Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

European Processor Initiative & RISC-V

Prof. Mateo Valero

BSC Director

erc

9/May/2018

RISC-V Workshop, Barcelona

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC-CNS objectives

Supercomputing services to Spanish and EU researchers

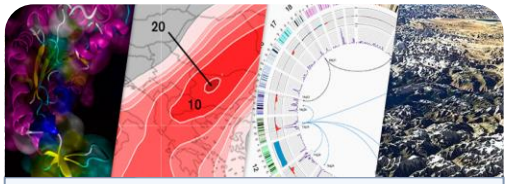

R&D in Computer, Life, Earth and Engineering Sciences

PhD programme, technology transfer, public engagement

BSC-CNS is a consortium that includes

Spanish Government	60%	GOBIERNO DE ESPAÑA · MINISTERIO DE ECONOMÍA Y COMPETITIVIDAD
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 and Resources

Data as of 31st of December 2017

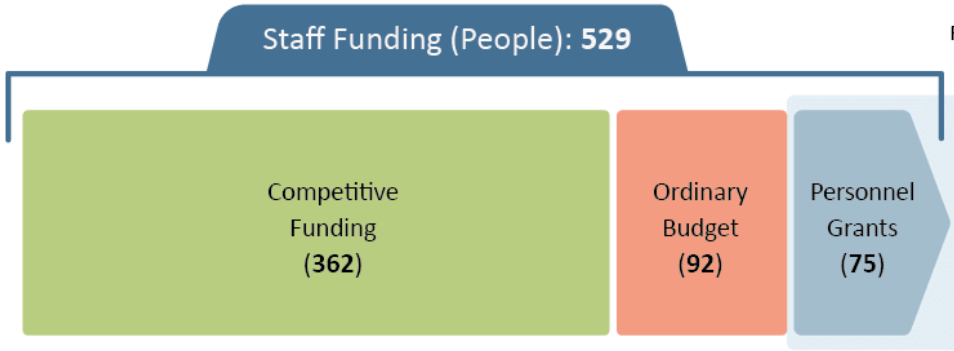

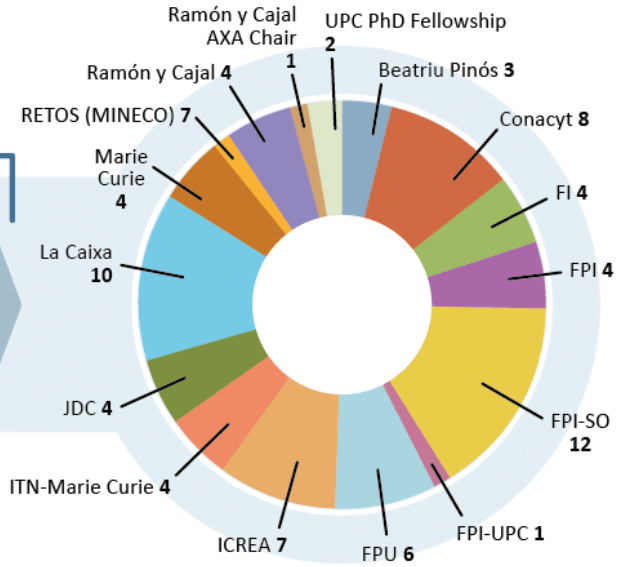

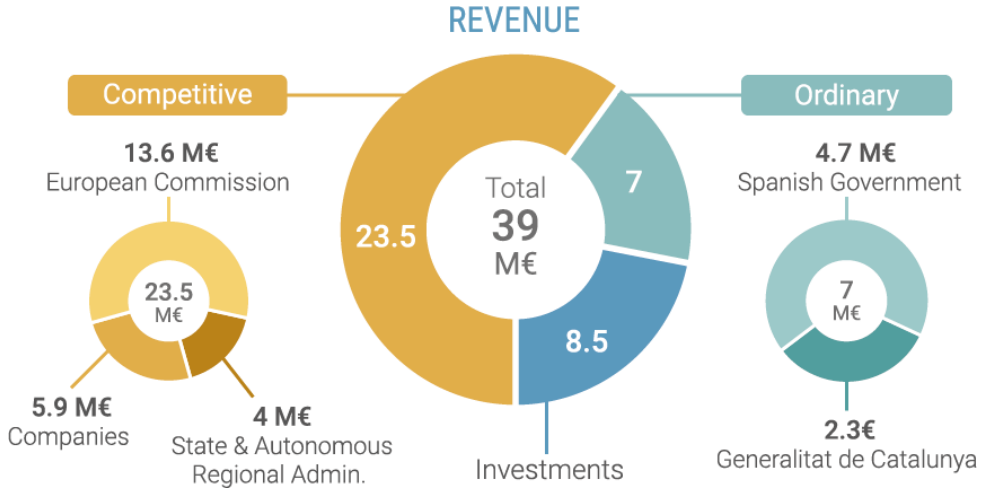

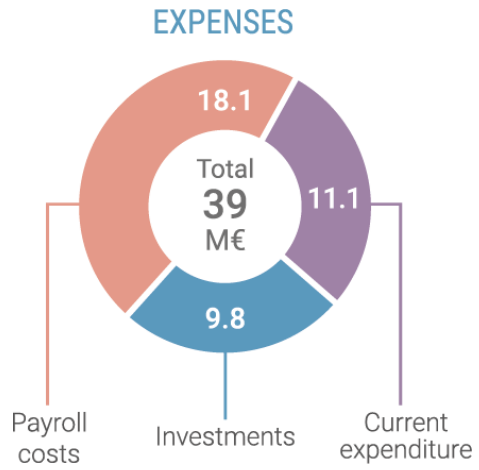

Context: The international Exascale challenge

- Sustained real life application performances, not just Linpack…
- Exascale will not just allow present solutions to run faster, but will enable new solutions not affordable with today HPC technology
- From simulation to high predictability for precise medicine, energy, climate change, autonomous driven vehicles…
- The International context (US, China, Japan and EU…)
- The European HPC programme
- The European Processor initiative
- BSC role

Top5 TOP500 SUPERCOMPUTER SITES — November 2017 list (compact view)

Rank	Site	Computer	Procs	Rmax	Rpeak	Mflops/Watt
1	Wuxi, China	Sunway SW26010 260C	10.649.600	93.015	125.436	6.051
2	Guangzhou, China	Xeon E5-2692+Phi	3.120.000 2736000	33.863	54.902	1.902
3	CSCS, Switzerland	Cray XC50, Xeon E2690 12C+P100	361.760 297920	19.590	25.326	8.622
4	JAMEST, Japan	ZettaScaler Xeon D-1571, PEZY-SC2	19860000 19840000	19.135	28.192	14.173
5	DOE/SC/Oak Ridge, US	Cray XK7, Opteron 16C+K20	560.640 261632	17.590	27.113	2.143
6	DOE/NNSA/LLNL, US	BlueGene/Q, BQC 16C	1.572.864	17.173	20.133	2.177
7	DOE/NNSA/LANL/SNL, US	Cray XC40, IXeon Phi 7250 (KNL)	979.968	14.137	43.902	3.678
8	DOE/SC/LBNL/NERSC, US	Cray XC40, Xeon Phi 7250 (KNL)	622.336	14.015	27.881	3.558
9	JCA, Japan	PRIMERGY, Xeon Phi 7250 (KNL)	556.104	13.555	24.913	4.896
10	Riken/AICS, Japan	K computer, SPARC64	705.024	10.510	11.280	830
11	DOE/SC/Argonne, US	BlueGene/Q, Power BQC 16C	786.432	8.587	10.066	2.177
12	TACC Texas, US	PowerEdge, Xeon Phi 7250 (KNL),	368.928	8.318	18.216	N/A
13	GSIC-TIT, Japan	SGI, Xeon E5,Tesla P100	135.828 120736	8.125	12.127	10.258
14	CINECA, Italy	Lenovo, Xeon Phi 7250 68C, KNL	314.384	7.471	15.372	N/A
15	UKMET, UK	Cray XC40, Xeon E2695 18C	241.920	7.039	8.129	N/A
16	BSC, Spain	Lenovo, Xeon Platinum 8160 24C	153.216	6.471	10.296	3.965

Top5

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

According to HPL

Rank	Name	Country	Rmax	Rpeak	% Efficiency
1	Sunway TaihuLight	China	93,015	125,436	74.15 %
2	Tianhe-2	China	33,863	54,902	61.68%
3	Piz Daint	Switzerland	19,590	25,326	77.35%
4	Gyoukou	Japan	19,135	28,192	67.87%
5	Titan	US	17,590	27,113	64.87%
6	Sequoia	US	17,173	20,133	85.30%
7	Trinity	US	14,137	43,902	32.20%
8	Cori	US	14,015	27,881	50.27%
9	Oakforest-PACS	Japan	13,555	24,913	54.41%
10	K Computer	Japan	10,510	11,280	93.17%

