Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

RISC-V y hardware abierto: una oportunidad y un reto para la UE

Prof. Mateo Valero

LOCA
European Laboratory for Open Computer Architecture

06/01/2020

RISC V Workshop. Universidad Autónoma Barcelona

Technological Achievements

- **Transistor (Bell Labs, 1947)**
  - **DEC PDP-1 (1957)**
  - **IBM 7090 (1960)**
- **Integrated circuit (1958)**
  - **IBM System 360 (1965)**
  - **DEC PDP-8 (1965)**
- **Microprocessor (1971)**
  - **Intel 4004**

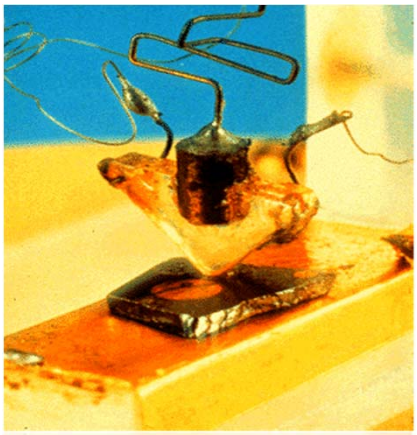

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

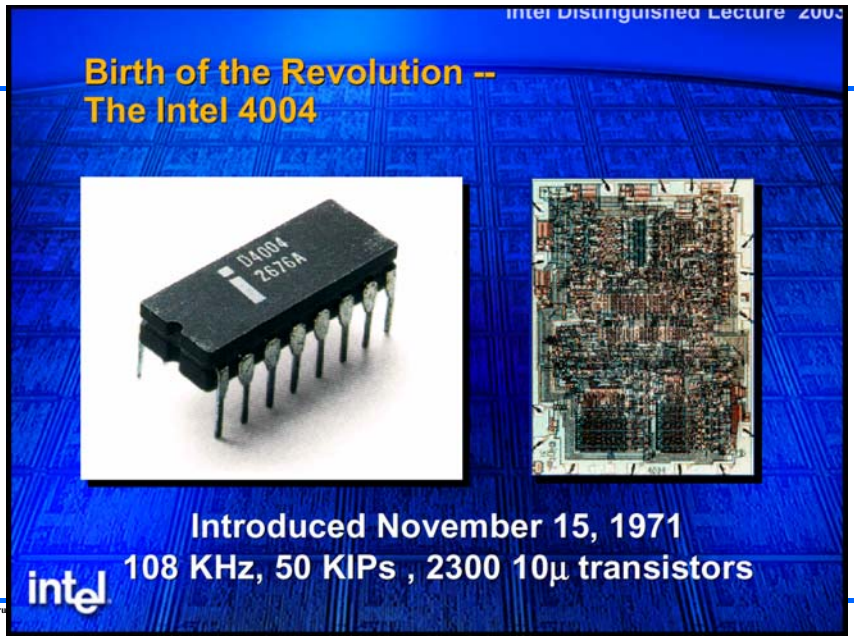
Intel Distinguished Lecture 2003
Birth of the Revolution --
The Intel 4004
Introduced November 15, 1971
108 KHz, 50 KIPs , 2300 10µ transistors
intel


ANNOUNCING TESLA V100
GIANT LEAP FOR AI & HPC
VOLTA WITH NEW TENSOR CORE
21B xtors | TSMC 12nm FFN | 815mm²
5,120 CUDA cores
7.5 FP64 TFLOPS | 15 FP32 TFLOPS
NEW 120 Tensor TFLOPS
20MB SM RF | 16MB Cache
16GB HBM2 @ 900 GB/s
300 GB/s NVLink

Professor Tomas Lang

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

OAK RIDGE National Laboratory

summit

System Overview

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

6

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC-CNS objectives

Supercomputing services to Spanish and EU researchers

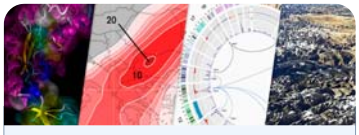

R&D in Computer, Life, Earth and Engineering Sciences

PhD programme, technology transfer, public engagement

BSC-CNS is a consortium that includes

Spanish Government	60%	
Catalan Government	30%	Generalitat de Catalunya Departament d'Empresa i Coneixement
Univ. Politècnica de Catalunya (UPC)	10%	

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

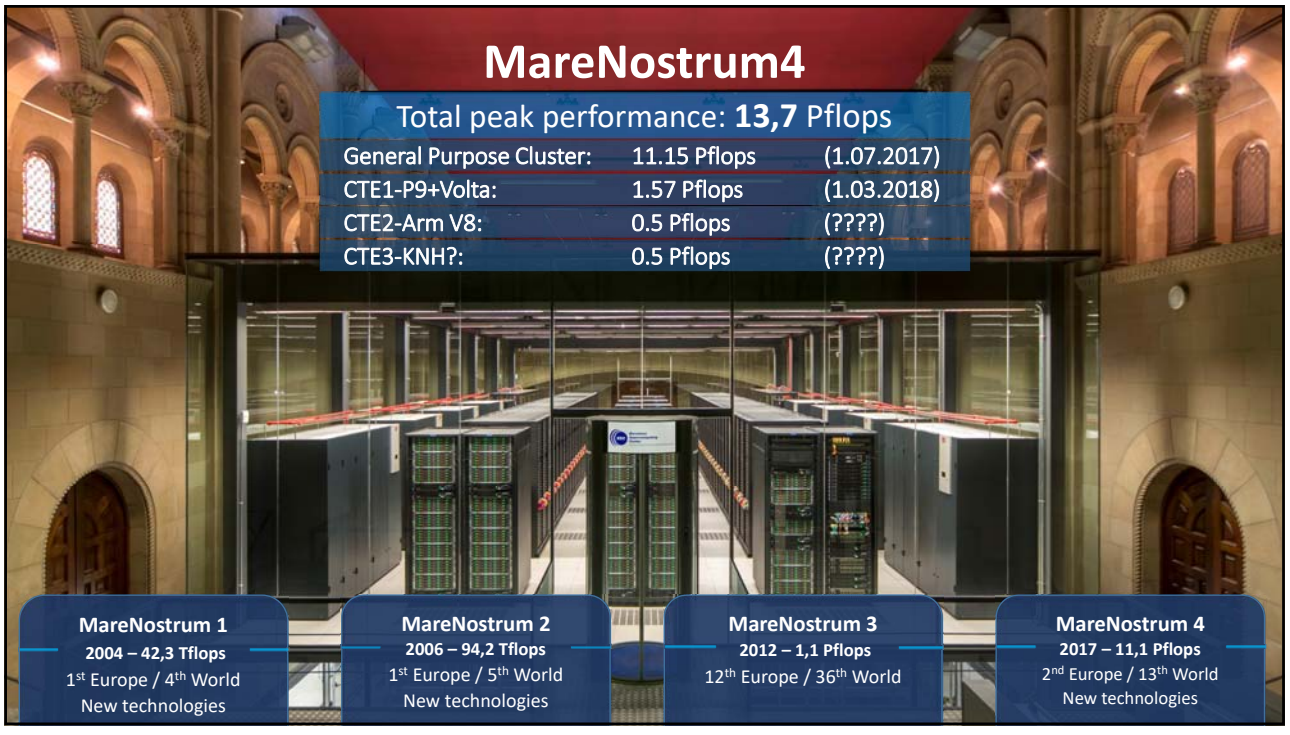

People

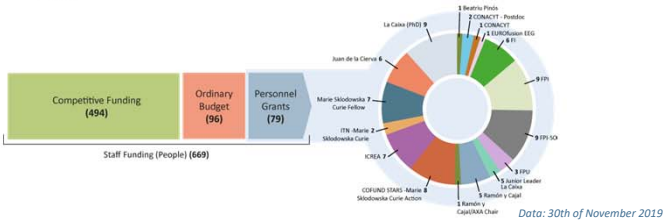

Data: 30th of November 2019

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

TOP-10 Spanish Organizations in Horizon 2020

Legal name	EU Contribution (€)	Project Participations
Agencia Estatal Consejo Superior de Investigaciones Científicas	238.259.045€	555
Fundación Tecnalia Research & Innovation	112.240.776€	248
Barcelona Supercomputing Center	**75.772.091€**	**137**
ATOS Spain Sa	61.610.810€	156
Universitat Politècnica de Catalunya	60.807.189€	166
Universidad Politécnica de Madrid	60.201.700€	163
Universitat Autònoma de Barcelona	60.101.687€	125
Fundació Institut de Ciències Fotòniques	57.742.386€	83
Universitat Pompeu Fabra	56.641.631€	113
Centro de Investigaciones energéticas, medioambientales y tecnologías	55.616.305€	73

Source:
European Commission, Participant Portal H2020 Projects
Updated Dec 2019

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

MareNostrum 5
A European pre-exascale supercomputer

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

Hosting Consortium:

Spain Portugal Turkey Croatia Ireland


Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EuroHPC

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 guest 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

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

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?
  - Inhibits opportunities in holistic co-design
    - Facing barrier to innovation
    - Being able to have a conversation or not

