



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación





18/09/2019



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



EXCELENCIA
SEVERO
OCHOA

El futuro MareNostrum 5 y el procesador europeo

Prof. Mateo Valero

Director



European Research Council

Supporting top researchers
from anywhere in the world



RES User Conference
18th and 19th of
September 2019

Patio de la Infanta,
Zaragoza



RED ESPAÑOLA DE
SUPERCOMPUTACIÓN



Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación

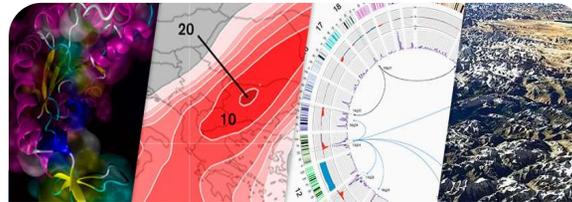
Ministerio de Ciencia, Innovación y Universidades
Departamento de Informática y Física
de la Universidad Zaragoza
Universidad Zaragoza

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%



Catalan Government

30%



Univ. Politècnica de Catalunya (UPC)

10%



Topic: High-Performance Computing in the **world** from 1980

• #	Institution	Count	Faculty
• 1	▶ Ohio State University	36.3	11
• 2	▶ Univ. of Illinois at Urbana-Champaign	33.7	19
• 3	▶ Polytechnic University of Catalonia	27.7	18
• 4	▶ Georgia Institute of Technology	26.8	22
• 5	▶ University of Minnesota	26.5	10
• 6	▶ University of Chicago	25.9	7
• 7	▶ Purdue University	22.5	15
• 8	▶ Indiana University	22.3	10
• 9	▶ ETH Zurich	17.9	5
• 10	▶ University of California - Berkeley	17.7	11

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

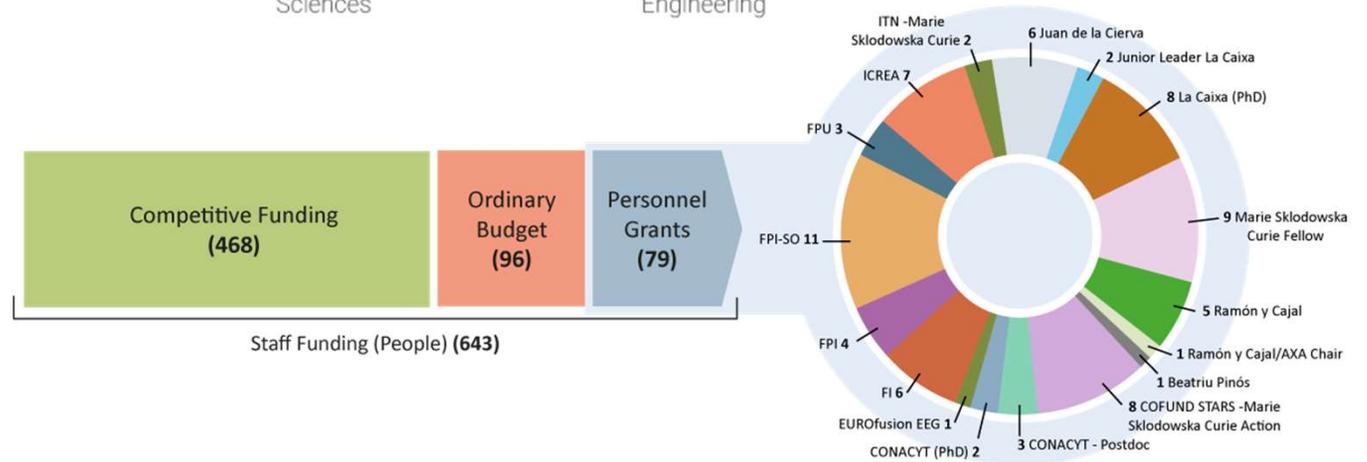
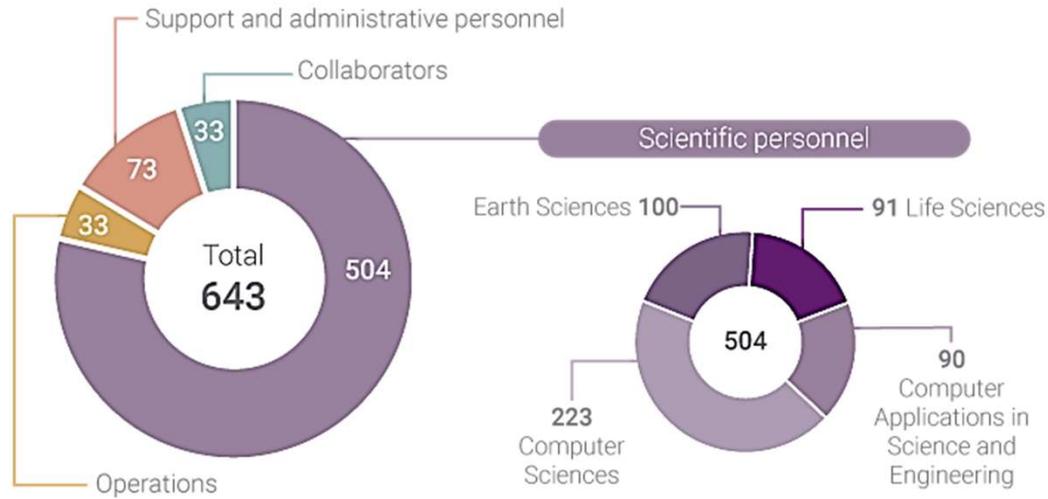
New technologies

