CINECA SCAI SuperComputing Application & Innovation

EuroHPC declaration

Slides stolen from several people presentations

Sanzio Bassini – May 2018

Why invest in HPC?

HPC is at the core of major advances and innovations in the digital age

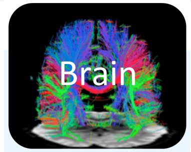

Strategic value for science

HPC enables breakthrough science

disease treatment; new therapies; brain; climate; chemistry; new materials; cosmology, astrophysics; high-energy physics; environment; transportation, earthquakes, etc.,

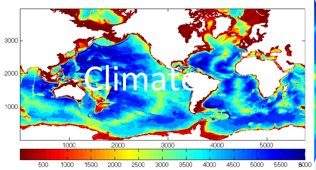

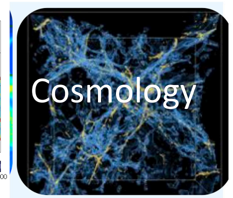

Strategic value for Industry

Market potential: new products, design and production cycles, decision processes, costs, resource efficiency, etc.

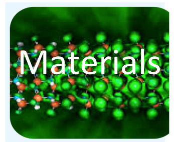

National security and defense

Complex encryption technologies, terrorism, forensics cyber attacks, nuclear simulations

High Energy physics

LHC Roadmap

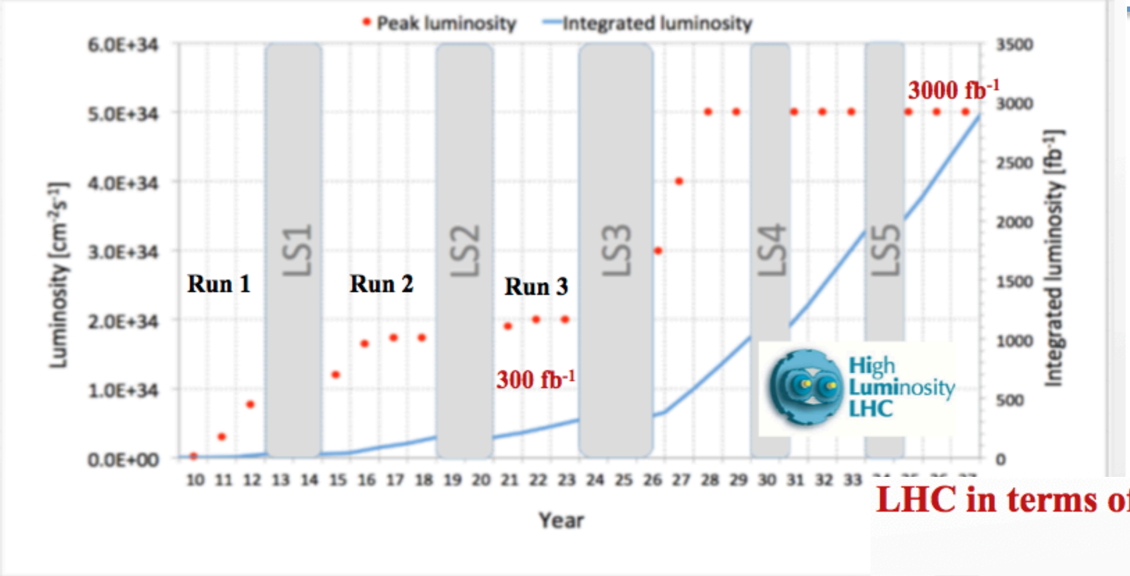

LHC in terms of data

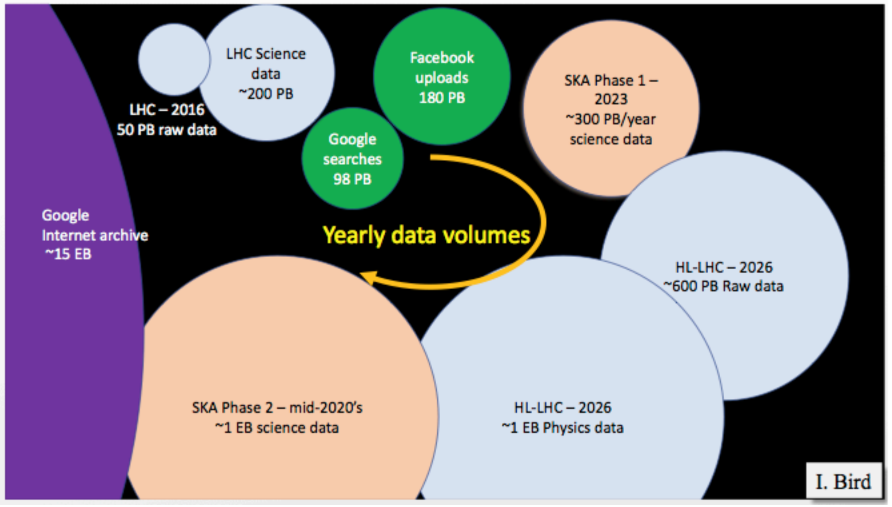

Resources Needs

- In principle computing resources increase ~ linearly with data.
- Physicists already adopted several improvements, at the moment it is needed:

1) factor ~ 4 of more CPU respect to the "20% per year growth"

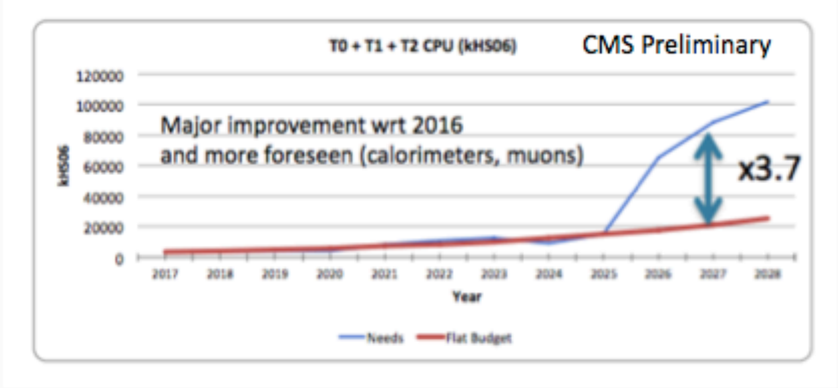

2) factor ~5 of more storage, cold storage (inactive data rarely used or accessed) concept already embedded

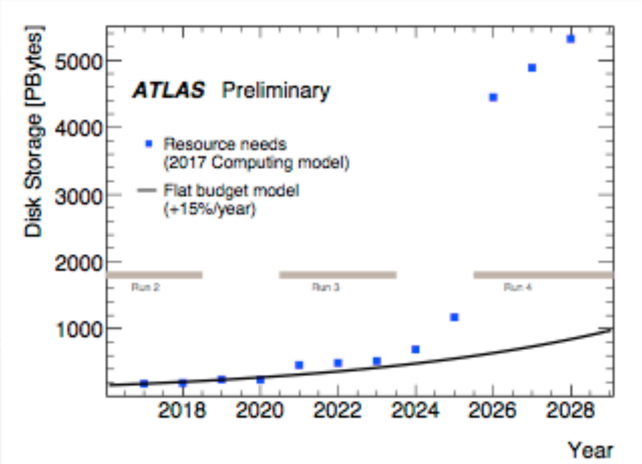

Nuclear Fusion Advanced Computing

Nuclear Fusion reactor with Magnetic Confinement Plasmas

Plasma Physics:

- Turbolence (Gyro-kinetic codes), Edge
- MHD (Equilibrium, Transport, Instabilities)
- Heating, Fast particles

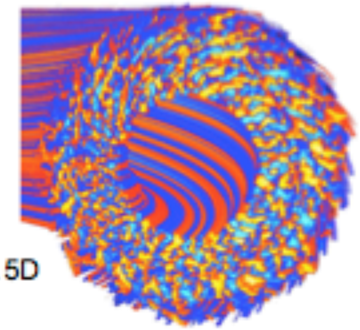

GYSELA 5D

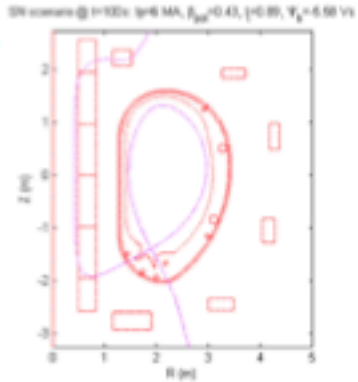

DTT equilibrium

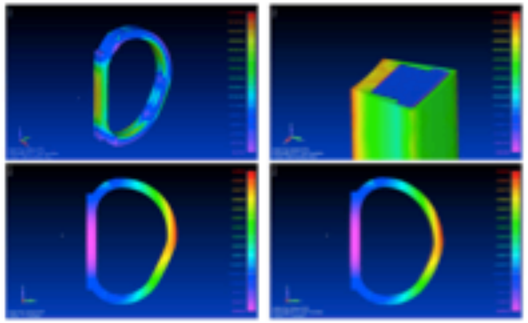

DTT: 3D stress analysis

Reactor technologies:

- Neutron Transport: MCNP: Montecarlo N-Particle Transport
- Materials: DFT: Density Functional Theory for Radiation Damage
- Structural Analysis : FEM (Ansys,Comsol Multiphisics)