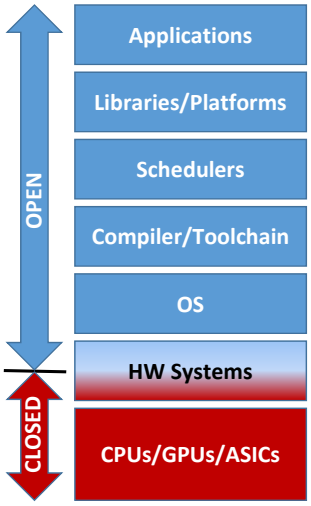

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

ARM-based prototypes at BSC

2011 Tibidabo	2012 KAYLA	2013 Pedraforca	2014 Mont-Blanc
ARM multicore	ARM + GPU CUDA on ARM	ARM + GPU Inifinband RDMA	Single chip ARM+GPU OpenCL on ARM GPU

PRACE PRACE PRACE MONT BLANC

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Mont-Blanc HPC Stack for ARM

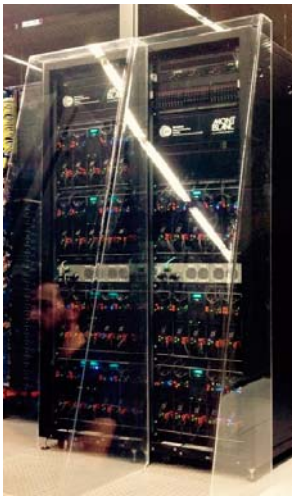

Industrial applications

Pharmacelera, MUREX, DASSAULT AVIATION, AVL, Fluids, Cenaero, Rolls-Royce

Applications

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

System software

ARM, ETH zürich, Inria, Barcelona Supercomputing Center, JÜLICH, Bull atos technologies

Hardware

ARM, Barcelona Supercomputing Center, Bull atos technologies, CEA, UC

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

The Exascale Race – The Japanese example

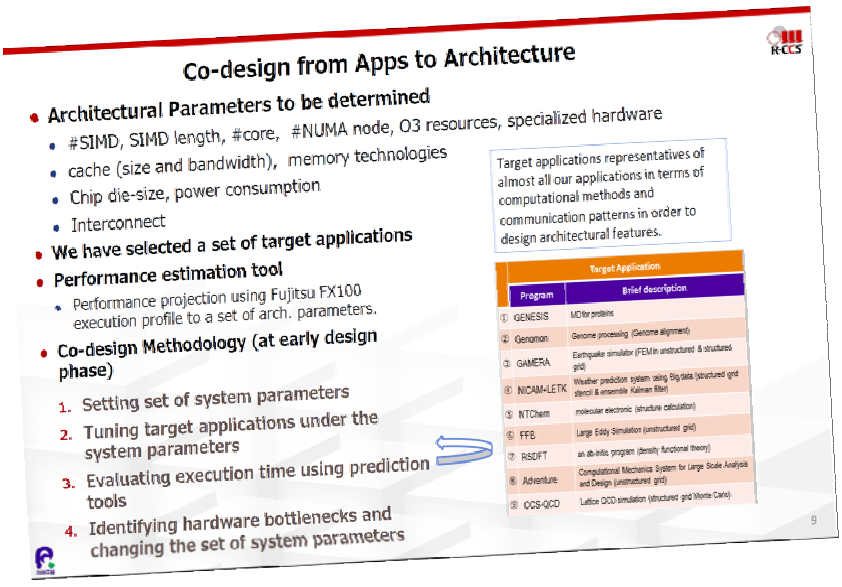

Co-design from Apps to Architecture

- Architectural Parameters to be determined
  - #SIMD, SIMD length, #core, #NUMA node, O3 resources, specialized hardware
  - cache (size and bandwidth), memory technologies
  - Chip die-size, power consumption
  - Interconnect
- We have selected a set of target applications
- Performance estimation tool
  - Performance projection using Fujitsu FX100 execution profile to a set of arch. parameters.
- Co-design Methodology (at early design phase)
  1. Setting set of system parameters
  2. Tuning target applications under the system parameters
  3. Evaluating execution time using prediction tools
  4. Identifying hardware bottlenecks and changing the set of system parameters

Target applications representatives of almost all our applications in terms of computational methods and communication patterns in order to design architectural features.

Target Application

Program	Brief description
① GENESIS	MD for proteins
② Genomon	Genome processing (Genome alignment)
③ GAMERA	Earthquake simulator (FEM in unstructured & structured grid)
④ NICAM+LETKF	Weather prediction system using Big data (structured grid stencil & ensemble Kalman filter)
⑤ NTChem	molecular electronic (structure calculation)
⑥ FFB	Large Eddy Simulation (unstructured grid)
⑦ RSDFT	an ab-initio program (density functional theory)
⑧ Adventure	Computational Mechanics System for Large Scale Analysis and Design (unstructured grid)
⑨ CCS-QCD	Lattice QCD simulation (structured grid Monte-Carlo)

The Exascale Race – The Japanese example

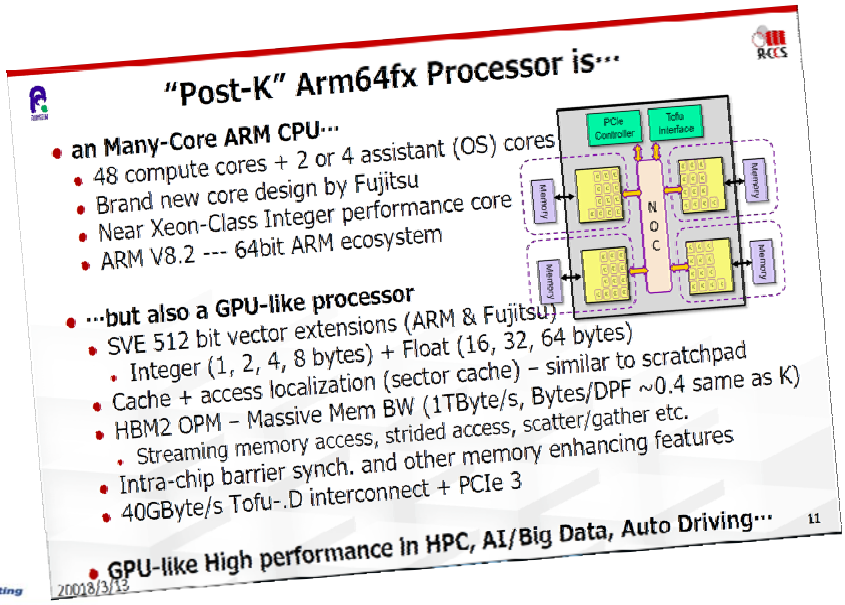

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

#	Site	Manufacturer	Computer	Country	Cores	Rmax [Pflops]	Power [MW]
1	Oak Ridge National Laboratory	IBM	Summit IBM Power System, P9 22C 3.07GHz, Mellanox EDR, NVIDIA GV100	USA	2,414,592	148.6	10.1
2	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1GHz, Mellanox EDR, NVIDIA GV100	USA	1,572,480	94.6	7.4
3	National Supercomputing Center in Wuxi	NRCPC	Sunway TaihuLight NRCPC Sunway SW26010, 260C 1.45GHz	China	10,649,600	93.0	15.4
4	National University of Defense Technology	NUDT	Tianhe-2A ANUDT TH-IVB-FEP, Xeon 12C 2.2GHz, Matrix-2000	China	4,981,760	61.4	18.5
5	Texas Advanced Computing Center / Univ. of Texas	Dell	Frontera Dell C6420, Xeon Platinum 8280 28C 2.7GHz, Mellanox HDR	USA	448,448	23.5	
6	Swiss National Supercomputing Centre (CSCS)	Cray	Piz Daint Cray XC50, Xeon E5 12C 2.6GHz, Aries, NVIDIA Tesla P100	Switzerland	387,872	21.2	2.38
7	Los Alamos NL / Sandia NL	Cray	Trinity Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries	USA	979,072	20.2	7.58
8	National Institute of Advanced Industrial Science and Technology	Fujitsu	AI Bridging Cloud Infrastructure (ABCI) PRIMERGY CX2550 M4, Xeon Gold 20C 2.4GHz, IB-EDR, NVIDIA V100	Japan	391,680	19.9	1.65
9	Leibniz Rechenzentrum	Lenovo	SuperMUC-NG ThinkSystem SD530, Xeon Platinum 8174 24C 3.1GHz, Intel Omni-Path	Germany	305,856	19.5	
10	Lawrence Livermore National Laboratory	IBM	Lassen IBM Power System, P9 22C 3.1GHz, Mellanox EDR, NVIDIA Tesla V100	USA	288,288	18.2	

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

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

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

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

Just an example: RISC-V in EPI

European Processor Initiative

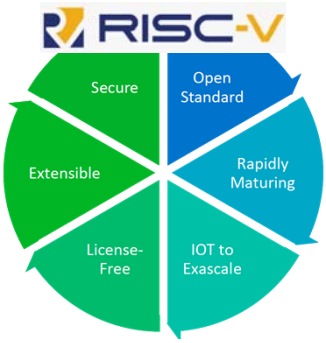

- More and more global IT actors are adopting RISC-V architectures to be vendor independent
- And of course the entire IoT ecosystem for lower performance, lower energy applications.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

EPI Partners

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

RISC-V in safety-critical systems in Europe

- **Safety-critical industry** in Europe gaining interest in RISC-V technology
  - **Space**: ESA needs microprocessors with no export restrictions
    - No SPARCv8, no PowerPC, no ISA with an owner
    - Now funding RISC-V initiatives
  - **Avionics** industry already expressed explicit interest, **automotive** industry now learning about RISC-V
- H2020 De-RISC: **TRL8 RISC-V space multicore** by early 2022
  - **Gaisler**, **Thales**, fentISS, and **BSC**
- H2020 SELENE: **TRL5 RISC-V space/avionics/automotive networked multicore** by end of 2022
  - **Gaisler**, **Airbus** (DE), **Airbus** (FR), **Siemens** (DE), **Siemens** (AT), **Virtual Vehicle**,… and **BSC**
- ECSEL FRACTAL: **Highly-adaptive RISC-V based prototype** for edge computing by mid 2023
  - **23 European partners (mostly industrial)**,… and **BSC**