People



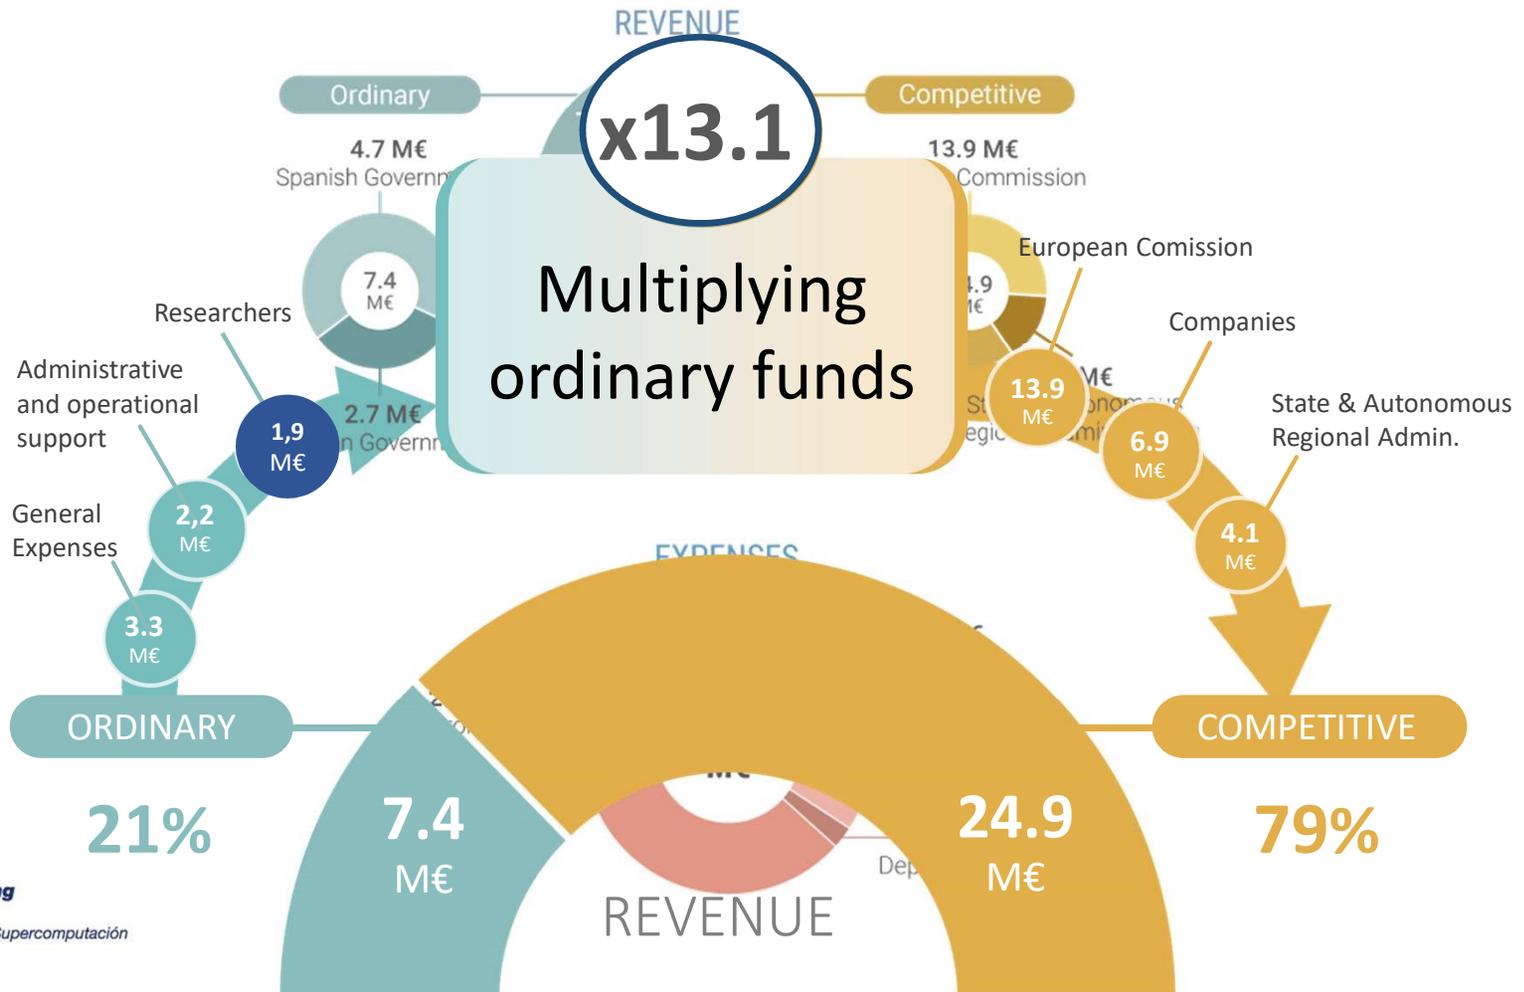
HR EXCELLENCE IN RESEARCH

Data as June 30, 2019



BSC Resources

2018 executed budget



TOP-10 Spanish Organizations in Horizon 2020

Legal name	EU Contribution (€)	Project Participations
CSIC	230,434,008 €	536
Tecnalia	106,426,784 €	239
Barcelona Supercomputing Center	76,524.698 €	132
Universitat Politècnica de Catalunya	59,475,312 €	158
Universitat Pompeu Fabra	56,816,732 €	109
ICFO	56,517,896 €	78
Universitat Autònoma de Barcelona	56,322,646 €	117
Universidad Politécnica de Madrid	55,004.745 €	155
Universitat Politècnica de València	53.806.967	139
ATOS Spain	52,902,517 €	148

La carrera mundial de los supercomputadores



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

Top 10, June 2019

Computer	Cores	Accelerators	Rmax [PFlop/s]	Rpeak [PFlop/s]	Power (MW)	Effeciency [GFlops/Watts]
IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband	2.414.592	2.211.840	148,6	200,8	10,1	14,7
IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband	1.572.480	1.382.400	94,6	125,7	7,4	12,7
Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway	10.649.600		93,0	125,4	15,4	6,1
TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000	4.981.760	4.554.752	61,4	100,7	18,5	3,3
Dell C6420, Xeon Platinum 8280 28C 2.7GHz, Mellanox InfiniBand HDR	448.448		23,5	38,7		
Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100	387.872	319.424	21,2	27,2	23,8	8,9
Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect	979.072		20,2	41,5	7,6	2,7
PRIMERGY CX2570 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, Infiniband EDR	391.680	348.160	19,9	32,6	1,6	12,1
ThinkSystem SD650, Xeon Platinum 8174 24C 3.1GHz, Intel Omni-Path	305.856		19,5	26,9		
IBM Power System S922LC, IBM POWER9 22C 3.1GHz, Dual-rail Mellanox EDR Infiniband, NVIDIA Tesla V100	288.288	253.440	18,2	23,0		

System Performance

- Peak performance of 200 petaflops for modeling & simulation
- Peak of 3.3 ExaOps for data analytics and artificial intelligence