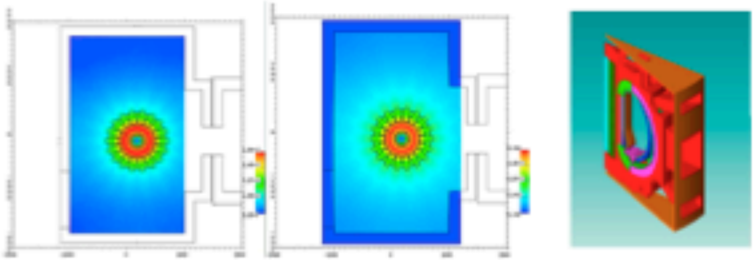

DTT: 3D MCNP analysis

ENEA F.Iannone et al.
EU Fusion Initiative

6

Carbon nanotubes as excitonic insulators

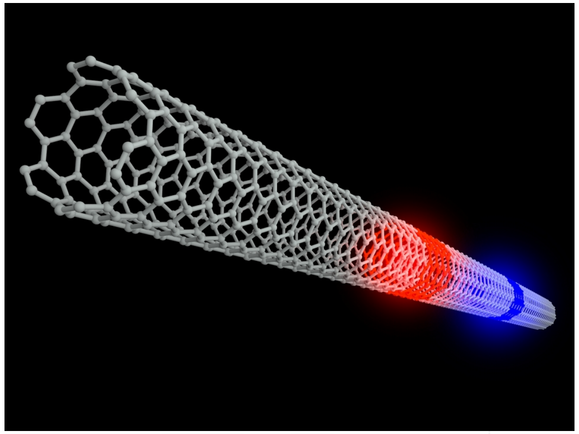

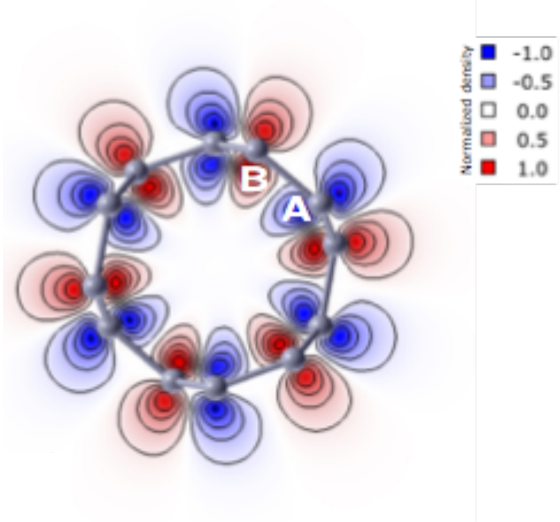

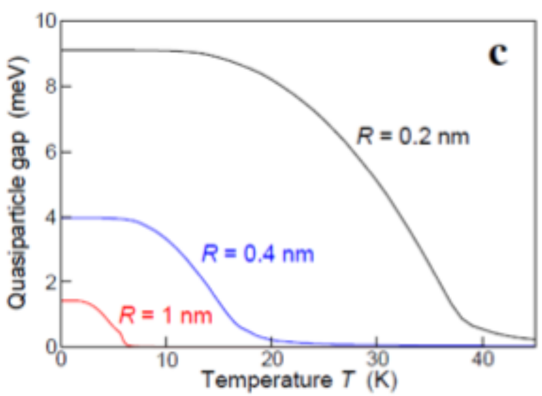

The **excitonic insulator** phase (EI), i.e. the instability of a zero gap semiconductor against the tendency of mutually attracting electrons and holes to form bound pairs, was speculated by W. Kohn in 1968 since then, the observation of the EI has remained elusive.

Here we have proved by means state of the art ab initio calculations
through **many-body perturbation theory** as well as **Quantum Monte Carlo** that the excitonic insulator is realised in zero gap **carbon nanotubes** (CNT). The excitonic order modulates the charge between the two carbon sublattices opening an experimentally observable gap.

By means of mean field models using parameters provided from the ab initio calculations we observed that below a critical temperature the exciton phase is present in all the armchair family of CNT with an electronic gap scaling approximatively as the inverse of the tube radius.

Carbon nanotubes as excitonic insulators; D. Varsano, S. Sorella, D. Sangalli, M. Barborini, S. Corni, E. Molinari and M. Rontani; Nature Communications 8, 1461 (2017)

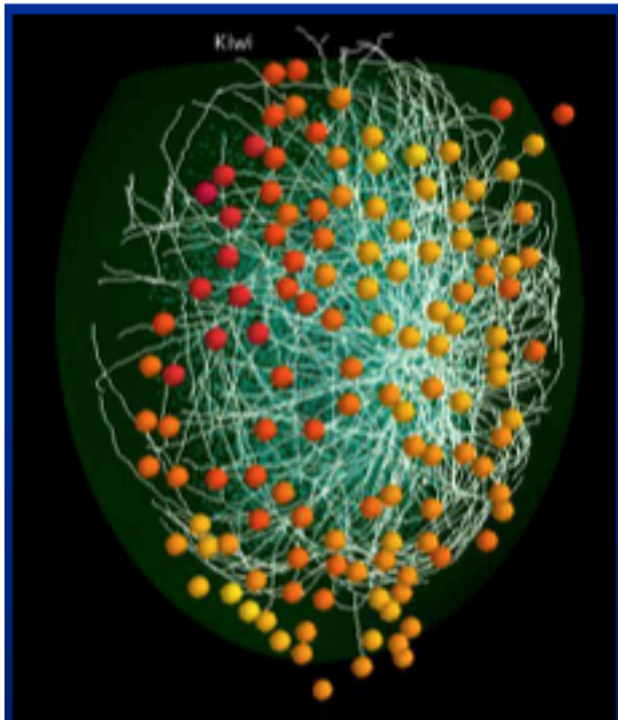

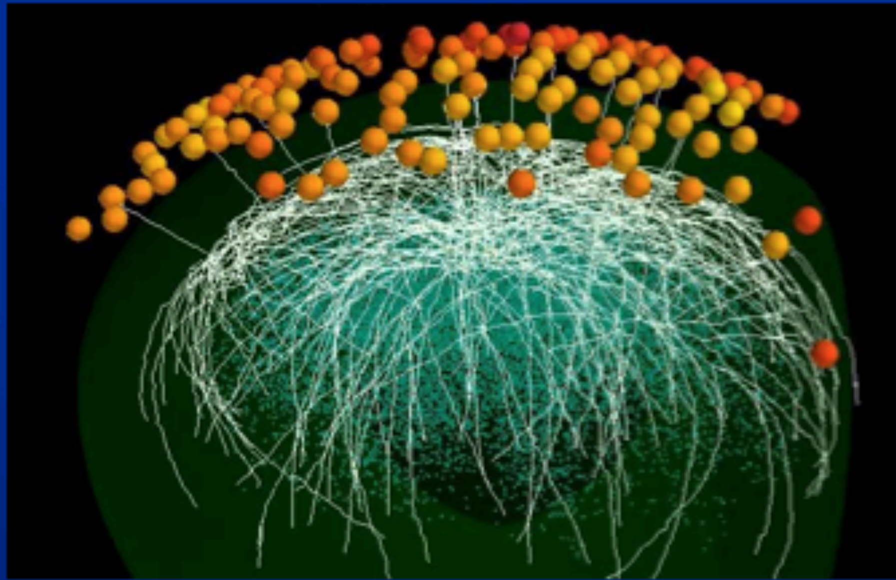

635 mitral cells
100K granule cells
$7 \cdot 10^5$ synapses
(1/20 of the real system area
32,000,000 nonlinear ODEs)

Table 2 | Model parameters and execution times for a typical simulation.

	Seg (min-max)	States (min-max) (v, channels, and syn. gates)		Syn (min-max)
MC (n = 635)	380,748 (169–1433)	5,259,735 (2536–20,020)		707,218 (300–2799)
GC (n = 68013)	4,344,724 (33–257)	26,892,317 (261–868)		707,216 (1–62)
Total	4,725,472	32,152,052		
	Computation time	Comm. time (spike exchange)	Comm. time (multisplit)	Total run time (2048 procs)
Average (sec)	27,149.35	68.53	555.94	32,552.86
Max (sec)	27,756.25	813.44	1453.96	

Currently installed on

CINECA Marconi

JSC JURECA, JUQUEEN

Typical 40 sec of sim. on 2048 processors, fully integrated NEURON+python implementation, $750 \cdot 10^6$ spikes: 9 hours, 10 Gb output, 99% eff.

Numerical Models to simulate climate change in different scenarious: Regionalization by dynamical or statistical downscaling