According to HPCG benchmark

Rank	Name	Rmax	Rpeak	% Efficiency
1	K Computer	603	11,280	5.34%
2	Tianhe-2	580	54,902	1.06%
3	Trinity	546	43,902	1.24%
4	Piz Daint	486	25,326	1.92%
5	Sunway TaihuLight	481	125,436	0.38%
6	Oakforest-PACS	385	24,913	1.54%
7	Cori	355	27,881	1.27%
8	Sequoia	330	20,133	1.64%
9	Titan	322	27,113	1.19%
10	Mira	167	10,066	1.66%

Top5

TOP500 SUPERCOMPUTER SITES

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

According to HPL

Rank	Name	Country	Rmax	Rpeak	% Efficiency
1	Sunway TaihuLight	China	93,015	125,436	74.15 %
2	Tianhe-2	China	33,863	54,902	61.68%
3	Piz Daint	Switzerland	19,590	25,326	77.35%
4	Gyoukou	Japan	19,135	28,192	67.87%
5	Titan	US	17,590	27,113	64.87%
6	Sequoia	US	17,173	20,133	85.30%
7	Trinity	US	14,137	43,902	32.20%
8	Cori	US	14,015	27,881	50.27%
9	Oakforest-PACS	Japan	13,555	24,913	54.41%
10	K Computer	Japan	10,510	11,280	93.17%
16	Mare nostrum	Spain	6,471	10,296	62,85%

According to HPCG benchmark

Rank	Name	Rmax	Rpeak	% Efficiency
1	K Computer	603	11,280	5.34%
2	Tianhe-2	580	54,902	1.06%
3	Trinity	546	43,902	1.24%
4	Piz Daint	486	25,326	1.92%
5	Sunway TaihuLight	481	125,436	0.38%
6	Oakforest-PACS	385	24,913	1.54%
7	Cori	355	27,881	1.27%
8	Sequoia	330	20,133	1.64%
9	Titan	322	27,113	1.19%
10	Mira	167	10,066	1.66%
15	Mare Nostrum	122	10,296	1,18%

Rank	Previous rank	Machine	Country	Number of cores	GTEPS
1	1	K computer	Japan	663,552	38,621
2	2	Sunway TaihuLight	China	10,599,680	23,755
3	3	DOE/NNSA/LLNL Sequoia	USA	1,572,864	23,751
4	4	DOE/SC/Argonne National Laboratory Mira	USA	786,432	14,982
5	5	JUQUEEN	Germany	262,144	5,848
6	new	ALCF Mira - 8192 partition	United States	131,072	4,212
7	6	ALCF Mira - 8192 partition	USA	131,072	3,556
8	7	Fermi	Italy	131,072	2,567
9	new	ALCF Mira - 4096 partition	United States	65,536	2,348
10	8	Tianhe-2 (MilkyWay-2)	China	196,608	2,061

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	259	**Shoubu system B** - PEZY Computing RIKEN -Japan	794,400	842.0	50	16.84
2	307	**Suiren2** - PEZY Computing KEK -Japan	762,624	788.2	47	16.77
3	276	**Sakura** - PEZY Computing PEZY Computing K.K. -Japan	794,400	824.7	50	16.49
4	149	**DGX SaturnV Volta** - NVIDIA Tesla V100 NVIDIA Corporation -United States	22,440	1,070.0	97	11.03
5	4	**Gyoukou** - PEZY-SC2 700Mhz Japan	19,860,000	19,135.8	1,350	14.17
6	13	**TSUBAME3.0** - NVIDIA Tesla P100 SXM2 Japan	135,828	8,125.0	792	10.26
7	195	**AIST AI Cloud** - NVIDIA Tesla P100 SXM2 Japan	23,400	961.0	76	12.64
8	419	**RAIDEN GPU subsystem** - NVIDIA Tesla P100 Japan	11,712	635.1	60	10.59
9	115	**Wilkes-2** - NVIDIA Tesla P100 University of Cambridge - United Kingdom	21,240	1,193.0	114	10.46
10	3	**Piz Daint -** NVIDIA Tesla P100 Switzerland	361,760	19,590.0	2,272	8.62
33	16	**MareNostrum**- Lenovo SD530 Barcelona Supercomputing Center Spain	153,216	6,470.8	1,632	3.97

Application processor performance MN3-MN4

Application	Cores	Performance
WRF	256	1.37
	128	1.06
GROMACS	1024	
	192	1.19
NAMD	2048	1.31
	1024	1.17
	728	1.25
	512	1.20
VASP	64	2.2
	32	2.0
HPL	96	2.24
	48	2.21

From MN3 to MN4

MareNostrum 4, chosen as the most beautiful data centre in the world

11 December 2017

The award, organised by DataCenter Dynamics, has been granted by popular vote.

MareNostrum 4 supercomputer has been the winner of the Most Beautiful Data Center in the world Prize, hosted by the Datacenter Dynamics (DCD) Company.

There are 15 prizes in different categories, besides the prize for the most beautiful data centre, which is elected by popular vote. MareNostrum 4 competed with such impressive facilities as the Switch Pyramid in Michigan, the Bahnhof Pionen in Stockholm or the Norwegian Green Mountain. BSC supercomputer has prevailed for its particular location, inside the chapel of Torre Girona, located in the North Campus of the Universitat Politècnica de Catalunya (UPC).

The awards ceremony took place on December 7th in London and both Mateo Valero, BSC Director, and Sergi Girona, Operations department Director, received the prize.

About MareNostrum 4

MareNostrum4

Total peak performance: **13,7** Pflops		
General Purpose Cluster:	11.15 Pflops	(1.07.2017)
CTE1-P9+Volta:	1.57 Pflops	(1.03.2018)
CTE2-Arm V8:	0.5 Pflops	(????)
CTE3-KNH?:	0.5 Pflops	(????)