Each node has

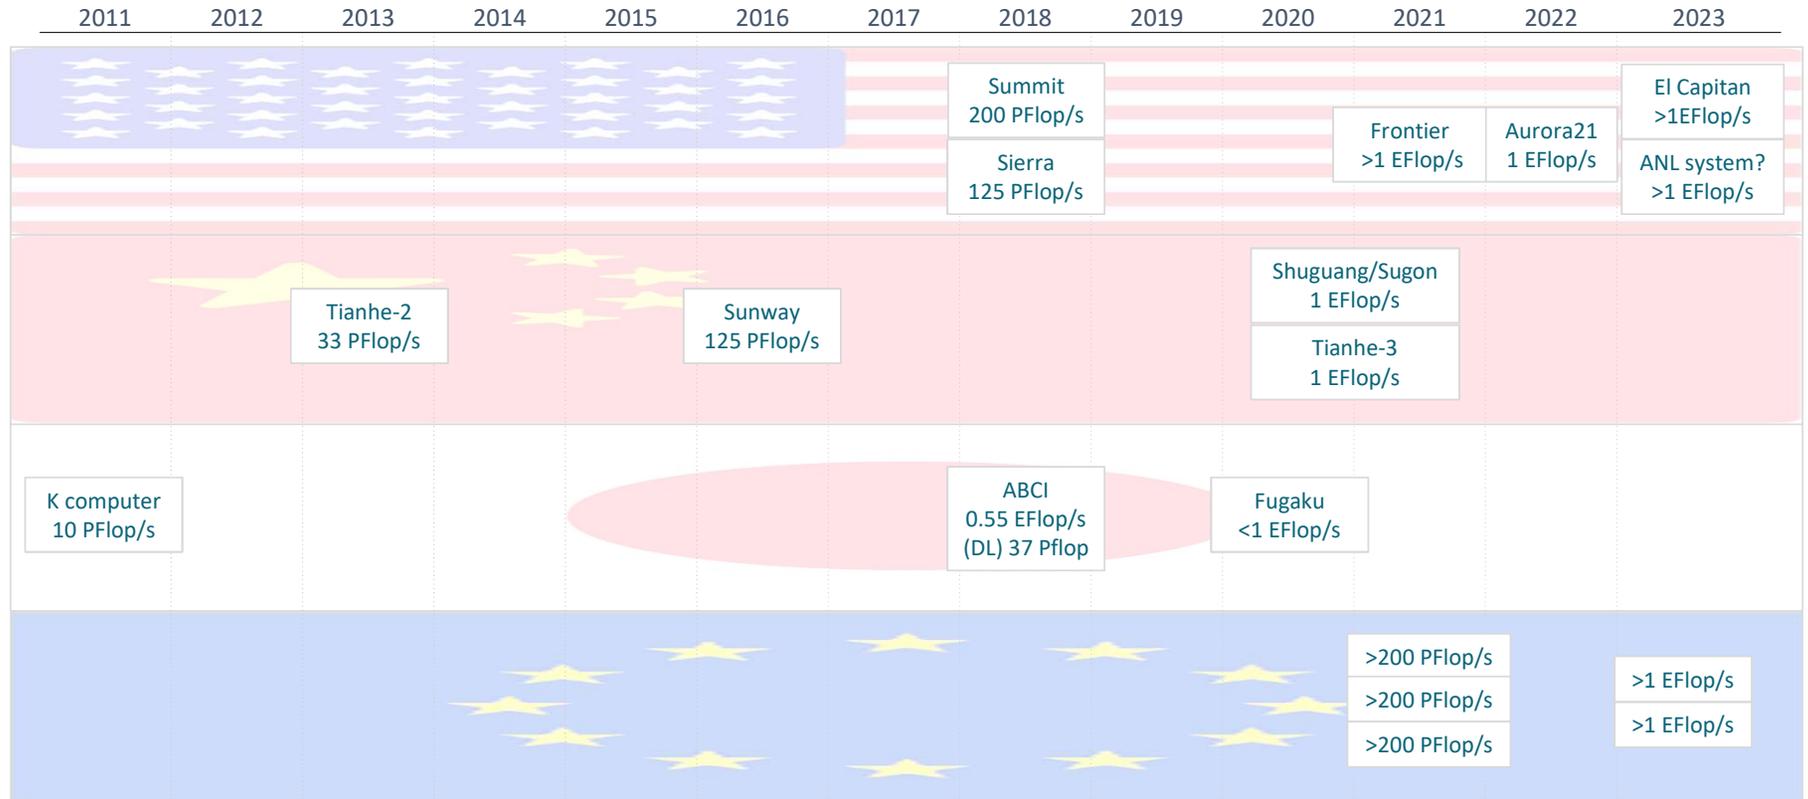
- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory
- 1.6 TB of NVMe memory

The system includes

- 4608 nodes
- Dual-rail Mellanox EDR InfiniBand network
- 250 PB IBM Spectrum Scale file system transferring data at 2.5 TB/s



