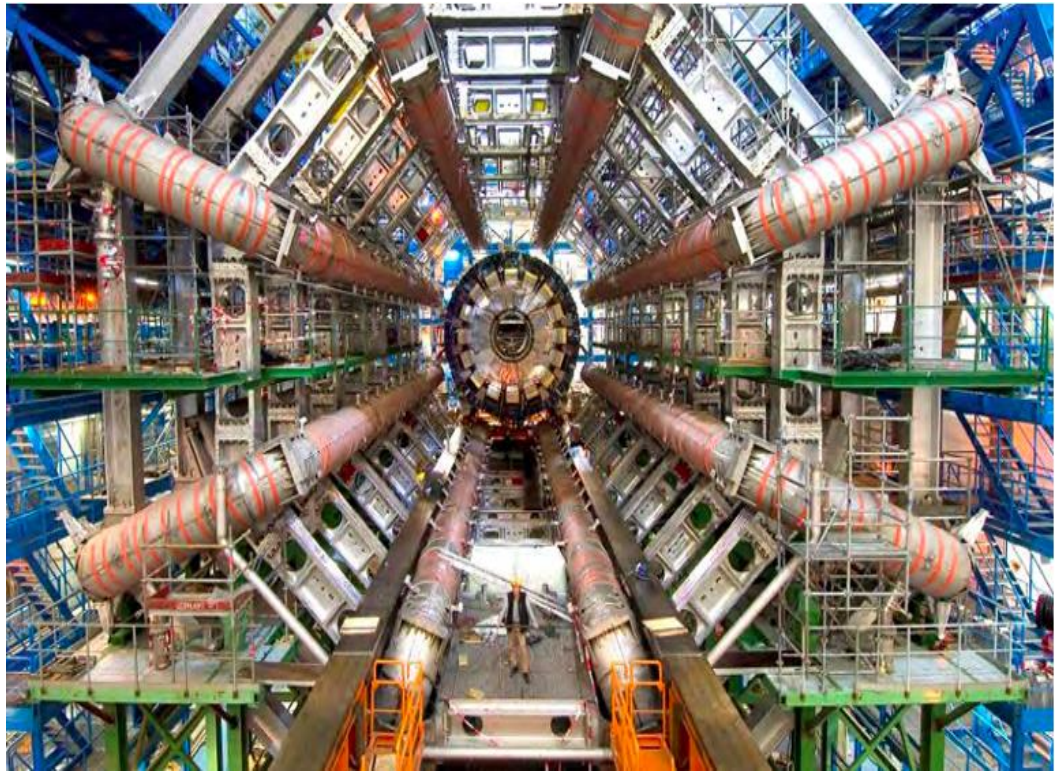

Big Data Association partnership - BDA

- Governance
- Inclusiveness
- National & International Networking
- Awareness
- Funds raise

