
BSC
Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

MareNostrum 5

Dr. Sergi Girona
Operations Director

October 2022

MareNostrum5 consortium @ Istanbul and Ankara

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC-CNS objectives

Supercomputing services
to Spanish and EU researchers

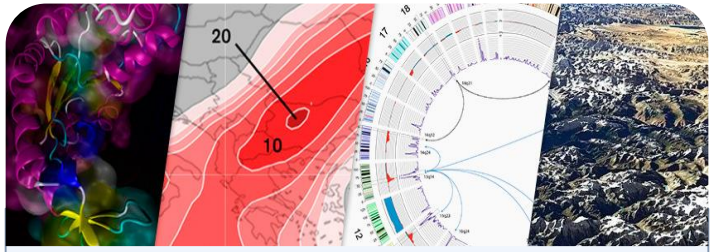

R&D in Computer, Life, Earth and
Engineering Sciences

PhD programme, technology
transfer, public engagement

BSC-CNS is a consortium that includes

Spanish Government	**60%**	Gobierno de España · Ministerio de Ciencia, Innovación y Universidades
Catalan Government	**30%**	Generalitat de Catalunya · Departament d'Empresa i Coneixement
Univ. Politècnica de Catalunya (UPC)	**10%**	Universitat Politècnica de Catalunya · BarcelonaTech

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

People

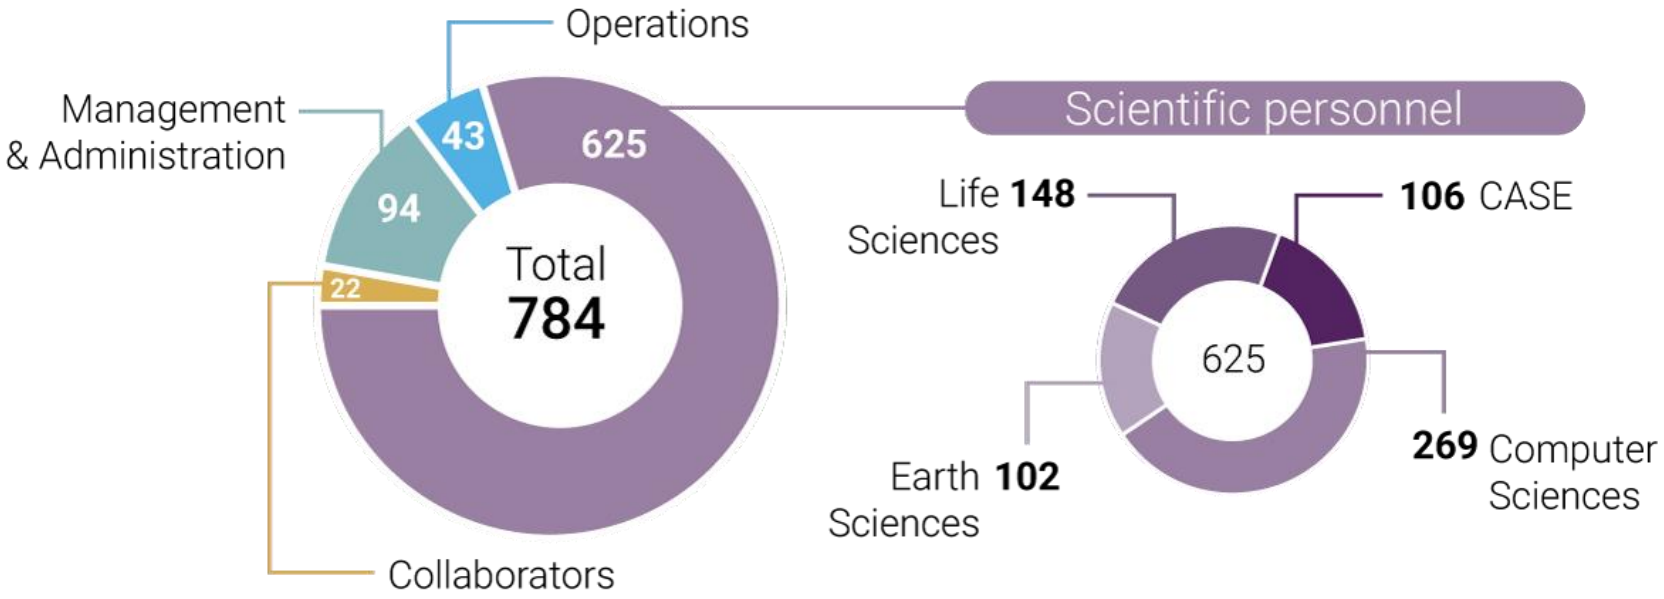

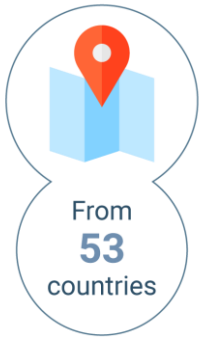

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC Resources

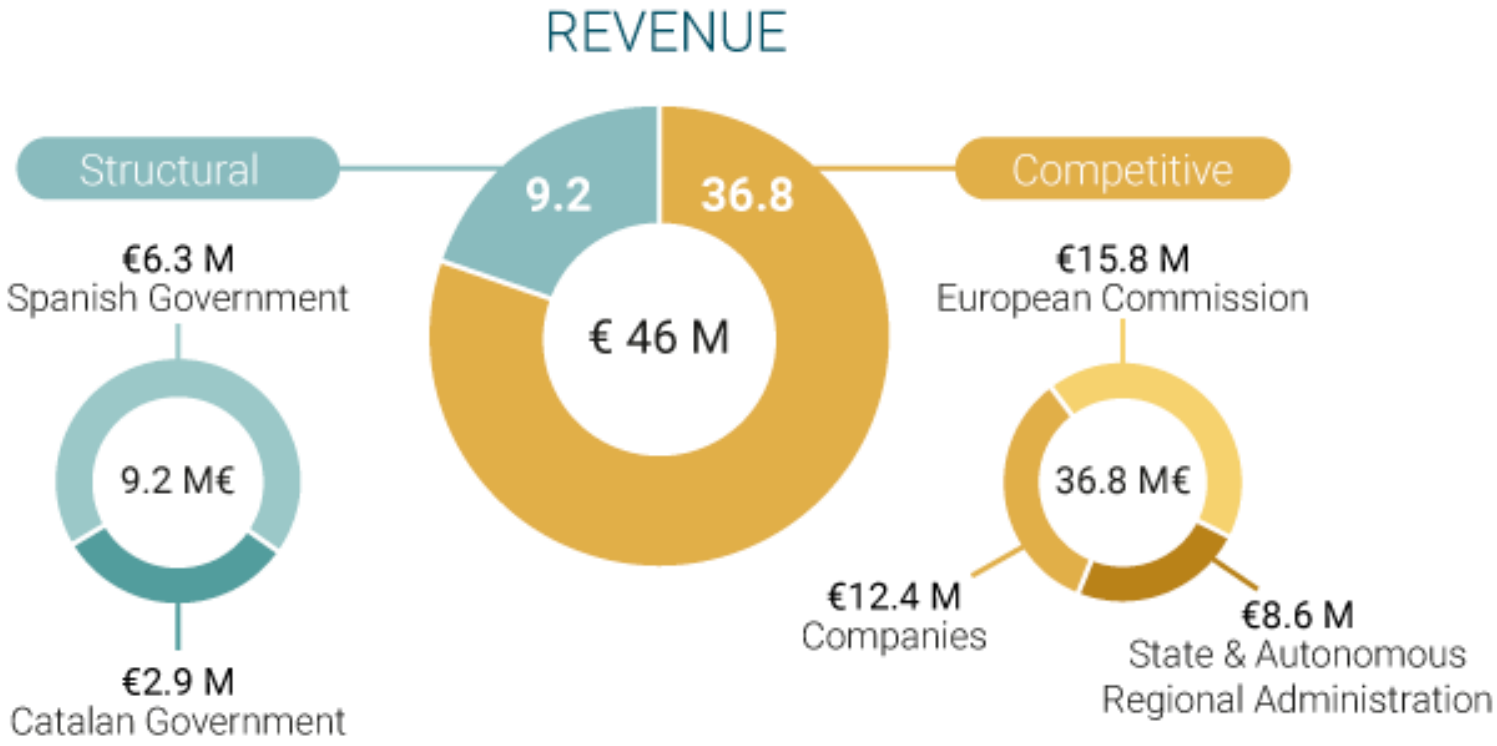

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

UPC contributes, in terms of assigned staff and room space, the equivalent of 10% of regular income

2021 executed budget

Mission of BSC Scientific Departments

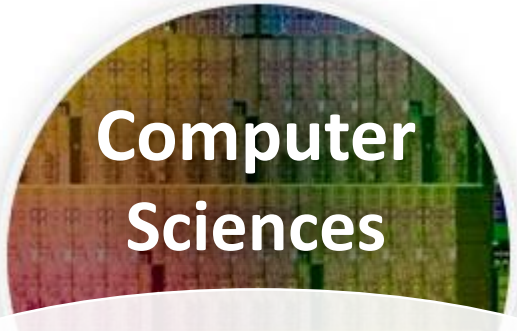

To influence the way machines are built, programmed and used: programming models, performance tools, Big Data, Artificial Intelligence , computer architecture, energy efficiency