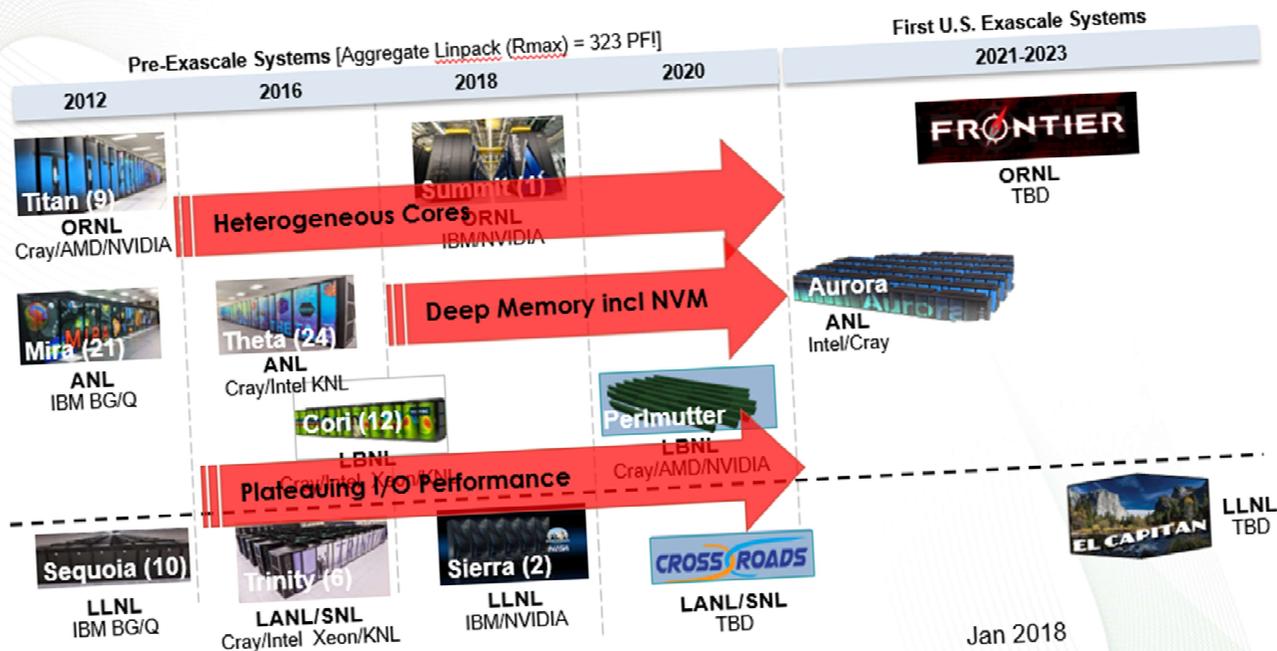
The race towards exascale



The Exascale Race – The US example

US Department of Energy (DOE) Roadmap to Exascale Systems

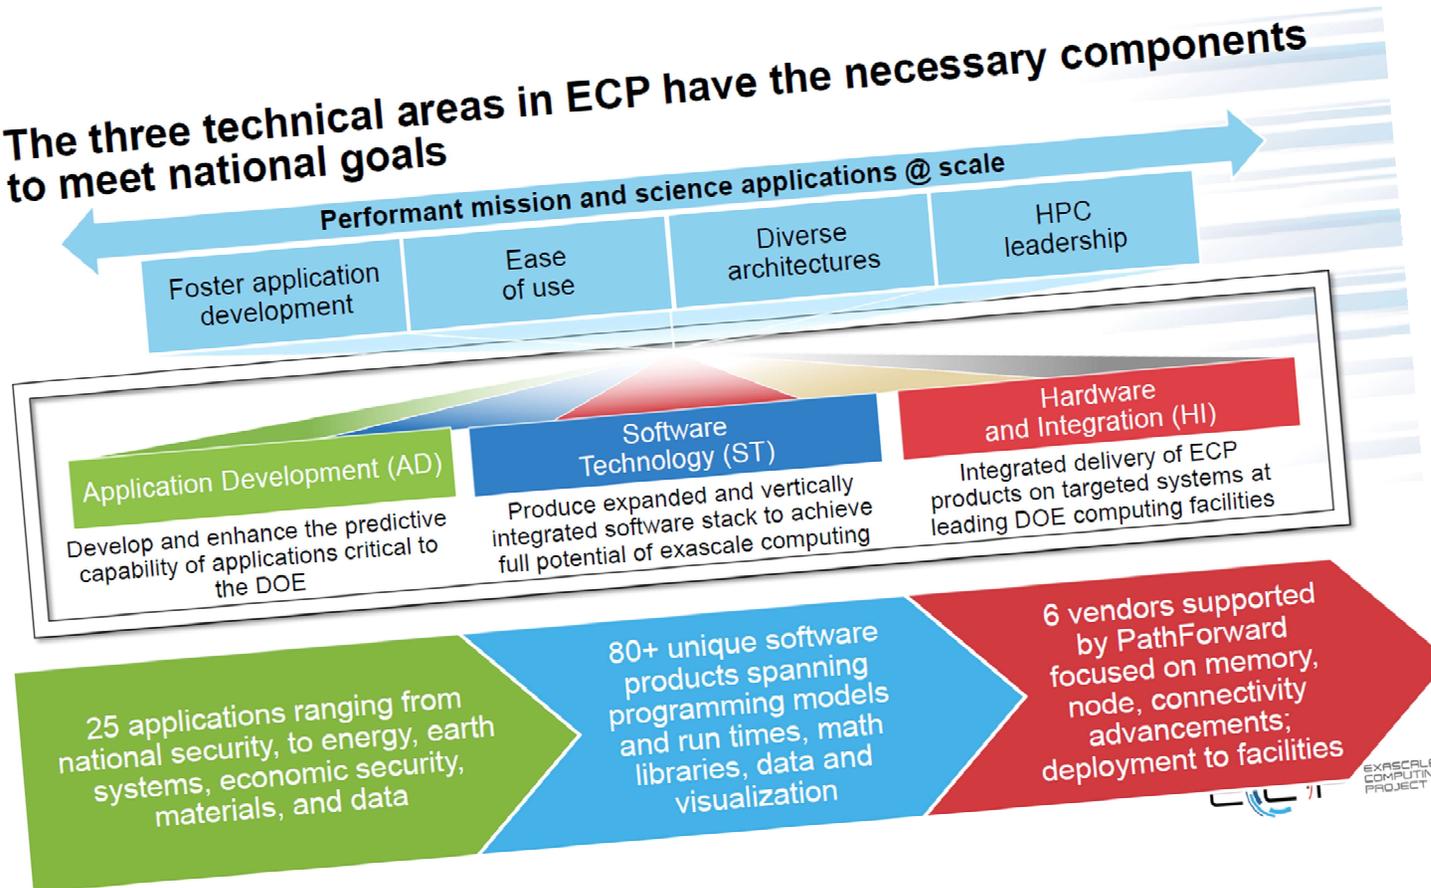
An impressive, productive lineup of *accelerated node* systems supporting DOE's mission



Jan 2018

The Exascale Race – The US example

The three technical areas in ECP have the necessary components to meet national goals



EuroHPC: Unifying European HPC technologies



EuroHPC-JU members:

Austria, Belgium, Bulgaria, Croatia, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, the Netherlands, Norway, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, Sweden, Switzerland 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

June 7: The new MareNostrum 5!!

La UE instalará en Barcelona uno de los tres superordenadores más veloces del continente

Apuesta por la ciencia

Llega una nueva generación de superordenadores a BCN

El Barcelona Supercomputing Center albergará esta joya de la tecnología a partir del 2020 || La Unión Europea aportará 100 millones de euros para el proyecto del MareNostrum 5

Barcelona recibe el espaldarazo de Europa con la adjudicación del nuevo supercomputador

La capital catalana acogerá en el BSC uno de los megaordenadores más potentes del continente, con 200.000 billones de operaciones por segundo

La Comisión Europea anunció ayer la adjudicación de uno de los tres nuevos supercomputadores de la UE a Barcelona, que se consolida así como una ciudad referente. El nuevo MareNostrum 5 será uno de los megaordenadores más potentes del continente. Capaz de procesar 200.000 billones de operaciones por segundo, será 18 veces más rápido que el actual MareNostrum 4. **TENDENCIAS 28, 29 Y EDITORIAL**



El superordenador más rápido de Europa recalará en Barcelona

La ciudad recibirá una inversión de la UE de 100 millones por albergar el supercomputador

Barcelona gana el supercomputador

Europa apuesta por la capital catalana para acoger uno de los megaordenadores más potentes del continente

Europa confía en España para «supercompetir» con EE.UU. y Asia

► La Comisión Europea elige Barcelona para instalar uno de los nuevos superordenadores, que contará con una inversión de 200 millones de euros