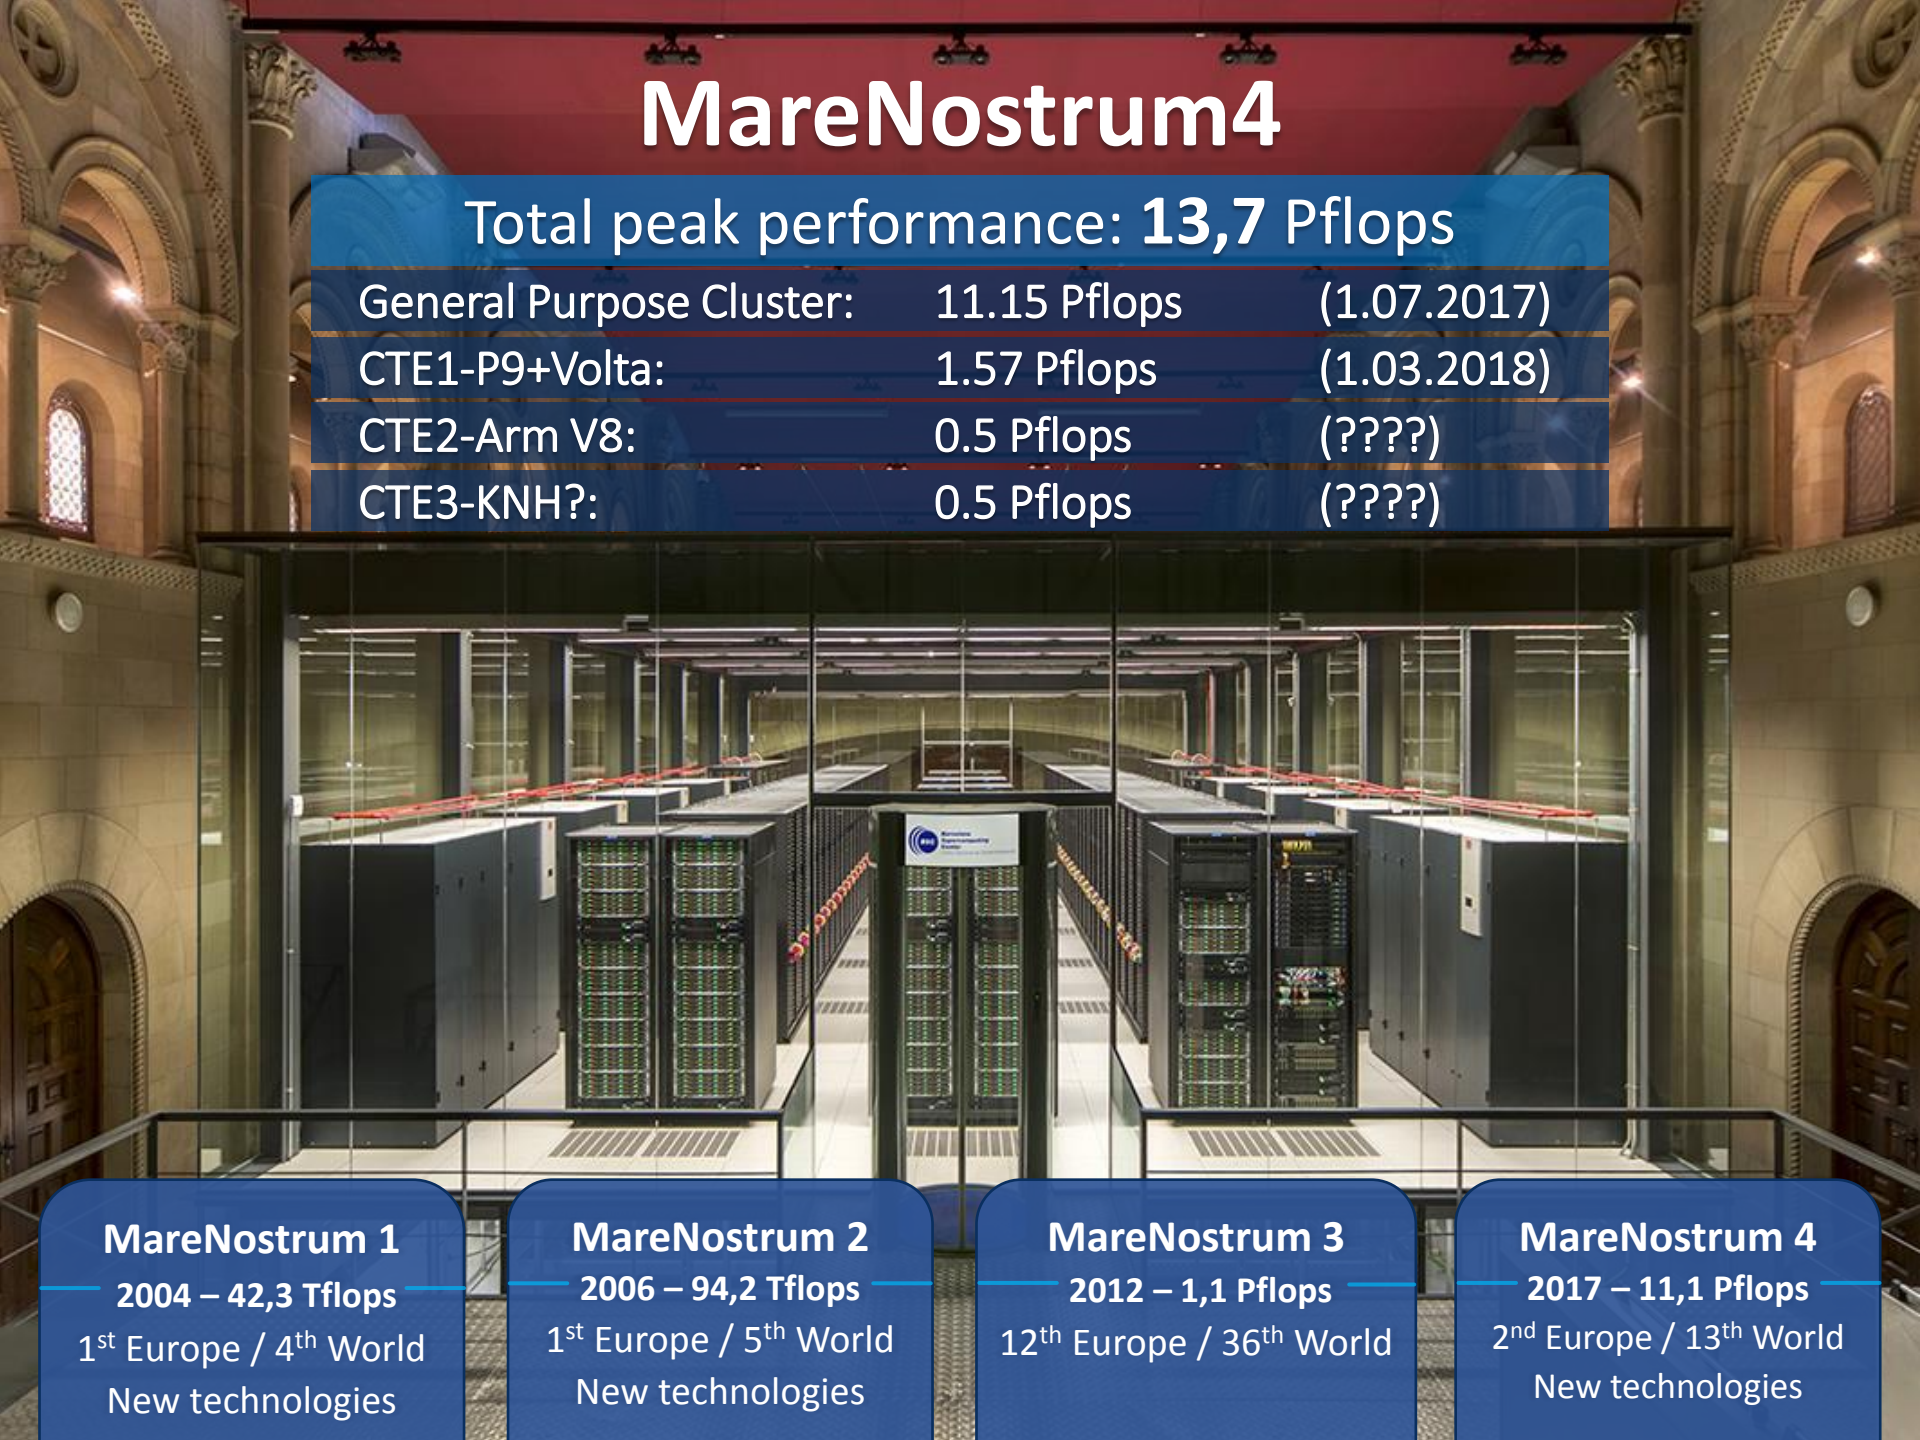

MareNostrum 1
2004 – 42,3 Tflops
1st Europe / 4th World
New technologies

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

Worldwide HPC roadmaps

The White House
Office of the Press Secretary

For Immediate Release July 29, 2015

Executive Order -- Creating a National Strategic Computing Initiative

EXECUTIVE ORDER

CREATING A NATIONAL STRATEGIC COMPUTING INITIATIVE

By the authority vested in me as President by the Constitution and the laws of the United States of America, and to maximize benefits of high-performance computing (HPC) research, development, and deployment, it is hereby ordered as follows:

with domestic technology.

From K computer…

… to Post K

with domestic technology.

From Tianhe-2..

…to Tianhe-2A

with domestic technology.

IPCEI on HPC

From the PPP for HPC…

to future PRACE systems…

…with domestic technology ?

US launched RFP for Exascale (April 2018)

- To develop at least two new exascale supercomputers for the DOE at a cost of up to $1.8 billion
- The deployment timeline for these new systems begins in the third quarter of 2021, with ORNL's exascale supercomputer, followed by a third quarter 2022 system installation at LLNL. The ANL addition or upgrade, if it happens, will also take place in the third quarter of 2022.
- The new systems can't exceed 40 MW, with the preferred power draw in the 20 to 30 MW (including exascale, counting storage, cooling and any other auxiliary equipment)
- The other critical requirement is that the ORNL and ANL systems are architecturally diverse from one other
- Proposals are due in May, the bidders will be selected before the end of the Q2
- Each system is expected to cost between $400 to $600 million second quarter.

Worldwide HPC roadmaps

The White House
Office of the Press Secretary

For Immediate Release July 29, 2015

Executive Order -- Creating a National Strategic Computing Initiative

EXECUTIVE ORDER

CREATING A NATIONAL STRATEGIC COMPUTING INITIATIVE

By the authority vested in me as President by the Constitution and the laws of the United States of America, and to maximize benefits of high-performance computing (HPC) research, development, and deployment, it is hereby ordered as follows:

with domestic technology.

From K computer…

… to Post K

with domestic technology.

From Tianhe-2..

…to Tianhe-2A

with domestic technology.

IPCEI on HPC

From the PPP for HPC…

to future PRACE systems…

…with domestic technology ?