Supercomputing Unified Platform-ER

I4.0, ML, IA

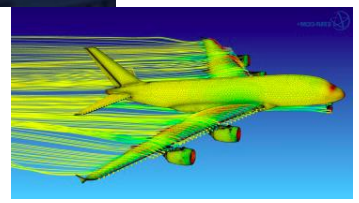

New material

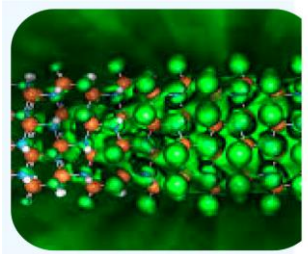

Extreme earth

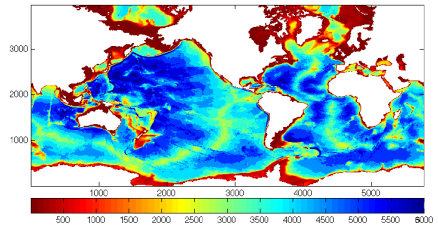

Life science, genomic, biobanking

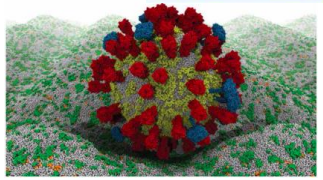

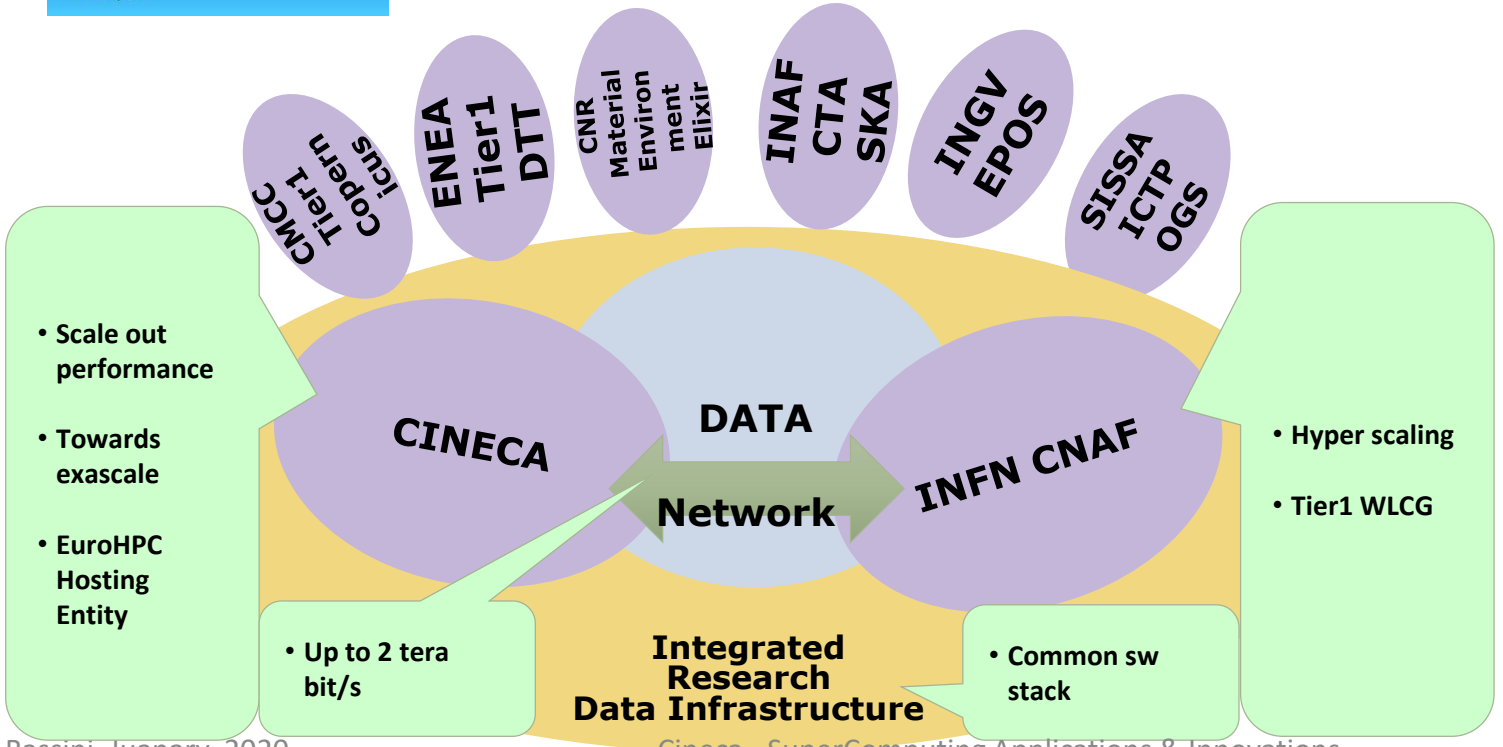

- Federated Infrastructure
- Federated National User ID Management
- Federated National Multi Tier storage services
- National infrastructure model for European Open Science Cloud Initiative

CINECA

Enabling federeted infrasctructure, data repository Test and experiment facility for