MareNostrum 5

A European pre-exascale supercomputer

- **200 Petaflops** peak performance (200×10^{15})
- **Experimental platform** to create supercomputing technologies “made in Europe”

Hosting Consortium:

Spain Portugal Turkey Croatia Ireland



|



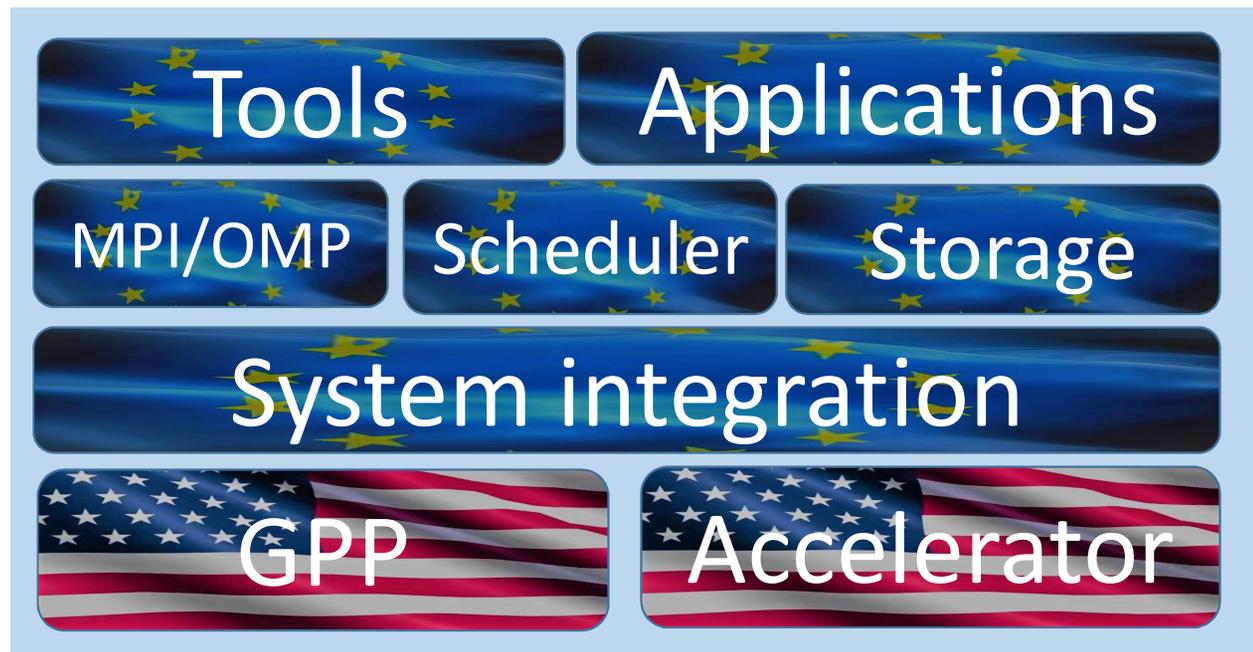
La carrera por la tecnología



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

Where Europe needs to be stronger

- ➔ Only 1 of the 10 most powerful HPC systems is in the EU
- ➔ HPC codes must be upgraded
- ➔ Vital HPC hardware elements are missing: general purpose processor and accelerators
- ➔ EU needs its own source of as many of the system elements as possible



BSC and the EC



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 needs to develop an entire domestic exascale stack from the processor all the way to the system and application software",
Mateo Valero,
Director of
Barcelona
Supercomputing
Center*

Director of Barcelona Supercomputing Center, Mateo Valero, makes a pledge for developing a strong HPC ecosystem.

Published on 12/04/2016

Europe has the competence and skills to engage in the global competition towards Exascale Supercomputing. To fully benefit from the opportunities of the digital single market, Europe must strengthen the fundamental research on which digital transformation is based and build a stronger European High Performance Computing (HPC) ecosystem.

In a [quest blog post](#) on Commissioner Günther Oettinger's [website](#) Mateo Valero stresses the need for Europe to join the race towards Exascale supercomputing. According to him, there is an open window of opportunity for the High Performance Computing (HPC) development that would stimulate scientific breakthroughs and have tremendous impact on society and industry.



Share

ARM-based prototypes at BSC



2011
Tibidabo

ARM multicore

2012
KAYLA

ARM + GPU
CUDA on ARM

2013
Pedraforca

ARM + GPU
Infiniband
RDMA

2014
Mont-Blanc

Single chip ARM+GPU
OpenCL on ARM GPU



Mont-Blanc HPC Stack for ARM



Industrial applications



Applications



System software



Hardware



Why Europe needs its own processor

- ➔ Processors now control almost every aspect of our lives
- ➔ **Security** (back doors, etc.)
- ➔ Possible **future restrictions on exports to EU** due to increasing protectionism
- ➔ **A competitive EU supply chain** for HPC technologies will create jobs and growth in Europe

A group of researchers showed how a Tesla Model S can be hacked and stolen in seconds using only \$600 worth of equipment



NSA May Have Backdoors Built Into Intel And AMD Processors



A jet sale to Egypt is being blocked by a US regulation, and France is over it



Google 'suspends some business with Huawei' after US blacklist



Amazon and Super Micro urge Bloomberg to retract 'unsupported' spy chip report



Car Hacking Remains a Very Real Threat as Autos Become Ever More Loaded With Tech