EU HPC Ecosystem

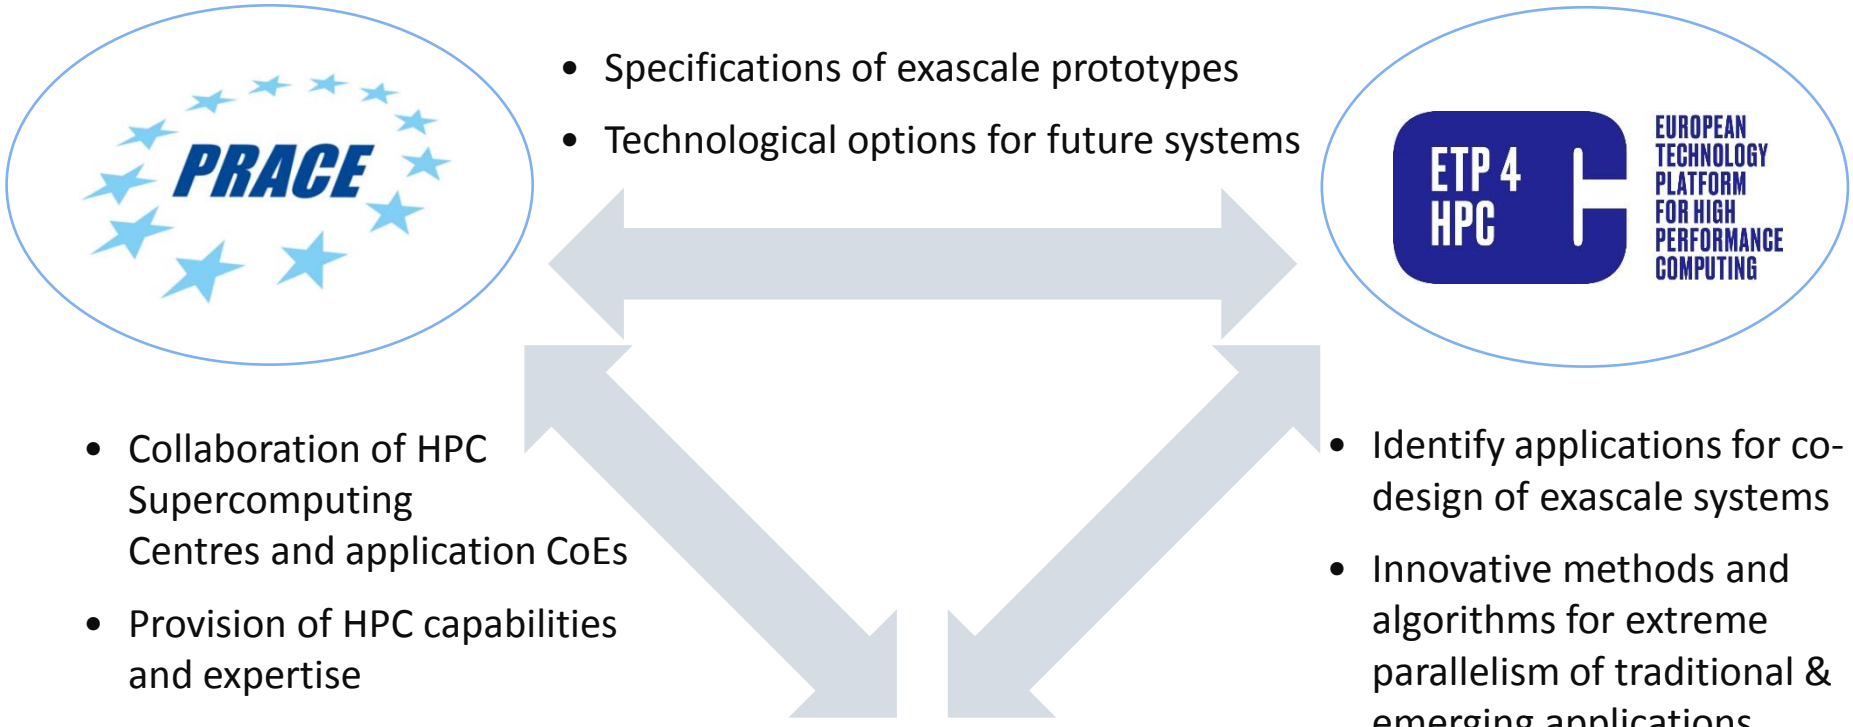

- Specifications of exascale prototypes
- Technological options for future systems

- Collaboration of HPC Supercomputing Centres and application CoEs
- Provision of HPC capabilities and expertise

- Identify applications for co-design of exascale systems
- Innovative methods and algorithms for extreme parallelism of traditional & emerging applications

Centers of Excellence in **HPC applications**

bioexcel
NOMAD NOVEL MATERIALS DISCOVERY
E CAM
POP
EoCoE
MaX
esiwace CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE
CompBioMed
CoeGSS Centre of excellence for GLOBAL SYSTEMS SCIENCE

BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación


SAGE
ExaHYPE
Memory & storage
Mathematics
NEXTGenIO
NLAFET
ExaFLOW
ExCAPE
ComPat
ExaNEST
Algorithms
READEX
ESCAPE
ALLScale
Interconnect
Programming tools
INTERTWINE
HPC
Ecosystem
EXDCI
Eurolab-4-HPC
ECOSCALE
ANTAREX
ExaNoDe
Compute
Data-intensive real-time
Mont-Blanc 3
Mango
EXTRA
GreenFLASH
Centres of Excellence
BioExcel
Biomolecular
COEGSS
Global systems
CompBioMed
Computational Biomedicine
EoCoE
Energy
E-CAM
Simulation Modelling
ESiWACE
Weather Climate
MAX
Materials
NOMAD
PoP
Performance optimisation