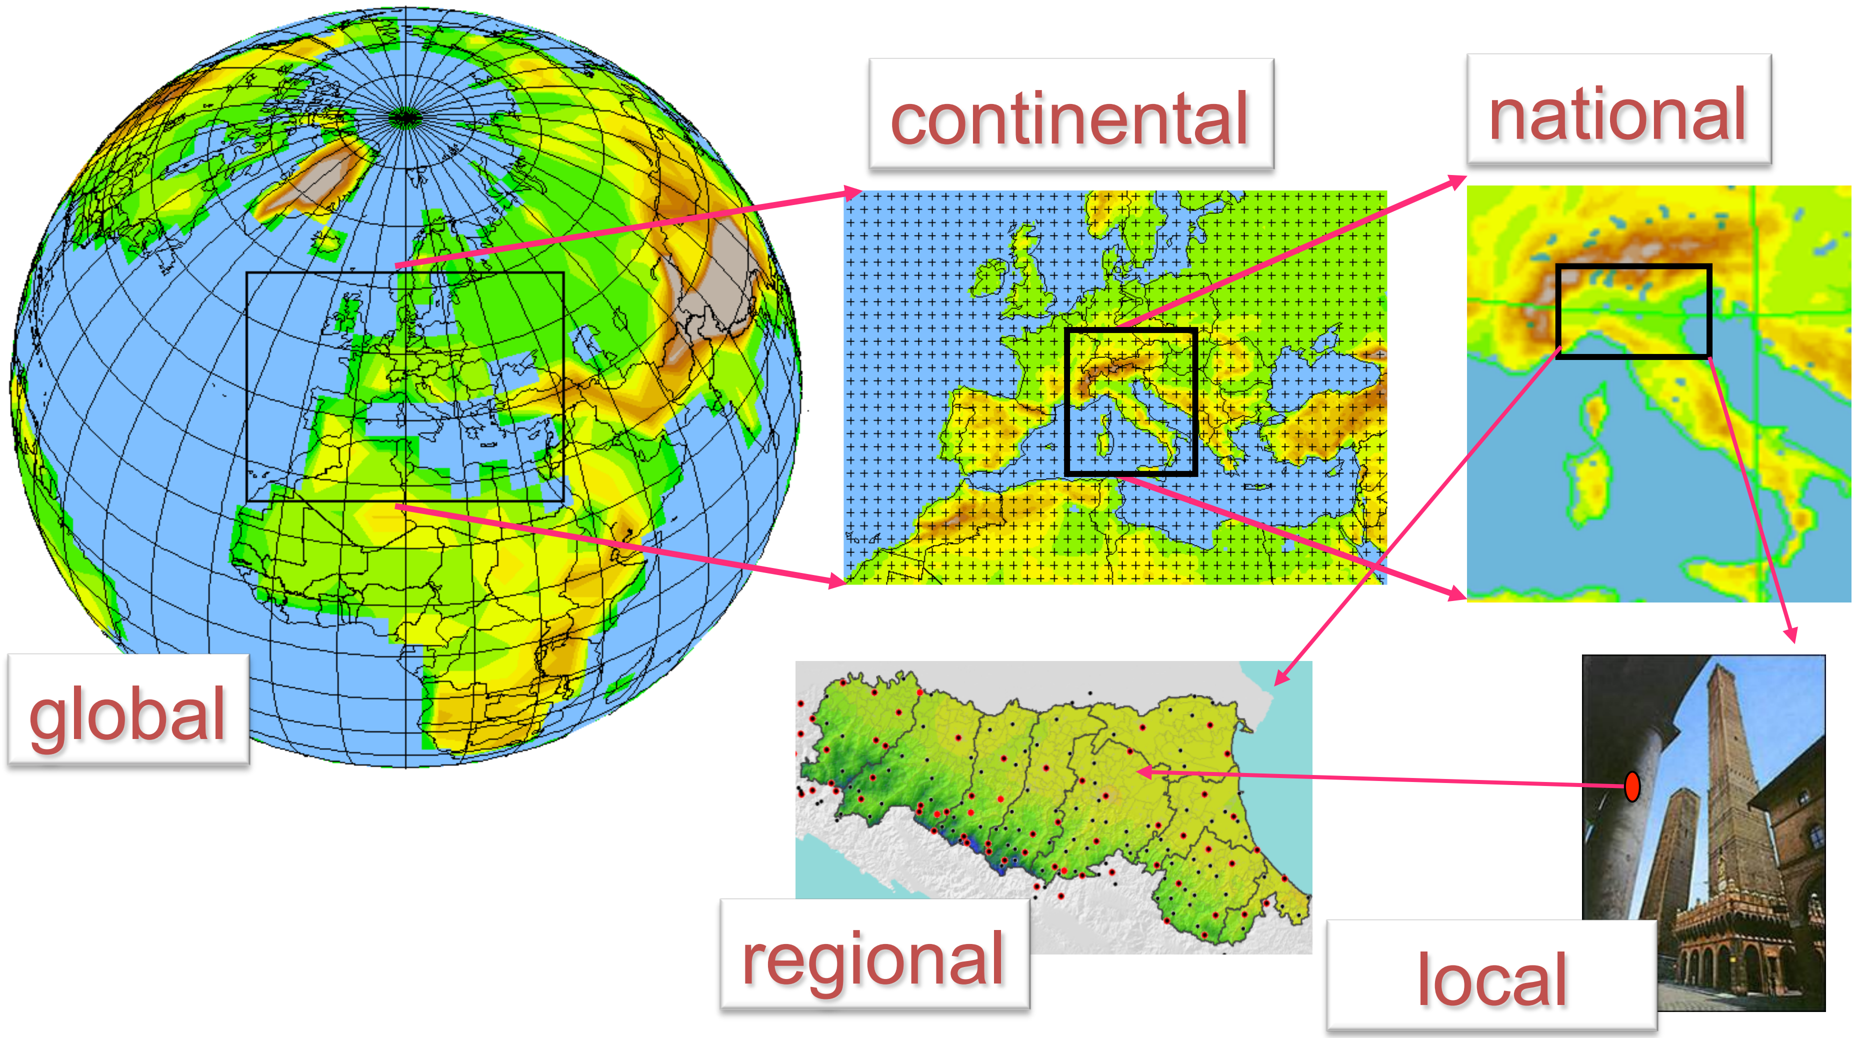

Model resolution – computing
Need of HPC

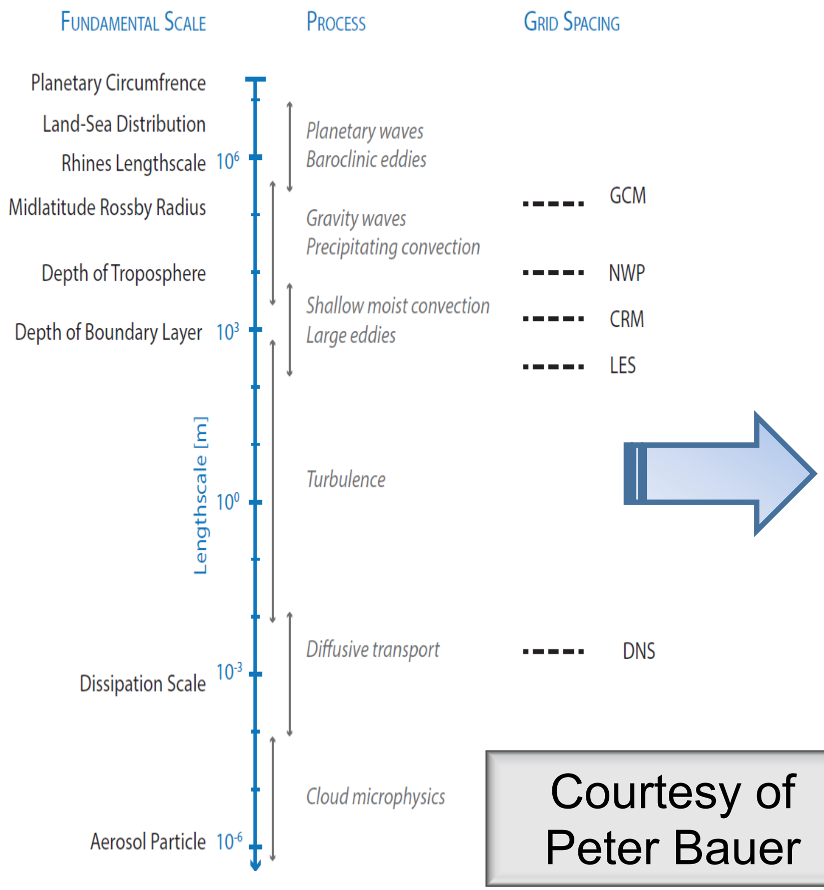

Courtesy of Peter Bauer

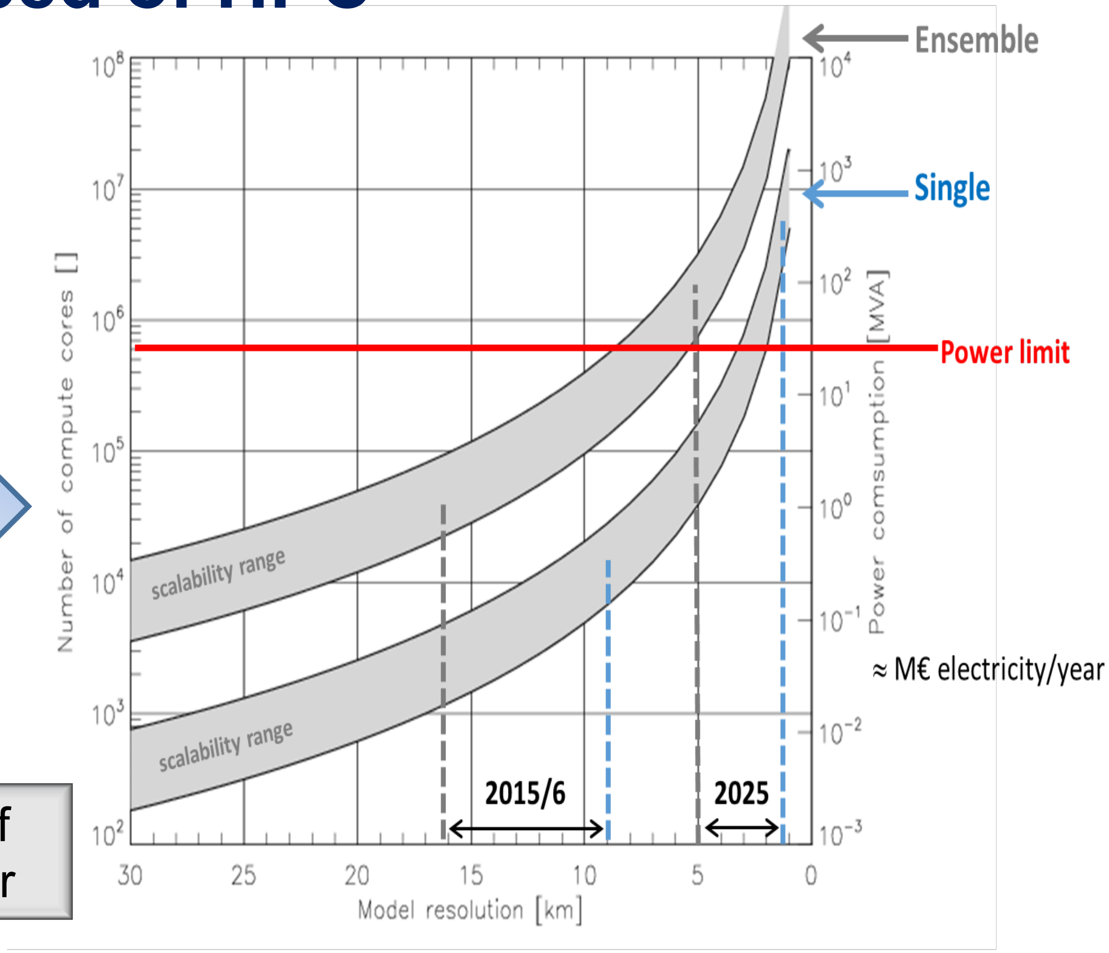

(Bauer et al. 2015)