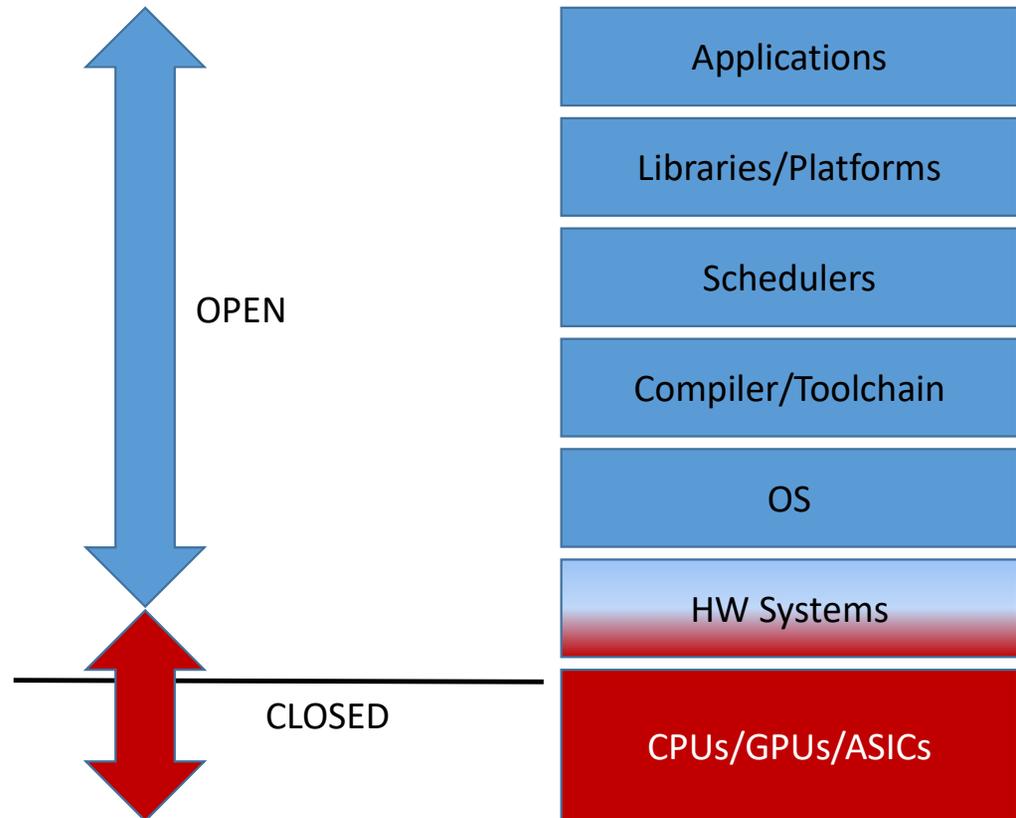


The US Cloud Act v The EU's GDPR - Data Privacy & Security

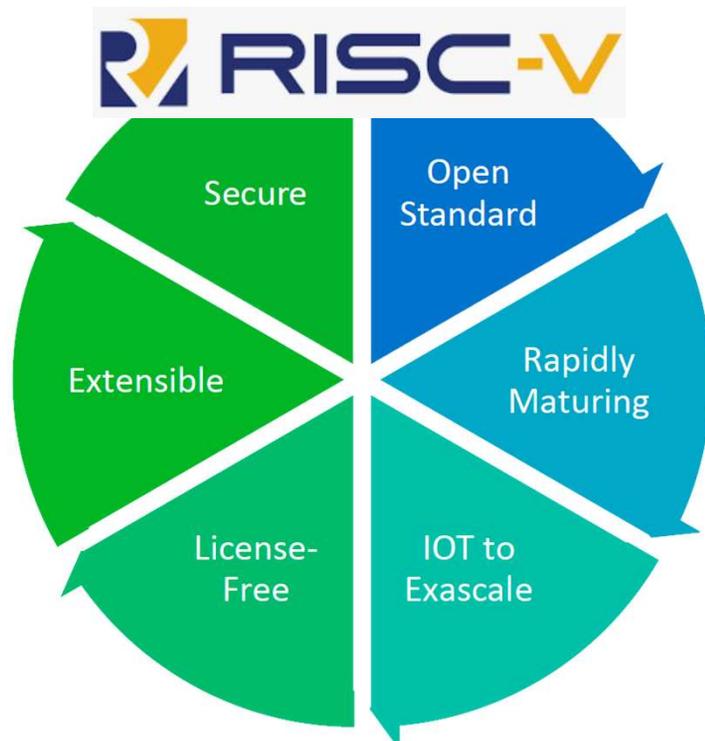


HPC Today

- Europe has led the way in defining a common open HPC software ecosystem
- **Linux** is the de facto standard OS despite proprietary alternatives
- Software landscape from Cloud to IoT already enjoys the benefit of open source
- Open source provides:
 - A common platform, specification and interface
 - Accelerates building new functionality by leveraging existing components
 - Lowers the entry barrier for others to contribute new components
 - Crowd-sources solutions for small and larger problems
- **What about Hardware and in particular, the CPU?**



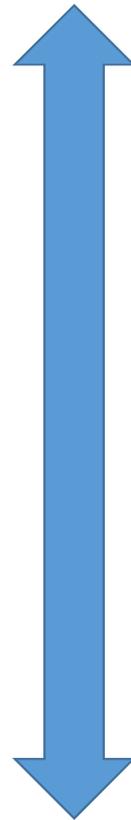
RISC-V is democratising chip-design



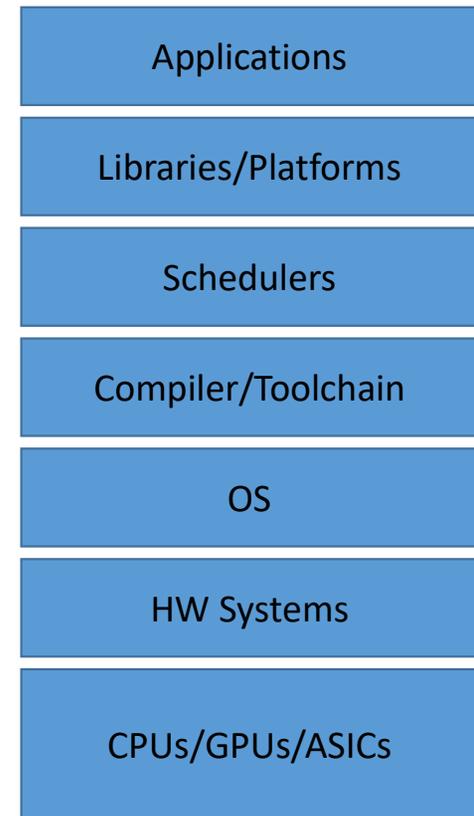
- ➔ More and more global IT actors are adopting RISC-V architectures to be vendor independent
 - ➔ Google
 - ➔ Amazon
 - ➔ Western Digital
 - ➔ Alibaba
- ➔ And of course the entire IoT ecosystem for lower performance, lower energy applications.
- ➔ Major opportunity for ICT industry also in Spain

HPC Tomorrow

- Europe can lead the way to a completely **open SW/HW stack for the world**
- RISC-V provides the open source hardware alternative to dominating proprietary non-EU solutions
- Europe can achieve complete technology independence with these foundational building blocks
- Currently at the same early stage in HW as we were with SW when Linux was adopted many years ago
- **RISC-V can unify, focus, and build a new microelectronics industry in Europe.**



OPEN



The European Processor Initiative

- In the same way BSC led the development of ARM processors for HPC in the various MontBlanc projects, now it leads the RISC-V HPC accelerator development in EPI
- EPI is a 100% funded EuroHPC project (120 M€) to develop European processor technology by 2022
- BSC was the original initiator of EPI and most active proponent in the scientific and technical community
- EPI is led by Atos/Bull with 28 partners from leading HPC industrial and academic centres