A big challenge, and a huge opportunity for Europe

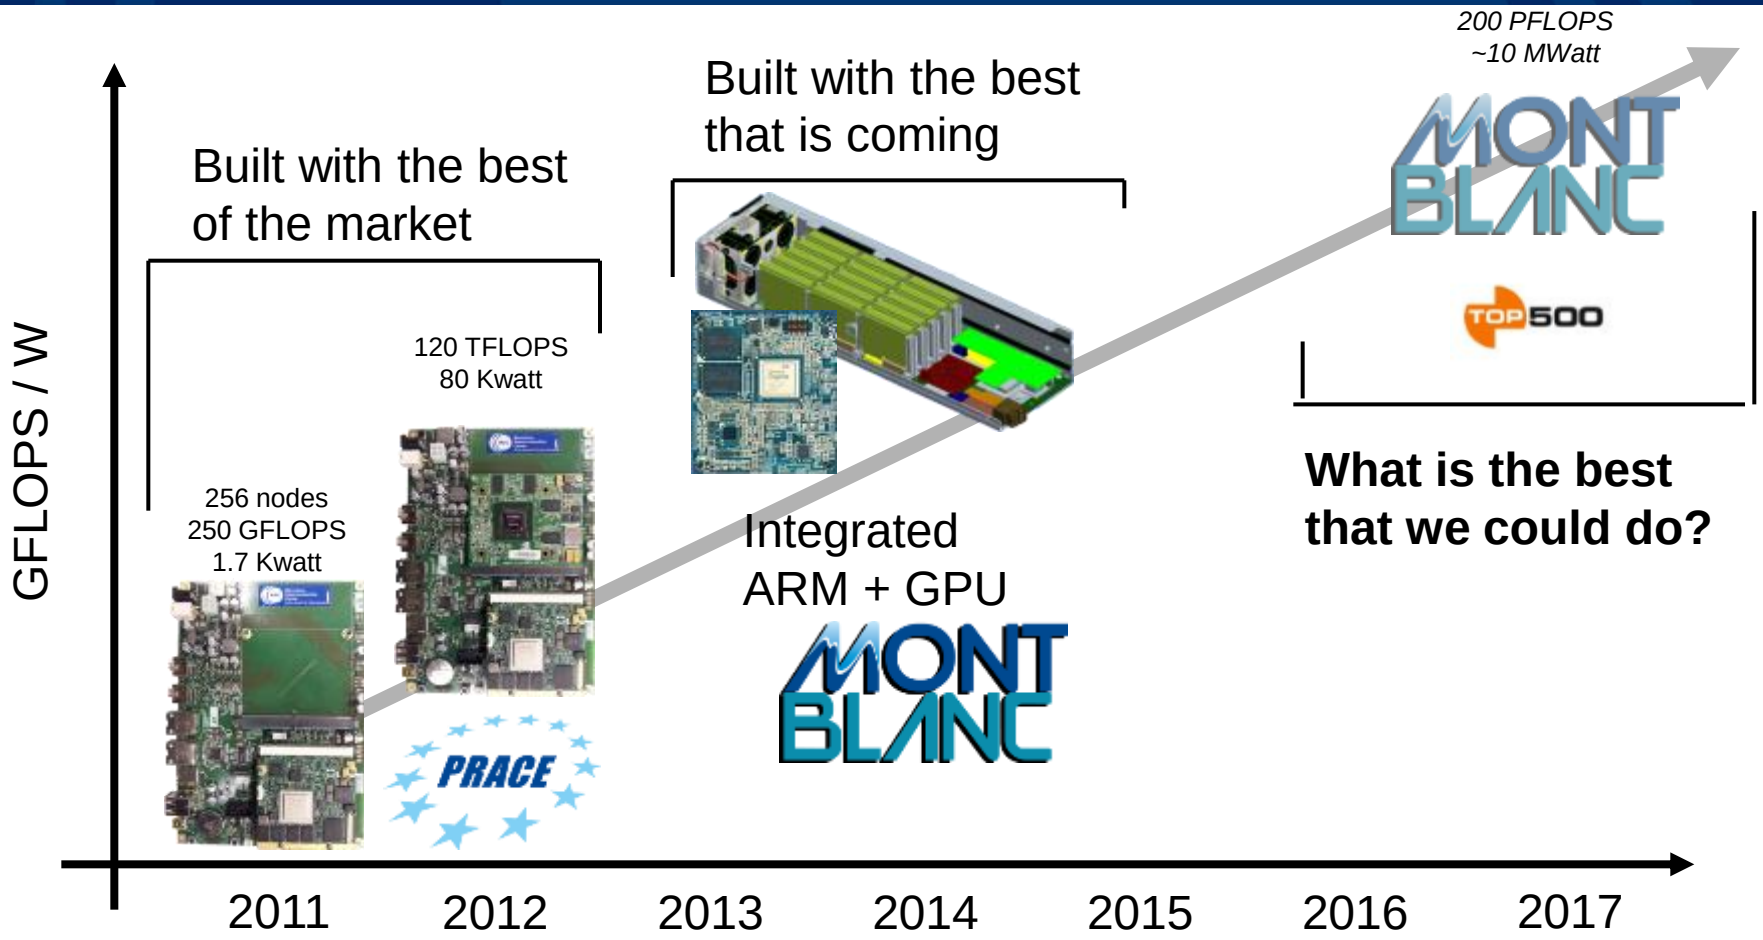

- Extend current mobile chips with the needed HPC features
  - Explore the use **vector architectures** in mobile accelerators (vector processor ARM-based, 15+ Teraflops chip, 150 watts)… unique opportunity for Europe
  - One design for all market segments: mobile, data centers, supercomputers

Mont-Blanc HPC Stack for ARM

Industrial applications

Pharmacelera, MUREX, DASSAULT AVIATION, AVL, Termo Fluids, Cenaero, ROLLS ROYCE

Applications

GENCI, CINECA, UNI GRAZ, HLRS, University of BRISTOL

System software

ARM, ETH zürich, Inria, BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación, JÜLICH FORSCHUNGSZENTRUM, Bull atos technologies

Hardware

ARM, BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación, Bull atos technologies, cea, UC UNIVERSIDAD DE CANTABRIA

BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación

World Top 20 machines (status November 2017)

EU not in HPC world leaders

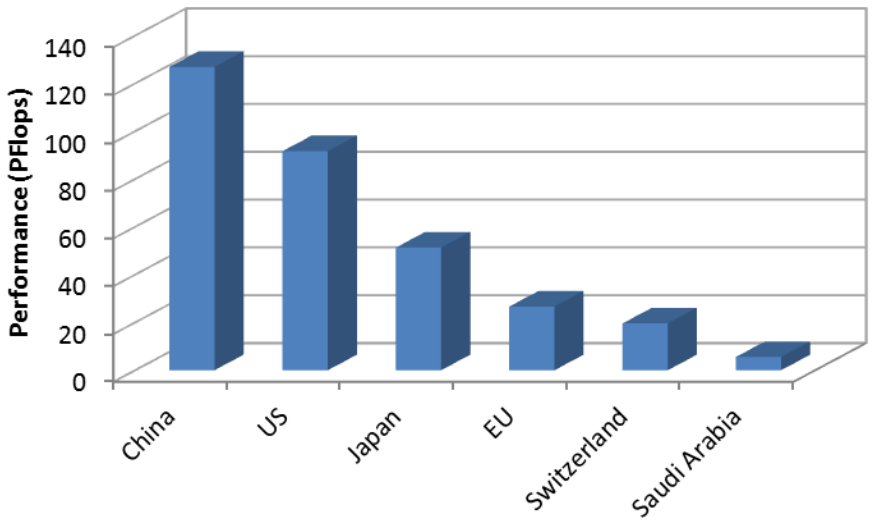

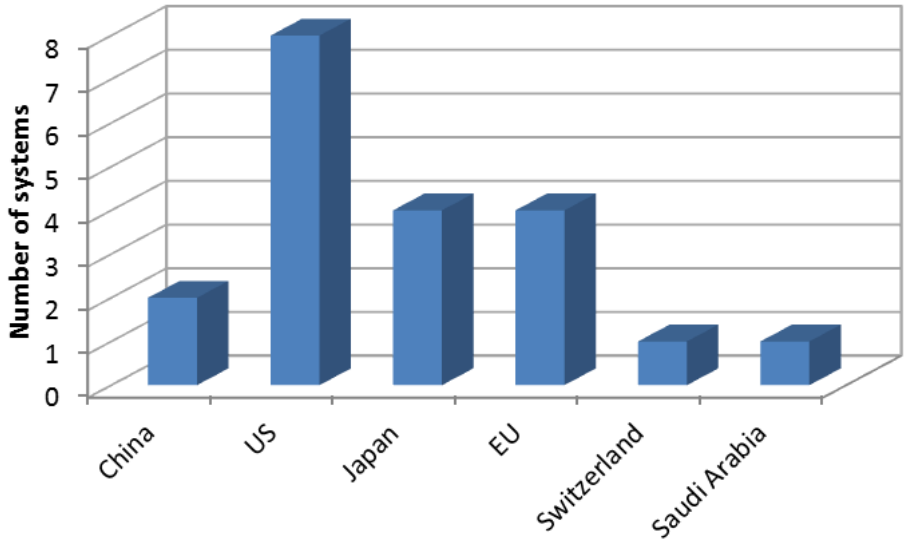

Europe has only 4 machines in world top 20

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

BSC and the European Commission

Final plenary panel at ICT - Innovate, Connect, Transform conference, 22 October 2015 Lisbon, Portugal.

The transformational impact of excellent science in research and innovation

"Europe can develop an exascale machine with ARM technology. Maybe we need an consortium for HPC and Big Data".

Seymour Cray Award Ceremony Nov. 2015
Mateo Valero

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

The European Commission and HPC

European Commission President Jean-Claude Juncker

"Our ambition is for Europe to become one of the top 3 world leaders in high-performance computing by 2020"

Paris, 27 October 2015

Vice-President Andrus Ansip

"I encourage even more EU countries to engage in this ambitious endeavour"

- Ministers from seven MS (France, Germany, Italy, Luxembourg, Netherlands, Portugal and Spain) sign a declaration to support the next generation of computing and data infrastructures

Digital Day Rome, 23 March 2017

The EuroHPC Declaration

Declaration signed in Rome, March 23rd, 2017 by:

France | Germany | Italy | Luxembourg | Netherlands | Portugal | Spain

Six more countries signed the Declaration:

Belgium | Slovenia | Bulgaria | Switzerland | Greece | Croatia

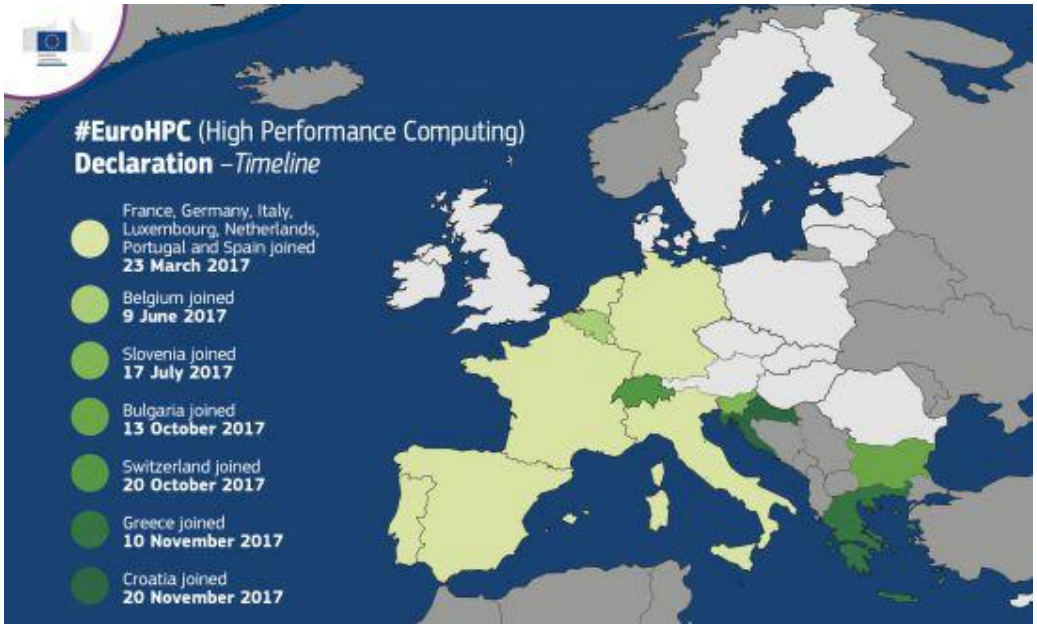

Agree to work towards the establishment of a **cooperation framework** - EuroHPC - for **acquiring and deploying an integrated exascale supercomputing infrastructure** that will be **available across the EU** for scientific communities as well as public and private partners