→ Science community agrees that at very high resolution models will make qualitative jump in accuracy, but this comes at a very high computing & data cost, and that we estimate to be a factor of 1000 short with current system

Genome-wide molecular docking

Approaching disease-specific poly-pharmacology by connecting drugs with biological targets and diseases

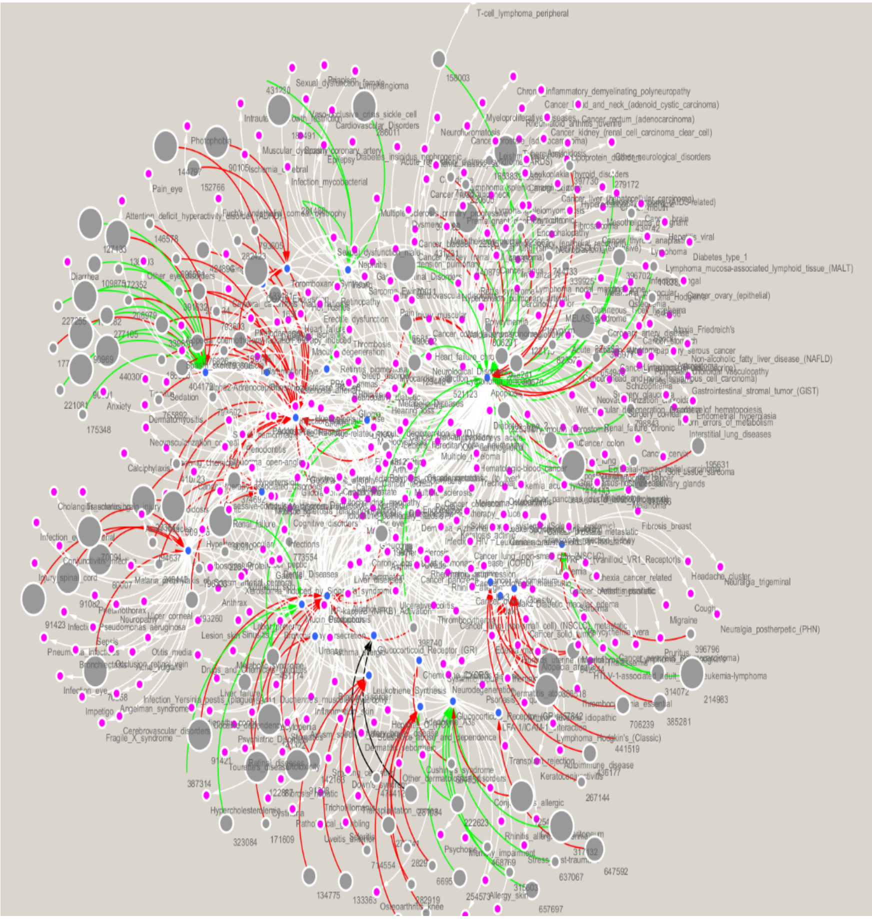

Genome-wide molecular docking simulations will predict

- Drug Efficacy
- Drug Safety
- Averse Effects
- Novel Use of Known Drugs

LiGen™ Push to the limit software optimization on standard hardware

- Develop energy and resource efficient algorithms
- Use self-functionalities to adapt and scale-out the application

As Is	To Be
Ligand: 100 Million	Ligand: 1 Billion
Targets: 1	Targets: 10
Docking time: 30 sec	Docking time: 3 sec
2048 CPU cores	1 Million CPU / GPU / MIC cores
UNICODE / Flat file	Binary / Database
Wall Time: 407 hours (17 Days)	Wall Time: 8 hours
Total Cost*: 83K	Total Cost*: 833K

Total Cost = 100 M Ligands X 1 Targets X 30/sec for each Dock X 0.1 Euro /hour

Development Roadmap of LiGen™

- Better Database Architectures
- Faster I/O
- Deeper Hardware / Software integration
- Ad Hoc Hardware
- New Evaluation Functions Based on Artificial Intelligence

DATA CENTER
21 PetaByte

Data
25.000.000 KM

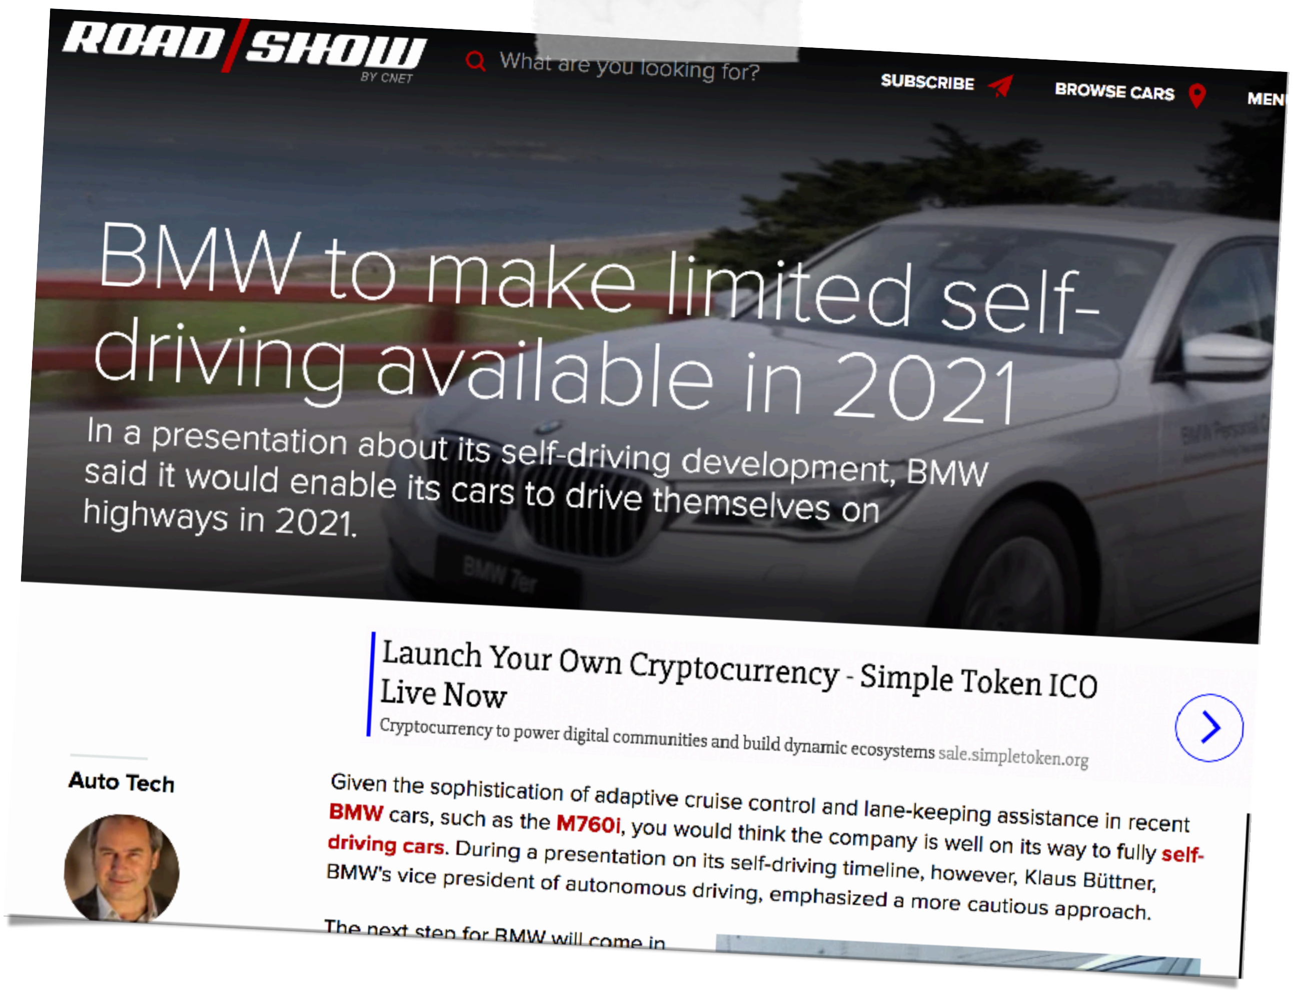

Weaknesses in HPC implementation

EU has no top ranked supercomputers and depends on non-EU technology

Funding Gap wrt USA, JP, CN

Weak EU supply chain

Weak integration of EU technology in HPC machines

HPC strategy implementation by EC is insufficient

Insufficient coordination of national investments

Demand is not met

Weaknesses in HPC implementation
Funding Gap

- The EU will invest ~€1 billion in HPC activities from 2014 to 2020
- US R&D investments: $1 to $2 billion per year
- China: over $1 billion per year
- Japan: $1.38 billion for 1 exascale system

No MS has the means to develop the necessary full HPC ecosystem on its own in competitive timeframes

