



**Barcelona
Supercomputing
Center**

Centro Nacional de Supercomputación



**EXCELENCIA
SEVERO
OCHOA**

Barcelona Supercomputing Center

The Future HPC will be Open

Prof. Mateo Valero

December 14, 2020

The 22nd IEEE International Conference on High Performance Computing and Communications (HPCC-2020)

14-16 December 2020, Fiji




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%



People evolution

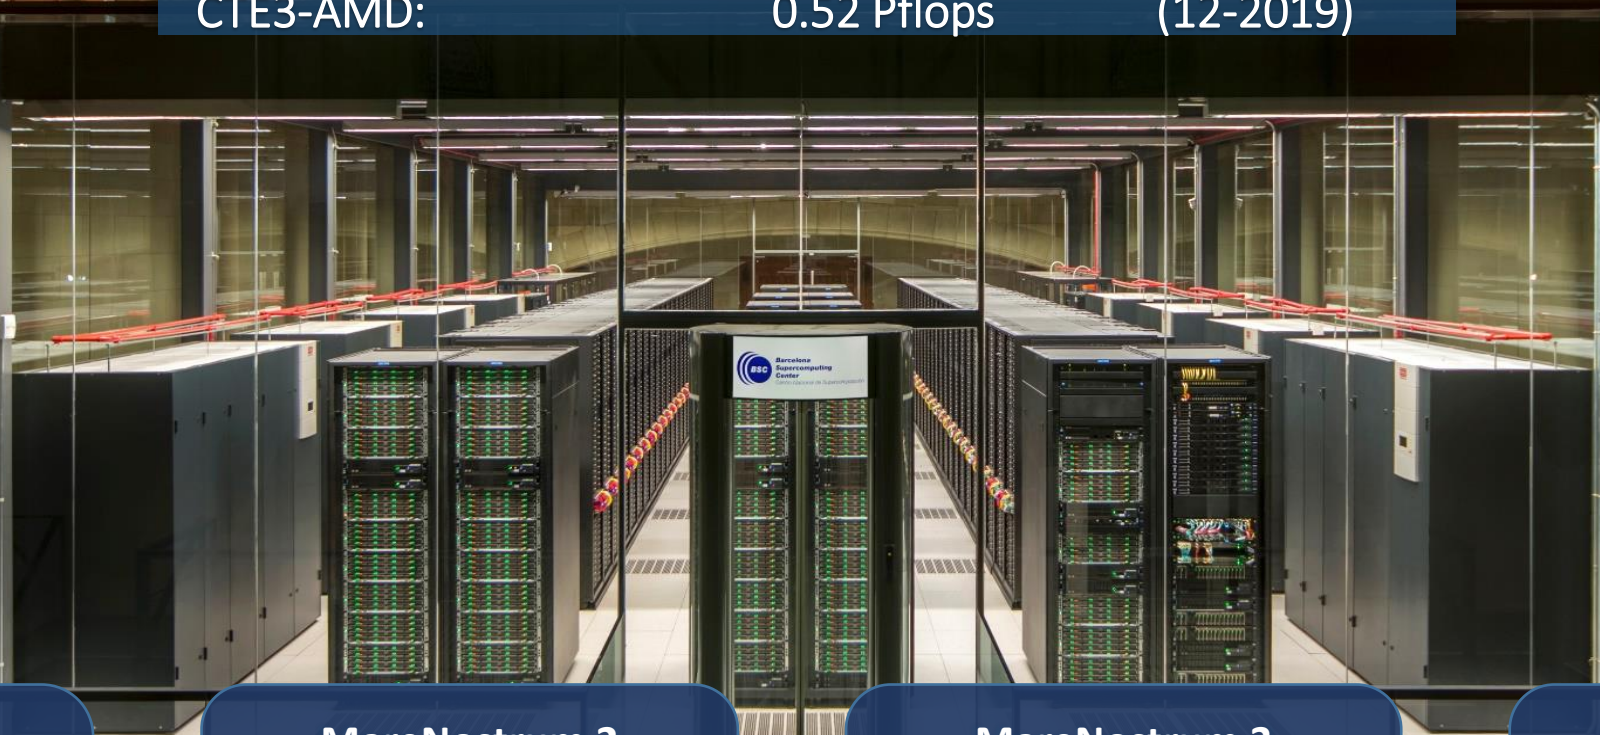


BSC Staff evolution 2005 - 2020

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

Today's technology trends




Massive penetration of Open Source Software

- IoT (Arduino),
- Mobile (Android),
- Enterprise (Linux),
- HPC (Linux, OpenMP, etc.)