To develop and implement global and regional state-of-the-art models for short-term air quality forecast and long-term climate applications

To understand living organisms by means of theoretical and computational methods (molecular modeling, genomics, proteomics)

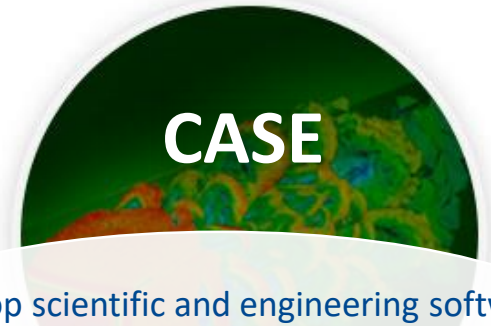

To develop scientific and engineering software to efficiently exploit super-computing capabilities (biomedical, geophysics, atmospheric, energy, social and economic simulations)

Open Project Portfolio July 2022

(includes HE projects in negotiation and national projects beginning soon)

284 running projects, total budget of 180M€*

139 H2020/HE projects running or in negotiation (18 as coordinator)	100M€
54 projects with Spanish and Catalan public funding	41M€
35 running contracts, mostly with companies	24M€
26 projects funded through other EU programmes (incl. Copernicus & CEF)	9M€
25 projects funded through other sources	4M€
5 strategic self-funded projects	2M€

Plus:

56 personnel grants (national funding)	5M€

Does not include ICREA, UPC, MSCA

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

* Budget for full term of project – normally multi-annual

Centers of excellence in HPC applications

BSC leads

- **ChEESE** — Centre of Excellence in Solid Earth
- **CoE Combustion** — Centre of Excellence in Combustion
- **PerMedCoE** — Centre of Excellence in Personalised Medicine

BSC participates

- **AISEE** — AI and Simulation-based Engineering at Exascale
- **bioexcel** — Centre of Excellence for Biomolecular Research
- **CompBioMed** — Centre of Excellence on Computational Biomedicine
- **esiwace** — Excellence in Simulation of Weather and Climate in Europe
- **EXCELLERAT** — European Centre of Excellence for Engineering Applications
- **MAX** — Materials design at the eXascale
- **NOMAD** — Novel Materials Discovery
- **eCAM** — A path to extreme-scale computing for industry and academia
- **HiDALGO** — HPC and Big Data Technologies for Global Systems
- **TREX** — Targeting Real chemical accuracy at the EXascale

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

November 2020

Scientific publications 2010-2019

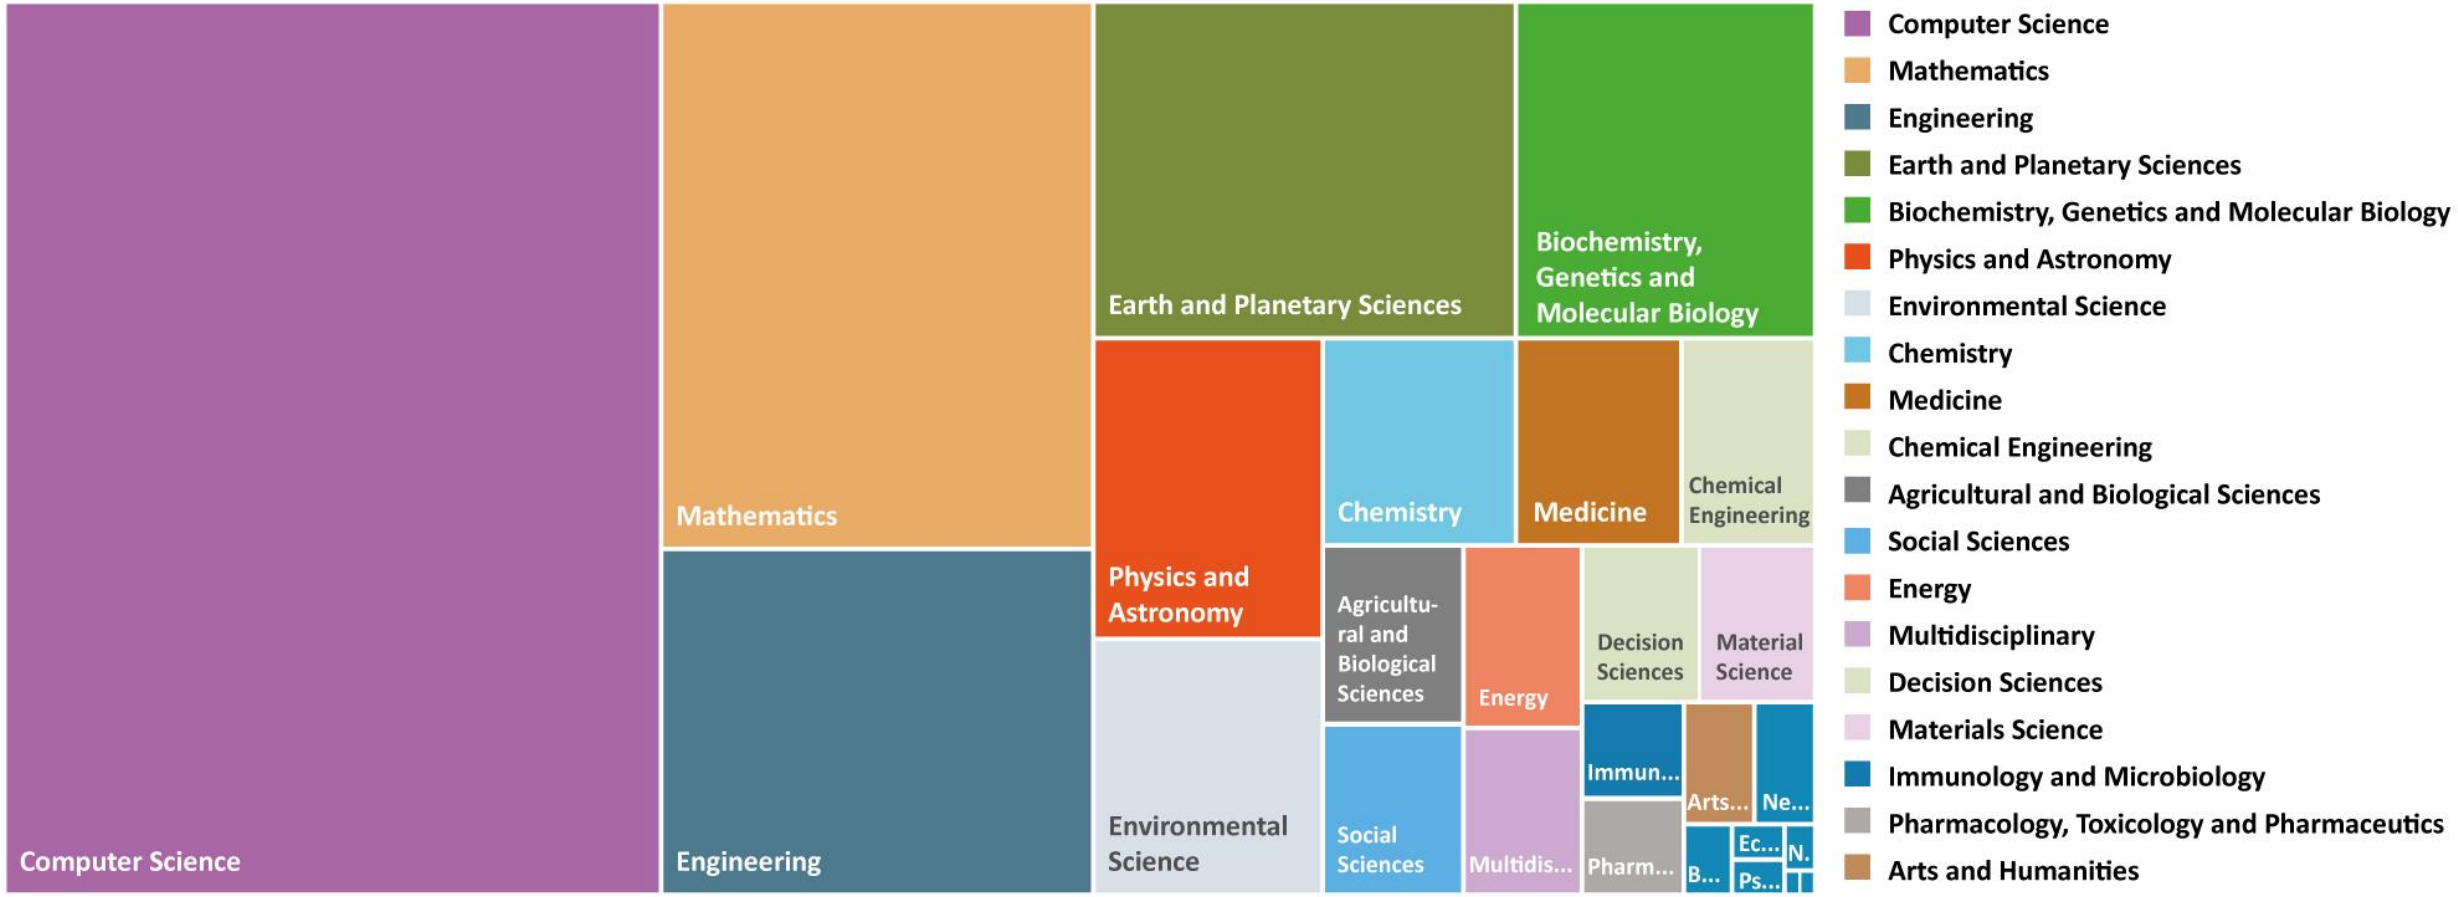

Total: 2,714

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Transferencia tecnológica en el BSC