EuroHPC latest news:

European Processor Initiative

- Europa portail: (January 2018)
  http://europa.eu/rapid/press-release_IP-18-64_en.htm

European Commission - Press release

Commission proposes to invest EUR 1 billion in world-class European supercomputers

Brussels, 11 January 2018

The European Commission unveiled today its plans to invest jointly with the Member States in building a world-class European supercomputers infrastructure.

Supercomputers are needed to process ever larger amounts of data and bring benefits tothe society in many areas from health care and renewable energy to car safety and cybersecurity.

HPC timeline in H2020 LEIT/FET (indicative)

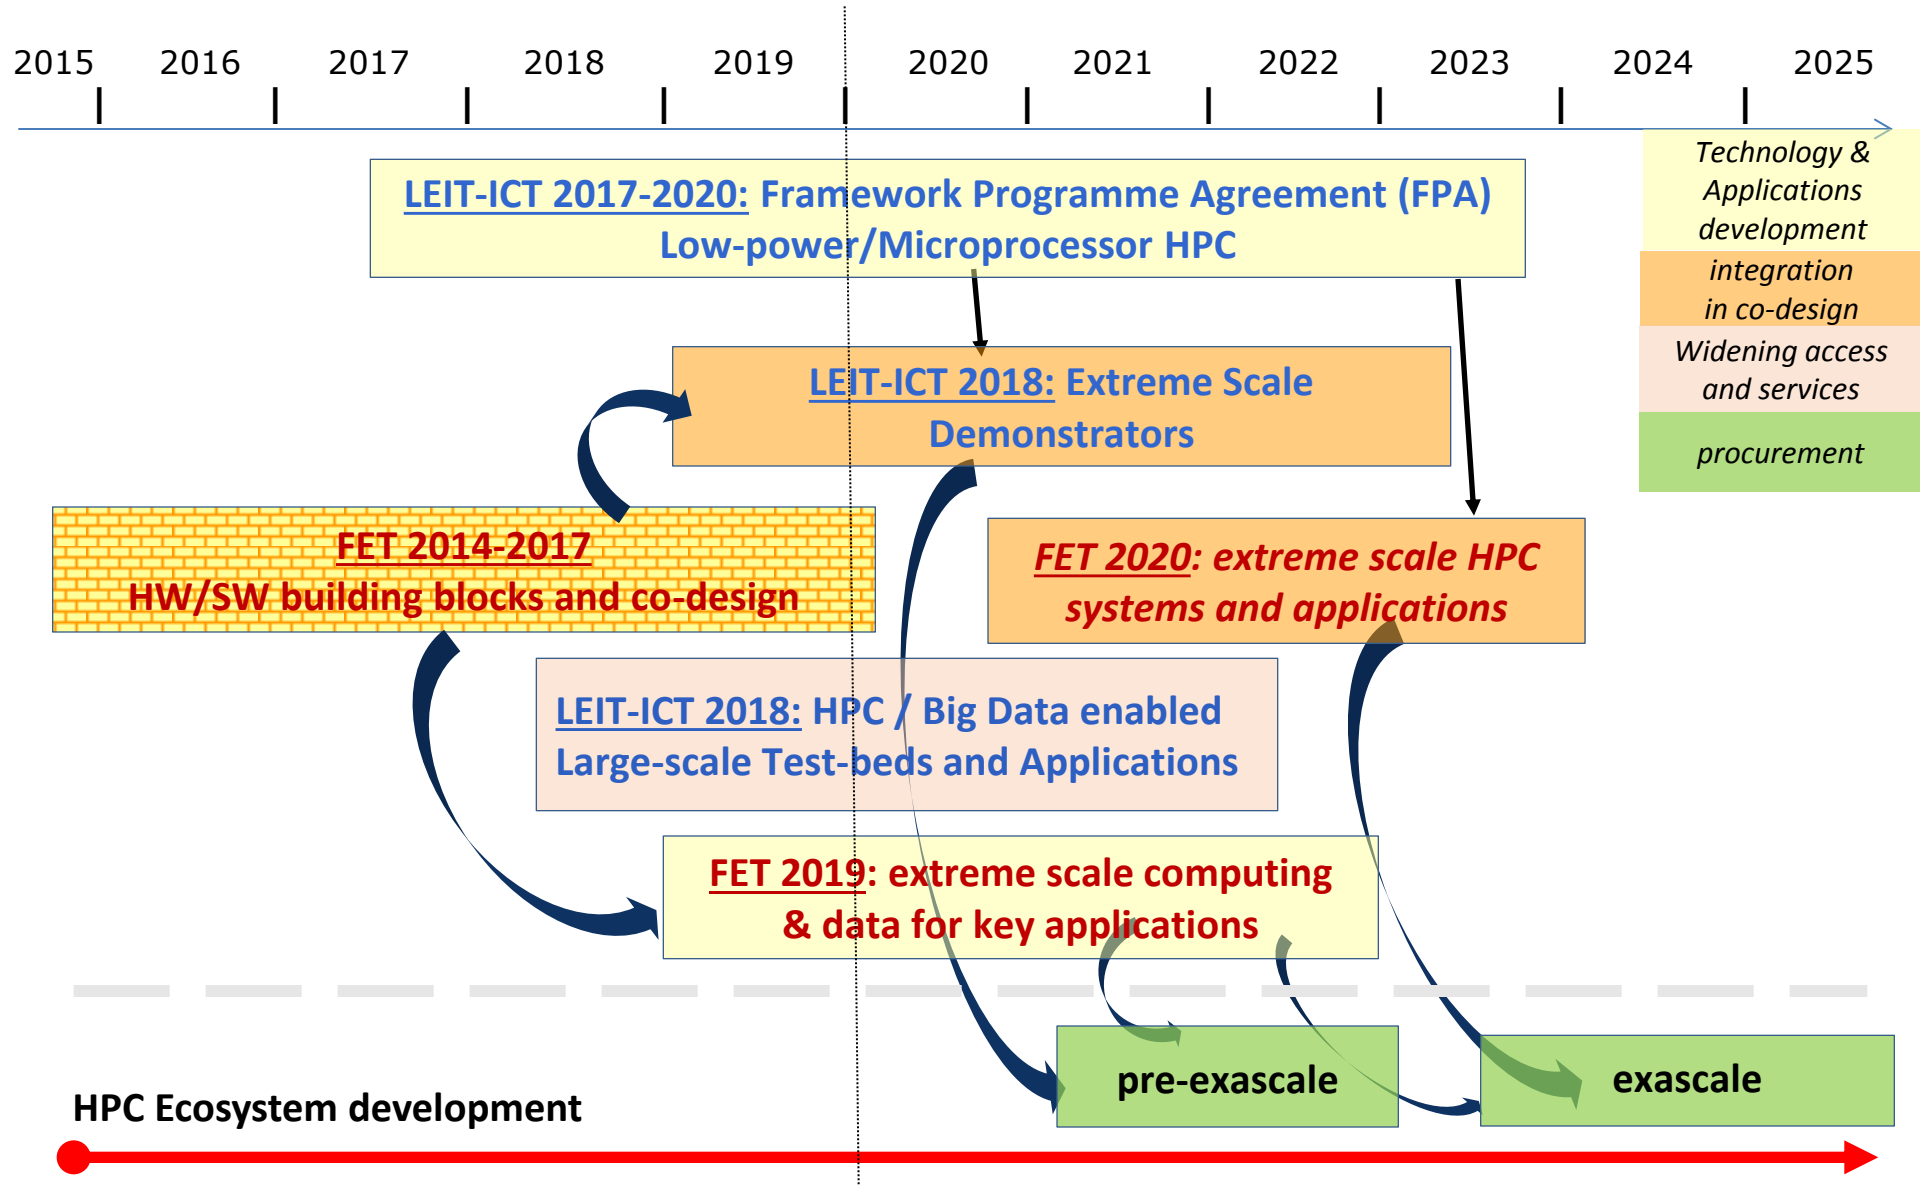

European Processor Initiative

EPI 23 partners, from research to industry from consortium to EU high tech fabless