EuroHPC Joint Undertaking

EOSC

ETP 4 HPC

EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

BDV BIG DATA VALUE ASSOCIATION

AI

HPC RoadMap

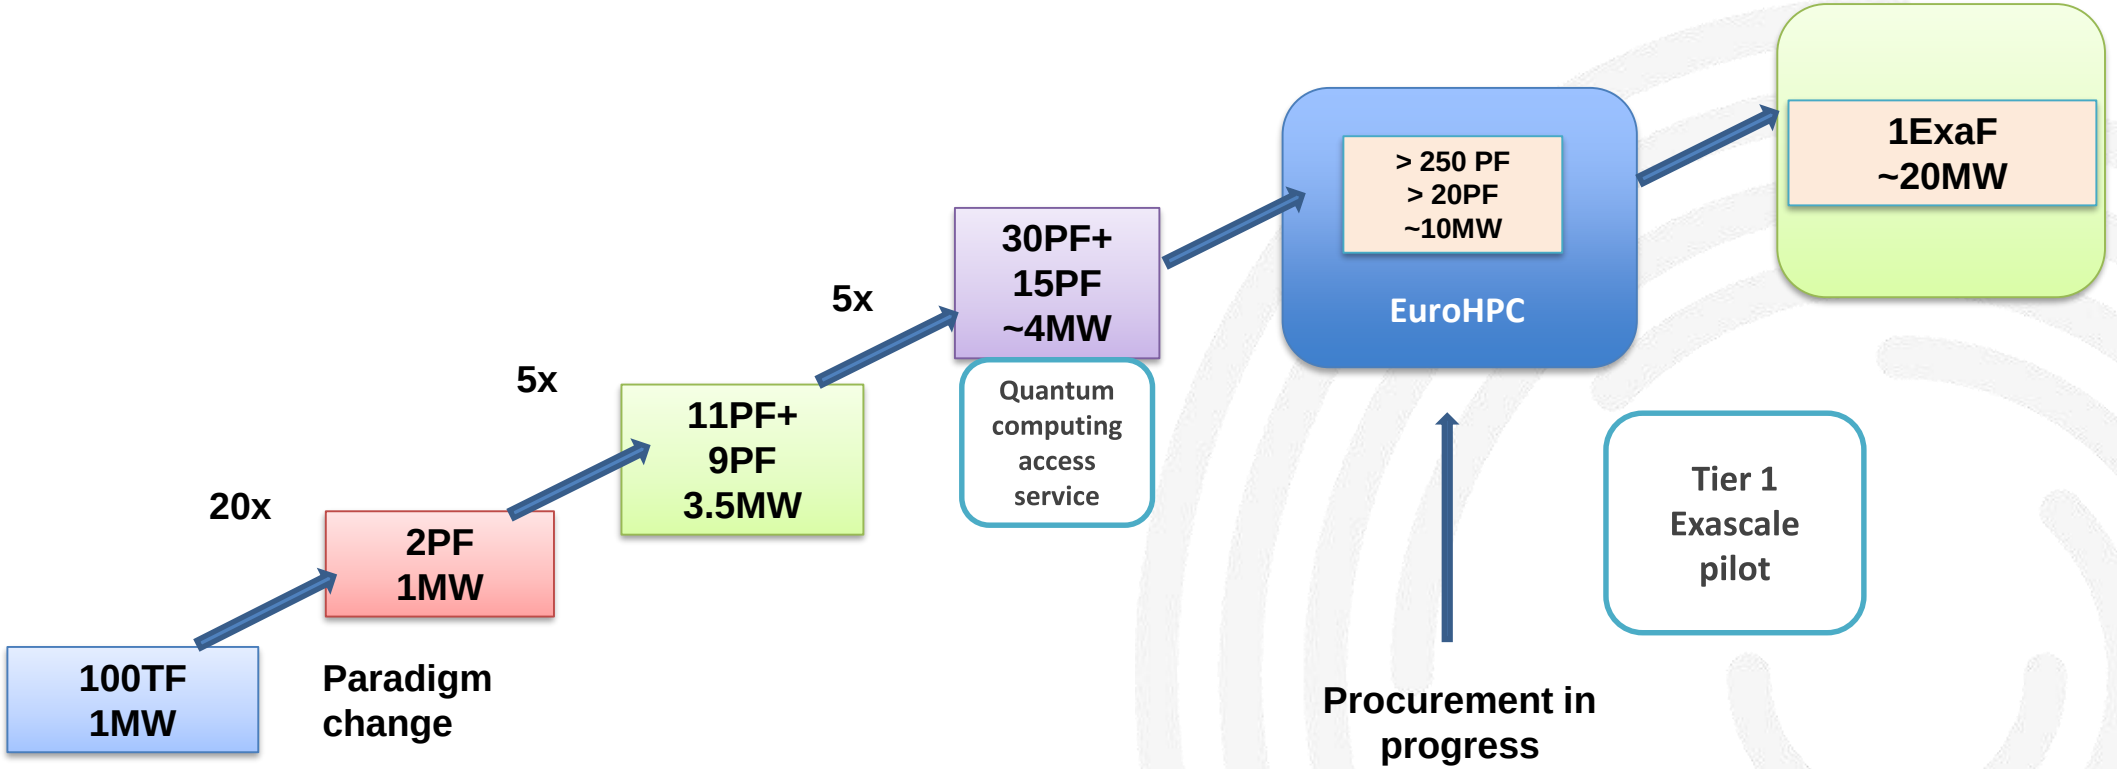

2009	2012	2016	2019 / 2020	2021	2025-2027
IBM SP6 Power6	Fermi IBM BGQ PowerA2	Marconi Lenovo Xeon+KNL	Marconi PPI4HPC ICEI - PPIHBP	Leonardo pre-Exascale with EuroHPC contribution	EuroHPC post Exascale/ National Full Exascale

CINECA

Ministero dell'Istruzione Università e Ricerca

MIUR

INFN
Istituto Nazionale di Fisica Nucleare

SISSA - ISAS

Regione Emilia-Romagna

REPUBLIKA SLOVENIJA