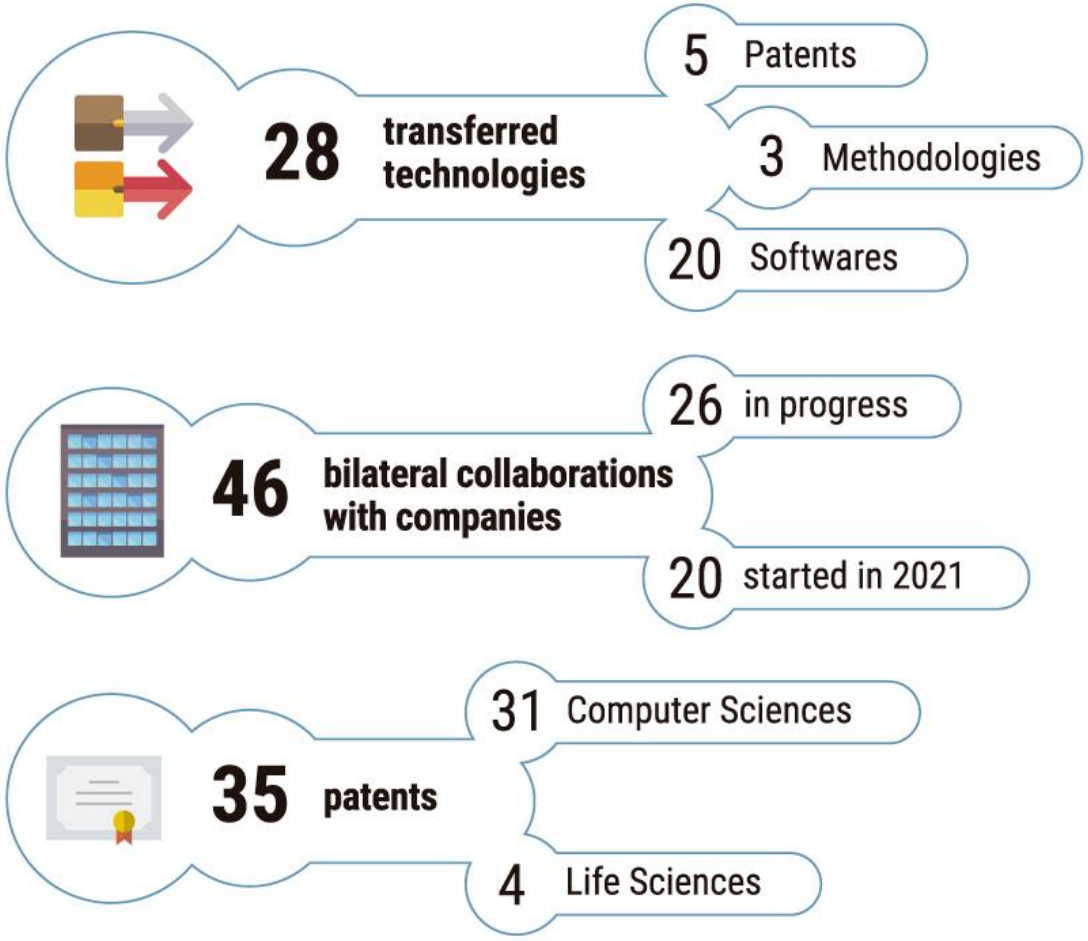

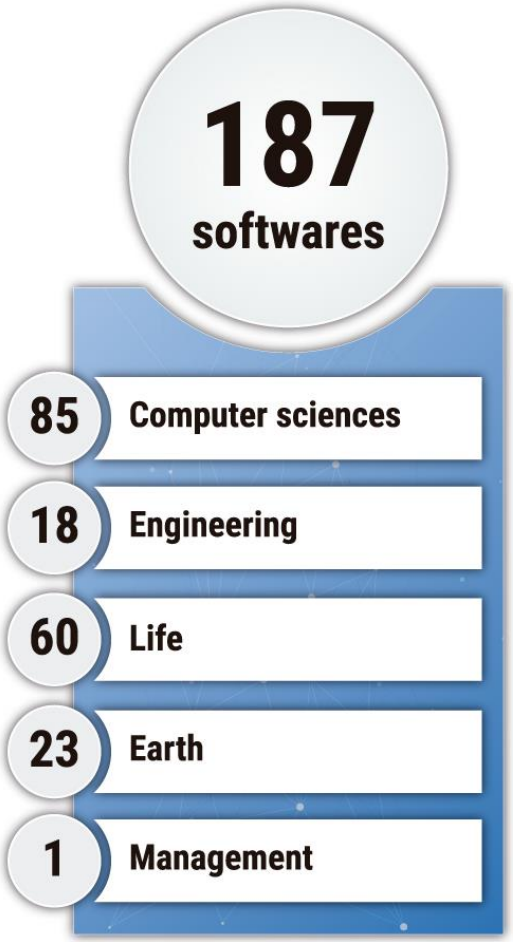

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Collaborations with Global IT industry 2020

Collaborations with Industry

REPSOL

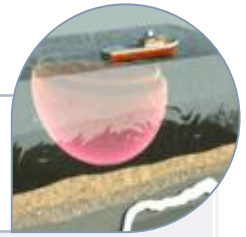

Research into advanced technologies for the exploration of hydrocarbons, subterranean and subsea reserve modelling and fluid flows

IBERDROLA

Research on wind farms optimization and wing energy production forecasts

CaixaBank

Collaboration agreement for the development of advanced systems of deep learning with applications to banking services

SEAT

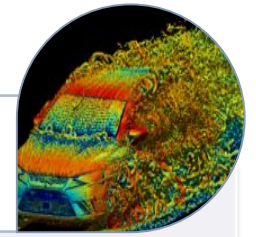

Simulations to improve the understanding of the rotating wheels flow physics and its impact over the aerodynamic performance

Aigües de Barcelona

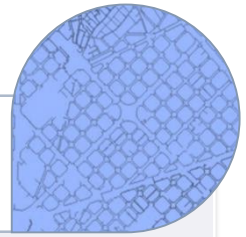

Advanced statistical methods to the optimization of maintenance, energy usage, and control of the city's water treatment and supply processes.

ikergune ETXETARgroup

Research on efficient data sensing, algorithms for analysis of industrial processes and visualization of large datasets of industrial data

ASEPEYO

Artificial Intelligence and Big Data techniques to improve the quality of care and personalized diagnosis

Rockwell Collins

BSC's dust storm forecast system licensed to be used to improve the safety of business flights.

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Spin-off del BSC

NBD INOSTRUM BIODISCOVERY

MITIGA SOLUTIONS

ELEM

NBYCOMP NearbyComputing

MASPA TECHNOLOGIES

Qbeast

QILIMANJARO QUANTUM·TECH

NEXTMOL

frontwave IMAGING

EAS

openchip

11	28	121	6 M€
Spin-Offs	Technologies transferred	Jobs created	Raised capital

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

MareNostrum 4

Total peak performance: **13.9** Pflops		
General Purpose Cluster:	11.15 Pflops	(1-07-2017)
CTE1-P9+Volta:	1.57 Pflops	(1-03-2018)
CTE2-Arm V8:	0.65 Pflops	(12-2019)
CTE3-AMD:	0.52 Pflops	(12-2019)

MareNostrum 1
2004 – 42.3 Tflops
1st Europe / 4th World
New technologies

MareNostrum 2
2006 – 94.2 Tflops
1st Europe / 5th World
New technologies

MareNostrum 3
2012 – 1.1 Pflops
12th Europe / 36th World

MareNostrum 4
2017 – 11.1 Pflops
2nd Europe / 13th World
New technologies

Spanish Supercomputing Network (RES)

RED ESPAÑOLA DE SUPERCOMPUTACIÓN

Infraestructuras Científicas y Técnicas Singulares

www.res.es

@RES_HPC

- Created 2006
- 14 institutions & 16 supercomputers
- HPC resources for scientific community
  - 12.000 Tflops
  - +600 million CPU hours/year
  - 3 calls/year
- Data management services available
  - 9,3 PB for multiannual projects
  - 1 call/year
- Support team to +1.000 regular users
- +200 scientific papers annually
- Member of Unique Scientific and Technical Infrastructure network (**ICTS**).
- Coordinated by **BSC-CNS**

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Updated: March 2021

Distributed supercomputing infrastructure

26 members, including **5 Hosting Members**
(Switzerland, France, Germany, Italy and Spain)

~ 220 PFlops/s of peak performance on **7 world-class systems**

> 30.000M core hours for research awarded

873 scientific projects **enabled**

> 17.000 people trained

> 65 companies supported

PRACE

Access prace-ri.eu/hpc-access

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Data as of July 2021

EuroHPC: towards European HPC technologies

EuroHPC-JU members

Austria, Belgium, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, the Netherlands, North Macedonia, Norway, Poland, Portugal, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden and Turkey.