Funding needs

First estimation: €4.7 – 5.2 billion

COM(2016) 178 of 19/4/2016

- 2 pre-exascale and 2 exascale machines, data and interconnection
- Technology development (processor, system, SW)
- Applications + wider access to HPC facilities for SMEs

- **Until end 2020**: €1B EC (H2020, CEF) + €1 B public & private sector
- **2021 – 2026+**: TBD

Overall envelope: ~ €8 – €9 billion

EuroHPC Declaration – Participating States

Declaration signed in Rome 23/03/2017 by:

For Italy signed the Ministries of Education University and Research and of Economic Development

- France
- Germany
- Italy
- Luxembourg
- Netherlands
- Portugal
- Spain

8 more countries signed the Declaration:

- Belgium
- Slovenia
- Bulgaria
- Switzerland
- Greece
- Croatia
- Czech Republic
- Cyprus

EuroHPC Mission and objectives

- **to provide scientists, industry and the public sector from the Union with latest HPC and Data Infrastructure** and support the development of its technologies and its applications across a wide range of fields.
- to provide a framework for **acquisition of an integrated world-class pre-exascale supercomputing** and data infrastructure in the Union;
- to provide Union level **coordination** and adequate **financial resources** to support the development and acquisition of such infrastructure, which will be accessible to users from the public and private sector primarily for research and innovation purposes;

Present EU Financial Framework **Next EU Financial Framework**

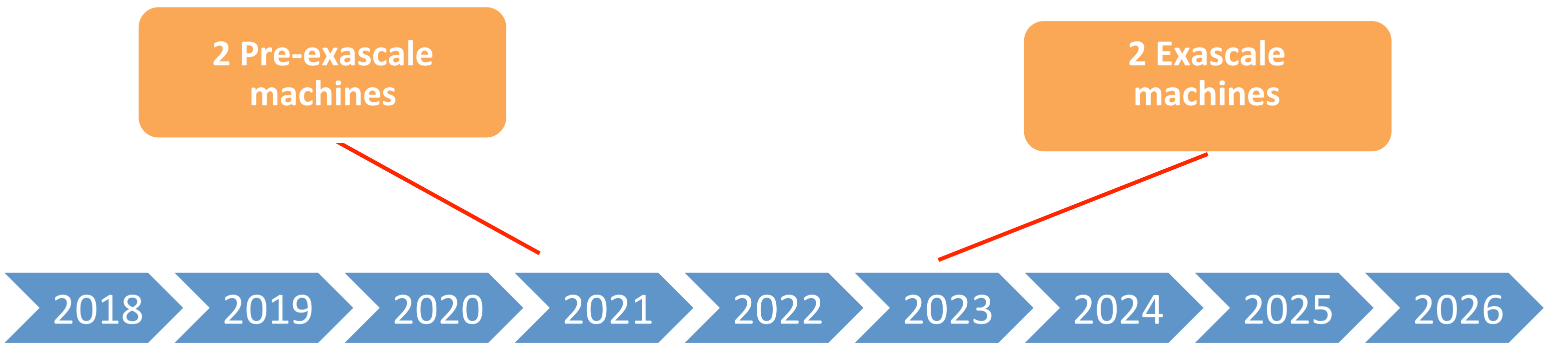

EuroHPC Mission and objectives

Joint Undertakings first period of agreement

2018 | 2019 | 2020 | 2021 | 2022 | 2023 | 2024 | 2025 | 2026

HPC ecosystem

Infrastructure Acquisition
Operating machines

- **Pillar 1:**
- Acquisition of infrastructure:
  - ➔ 2 pre-exascale machines
  - ➔ >= 2 Peta-scale machines
- Installation, deployment and operation via hosting entities + access to users

Research & Innovation
Applications & Skills

- **Pillar 2:**
- European exascale technologies and systems (incl. low power processor)
- Excellence in HPC applications; CoE; competence centres for industry (incl.SME); Training and Outreach

Activites and Funding

		EC	PS	Total (M€)	Private Mem.
Infrastructure Acquisition Operating machines	Pillar 1	270	290	560	
Research & Innovation Applications & Skills	Pillar 2	206	186	392	422
JU Admin/Running costs		10	10	20	

National priorities for the JU workprograms and first estimates of national contribution

HPC infrastructure design point

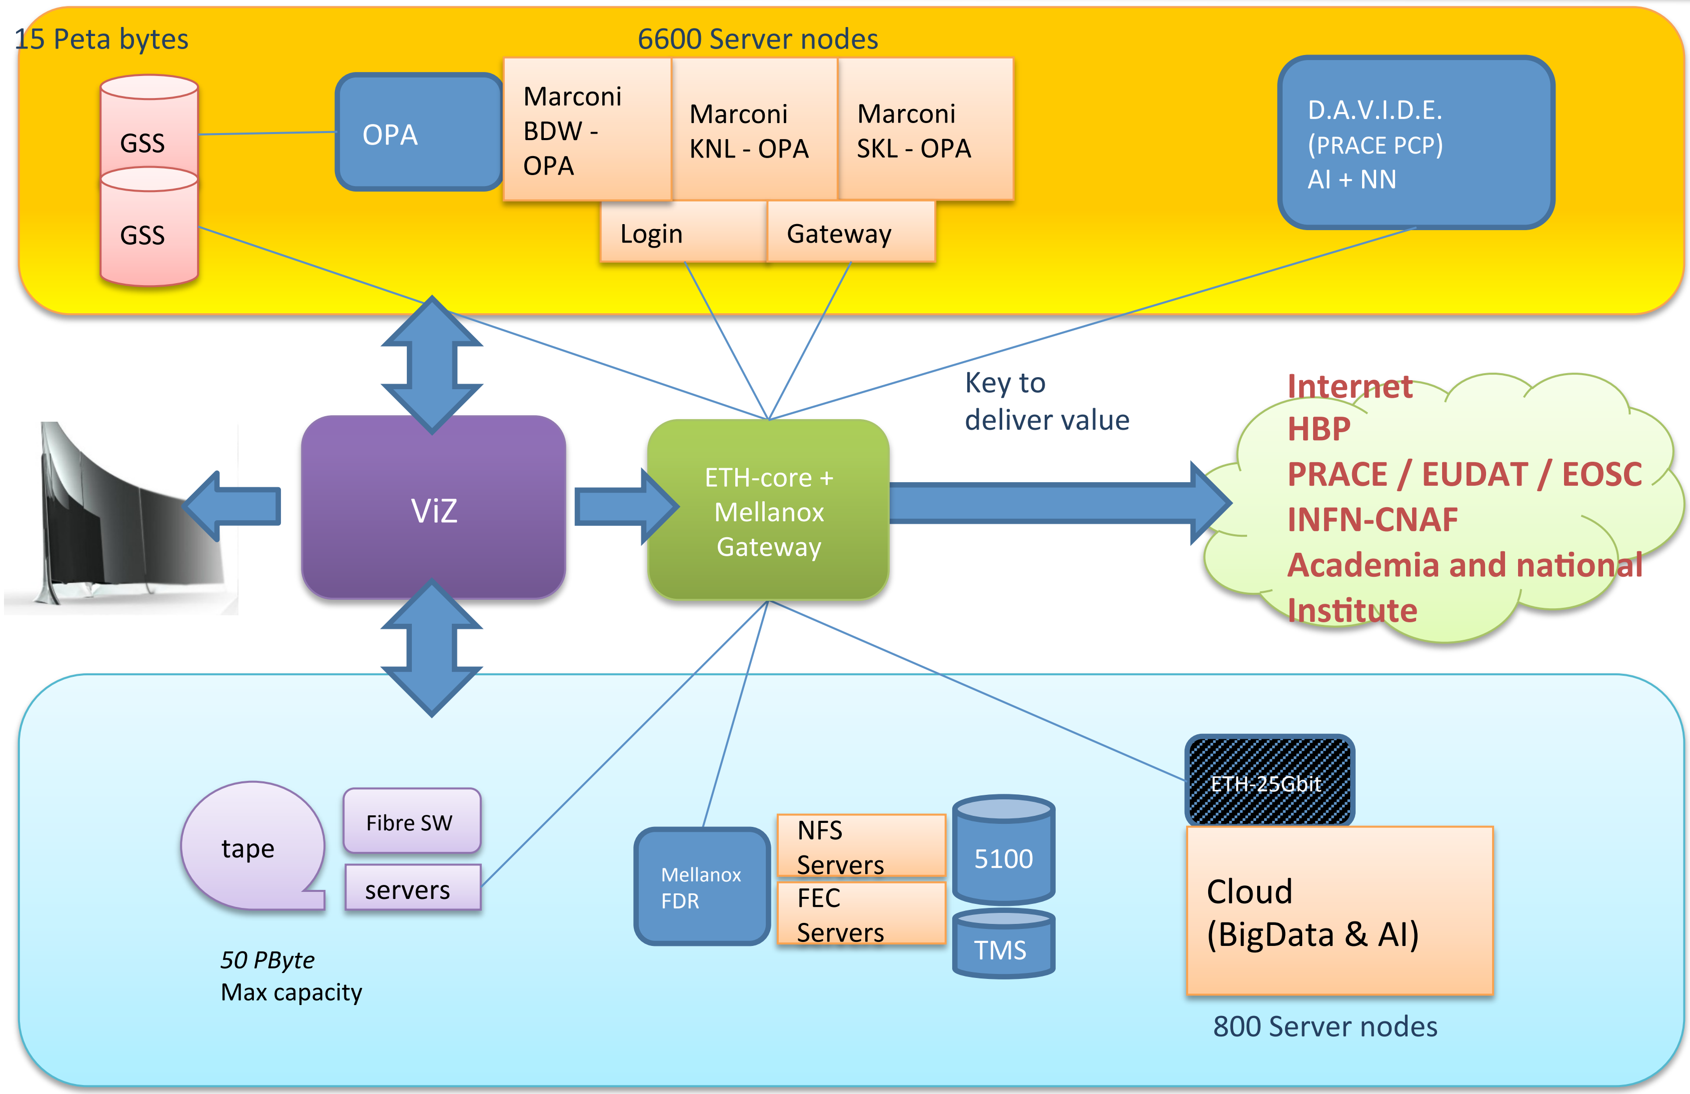

Top500 (November 2017)