REPUBLIC OF SLOVENIA
Ministrstvo za izobraževanje, znanost in šport
Ministry of education, science and sport
Masarykova cesta 16, SI - 1000 Ljubljana

MINISTRY
OF EDUCATION, SCIENCE,
RESEARCH AND SPORT
OF THE SLOVAK REPUBLIC

AUSTRIA
REPUBLIK ÖSTERREICH

Innovációs és Technológiai
Minisztérium

World class HPC European Hub

Bologna Big Data Science Park

A modern energy efficient new data center capable of hosting a full exascale system

Cooling equipment
3MW (2020) -> 5MW(2023)

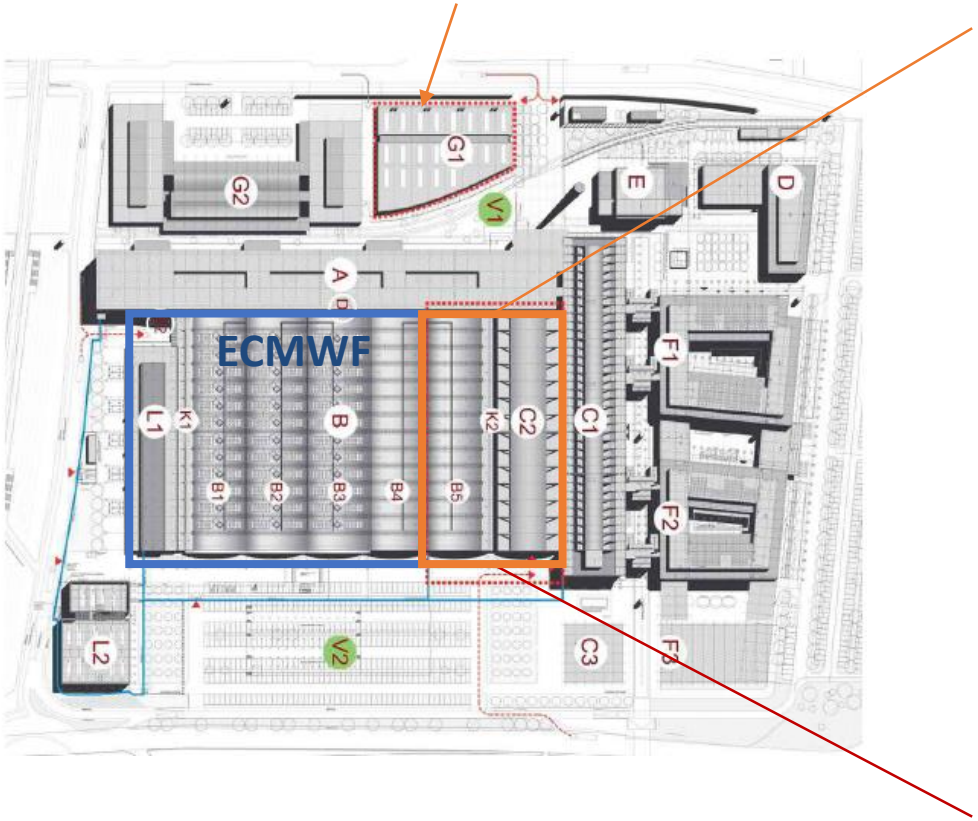

PUE < 1.1

Computer Rooms
10MW (2020) -> 20MW (2023)