"A new legal and funding structure – the EuroHPC Joint Undertaking – shall acquire, build and deploy across Europe a world-class High-Performance Computing (HPC) infrastructure.

It will also support a research and innovation programme to develop the technologies and machines (hardware) as well as the applications (software) that would run on these supercomputers."

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EuroHPC Joint Undertaking

October 2022

MareNostrum5 concept

Hosting Consortium:

Spain Portugal Turkey

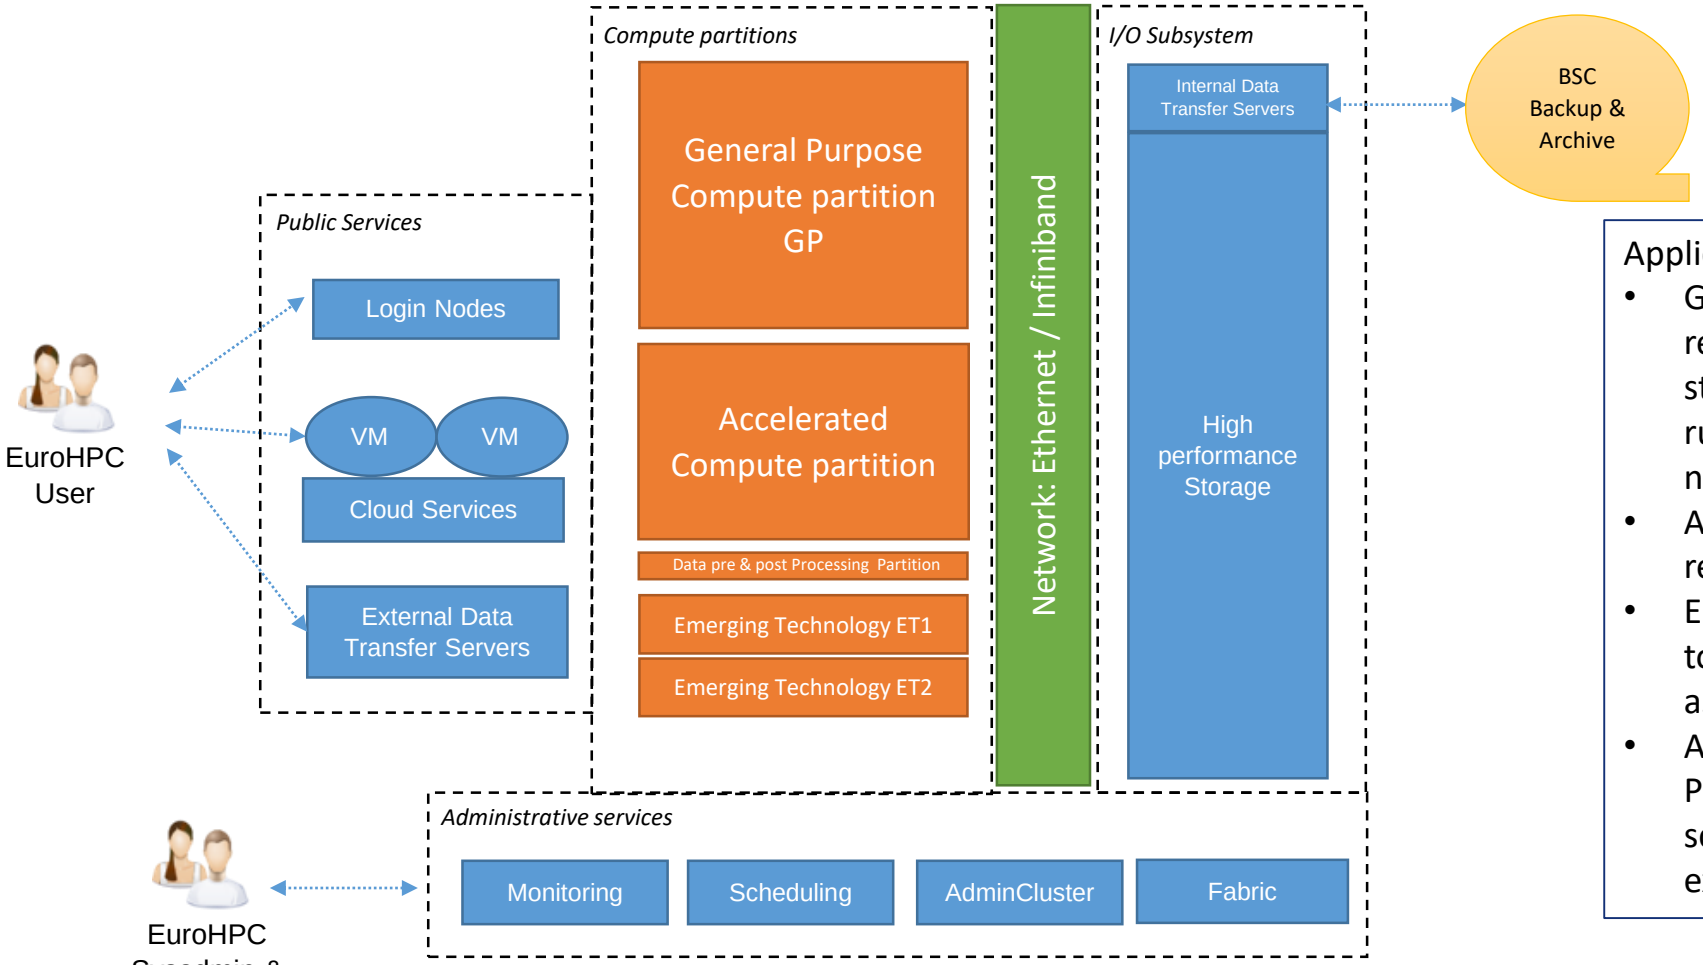

Applications:

- General purpose partition, open to all researchers with MPI, OpenMP codes, standard HPC codes. Scalable machine to run codes with high scalability, thousands of nodes.
- Accelerated partition: Any GPU application ready to scale to thousands of GPUs
- Emerging technologies: prepare workloads to exascale era, exascale technology assessment
- Any domain with workflows mixing General Purpose and GPU, e.g. Earth science, Life science, Engineering, AI and AI driven executions.

https://etendering.ted.europa.eu/cft/cft-display.html?cftId=9758

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union's Connecting Europe Facility and the Horizon 2020 research and innovation programme, as well as the Participating States Spain, Portugal, and Turkey

MareNostrum 5. A European pre-exascale supercomputer

- **200 Petaflops** peak performance (200 x 10^{15})*
- **Experimental platform** to create supercomputing technologies “made in Europe”
- **217 M€** of investment

EuroHPC Joint Undertaking

Hosting Consortium:

Spain **Portugal** **Turkey** **Croatia**

European Commission

* *At the time of call for HE, peak performance expected of 200 Petaflops*
- *At the time of tender publications, minimum aggregated sustained HPL of 205 Petaflops*
- *Contract signed on July 2022*

Generalitat de Catalunya
Departament d'Empresa i Coneixement

Gobierno de España
Ministerio de Ciencia, Innovación y Universidades

UPC Universitat Politècnica de Catalunya BarcelonaTech

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The acquisition and operation of the EuroHPC supercomputer is funded jointly by the EuroHPC Joint Undertaking, through the European Union’s Connecting Europe Facility and the Horizon 2020 research and innovation programme, as well as the Participating States Spain, Portugal, Croatia, and Turkey

EuroHPC Joint Undertaking

Data Center

- Space: 900 sqm
- Energy
  - **Green energy**
  - Up to 20 MW (<10 exp.)
    - Extendable to 40 MW
  - Emergency Line, 1+1 MW UPS
- Cooling
  - **Heat Reuse**
  - PUE below 1,08
  - 1500 m3/h, 28-38ºC
    - 1170 m3/h , 30-40ºC
    - 302 m3/h, 17-27ºC
    - 151 m3/h, 8-14 ºC
  - N+2 redundancy
- Data Center Monitoring & Control
  - **Water quality & treatment**
  - Fire detection & extinction