Moore's Law + Power =
Specialization (HW/SW Co-Design)

- More cost effective
- More performant
- Less Power



New Open Source Hardware
Momentum from IoT and the
Edge to HPC

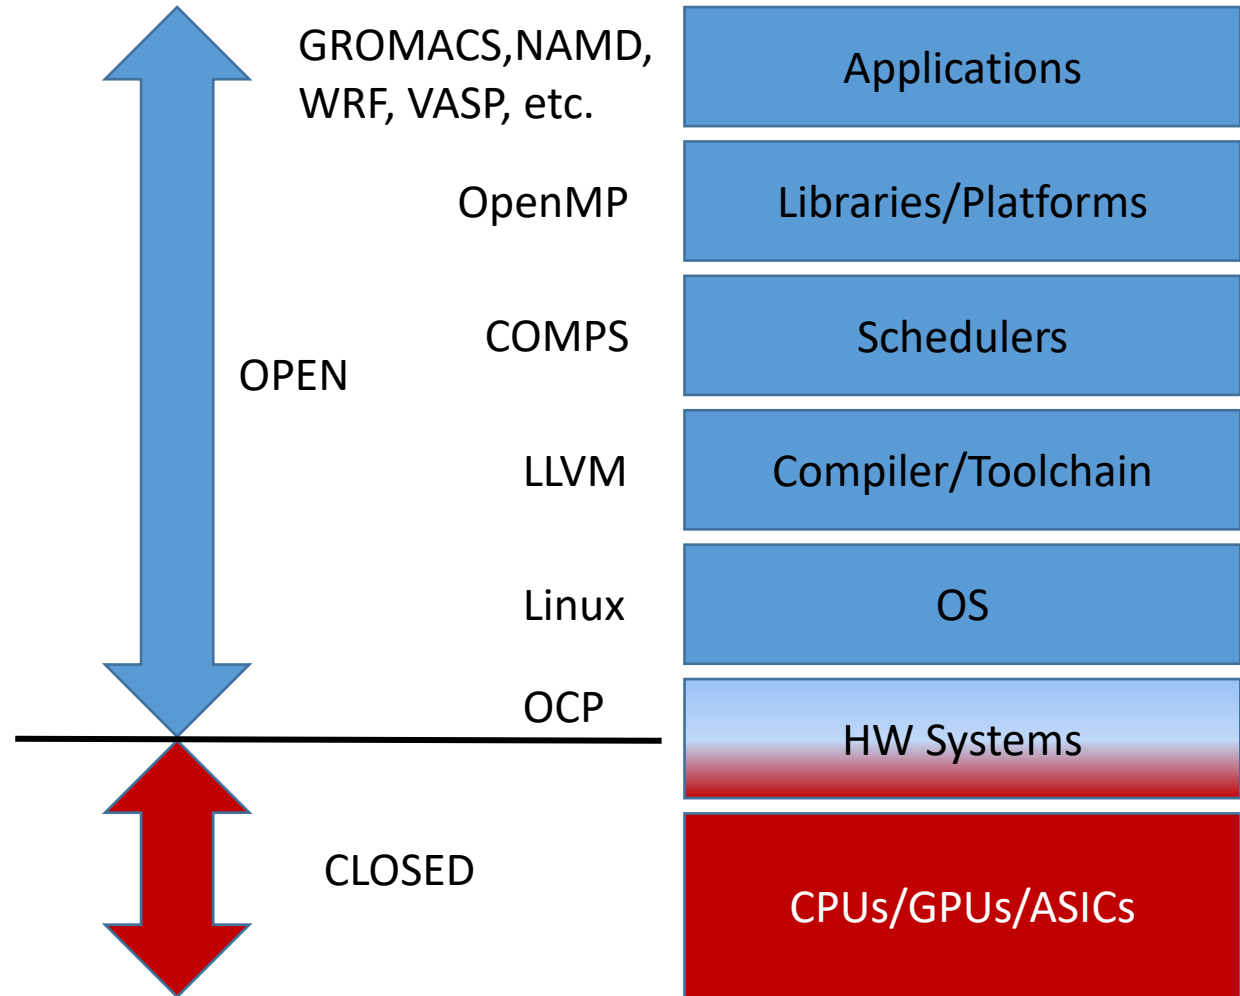
- RISC-V
- OpenPOWER
- MIPS

Linux History

- 1991: Started development, release
- 1992: X Windows released
- 1998: Adopted by many major companies
- 2004: 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

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?**



Mont-Blanc HPC Stack for ARM



Industrial applications



Applications



System software



Hardware



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+LETK | 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 | RSDFE | 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