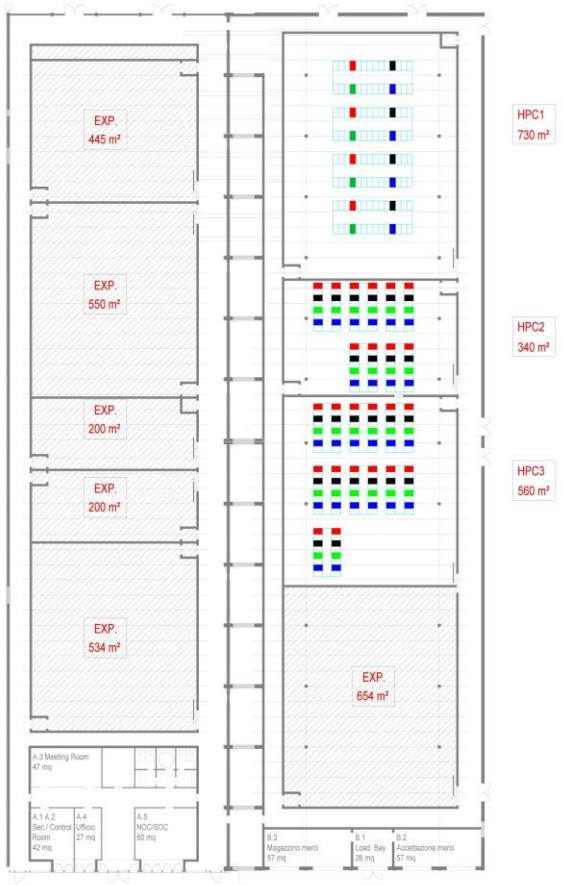

DATA ROOM STAGE 1: 1600 sqm
DATA ROOM STAGE 2: 2600 sqm
ANCILLARY SPACES: 900 sqm

8MW hot water DLC
Compute nodes

2MW AIR Cooled
Storage + Ancillary

Leonardo High Level description

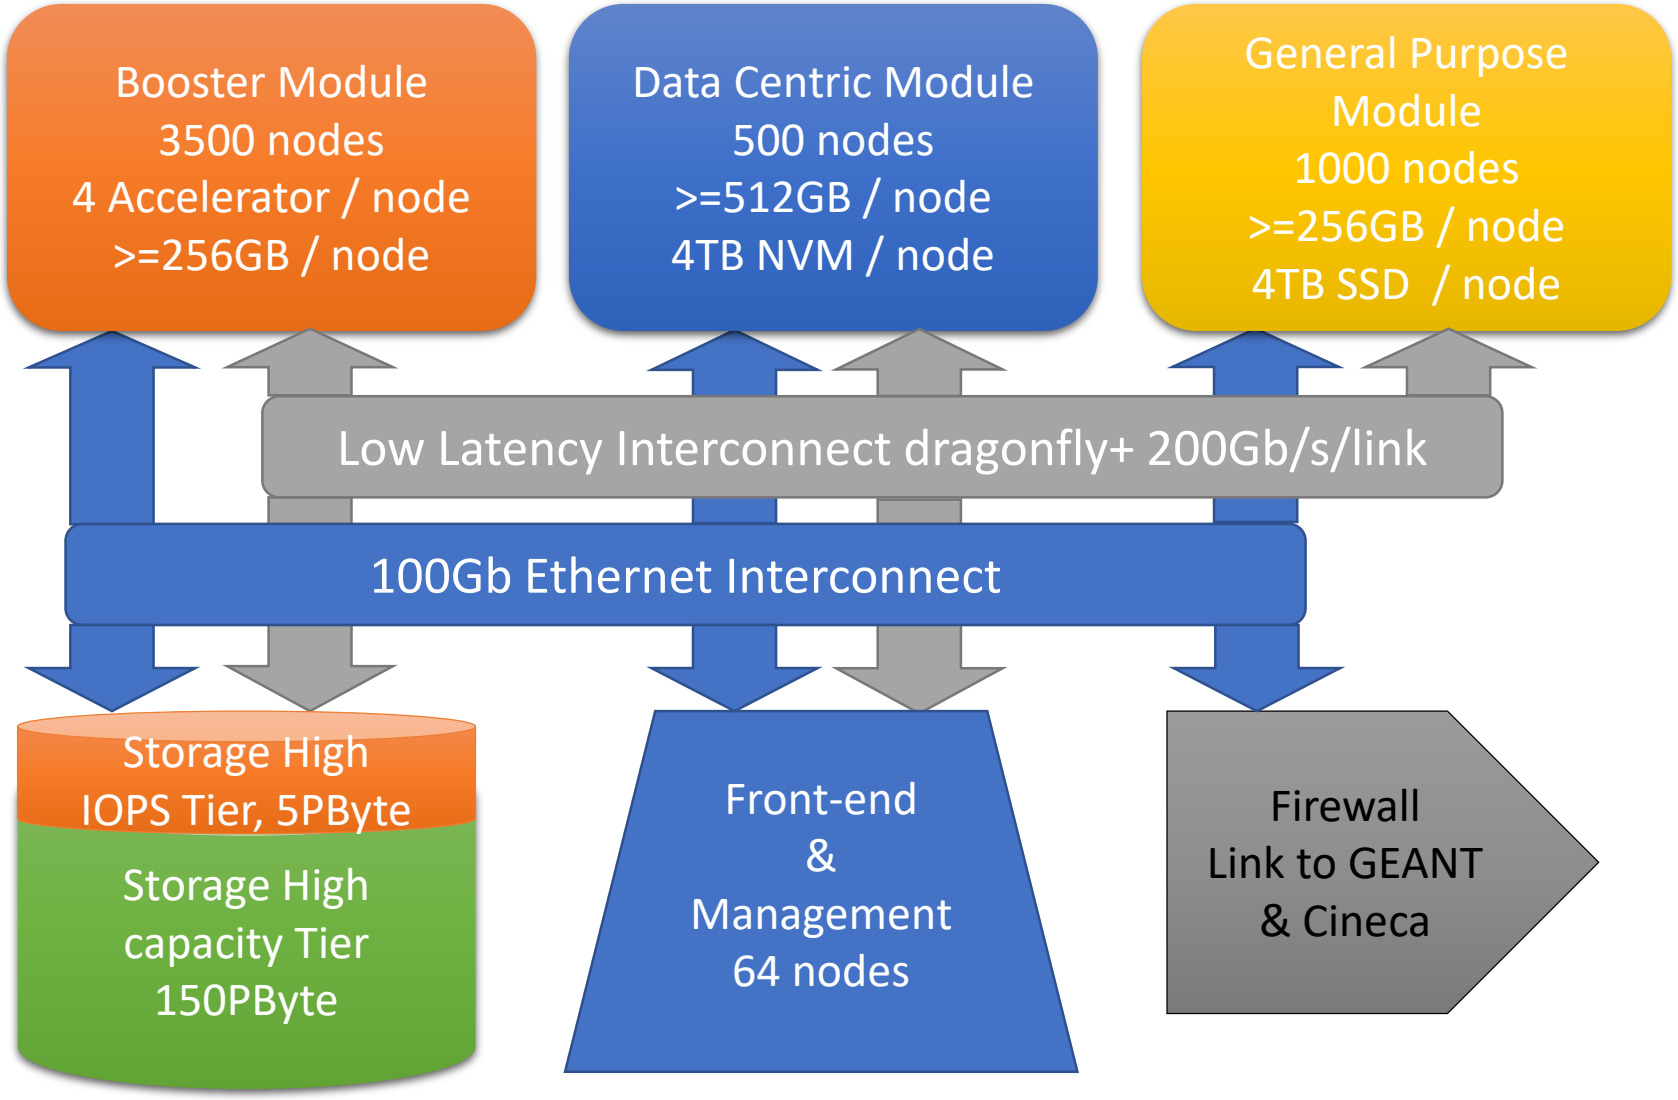

Leonardo will enable

Medical treatments tailored to the patient

internet-of-Things

New material from scratch

Simulate formation of galaxies

Source of gamma-ray bursts

Predict solar eruptions

Properties of elementary particles

Next-generation weather forecasting

Advanced quantum chromodynamics

Uncertainty quantification in predictions

Understanding of general relativity

New algorithms / non-traditional areas

New batteries

Direct numerical simulations of Navier-Stokes eqs.