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

BSC roadmap towards Exascale computing in Europe

Prof. Mateo Valero

LOCA
European Laboratory for Open Computer Architecture

1/22/2020

EPI SGA1

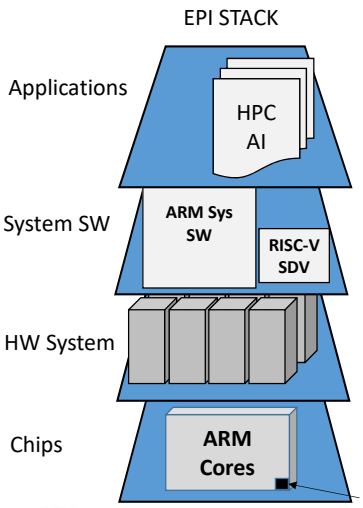

Industry	Research
• Existing Apps	• Accelerated HPC & AI Apps

• ARM SW Ecosystem (from Mont Blanc research)	• System Software
	  • RISC-V System SW
	  • Software Development Vehicles (SDV)

• Bull Sequana	• Test chip (22nm) and board

• ARM (and other IP providers)	• RISC-V Accelerators
	  • Vector, Stencil, AI, VPR
	  • L2+HN, NoC

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

Industry Chip Development Cycle

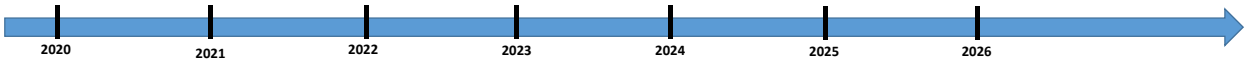

people

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

Mapping Chip Development onto Research Roadmap

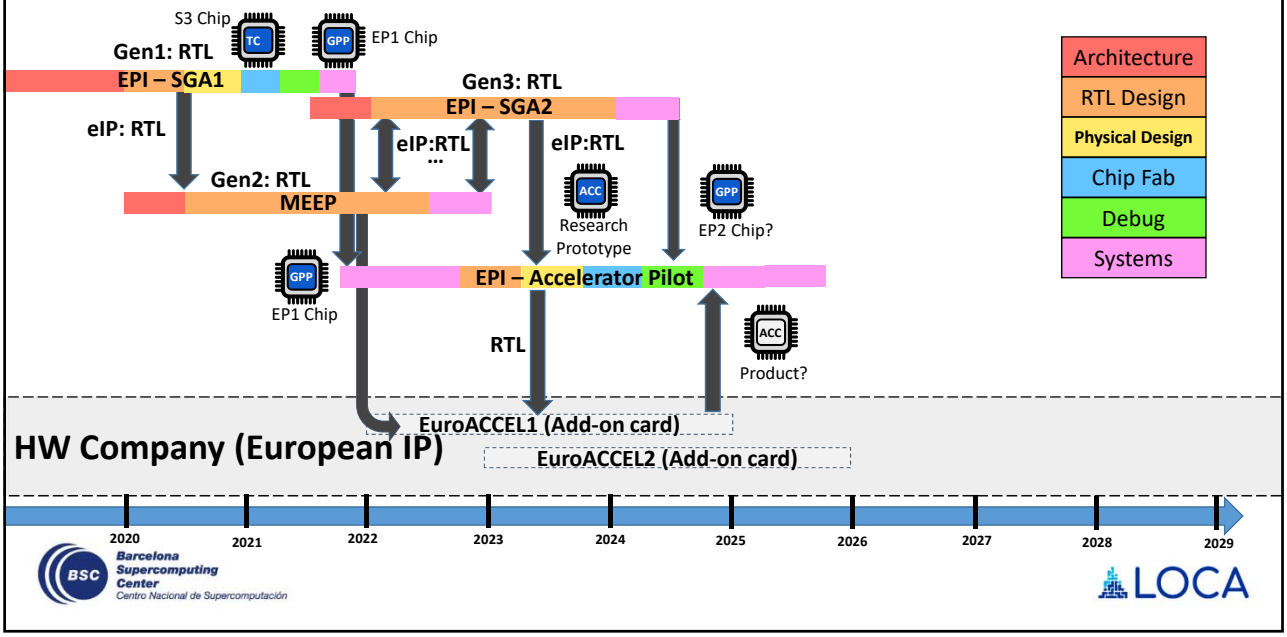

MareNostrum5 R&D -> MEEP

(MareNostrum Experimental Exascale Platform)

Large-scale Emulations for Exascale Systems

Software stack development vehicle
- Apps down to OS and tools

FPGA Emulation Platform
- Open Compute Accelerator Module (OAM) FPGA + HBM
  - RISC-V: Scalar+DLP accelerators+Memory

Develop & Improve Technology Readiness Level (TRL):
- RISC-V: Scalar + Vector (and other DLP accelerator) cores
- Memory hierarchy
- Software stack

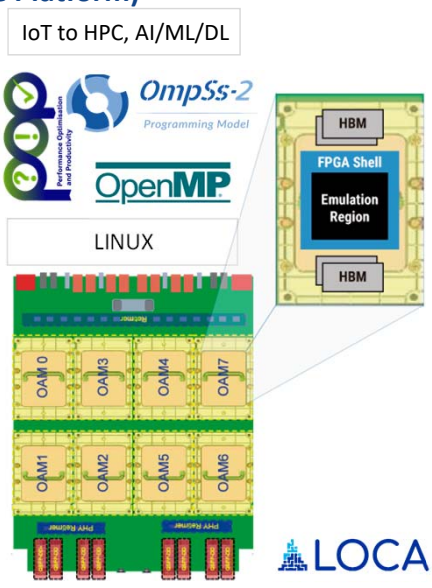

The BSC HPC Roadmap to Technology Independence

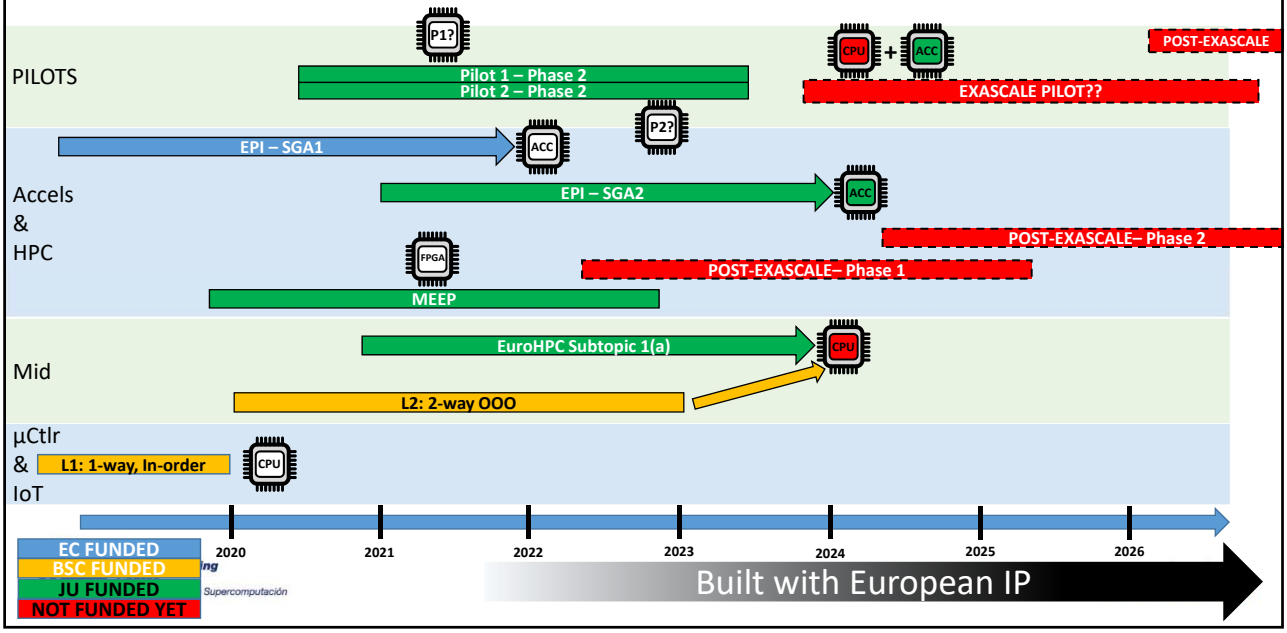

First Lagarto Tapeout

- 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)
  - Physicial design (IMB-CNM + BSC)
  - Tapeout and bringup (IMB-CNM + BSC)

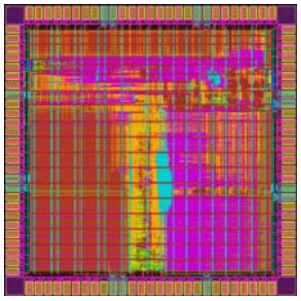

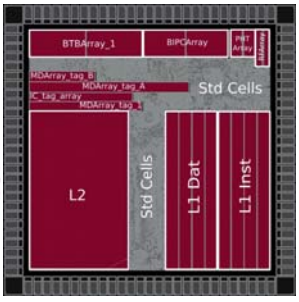

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

LOCA

Enabling HW/SW Co-Design for IoT to HPC

Name

Position

LOCA
European Laboratory for Open Computer Architecture

Today´s technology trends

Massive penetration of Open Source Software
- IoT (Arduino),
- Mobile (Android),
- Enterprise (Linux),
- HPC (Linux, OpenMP, etc.)