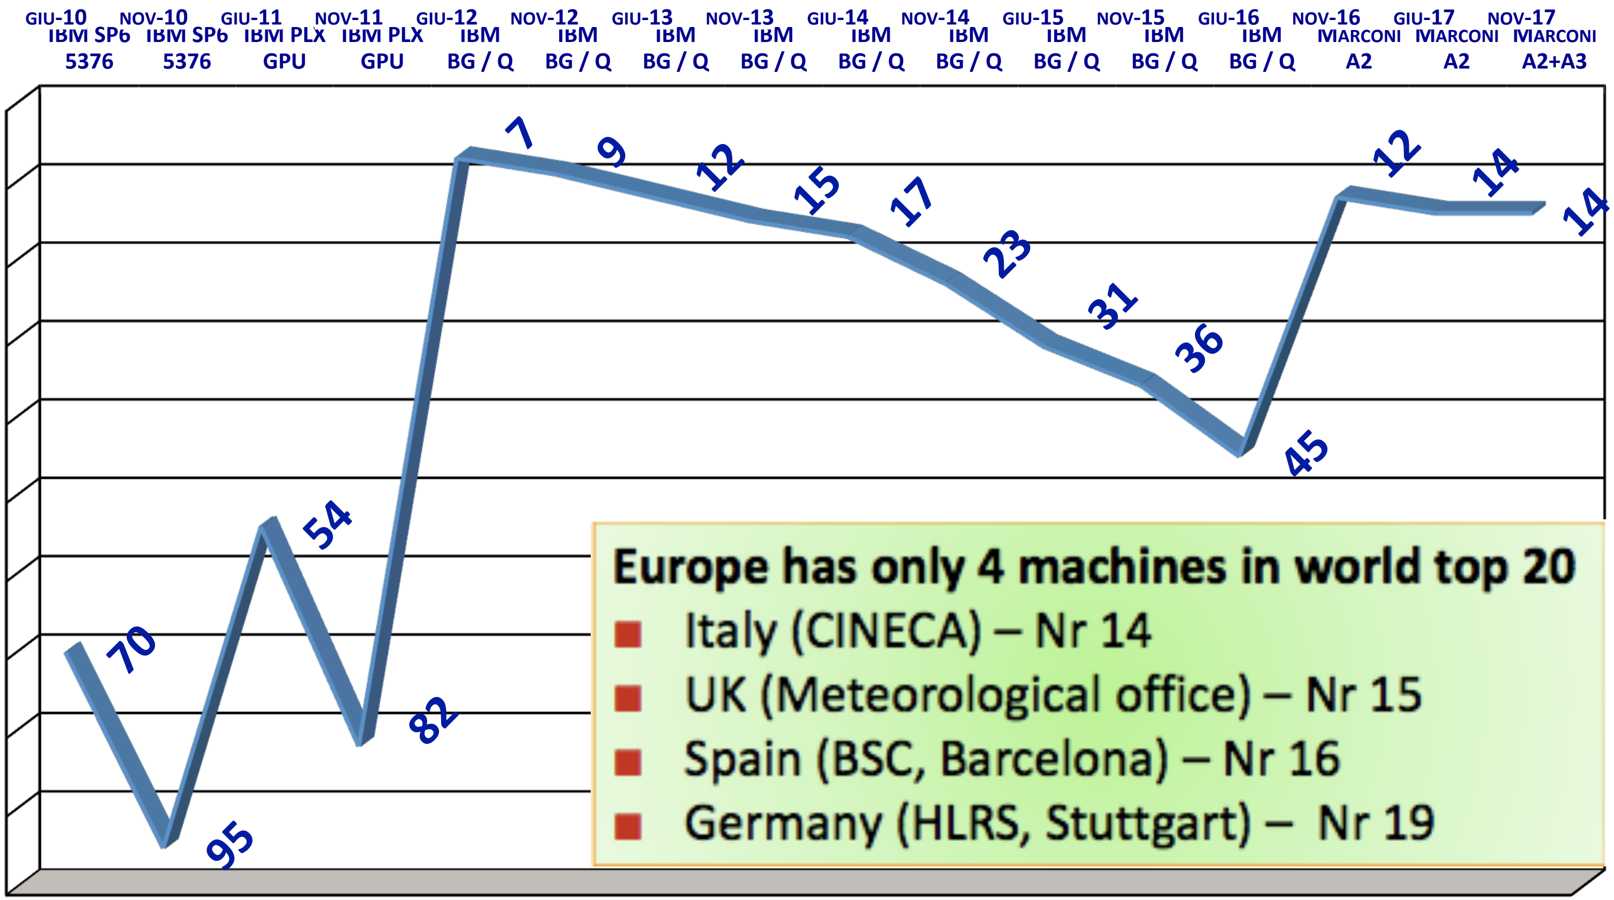

Europe has only 4 machines in world top 20

- Italy (CINECA) – Nr 14
- UK (Meteorological office) – Nr 15
- Spain (BSC, Barcelona) – Nr 16
- Germany (HLRS, Stuttgart) – Nr 19