Improve nuclear power, hydropower, wind turbines

Scientific Progress Without Computers
(two paradigms)

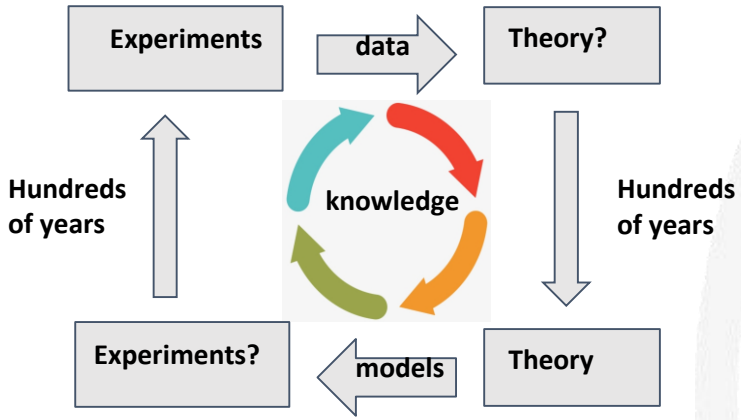

Scientific Progress With Computers
(four paradigms)

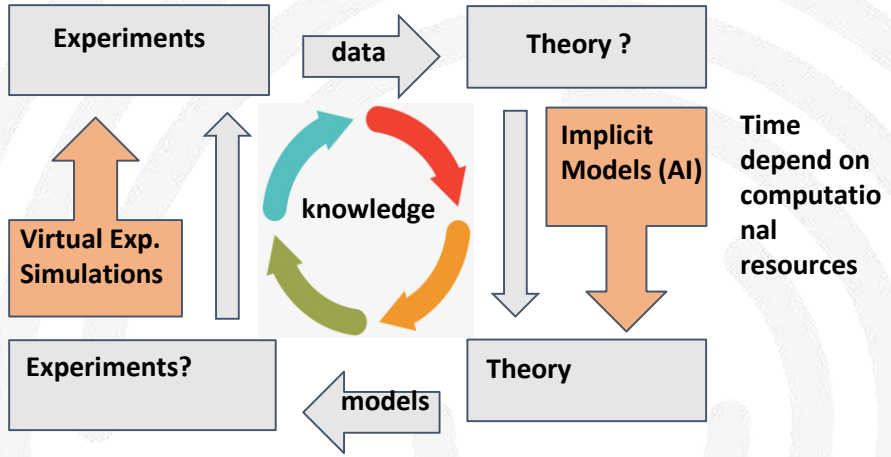

AI and HPC boost CADD

Computer Aided Drag Discovery

ANTAREX 10^18

FET-HPC

European Commission

Exscalate *smart* solution

EXaSCale smArt pLatform Against paThogEns

EXAMON *smart* management of HPC

EXAscale MONitoring Framework

POLITECNICO MILANO 1863

Dompé

CINECA

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

Exscalate platform combine **A.I.** and **HPC** to boost green Computer Aided Drug Discovery

EXAMON platform combine **A.I.** and **BigData** to boost green and smart HPC

new models algorithm optimization

faster greener

A.I.

Ultra fast Drug Discovery platform

HPC Simulation

Training Autotuning

Data Parameters

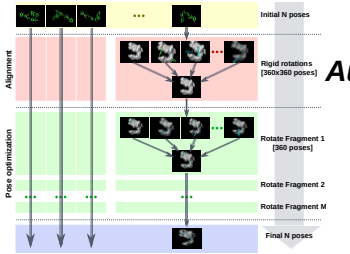

100x Faster!

System & Application Monitoring

Training Modelling

HPC Infrastructure Smart Management

A.I.

Run -time management

Power Reduction Anomaly Detection

20% Greener!