New Open Source Hardware Momentum from IoT and the Edge to HPC
- RISC-V
- OpenPOWER
- MIPS

Moore´s Law + Power = Specialization (HW/SW Co-Design)
- More cost effective
- More performant
- Less Power

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

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.

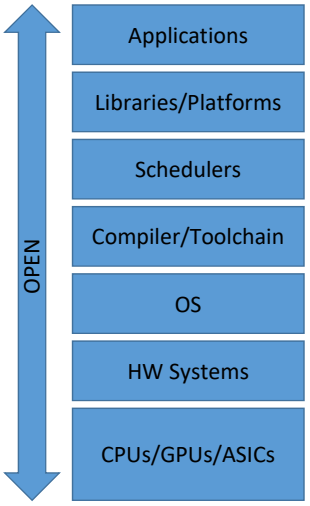

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

Open Source Beyond 2020

"**Open Source has become mainstream across all sectors of the software industry** during the past 10 years. To a large extent, **open software re-use has proven economically efficient**. The level of maturity of Open Source Hardware (OSH) remains far lower than that of Open Source Software (OSS). However, **business ecosystems for OSH are developing fast so that OSH could constitute a cornerstone of the future Internet of Things (IoT) and the future of computing**."

- DG Connect & DG IT Workshop, Brussels, Nov. 14-15, 2019

Source: https://ec.europa.eu/digital-single-market/en/news/workshop-about-future-open-source-software-and-open-source-hardware

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

b3
b4

BSC full stack

SW

HW

Software/Hardware Co-Design

HPC Applications

Specialization using HW/SW Co-Design

LOCA
European Laboratory for Open Computer Architecture

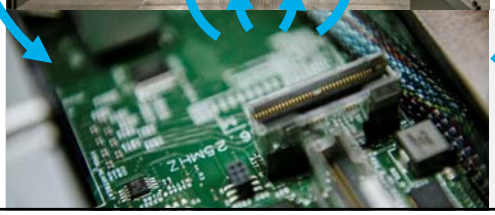

HPC Hardware

LOCA

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA Goals

- Mechanism to extend open source ecosystem to include H/W
  - Add H/W expertise to BSC and European partners, leverage existing S/W expertise
  - Open European IP repository → rapid implementation
  - Provide usable Open Source H/W
  - Intersection of academia and industry
  - Catalyst to reinvigorate European ICT industry Global collaboration and training center
  - Incubator for European IP

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

Diapositiva 35

b3 I would really love a pretty pciture of graphic of this.
bscuser; 25/10/2019

b4 Computer Science should blend into Apps and CASE blends into Apps, CS, and Ops.
bscuser; 25/10/2019

b8
b9

European Collaboration & Education

Casteller
(human tower)

- Traditional chip design is done in a Master/Apprentice environment
- LOCA recreates this environment by bringing in Masters from industry to collaborate with a variety of people, pushing beyond RTL
- Professors, students, and industry veterans all together
- Ideal sandbox for creative and innovative work
- Research and Design to chip fabrication and software systems

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

From IoT, Edge Computing, Clouds to Supercomputers

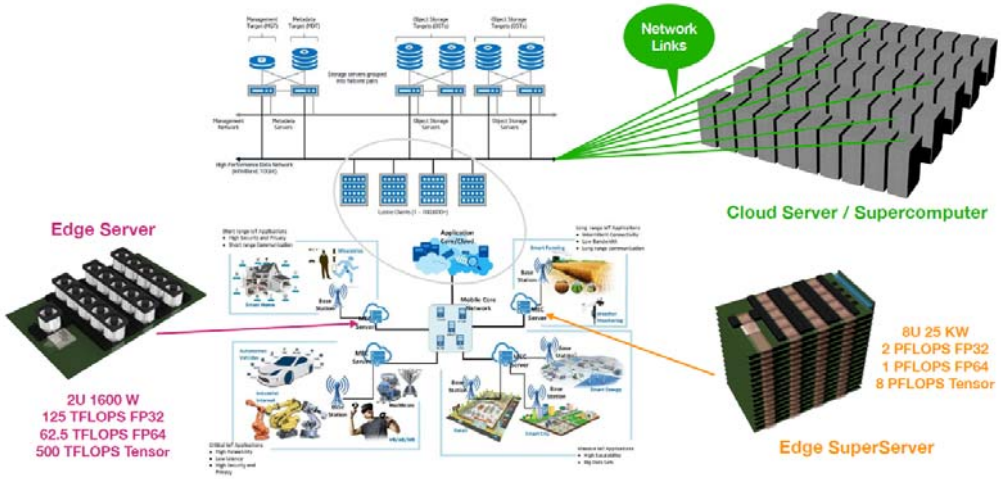

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

Diapositiva 37

b8 Can we add a graphic of people coming together? Maybe the Castellers?
bscuser; 25/10/2019

b9 Maybe a construction sandbox with masters and apprentice builders as a theme to this slide?
bscuser; 25/10/2019

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

- 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
- 6 nm initial, 5 and 3 nm follow-on designs

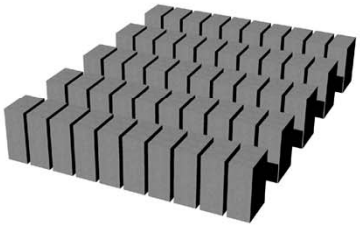

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

BSC RISC-V Accelerator

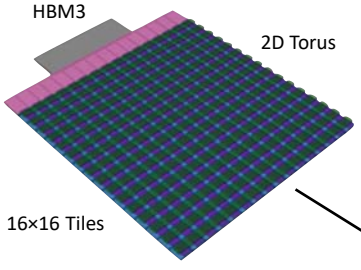

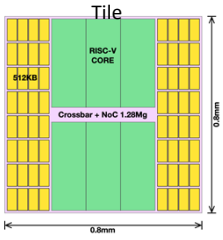

164 TFLOPS
1 TB HBM3
3000W

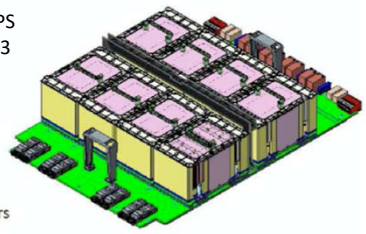

128MB Distributed Memory

1536 FMA
1.666 GHz
5.12 TFLOPS

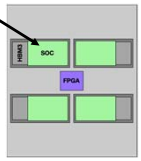

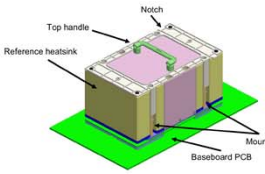

20.5 TFLOPS
128 GB HBM3
375W

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

P. Hsu 21 Jan. 2020

Focus, Freedom, and Forward

- BSC is embracing a Open Source Hardware to complement OSS
  - RISC-V has the momentum to succeed as the Open Source ISA, much like Linux
- BSC is building the infrastructure to support the future of computing
  - Combining applications, system software, hardware design and hardware
- BSC can organize and coordinate full stack efforts
  - More resources = faster time to success
  - More partners (Academia/Industry)
  - Research/Resource multiplier
- Europe and the world can unite around the BSC Vision to move forward faster, together

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

BSC is Made in Europe

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA makes Open Source Ecosystems a reality in the future

LOCA

Come to Barcelona, and learn more about LOCA

Casteller
(human tower)

LOCA
European Laboratory for Open Computer Architecture

→ European Open Hardware, Architecture Initiative

We must stand on the shoulders of giants to build great things.
We are assembling many giants and hope you can join us.

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

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

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

HR EXCELLENCE IN RESEARCH

LOCA

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

EuroHPC
Joint Undertaking

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

MareNostrum RISC-V inauguration 2021

MN6-RISC-V
2025???