HPC and Verticals

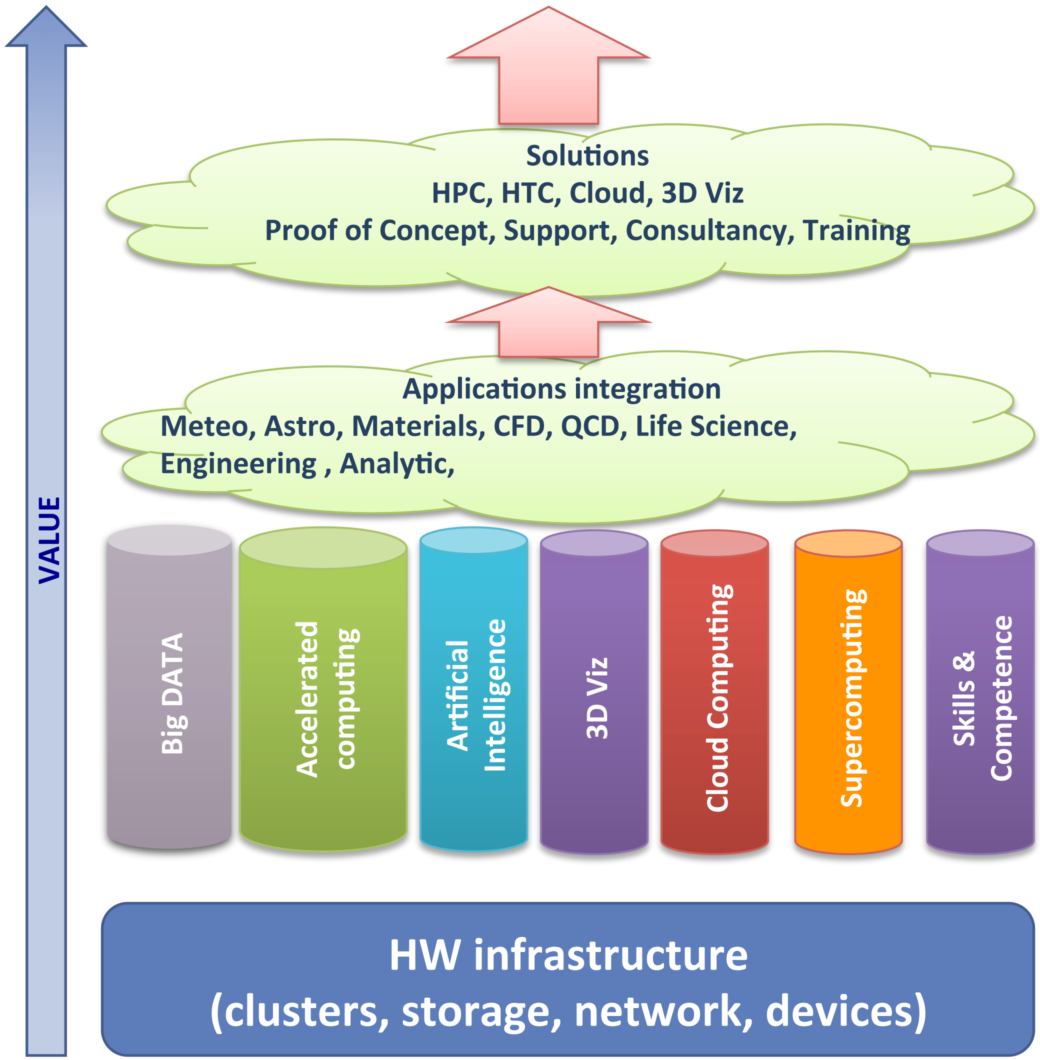

Roadmap toward exascale

CINECA

Peak Performance

10^18 Flops — Moore law

FPU Performance

10^9 Flops — Dennard law

→ Number of FPUs

10^9

- **10^5 FPUs in 10^4 servers**
- 10^4 FPUs in 10^5 servers

Working hypothesis

Co-design → Heterogeneus — Latency vs Throughput

Power projection

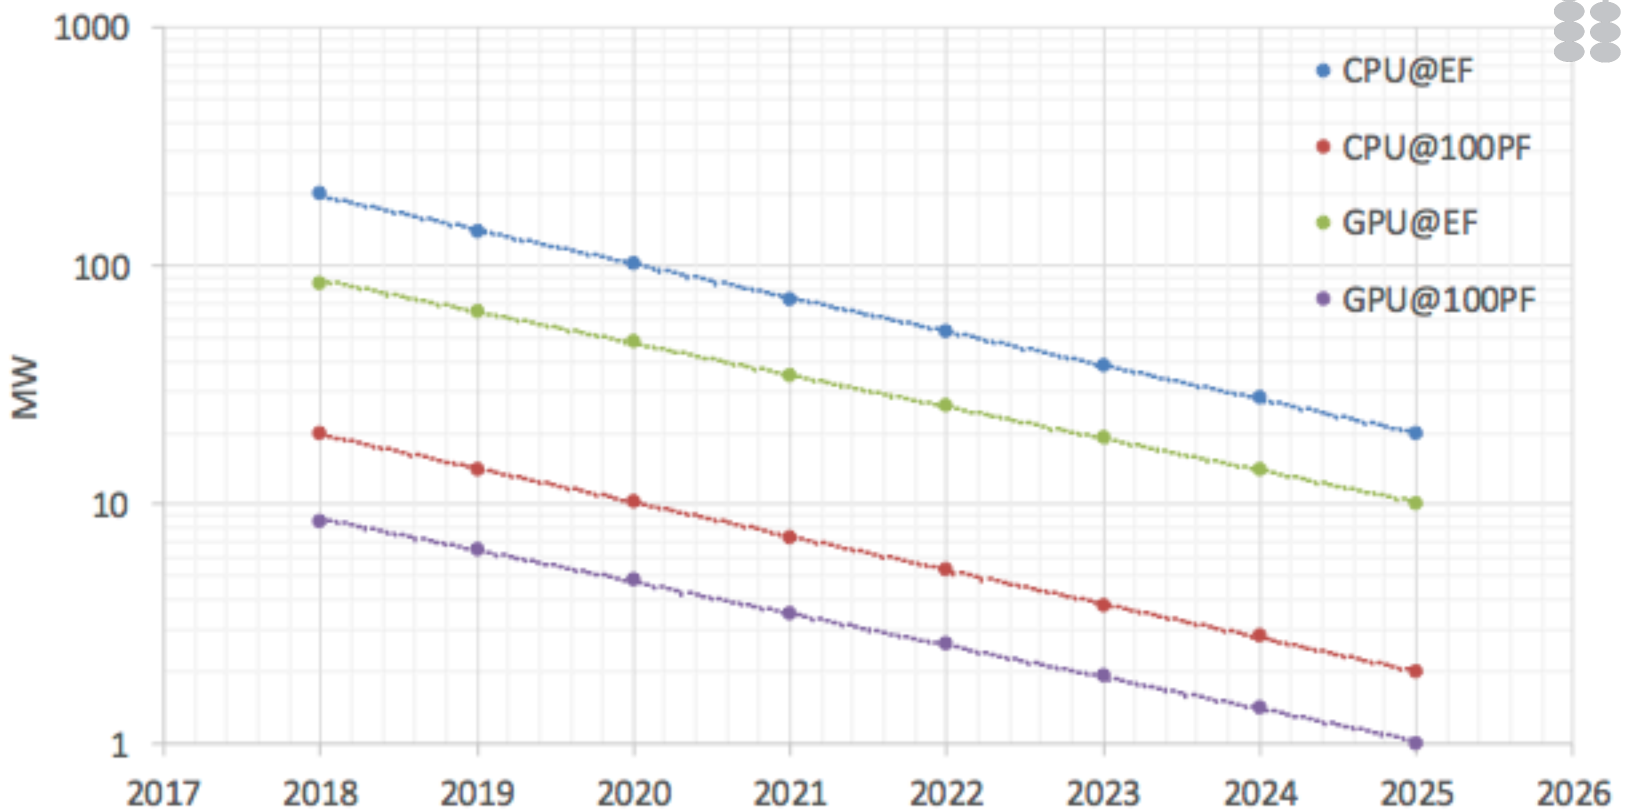

Peek Perf (DP) @ 10MW

	2018	2019	2020	2021	2022	2023	2024	2025	2026
CPU	50PF	70PF	100PF	140PF	200PF	250PF	330PF	500PF	750PF
GPU	125PF	166PF	200PF	300PF	385PF	525PF	715PF	1EF	1.3EF

Cineca Roadmap

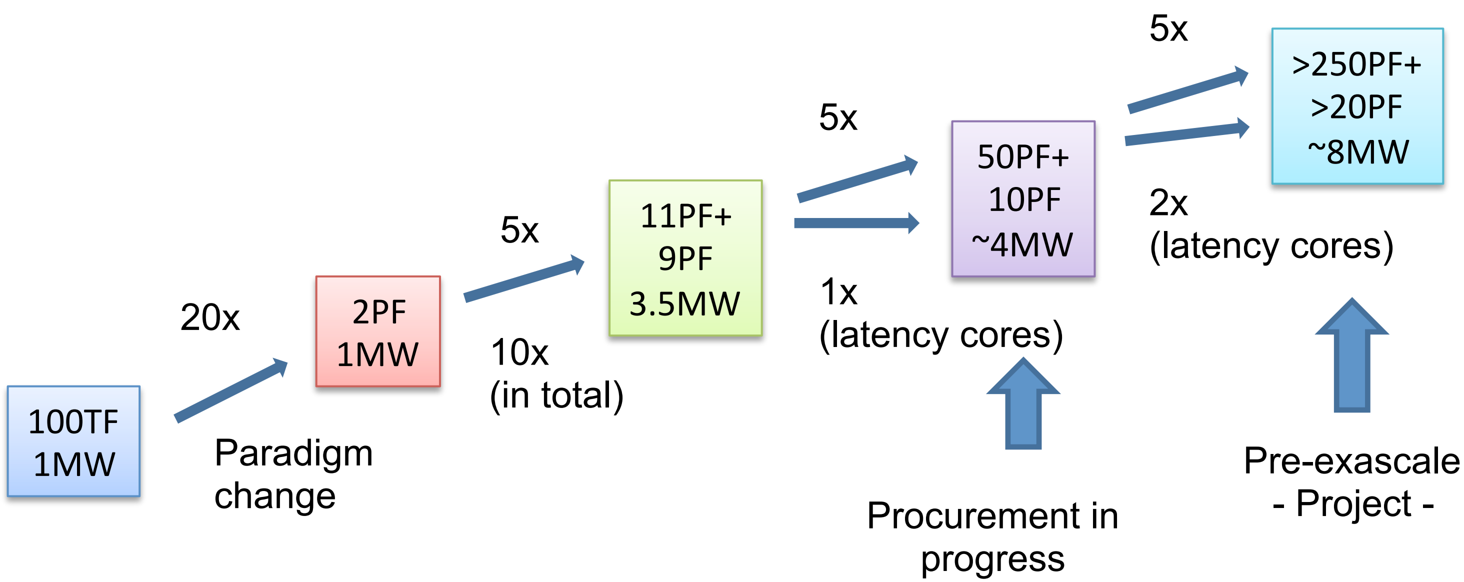

2009	2012	2016	2019	2021/2022
IBM SP6 Power6	Fermi IBM BGQ PowerA2	Marconi Lenovo Xeon+KNL	Marconi PPI4HPC ICEI - PPIHBP	EuroHPC