Big Data, AI for Smart HPC

EXAMON
EXAscale **MON**itoring
Framework

EXAMON

Common Interface

Machine Learning | Performance Analysis | Energy efficiency

Job Scheduling | Resources Management | Data Visualization

Reactive and Proactive Feedbacks

Heterogeneous Sensors

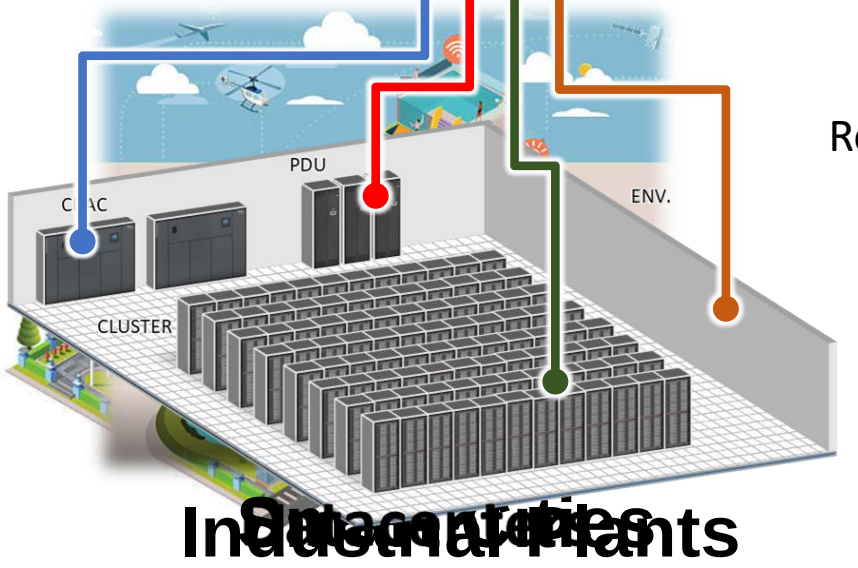

Industrial Plants

Multitherman

EU H2020 FETHPC project ANTAREX (g.a. 671623)

https://github.com/EEESlab/examon

F. Beneventi et al., “Continuous learning of HPC infrastructure models using big data analytics and in-memory processing tools”
A. Bartolini et al, “The DAVIDE Big-Data-Powered Fine-Grain Power and Performance Monitoring Support”

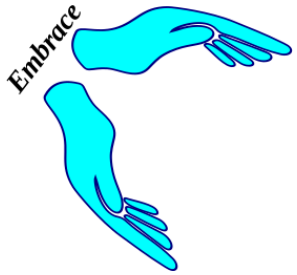

Glimpse of some more projects involving ML

Highlander: HPC to support smart land and agriculture services

Designing multidisciplinary protocols to investigate, using machine learning approaches, metrics originating from multiple domains (eg., climatic and genomic data) to establish inter-relations and to monitor any future variations of one due the other.

IMC: I-media cities to support smart digital cultural heritage

The innovative platform for automatic annotation of multimedia content items that have already been enriched with metadata and that will be further manually annotated. The annotations allow the efficient retrieval of information - at image, frame or shot level.

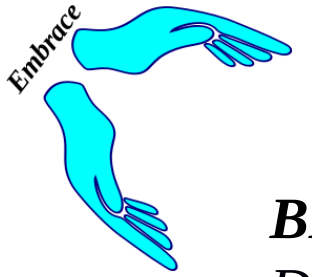

Glimpse of some projects involving ML

BELT: Big-data technologies and extreme scale analytics

Development of a novel multi-layered (Edge-Cloud-HPC) infrastructure which allows the prepossessing of large volume of data, enables extreme-scale analytics (deep learning and predictions), and automated decision-making.

ROBS: Robotic frameworks for biotarget screening

Robotization of smart frameworks for drug screening against biotargets with intelligent-decision making workflow to limit human interventions to the lowest possible.

…..To name a few

HPC Support For Data Science And Machine Learning

Containerized solutions for deep learning tools

Objectives

Testing Intel deep learning container (tensorflow, pytorch) on HPC

Debugging functionality & portability issues

Performance benchmarking

Benchmarking showcase DL models with publicly available datasets

Benchmarking use-case models on domain specific datasets (on demand)

Employing new technologies to address scale-up issues

Integrating other technologies (eg: horovod) with containers to achieve scalability with distributed training on large number of nodes

Thank you for your attention

www.cineca.it