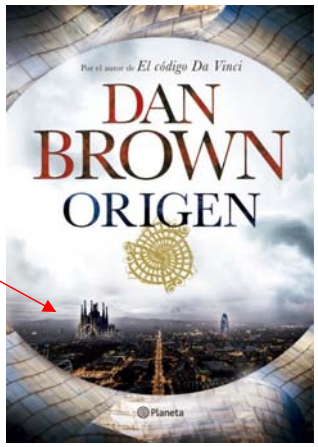

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

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

RISC-V

- 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**

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

Mapping Chip Development onto Research Roadmap

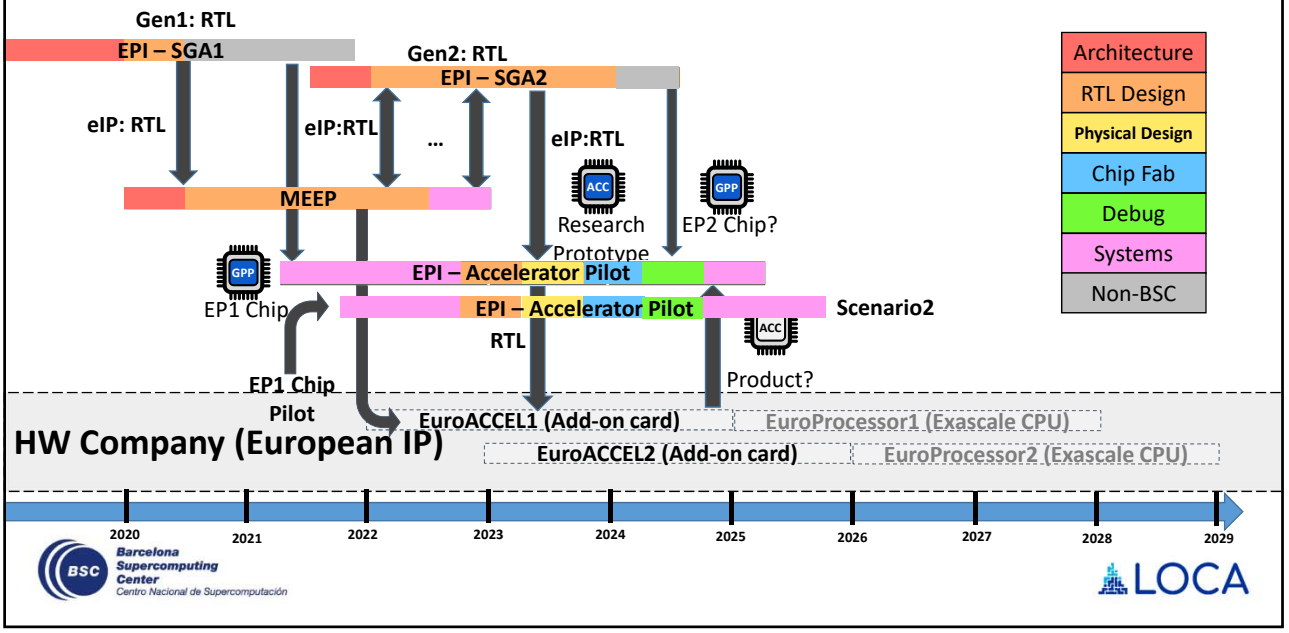

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

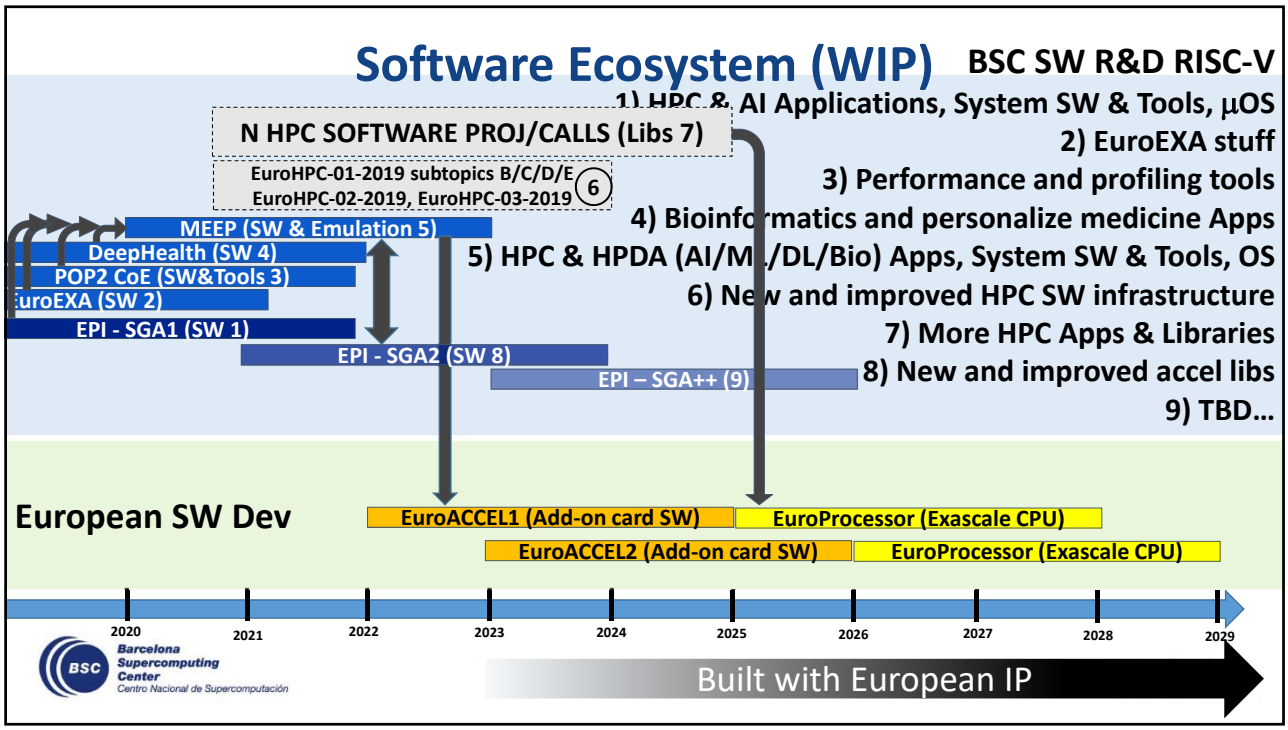
Software Ecosystem (WIP)
BSC SW R&D RISC-V
1) HPC & AI Applications, System SW & Tools, µOS
2) EuroEXA stuff
3) Performance and profiling tools
4) Bioinformatics and personalize medicine Apps
5) HPC & HPDA (AI/ML/DL/Bio) Apps, System SW & Tools, OS
6) New and improved HPC SW infrastructure
7) More HPC Apps & Libraries
8) New and improved accel libs
9) TBD…
N HPC SOFTWARE PROJ/CALLS (Libs 7)
EuroHPC-01-2019 subtopics B/C/D/E
EuroHPC-02-2019, EuroHPC-03-2019
6
MEEP (SW & Emulation 5)
DeepHealth (SW 4)
POP2 CoE (SW&Tools 3)
uroEXA (SW 2)
EPI - SGA1 (SW 1)
EPI - SGA2 (SW 8)
EPI – SGA++ (9)
European SW Dev
EuroACCEL1 (Add-on card SW)
EuroProcessor (Exascale CPU)
EuroACCEL2 (Add-on card SW)
EuroProcessor (Exascale CPU)
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
Barcelona Supercomputing Center
Centro Nacional de Supercomputación
Built with European IP

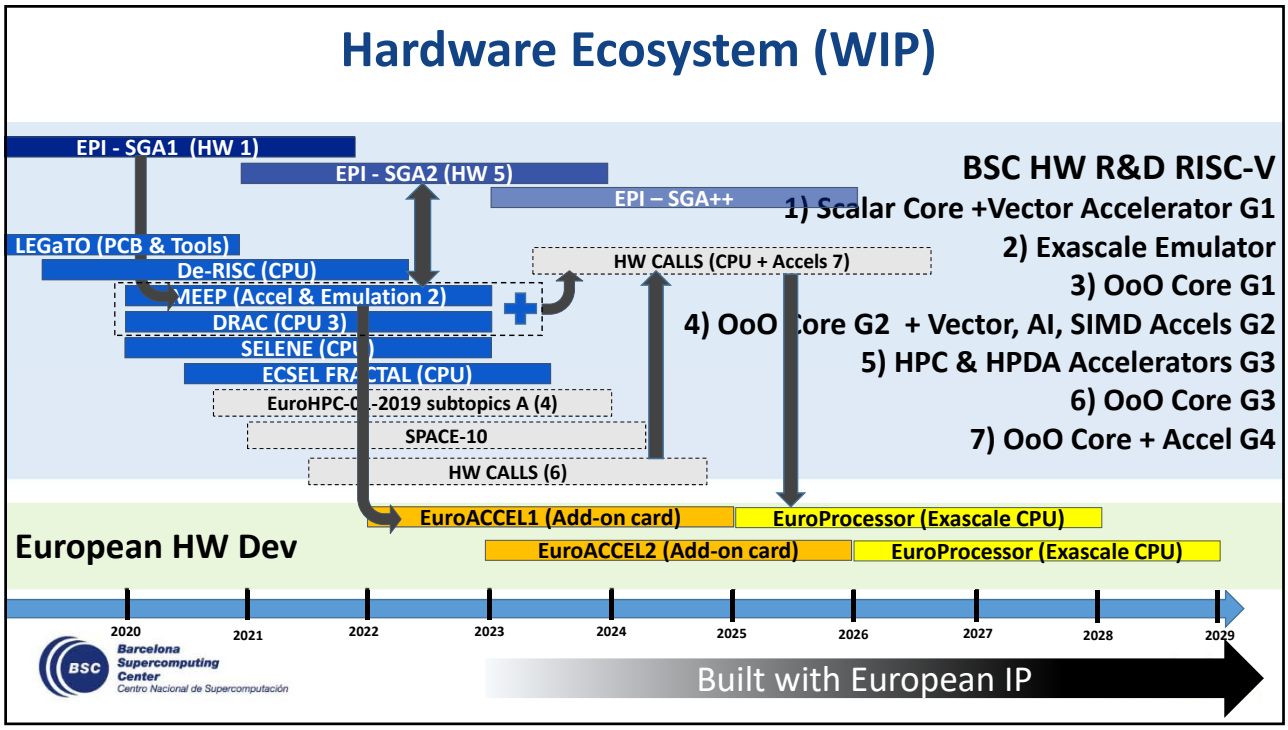
Hardware Ecosystem (WIP)
BSC HW R&D RISC-V
1) Scalar Core +Vector Accelerator G1
2) Exascale Emulator
3) OoO Core G1
4) OoO Core G2 + Vector, AI, SIMD Accels G2
5) HPC & HPDA Accelerators G3
6) OoO Core G3
7) OoO Core + Accel G4
EPI - SGA1 (HW 1)
EPI - SGA2 (HW 5)
EPI – SGA++
LEGaTO (PCB & Tools)
De-RISC (CPU)
MEEP (Accel & Emulation 2)
DRAC (CPU 3)
SELENE (CPU)
ECSEL FRACTAL (CPU)
EuroHPC-01-2019 subtopics A (4)
SPACE-10
HW CALLS (6)
HW CALLS (CPU + Accels 7)
European HW Dev
EuroACCEL1 (Add-on card)
EuroProcessor (Exascale CPU)
EuroACCEL2 (Add-on card)
EuroProcessor (Exascale CPU)
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
Barcelona Supercomputing Center
Centro Nacional de Supercomputación
Built with European IP