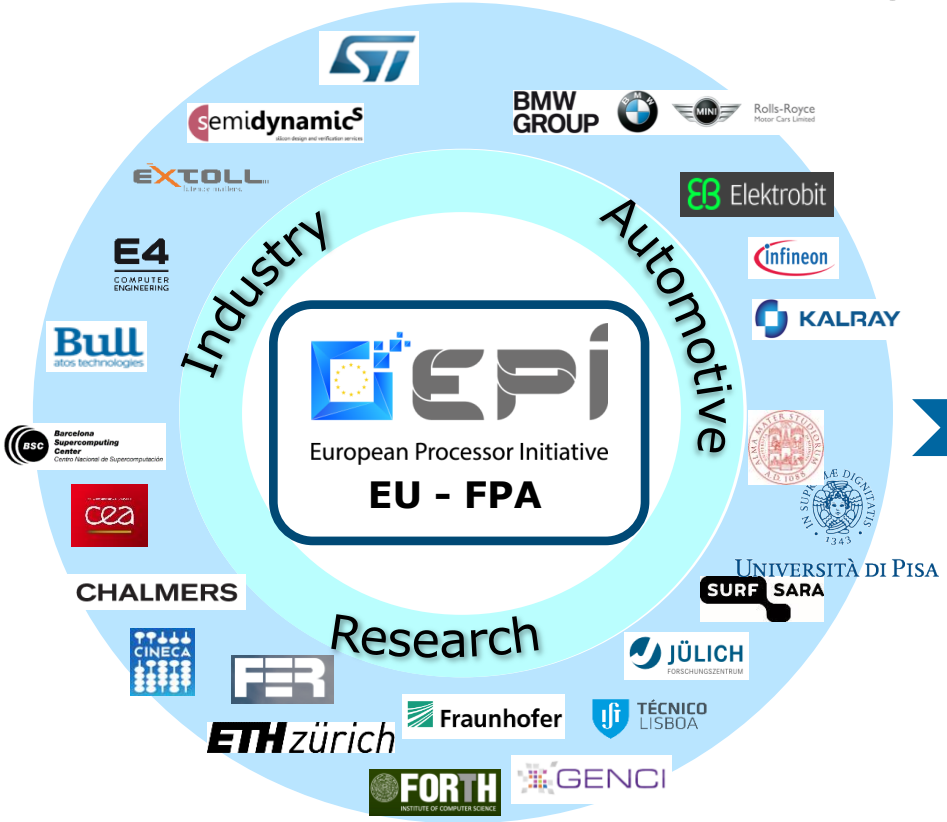

Semiconductor

Fabless company
Industrial hand of EPI
Incorporated by a couple EPI members and external investors

EPI Common Platform

1st EPI production

Three streams

- **General purpose and Common Platform**
  - ARM SVE or other candidates…
  - BULL: System integrator → chip integrator
- **Accelerator**
  - RISC-V
  - EU design: BSC, CEA, Chalmers, ETHZ, EXTOLL, E4, FORTH, Fraunhofer, IST, UNIBO, UNIZG, Semidynamics
- **Automotive**
  - Infineon, BMW…