EPI Partners



Rolls-Royce
Motor Cars Limited



ALMA MATER STUDIUM
UNIVERSITÀ DI BOLOGNA



CHALMERS



UNIVERSITÀ DI PISA



UNIVERSITY OF ZAGREB
FACULTY OF ELECTRICAL
ENGINEERING
AND COMPUTING



PROVE & RUN



Exascale supercomputing initiative at BSC

- ➔ Ground floor opportunity to design and build a European supercomputer at the best supercomputing center in Europe!
- ➔ The open-source hardware opportunity
- ➔ RISC-V HPC accelerator: from concept to implementation
- ➔ Latest silicon technologies: 7nm, 5nm and 3 nm
- ➔ Working with industrial and academic partners
- ➔ HPC, automotive, bio, meteorological and other workloads

Barcelona desarrollará el chip de los superordenadores europeos

La CE financia una tecnología clave para la soberanía informática del continente

El proyecto del chip europeo estará liderado por Barcelona

LAVANGUARDIA

Barcelona desarrolla el chip de los futuros superordenadores europeos

El superordenador MareNostrum 5 se lanzará a la conquista de procesadores y chips 'made in Europe'

El MareNostrum 5 incluirá una plataforma para crear chips europeos

El próximo superordenador contribuirá al desarrollo de tecnologías íntegramente desarrolladas en Europa

El súperordenador presentará batalla en la fabricación de chips y procesadores europeos

The future is wide open!

- ➔ There is an urgent need, from mobile phones to supercomputers: more compute at lower power
- ➔ The RISC-V ecosystem is in the nascent period where it can become the de facto open hardware platform of the future
 - ➔ An opportunity for Europe to lead the charge to creating a full stack solution for everything, from supercomputers down to IoT devices
- ➔ Our main aim: create European chips that meet the needs of future European and global markets across HPC, cloud, automotive, mobile to IoT
- ➔ **This is the framework for the Exascale Supercomputing Initiative at BSC**



How to implement this “Open Future World”?

➔ The BSC launches **LOCA**, the new **European Laboratory for Open Computer Architecture**, a joint long-term initiative to promote a vibrant RISC-V ecosystem, HQ in Southern Europe, supported by:

➔ The European Commission



➔ The BSC trustees



➔ The European Academics

UPC, Cantabria, Chalmers, Rome Sapienza, Zagreb, Forth, ETH, EPFL.....

➔ The main IT worldwide companies

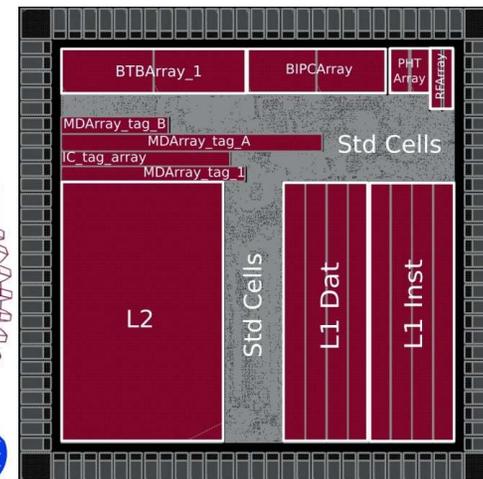
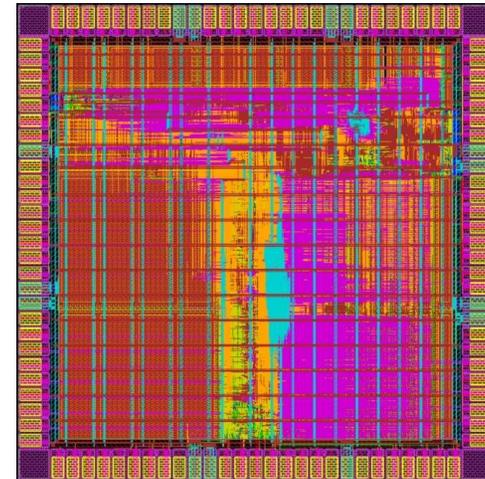


➔ The digital technology industry in Spain?



First Lagarto Tapeout

- Target design:
 - Simple in-order core with 5 stages, single issue
 - 16KB L1 caches, 64KB L2 cache, TLB
 - Memory controller on the FPGA side
 - FPGA – ASIC communication via packetizer
 - Debug ring via JTAG
 - Target technology: TSMC 65nm
 - Design fits in the total area budget of 2.5mm²
 - Submitted for fabrication in May 2019
- Collaborative project with different teams:
 - RTL Design: Lagarto (BSC + CIC-IPN)
 - Verification (BSC)
 - Logic Synthesis (UPC + BSC)
 - Physical design (IMB-CNM + BSC)
 - Tapeout and bringup (IMB-CNM + BSC)

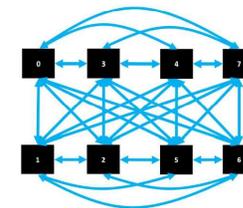
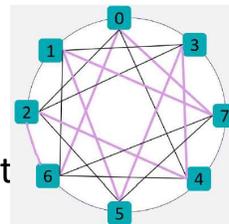
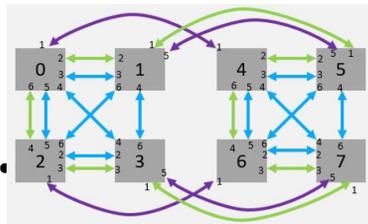


The HPC Future is Wide Open!

- Can open source hardware play a big role like open source software?
- How do we build flexible accelerators?
- Can we leverage commodity components and merge them with HPC systems?
- Can we jumpstart HPC hardware development in Europe?