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

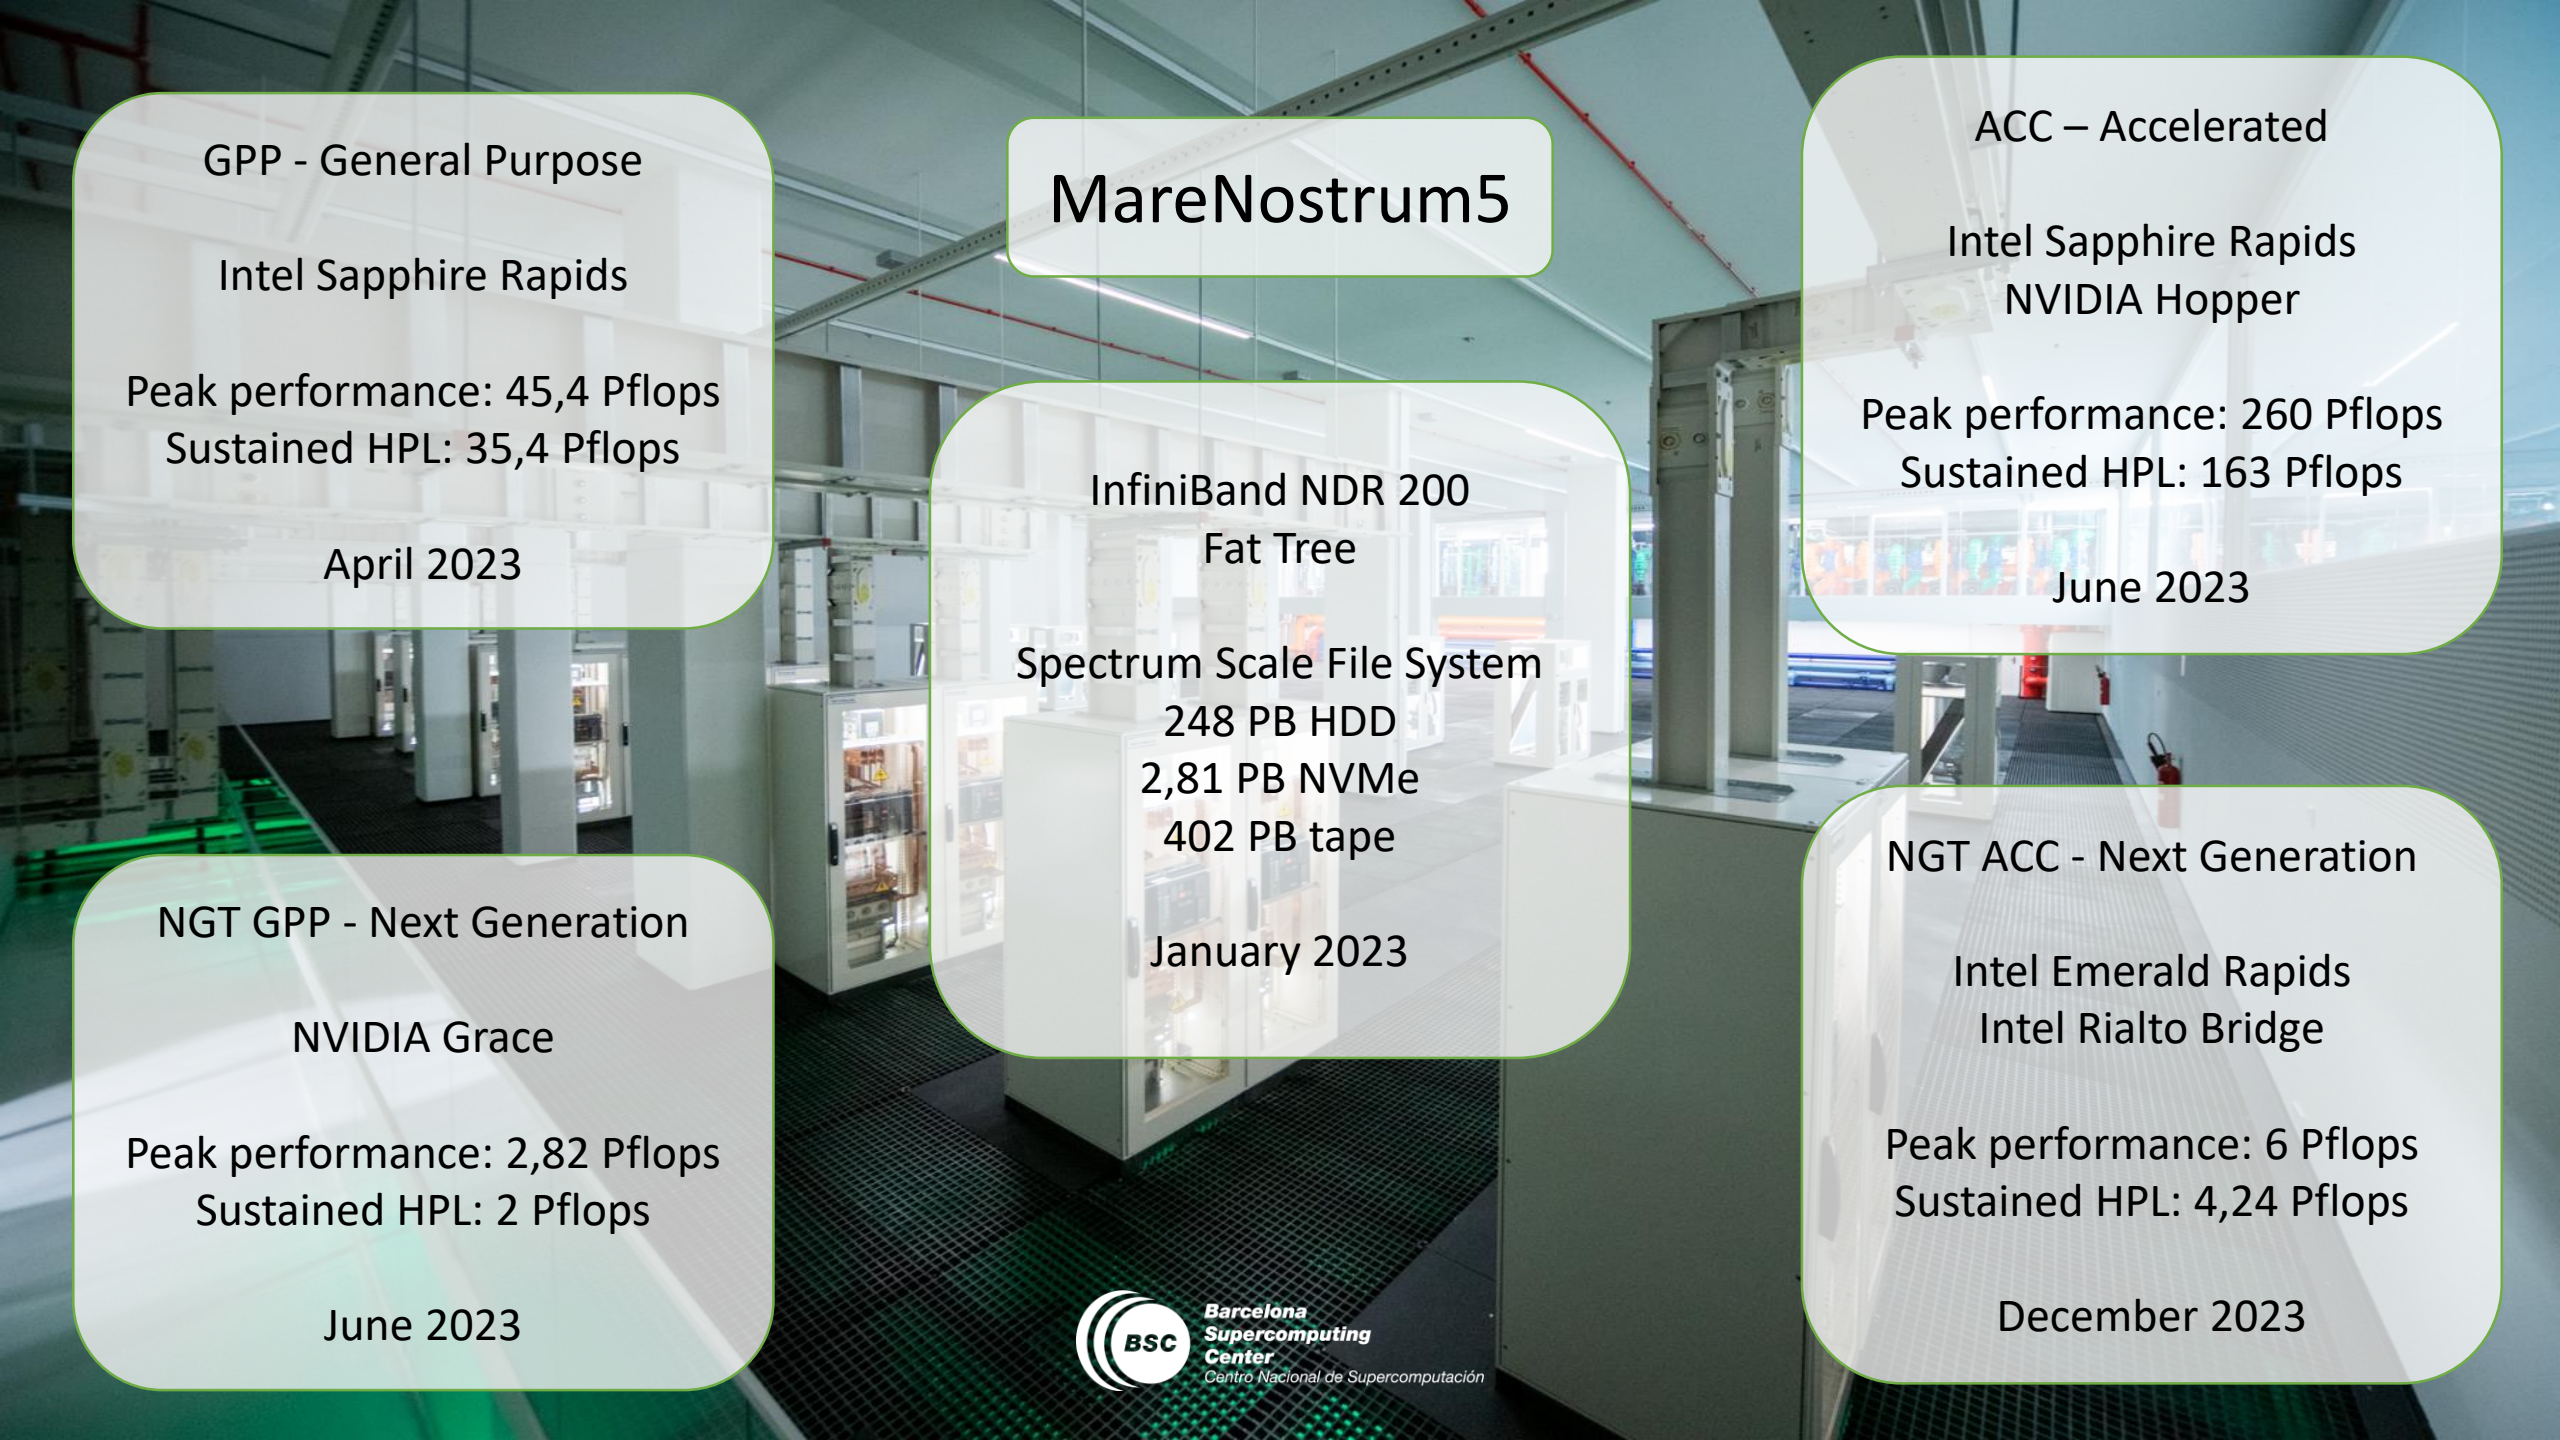

MareNostrum5

GPP - General Purpose

Intel Sapphire Rapids

Peak performance: 45,4 Pflops
Sustained HPL: 35,4 Pflops

April 2023

ACC – Accelerated

Intel Sapphire Rapids
NVIDIA Hopper

Peak performance: 260 Pflops
Sustained HPL: 163 Pflops

June 2023

InfiniBand NDR 200
Fat Tree

Spectrum Scale File System
248 PB HDD
2,81 PB NVMe
402 PB tape

January 2023

NGT GPP - Next Generation

NVIDIA Grace

Peak performance: 2,82 Pflops
Sustained HPL: 2 Pflops

June 2023

NGT ACC - Next Generation

Intel Emerald Rapids
Intel Rialto Bridge

Peak performance: 6 Pflops
Sustained HPL: 4,24 Pflops

December 2023

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Compute partitions overview

		Cooling	Nodes Total	Technology	Processor/Accelerator		Memory	PFlops (HPL)		Local Drive	High-Perf. Network
	General Purpose	DLC +RDHX	>6000	Lenovo	2x Intel Sapphire R.		>2GB/core 256GB DDR5	35.43	>205	960GB NVMe	1x NDR200 Shared by 2 nodes
			>200				>8GB/core 1024GB DDR5				
			>50		2x Intel Sapphire R. HBM		> 0.5GB HBM/core 128GB HBM + 32GB DDR5	0.34			
	Accelerated	DLC	> 1000	Atos	2x Intel Sapphire R.		512GB	163		480GB NVMe	4x NDR200
					4x Nvidia Hopper 64GB HBM						
Next Gen	General Purpose	AC +RDHX	> 400	Atos	Nvidia Grace	144c @ > 2.4GHz	240GB LPDDR5	2		128GB NVMe	1x NDR200
	Accelerated	DLC +RDHX		Lenovo	2x Intel Emerald R.		512GB DDR5	4.24		960GB NVMe	2x NDR
					4x Intel Rialto Bridge ≧128GB HBM						

Compute racks Infrastructure

	General Purpose	Accelerated	Next Gen General Purpose	Next Gen Accelerated
Size of the rack (H x W x D) (in cm)	201 x 60 x 160	225 x 90 x 135 (no doors)	~ 202 x 60 x 127	201 x 60 x 160