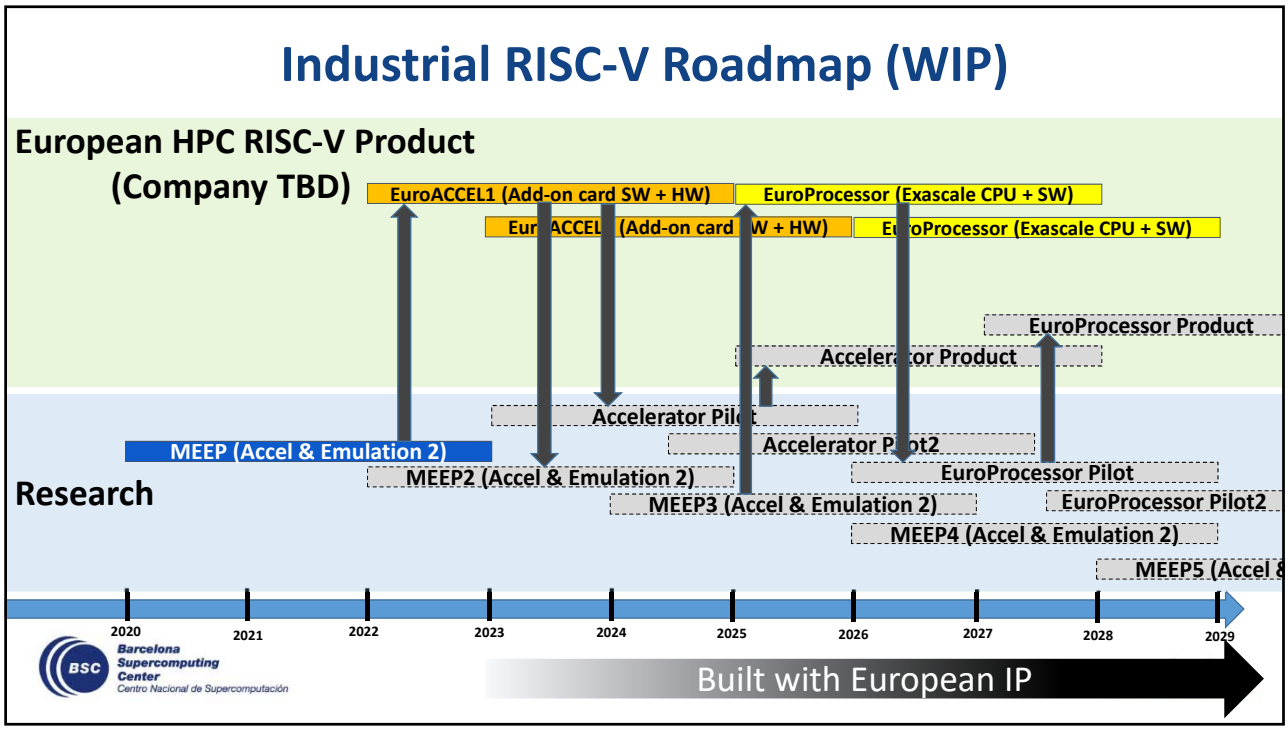
Industrial RISC-V Roadmap (WIP)
European HPC RISC-V Product
(Company TBD)
EuroACCEL1 (Add-on card SW + HW)
EuroProcessor (Exascale CPU + SW)
EuroProcessor Product
Accelerator Product
Accelerator Pilot
Accelerator Pilot2
EuroProcessor Pilot
EuroProcessor Pilot2
MEEP (Accel & Emulation 2)
MEEP2 (Accel & Emulation 2)
MEEP3 (Accel & Emulation 2)
MEEP4 (Accel & Emulation 2)
MEEP5 (Accel &
Research
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
Barcelona Supercomputing Center
Centro Nacional de Supercomputación
Built with European IP

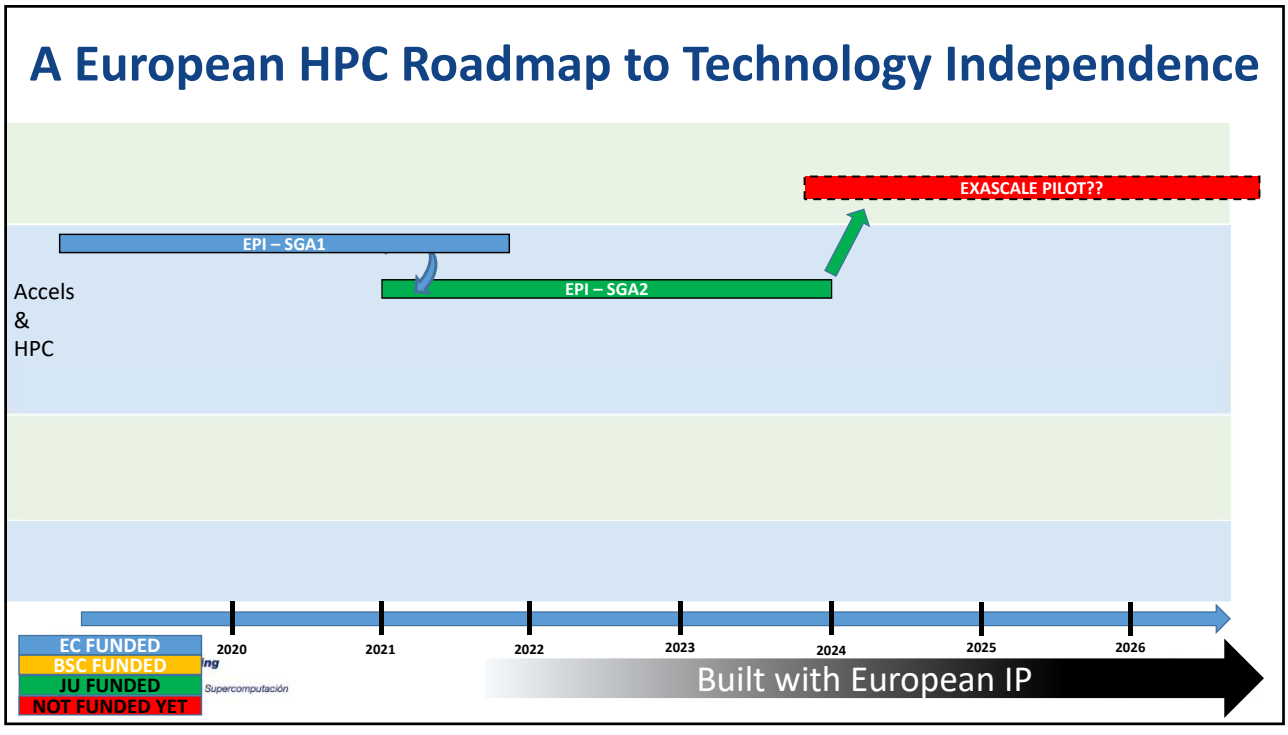
A European HPC Roadmap to Technology Independence
EXASCALE PILOT??
EPI – SGA1
EPI – SGA2
Accels
&
HPC
2020
2021
2022
2023
2024
2025
2026
EC FUNDED
BSC FUNDED
JU FUNDED
NOT FUNDED YET
Supercomputación
Built with European IP