EPI ROADMAP

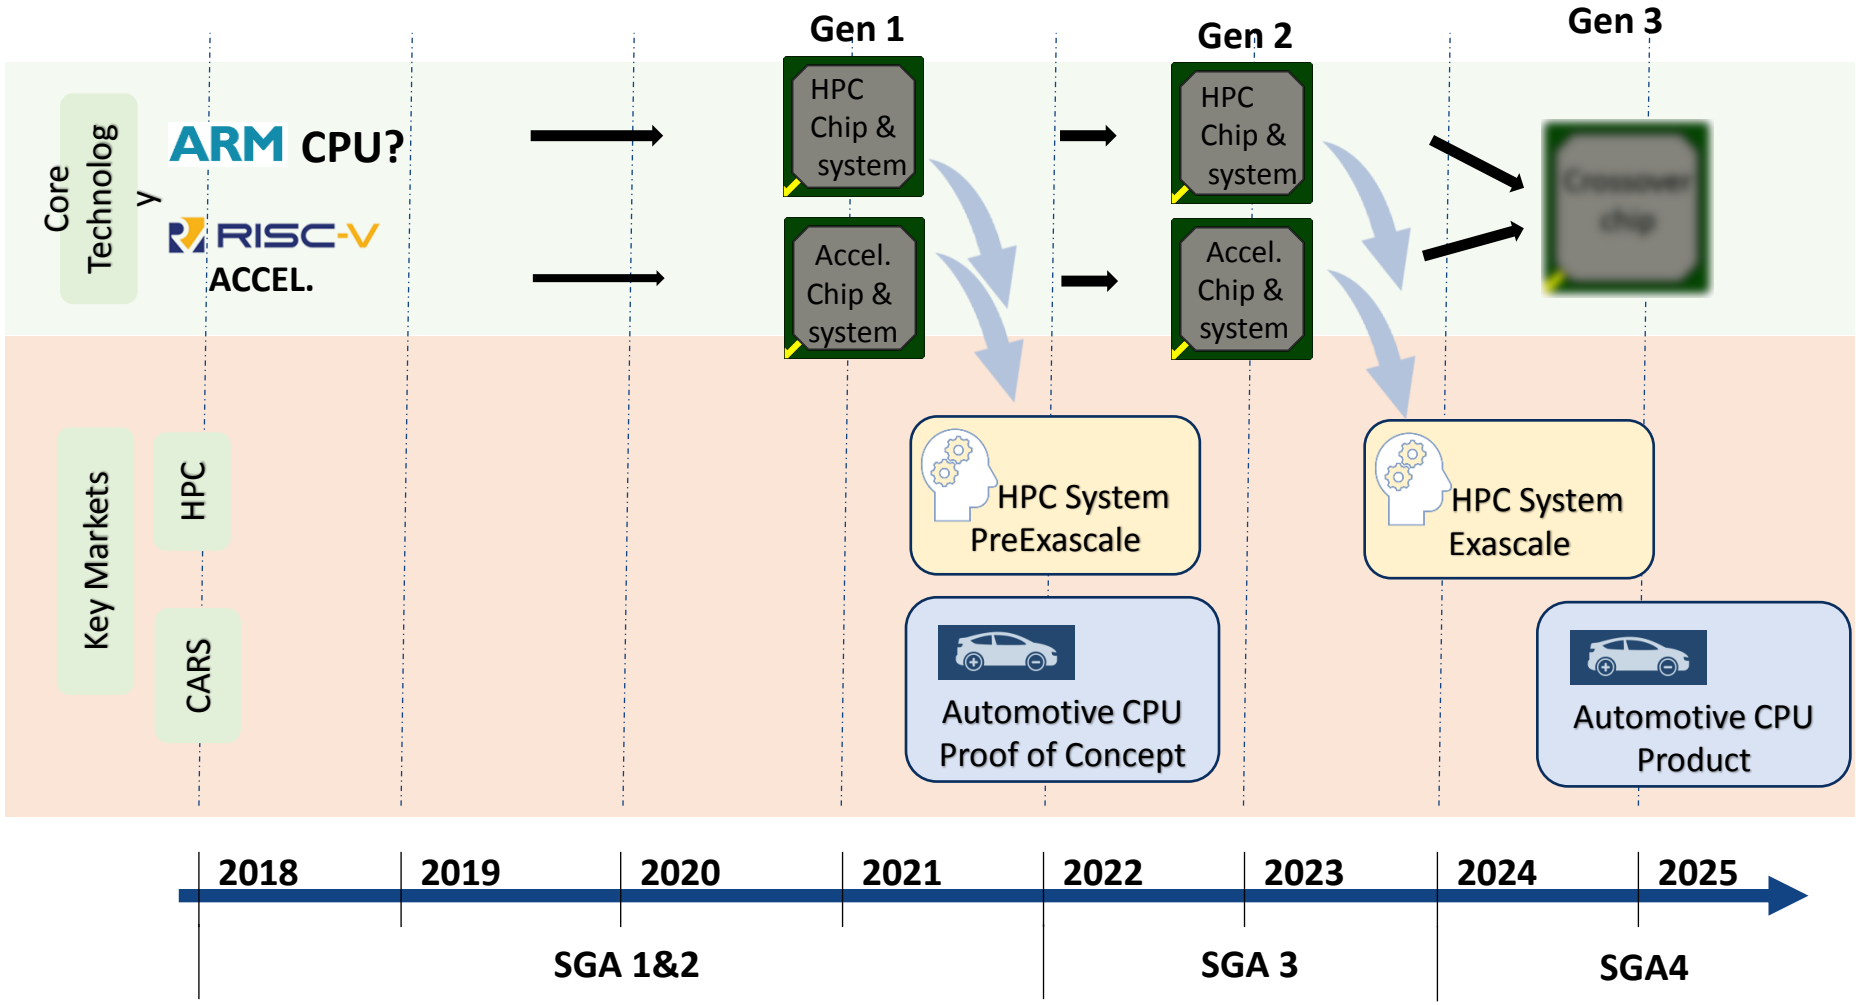

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

European Processor Initiative

RISC-V accelerator vision @ EPI

- **High throughput devices**
  - **Long Vectors** (a la Cray? A la Cyber205? ...)
    - Decouple Front end - Back end engines
    - Optimize memory throughput ([Command vector, 98])
    - Explicit locality management (long register file)
  - **ISA is important**
    - Decouple/hide again hardware details, reuse SW technologies (compilers, OS,…),
    - Specific instructions?
  - **“limited” number of control flows**
- **Hierarchical Acceleration**
  - **Nesting**
- **Low power**: ~ low voltage x ~ low frequency
- **MPI+OpenMP**
  - Task based, throughput oriented programming approach
  - Malleability in application + Dynamic resource (cores, power, BW) management
  - Intelligent runtimes & Runtime Aware Architectures
    - Architectural support for the runtime
- **Accelerator for ML**
  - Specialized “non Von-Neumann” compute and data motion engines (neural/stencil)
  - Tuned numerical precision

BSC and EPI

- EPI is a H2020 EU funded initiative restricted to the 23 original partners, selected according to EU rules
- EPI plans considering additional participants in future, provided resources will become available
- In EPI BSC is the leader of the Accelerator activities and contributor in the rest of the technical programme, including the Common Platform
- BSC will promote the EPI agenda within its vast academic network
- BSC is open to additional collaboration outside and within EPI to anyone in the world interested in producing RISC-V IP in Europe and especially in Barcelona
- Collaboration with the HPC global vendors will remain a key element of BSC strategy
- Everybody interested in RISC-V is welcome! Just come and talk to us…

BSC & The Global IT Industry 2018

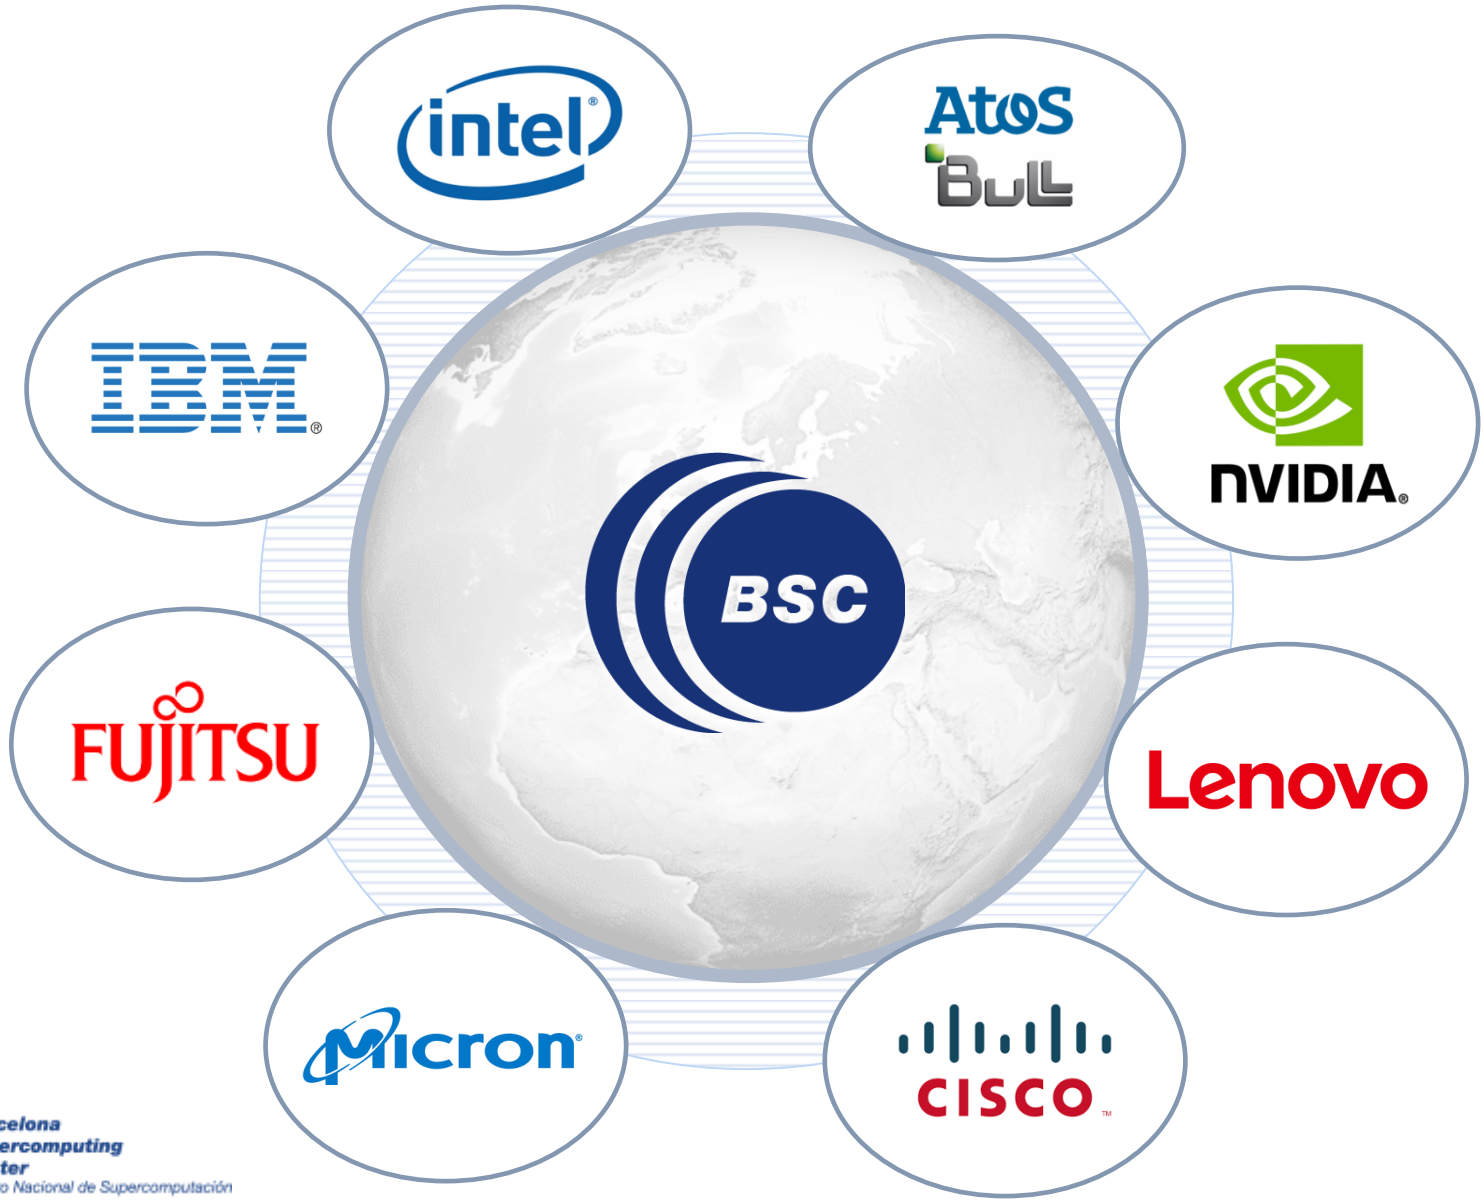

BSC is Hiring

European Processor Initiative

BSC is looking for talented and motivated professionals with expertise in the design and verification of IPs to be integrated into top-level HPC SoC designs. The immediate responsibilities of this group will be related to The European Processor Initiative.

Experienced professionals (Engineers and/or PhD holders) wanted for:

- **RTL/Microarchitecture**
- **Verification**
- **FPGA Design**

Find out more:

https://www.bsc.es/join-us/job-opportunities/103csrre

Or contact: rrhh@bsc.es

HR EXCELLENCE IN RESEARCH

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Mare Nostrum RISC-V inauguration 202X

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA



Thank you

mateo@bsc.es

Barna April 9th, 2018