Weight (kg / m2)	1637 kg/m²	< 2415 kg/m2	~ 1000 kg/m2	675 kg/m2
Compute nodes per rack	72	32	68	24
KW per rack (average) HPL	> 65 kW	> 110 kW	> 50 kW	> 70 kW
Cooling mechanism	Direct-To-Node Warm Water Cooling and Rear Door Heat Exchangers	Direct Warm Water Cooling	Rear Door Heat Exchangers	Direct-To-Node Warm Water Cooling and Rear Door Heat Exchangers
Residual heat to ambient (kW)	Room neutral	3.86kW	Room neutral	None

Other racks Infrastructure

	Management + Eth HUB + Pre-Post Compute	IB Network	Storage	Archive
Number of racks	3	3	25	26
Size of the rack (H x W x D) (in cm)	201 x 60 x 120 (w/o RDHx)	201 x 60 x 120 (w/o RDHx)	202 x 60 x 120	180 x 78 x 121
Weight (kg / m2)	1095 kg/m²	704.9 kg/m²	1679 - 1826 kg/m²	500 – 630 kg/m²
Linpack kW per rack (average)	10-11.2 kW	22.6 kW	16-22 kW	1.4 kW
Cooling mechanism	Rear Door Heat Exchanger	Rear Door Heat Exchanger	Rear Door Heat Exchanger	N/A
Residual heat to ambient (kW)	Room neutral	Room neutral	Room neutral	Room neutral

MareNostrum5 – Software stack

Software type	MN5
Operating system	Red Hat Enterprise Linux
Compiler Suite	Intel OneAPI HPC Toolkit Nvidia SDK (PGI)
Numerical libraries	Intel MKL Nvidia SDK
Debugging/profiler tools	BSC Performance tools ARM DDT Nvidia SDK Intel OneAPI HPC Toolkit (vtune, ...)
Resource and workload manager	SLURM Only one Slurm cluster, with different partitions
Energy Efficiency and Power Management	EAR

Applications

- Digital twins
  - Human body
  - Climate change
  - Smart cities

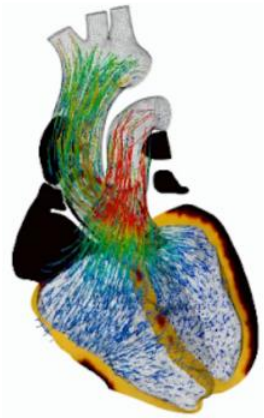

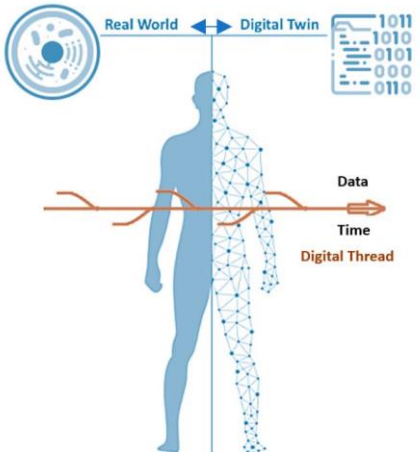

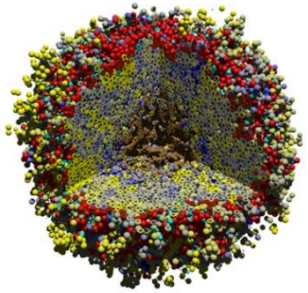

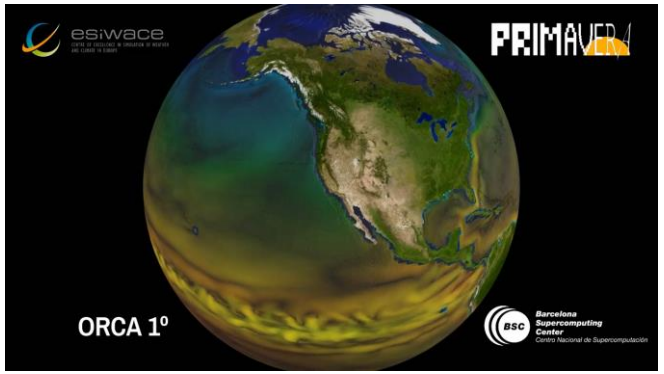

- Standard HPC codes
  - Climate research, engineering, material science and earth sciences, …