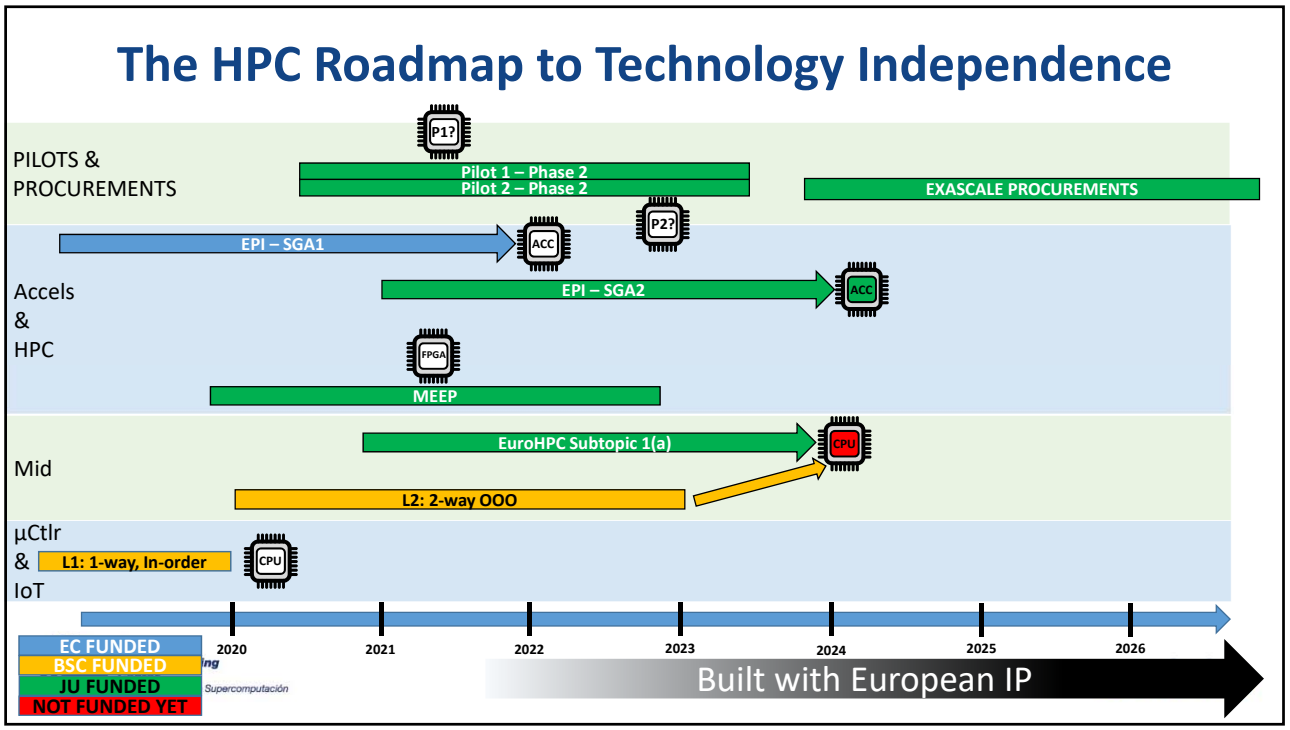
The HPC Roadmap to Technology Independence
PILOTS & PROCUREMENTS
P1?
Pilot 1 – Phase 2
Pilot 2 – Phase 2
EXASCALE PROCUREMENTS
P2?
EPI – SGA1
ACC
Accels & HPC
EPI – SGA2
ACC
FPGA
MEEP
Mid
EuroHPC Subtopic 1(a)
CPU
L2: 2-way OOO
µCtlr & IoT
L1: 1-way, In-order
CPU
2020
2021
2022
2023
2024
2025
2026
EC FUNDED
BSC FUNDED
JU FUNDED
NOT FUNDED YET
Supercomputación
Built with European IP

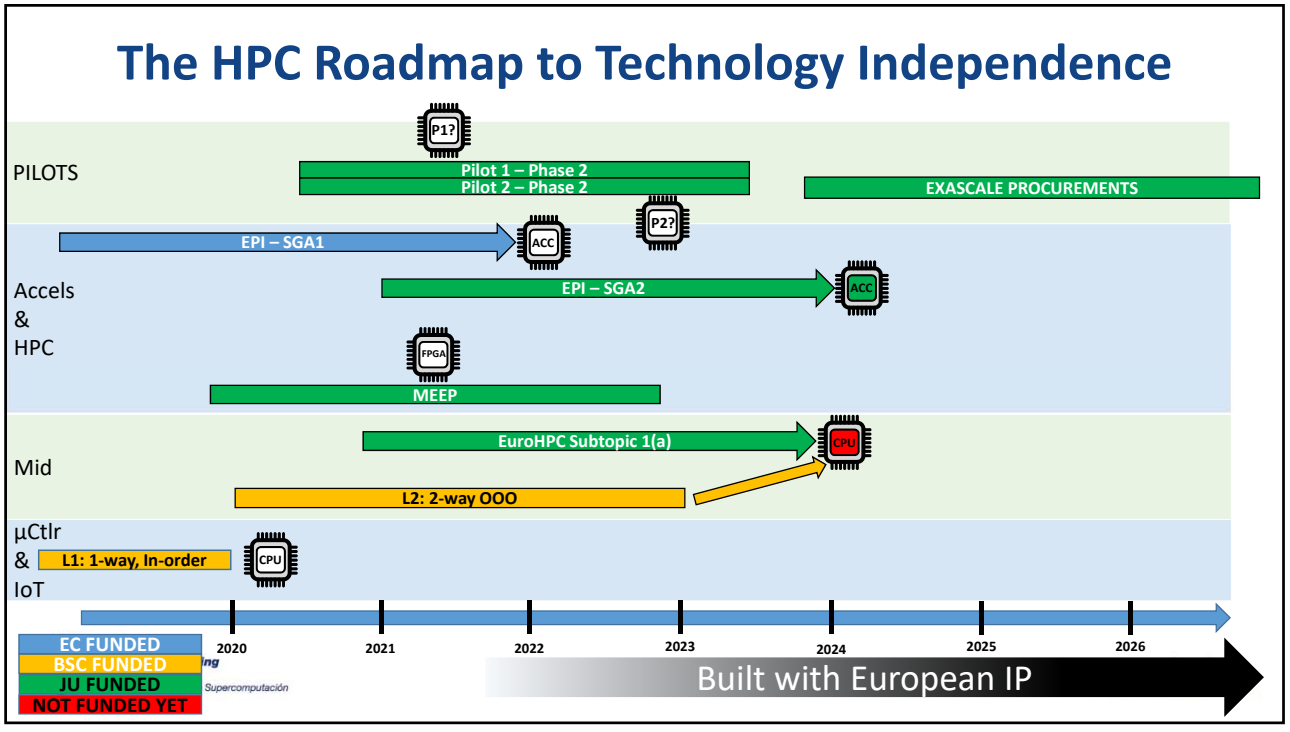
The HPC Roadmap to Technology Independence
PILOTS
P1?
Pilot 1 – Phase 2
Pilot 2 – Phase 2
EXASCALE PROCUREMENTS
P2?
EPI – SGA1
ACC
Accels & HPC
EPI – SGA2
ACC
FPGA
MEEP
Mid
EuroHPC Subtopic 1(a)
CPU
L2: 2-way OOO
µCtlr & IoT
L1: 1-way, In-order
CPU
2020
2021
2022
2023
2024
2025
2026
EC FUNDED
BSC FUNDED
JU FUNDED
NOT FUNDED YET
Supercomputación
Built with European IP

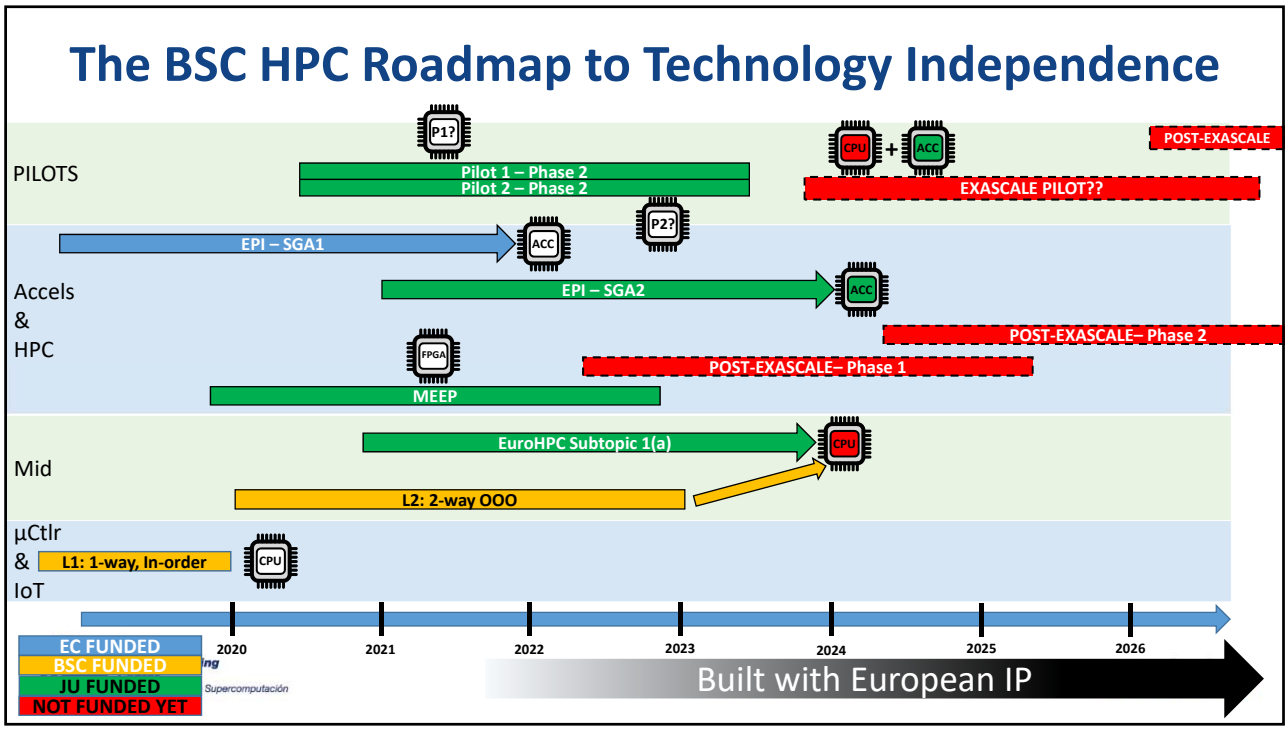
The BSC HPC Roadmap to Technology Independence
PILOTS
Accels & HPC
Mid
µCtlr & IoT
P1?
Pilot 1 – Phase 2
Pilot 2 – Phase 2
CPU
ACC
POST-EXASCALE
EXASCALE PILOT??
EPI – SGA1
P2?
EPI – SGA2
POST-EXASCALE– Phase 2
FPGA
POST-EXASCALE– Phase 1
MEEP
EuroHPC Subtopic 1(a)
L2: 2-way OOO
L1: 1-way, In-order
2020
2021
2022
2023
2024
2025
2026
EC FUNDED
BSC FUNDED
JU FUNDED
NOT FUNDED YET
Built with European IP