Towards Italian HE candidature

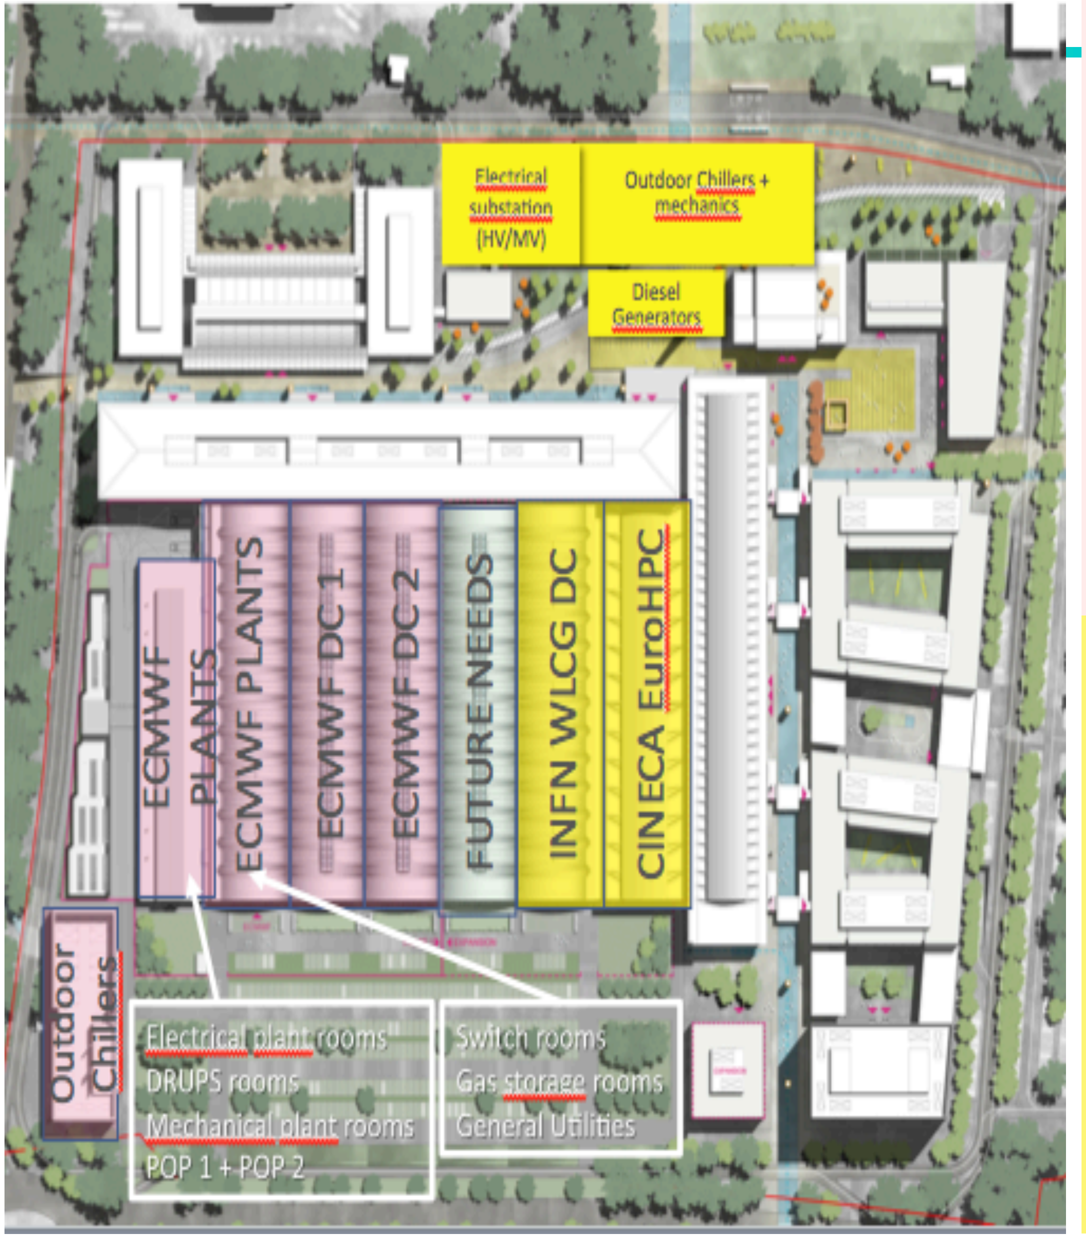

ECMWF DC main characteristics

- 2 power line up to 10 MW (one bck up of the other)
- Expansion to 20 MW
- Photovoltaic cells on the roofs (500 MWh/ year)
- Redundancy N+1 (mechanics and electrical)
- 5 x 2 MW DRUPS
- Cooling
  - 4 dry coolers (1850 kW each)
  - 4 groundwater welles
  - 5 refrigerator units (1400 kW each)
- Peak PUE 1.35 / Maximum annualized PUE 1.18

INFN – CINECA DC main characteristics

- up to 20 MW (one bck up of the other)
- Possible use of Combined Heat and Power Fuel Cells Technology
- Redundancy strategy
- Cooling
  - dry coolers
  - groundwater welles
  - refrigerator units
- PUE < 1.2 – 1.3 / Max Annualized < 1.2 / 1.17

The CINECA-INFN datacentre timeline

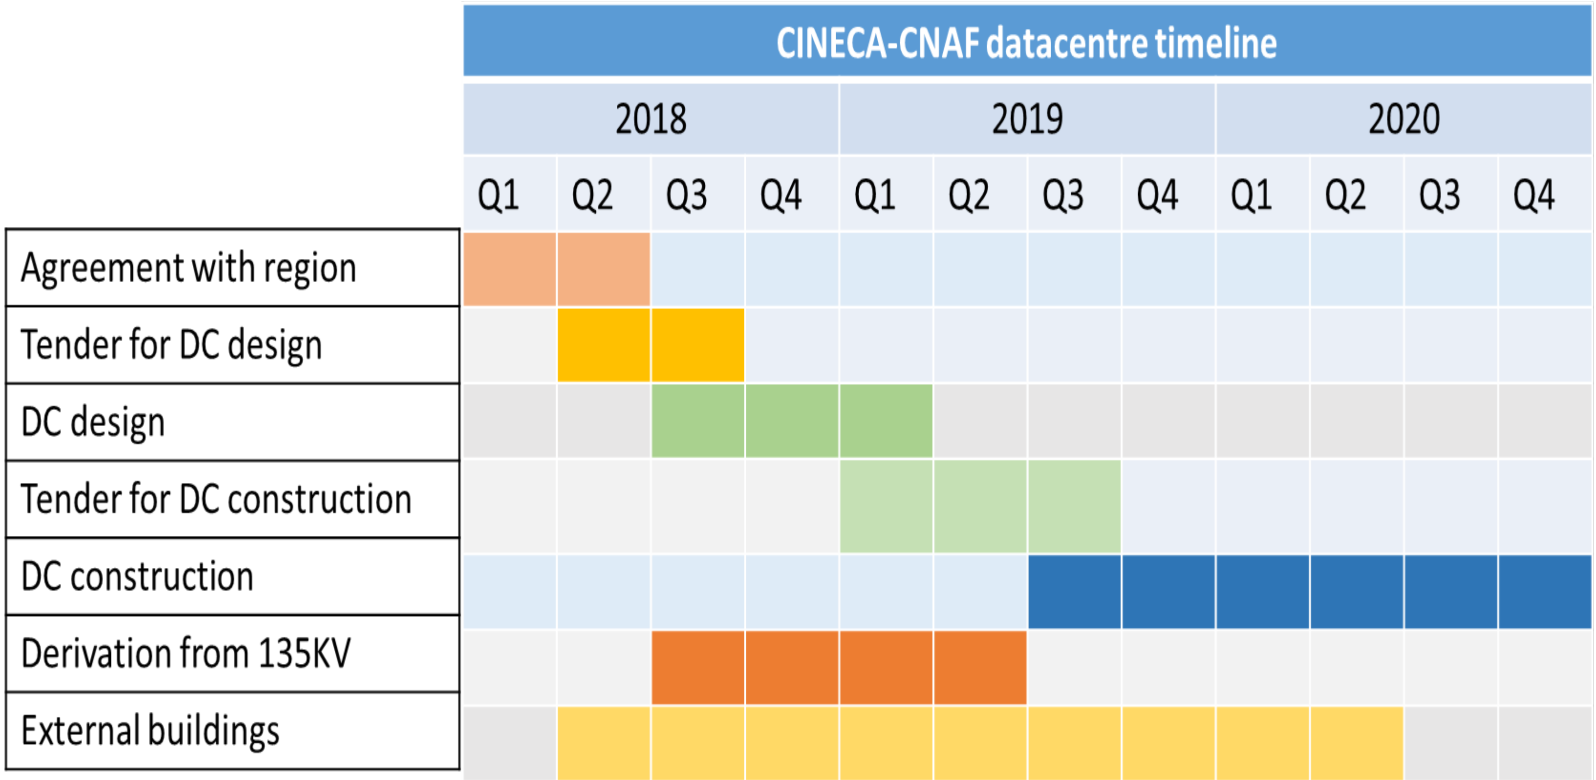

Italian Estimates contribution for JU Pillar 1

- Solid national CINECA-INFN strategy and road-map for next 10 years
- Great DC location, with cross-fertilization opportunities
  - CINECA EuroHPC, INFN WLCG, ENEA DTT & Eurofusion, ECMWF, Italia Meteo, etc…
  - R&D labs of private companies
- 50% of TCO in kind contribution commitment for a pre-exascale system
- Commitment to represent an HE including others PS

EU funding for R&I activities in 2019-2020

Priorities for JU	2018	2019	2020
µ-processor & co-design of machines	*80 M€ (H2020)*	*40 M€ (EuroHPC)*	*Up to 160 M€ Co-design -*
Support to SMEs		*8 M€ (EuroHPC)*	
Technology R&D		*68 M€ (EuroHPC)*	
Applications	*90 M€ (H2020)*		*? M€*

FPA LPP: ~100 M€
Codesign 2 machines: 60 M€

Focusing on exascale priorities

Italian National Priority JU
Pillar 2

Priorities for JU	Calls programmed 2019	Calls programmed 2020	Italian National priorities
µ-processor & co-design of machines	Framework Partnership Agreement in European low-power microprocessor technologies (Phase 2) **40 M€**	FETHPC-03-2020 : Co-design of extreme scale HPC systems and applications **160 M€**	1) Co-design for **computing accelerators** with standard interfaces: PCIe, OPENCAPI, ZENG 2) Co-design of micro architecture of a computing efficient, **exa-scalable direct network** 3) HW/SW support at micro architecture level to optimize the use of **Machine Learning** methods in scientific computing (either for HPC and HTC applications) 4) HW/SW co-design with **full applications**
Support to SMEs	INFRAINNOV-01-2019: Stimulate the innovation potential of SMEs **8 M€**		1) Public Private **HPC Cloud Service** 2) HPC **power efficiency** with innovative solutions 3) Academia and Research InstituteTechnology transfer initiatives to improve **3rd mission activities** 4) **Industria 4.0** funded by Ministry of Economic Development
Technology R&D	FETHPC-02-2019 : HPC - Extreme scale computing technologies, methods and algorithms for key applications and support to the HPC ecosystem **68 M€**		1) Toward ecosystems for energy efficient - **Power Management and Performance Monitoring** 2) **Advanced interconnect topologies** optimized for low latency and high throughput 3) **Advanced programming model** and software tools for parallel ExaScalable scientific application 5) Development of **exascale oriented algorithms** (eg mixed precision algorithms)
Applications			1) Increase funding on applications with respect to the **co-design processes** 2) development of **novel scientific and algorithm approaches** able to take full advantage olarge scale HPC machines 2) **LQCD** at extreme scale 3) **Bio-computing** and Complex system simulation 4) **Neural network simulation** and Neural network training for embedded solution 5) **Machine Learning** methods applied to off-line/on-line particle tracking at HL-LHC experiments 6) **Quantum computing** simulation and modeling