- AI applications and Big Data Processing

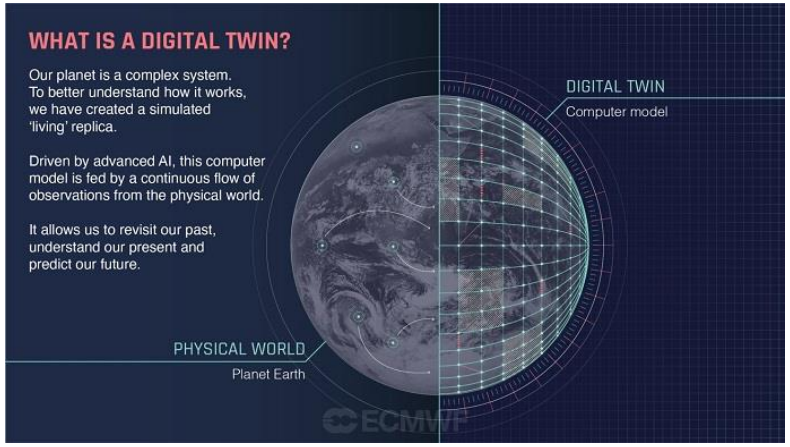

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

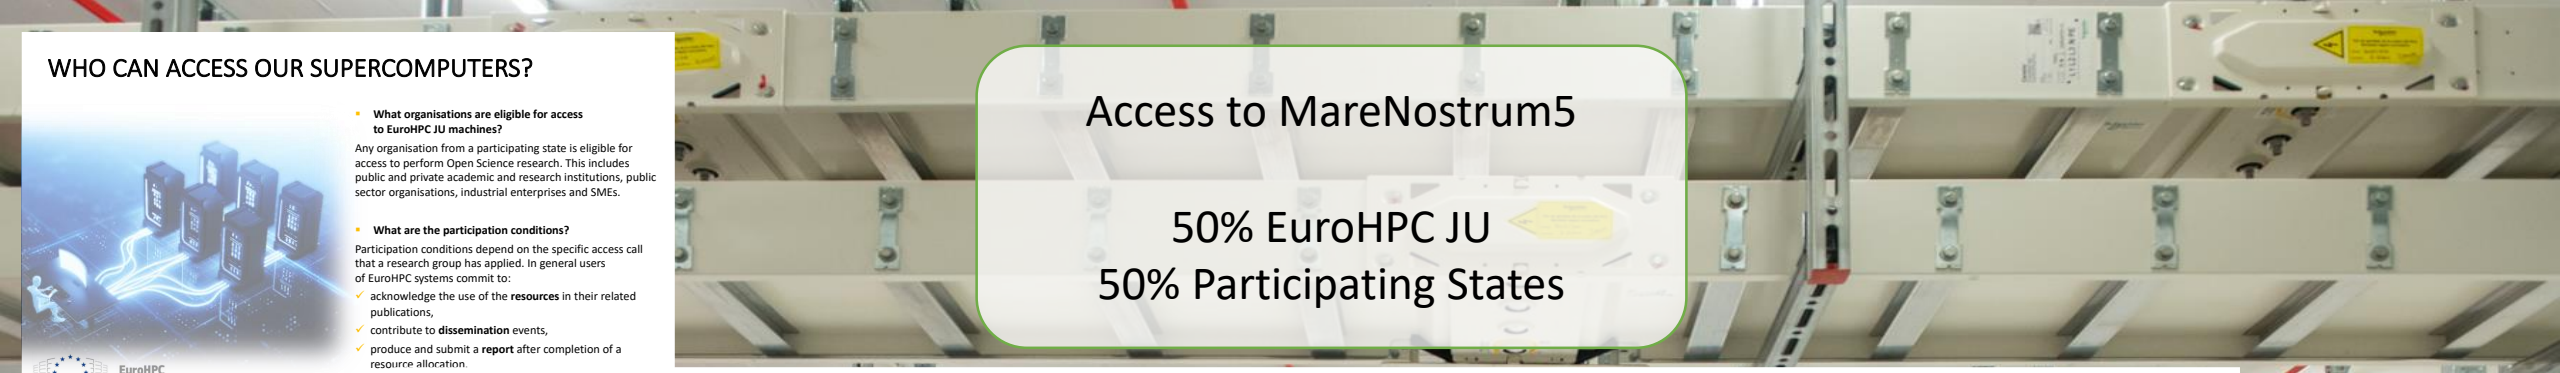

Access to MareNostrum5

50% EuroHPC JU
50% Participating States

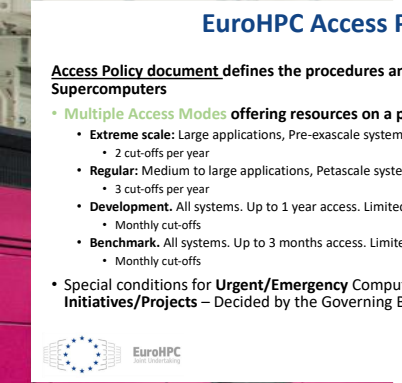

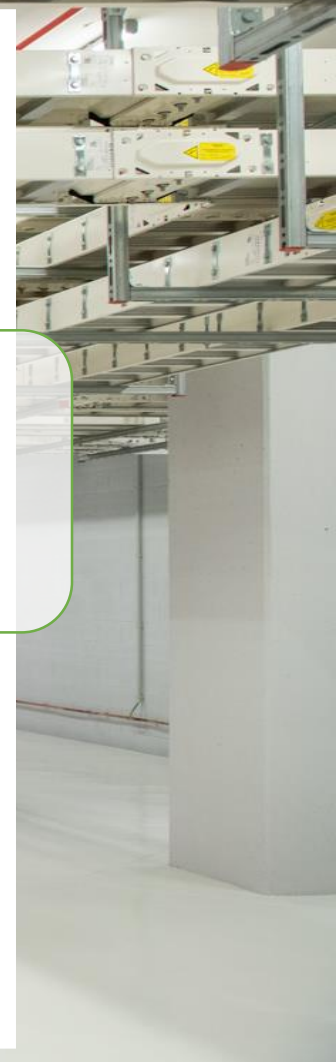

EuroHPC Access Policy

document defines the procedures and conditions for access the EuroHPC ers

cess M **riodic and continuously open call basis.**

cale: La Peer-reviewed

offs per y

Medium s. Peer-reviewed

offs per y

- **Development.** All esources.
  - Monthly cut-offs
- **Benchmark.** All systems. Up to 3 months access. Limited resources.
  - Monthly cut-offs

- Special conditions for **Urgent/Emergency** Computing & Access for **Strategic Initiatives/Projects** – Decided by the Governing Board.