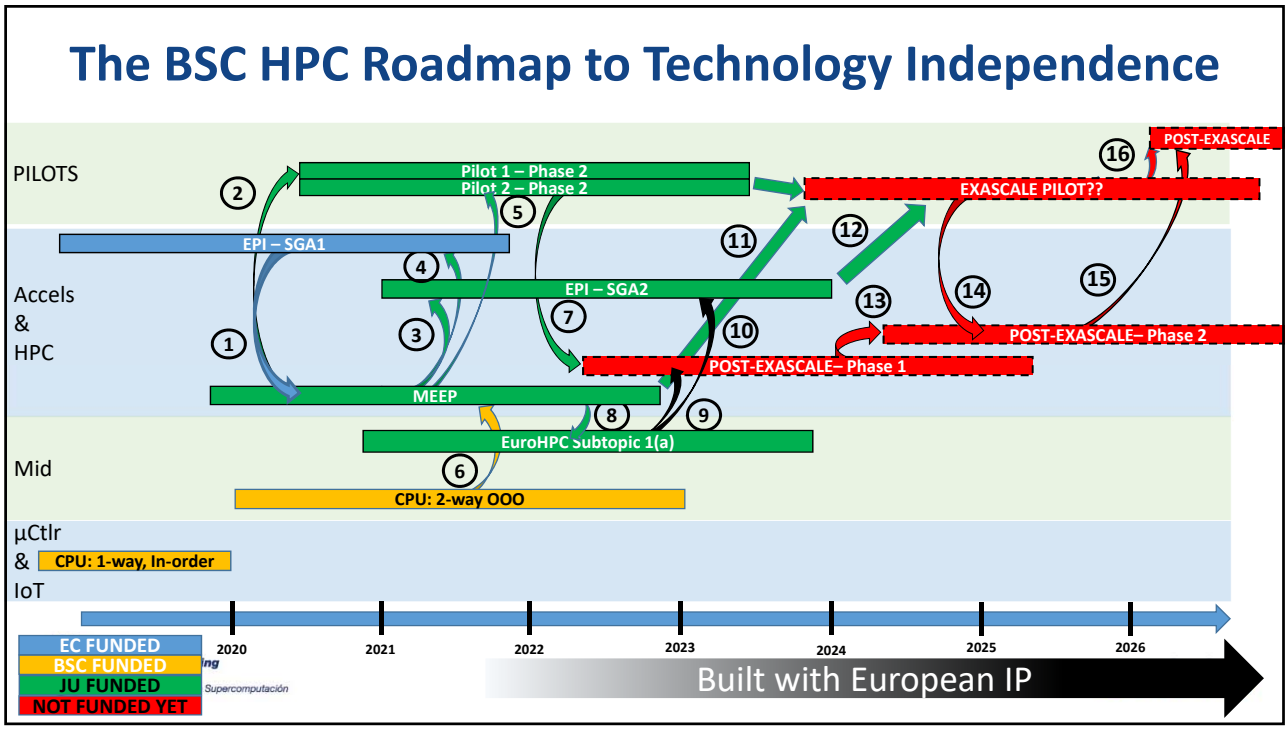
The BSC HPC Roadmap to Technology Independence
PILOTS
Accels & HPC
Mid
µCtlr & IoT
Pilot 1 – Phase 2
Pilot 2 – Phase 2
POST-EXASCALE
EXASCALE PILOT??
EPI – SGA1
EPI – SGA2
POST-EXASCALE– Phase 2
POST-EXASCALE– Phase 1
MEEP
EuroHPC Subtopic 1(a)
CPU: 2-way OOO
CPU: 1-way, In-order
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
2020
2021
2022
2023
2024
2025
2026
EC FUNDED
BSC FUNDED
JU FUNDED
NOT FUNDED YET
Built with European IP

EPAC architecture

- **Objective: Develop** & demonstrate **fully European** processor IPs based on the **RISC-V ISA.** Build on existing EU IP, leverage EU background and vision
- Provide a very **low power** and high computing throughput **accelerator for HPC & Emerging -> Automotive**

Main Contributors:

- **RISC-V in the Tile: 8 Vector cores, 8 STX, 1 VRP**
- **Vector Core :**
  - **Spain + Italy + Croatia**
  - **C: RISC-V Scalar core**
  - **V: RISC-V Vector core: 8 lanes**
- **STX**
  - **Switzerland + Germany**
  - **RISC-V NTX (AI/ML/DL) + Stencil**
- **VRP**
  - **(France)**
  - **RISCV + Extended precision FPU**

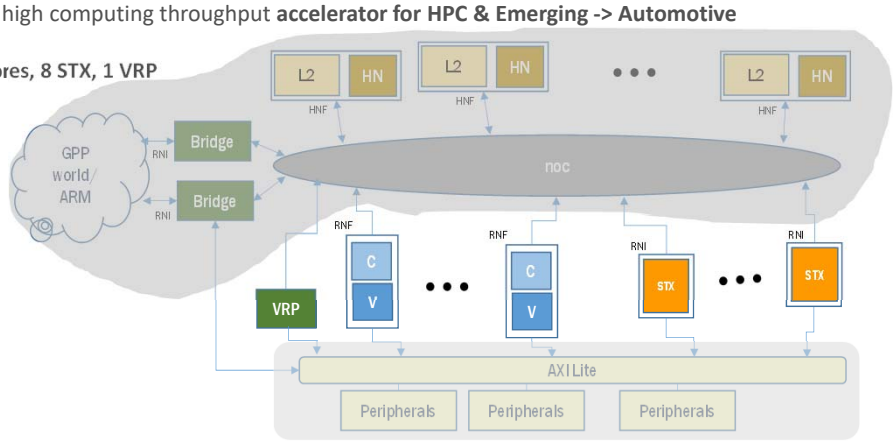

UNIBO | IST | CEA | Chalmers | FORTH | Semidynamics | BSC | ETH + FhG | EXTOLL | ETH

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

EPAC architecture

Vector

- **Performance**
  - @ 1.5 GHz → 24GF/core → 196GF/tile
  - State/GF: 80KB KB/core; +~640KB / tile
- **Area**
  - ~fit into budgeted slot
- **Power**
  - Expectation based on NTX
- **NTX:**
  - **20 Gflop/s, 108 Gflop/sW (FP32)**

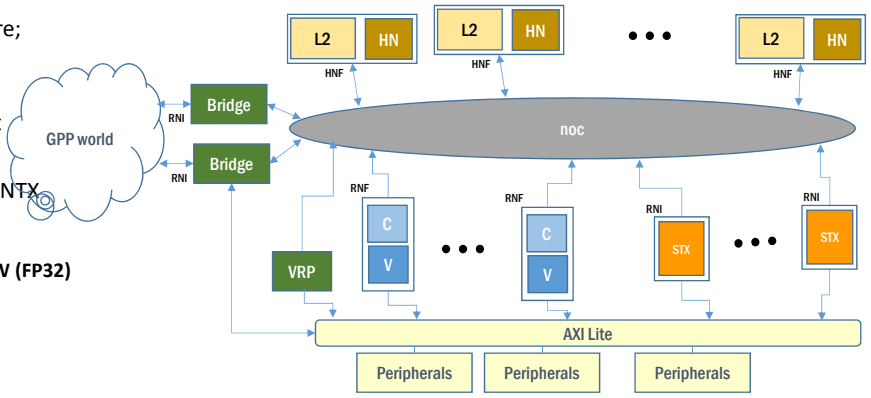

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Software History

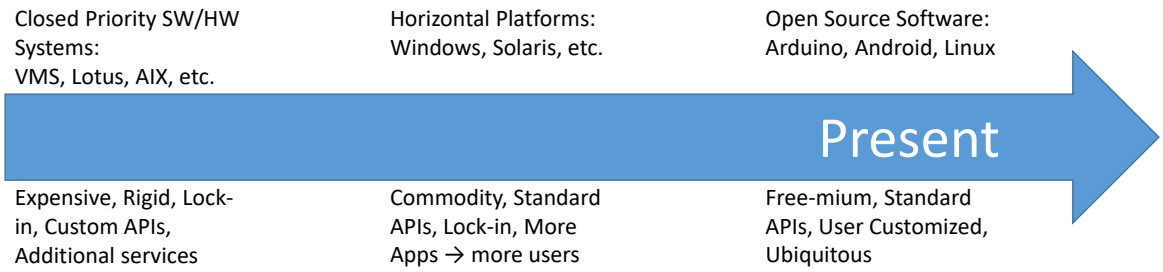

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

Linux History

- 1991: Started development, release
- 1992: X Windows released
- 1998: Adopted by many major companies
- 2004: MareNostrum BSC Supercomputer OS
- 2009: Basis for many new business systems and the cloud
- 2013: Android in 75% of the world smart phones
- 2015: De facto OS for IoT, mobile, cloud, and supercomputers
- 2019: 86,1% of the smartphones, 67% of public servers, 75% of embedded systems (*https://ec.europa.eu/digital-single-market/en/news/workshop-about-future-open-source-software-and-open-source-hardware)

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA

RISC-V History

- 2010: Started development and initial proposal
- 2015: RISC-V Foundation formed
- 2019: Adopted by many major companies
  - Starting in the embedded market with already over 1 Billion CPUs
  - RISC-V Foundation moves to Switzerland

The time is now to embrace and support RISC-V from IoT to HPC

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

LOCA