An Open Path to the Future

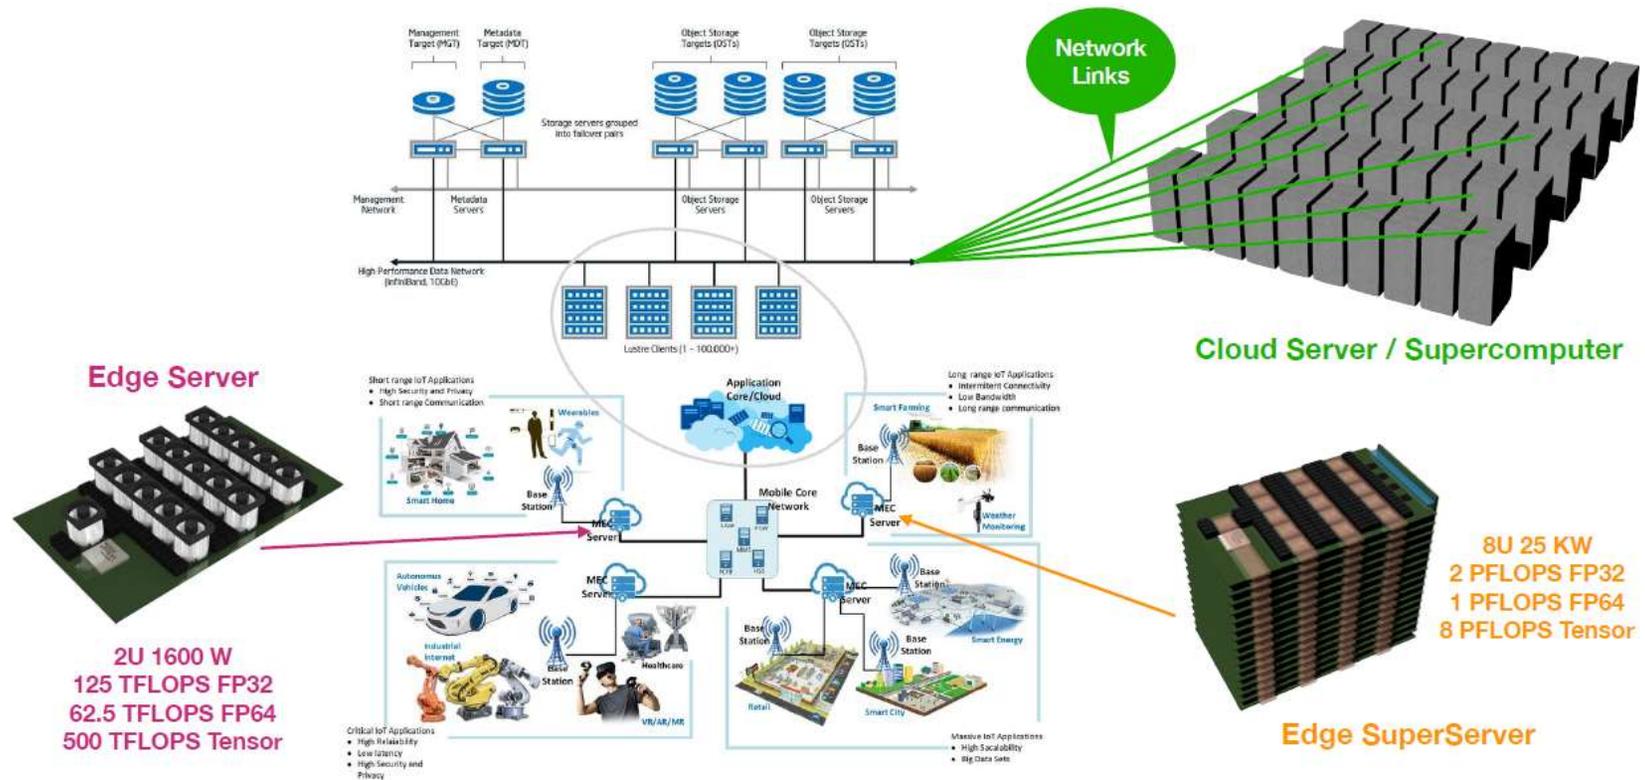
- We can change the balance of Host CPUs to Accelerators
 - OCP Accelerator Module (OAM): 1-8 or more accelerators per CPU
 - Partial to All-to-All OAM communication topologies:



enabling a variety of applications, such as HPC, AI, ML, DL, etc.

- FPGAs: Accelerators, prototypes and emulators
- MareNostrum Experimental Exascale Platform

From IoT, Edge Computing, Clouds to Supercomputers

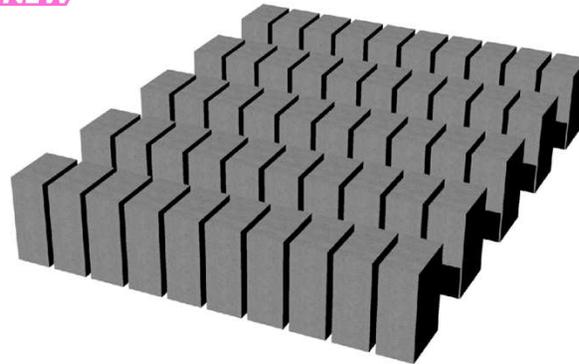


What does a 30 MW ExaFLOP SC look like?...

We have some ideas, come join the fun!

- 64 cabinets: 1.0 Exaflops
- Cabinet: 16 Petaflops, 400 KW (water cooled)
- 256 nodes, 24,576 cores
- 128 to 512 Terabytes DRAM
- 0.1 Byte/flop bandwidth ratio
- 40 Gflops/W efficiency
- 7nm initial, 5 and 3 nm follow-on designs

THE
NEW
NEW



BSC is hiring... Creating high value job opportunities in Spain

BSC is looking for talented and motivated professionals with expertise in the design and verification of IPs to be integrated into a European HPC accelerator. The design is based on a RISC-V architecture. This is a NEW project to build an energy efficient Exascale system.

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

- ➔ RTL / Microarchitecture
- ➔ Verification
- ➔ FPGA design
- ➔ Simulation
- ➔ Software: compilers/OS/RT

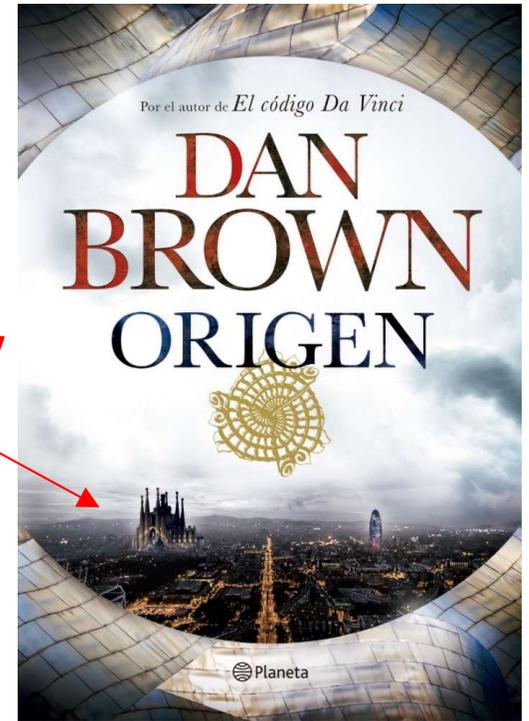
RISC-V has the opportunity to be like Linux. It would be global and go beyond Airbus and Galileo!



MareNostrum RISC-V inauguration 2021



MN6-RISC-V
2025???





**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



EXCELENCIA
SEVERO
OCHOA

Thank you

mateo.valero@bsc.es