Participating States

Spain **Portugal** **Turkey**

- According to contributions
- Own decision access methods

EuroHPC Joint Undertaking

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

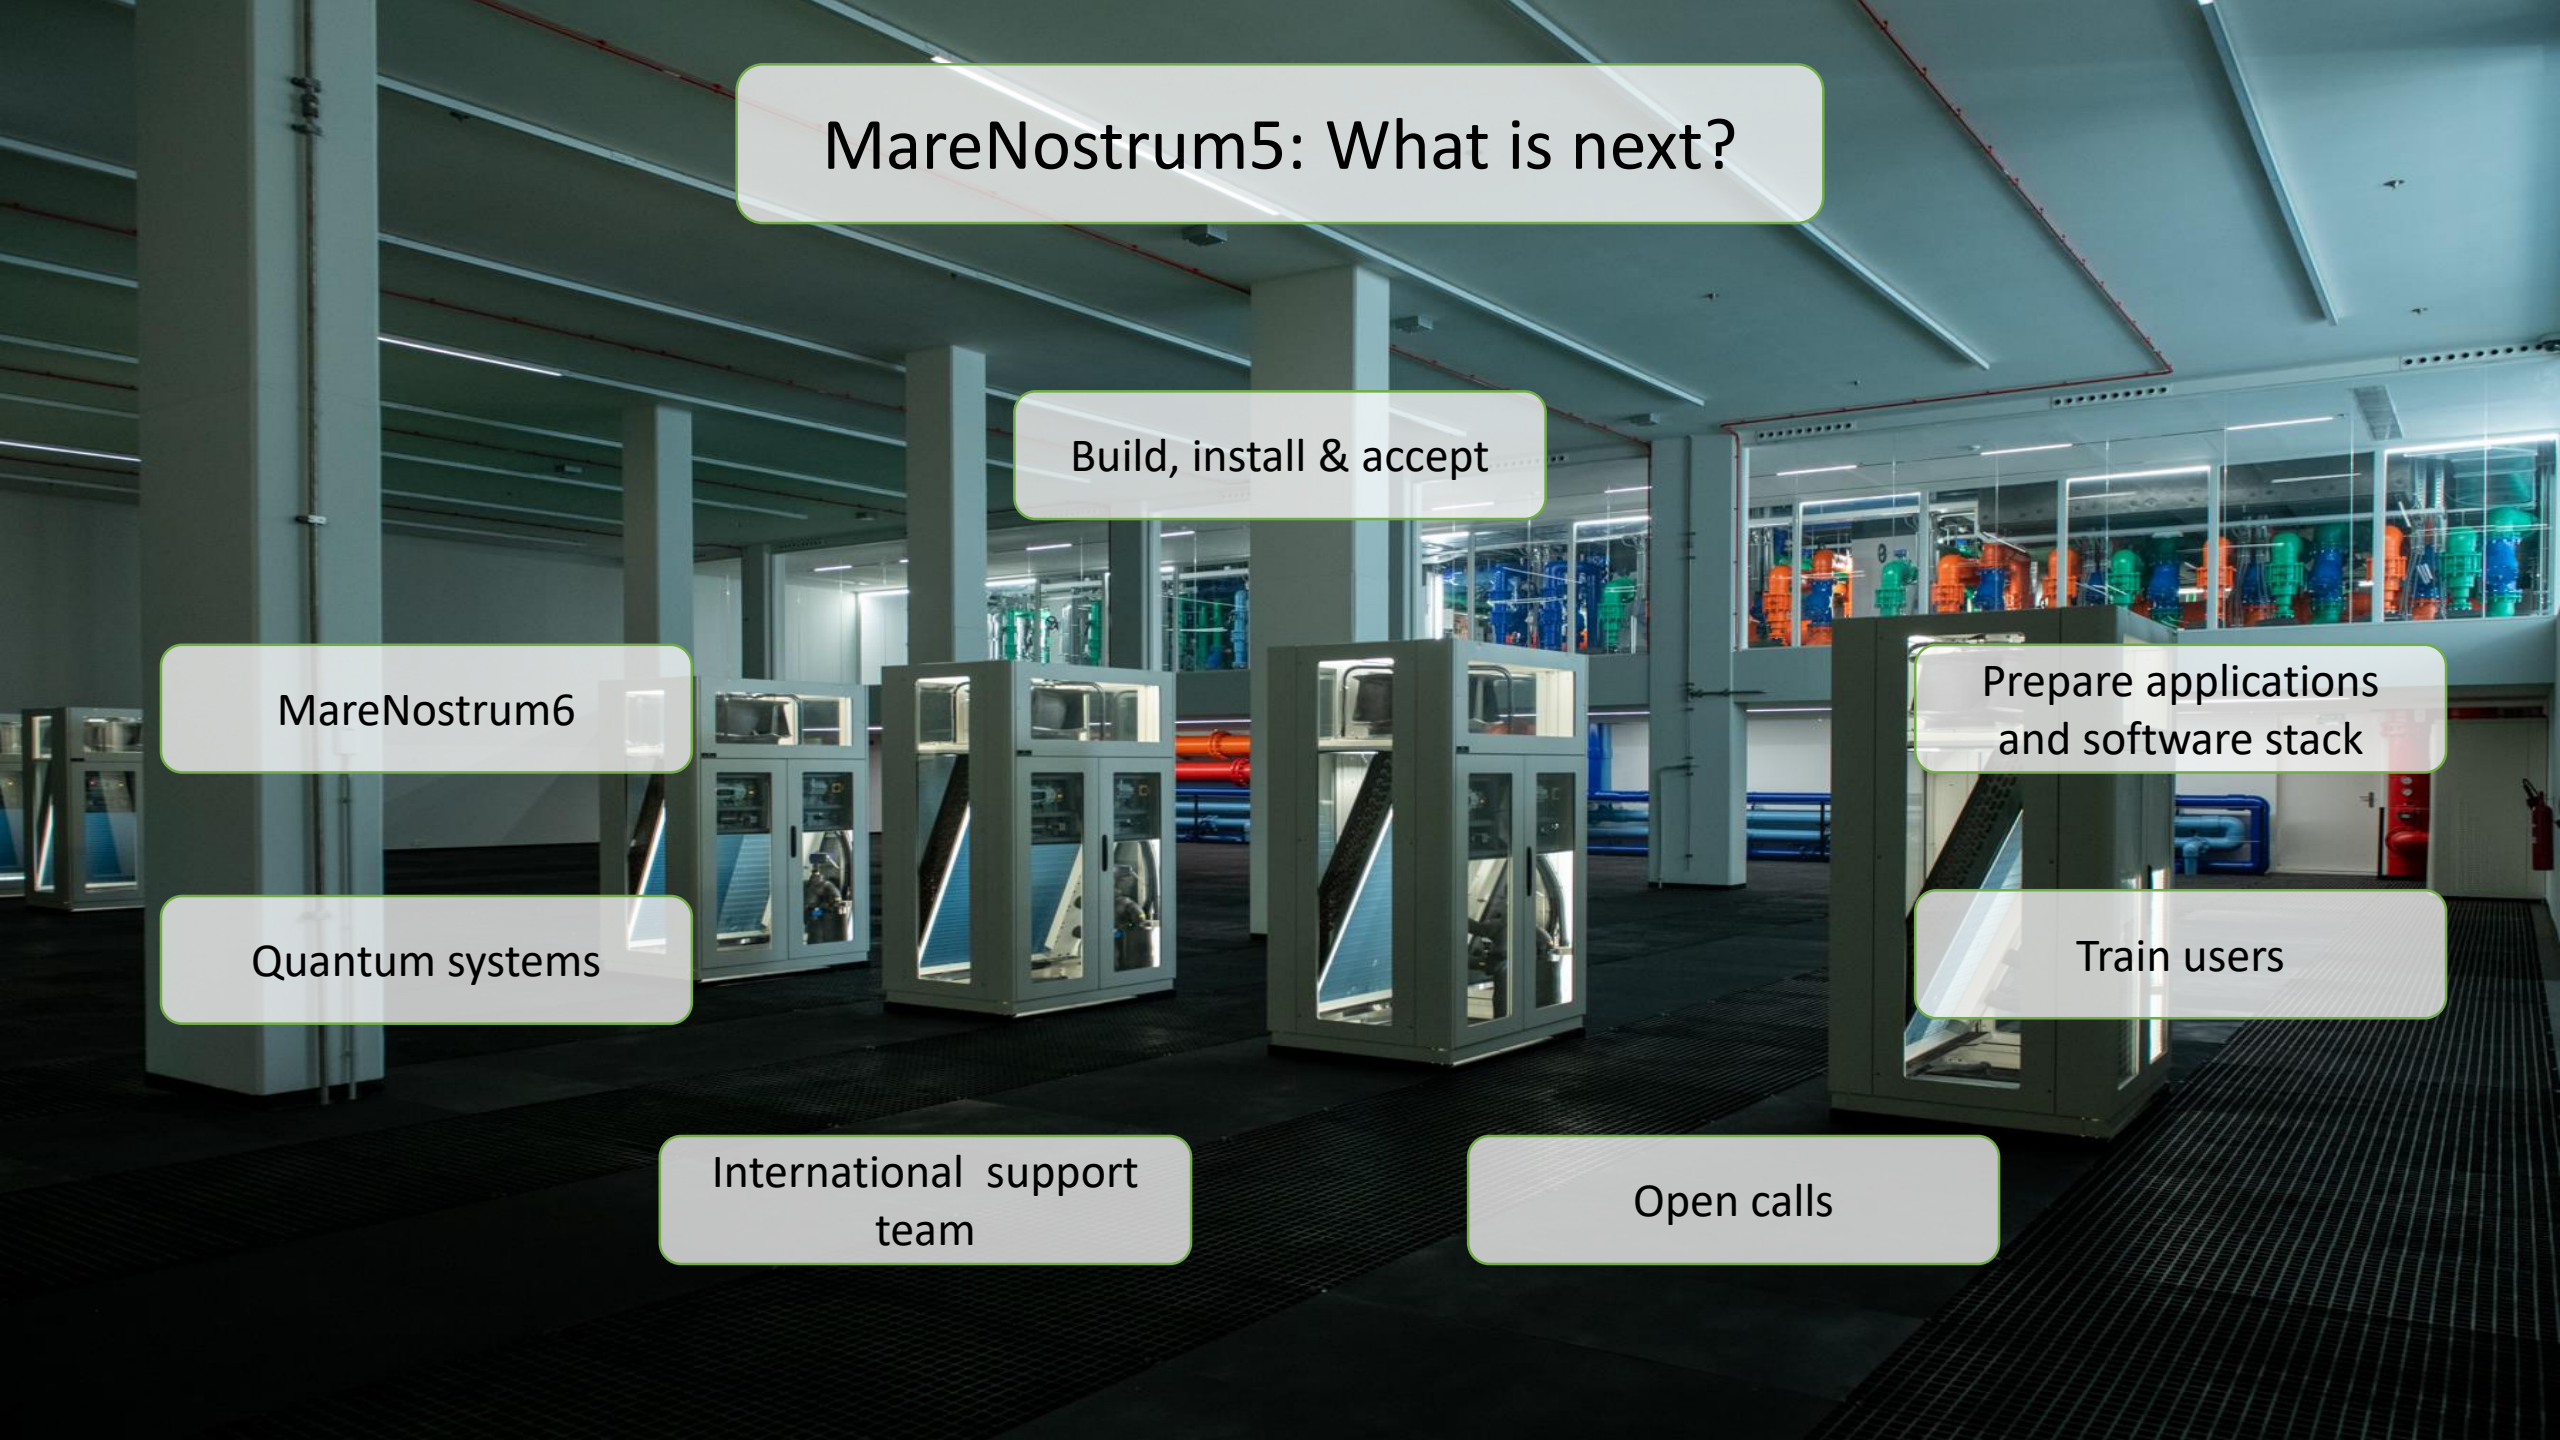
MareNostrum5: What is next?
Build, install & accept
MareNostrum6
Prepare applications and software stack
Quantum systems
Train users
International support team
Open calls

Thank you

sergi.girona@bsc.es

Barcelona Supercomputing Center
Centro Nacional de Supercomputación