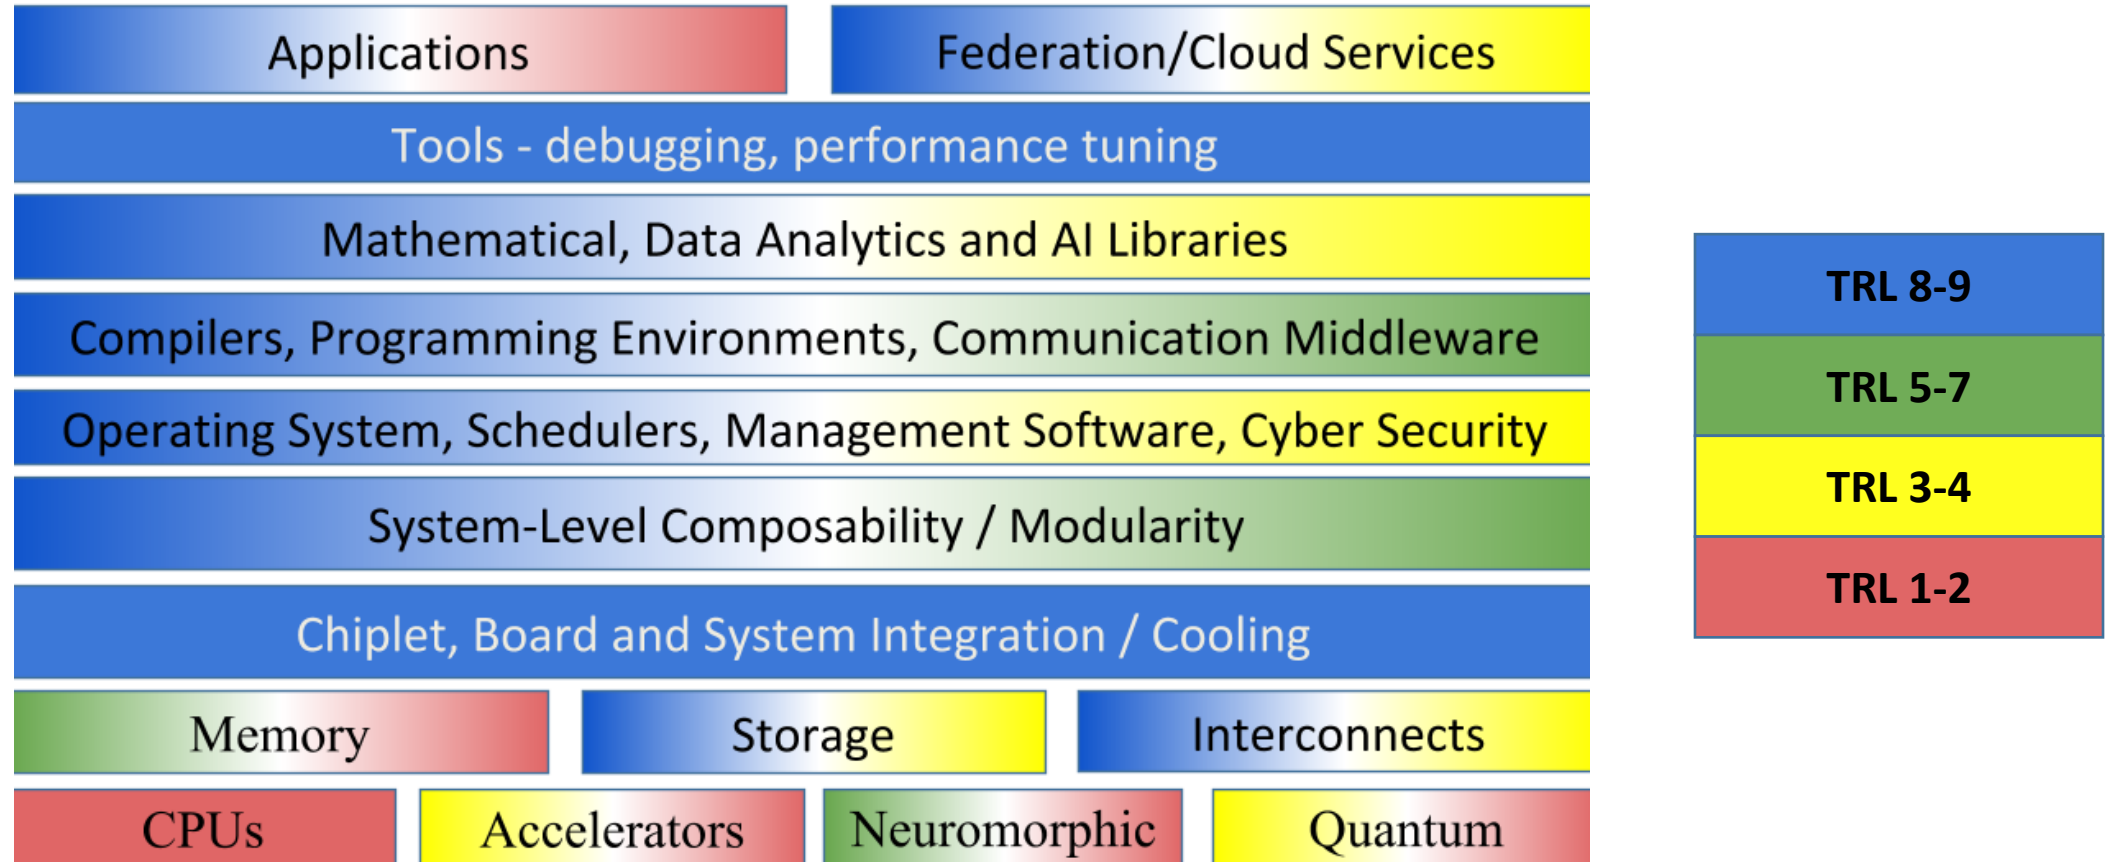
"Post-K" Arm64fx Processor is...

- **an Many-Core ARM CPU...**
 - 48 compute cores + 2 or 4 assistant (OS) cores
 - Brand new core design by Fujitsu
 - Near Xeon-Class Integer performance core
 - ARM V8.2 --- 64bit ARM ecosystem
- **...but also a GPU-like processor**
 - SVE 512 bit vector extensions (ARM & Fujitsu)
 - Integer (1, 2, 4, 8 bytes) + Float (16, 32, 64 bytes)
 - Cache + access localization (sector cache) – similar to scratchpad
 - HBM2 OPM – Massive Mem BW (1TByte/s, Bytes/DPF ~0.4 same as K)
 - Streaming memory access, strided access, scatter/gather etc.
 - Intra-chip barrier synch. and other memory enhancing features
 - 40GByte/s Tofu-.D interconnect + PCIe 3
- **GPU-like High performance in HPC, AI/Big Data, Auto Driving...**



20018/3/13

The European Open System Stack



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




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
- 2020
 - RISC-V Foundation moves to Switzerland

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

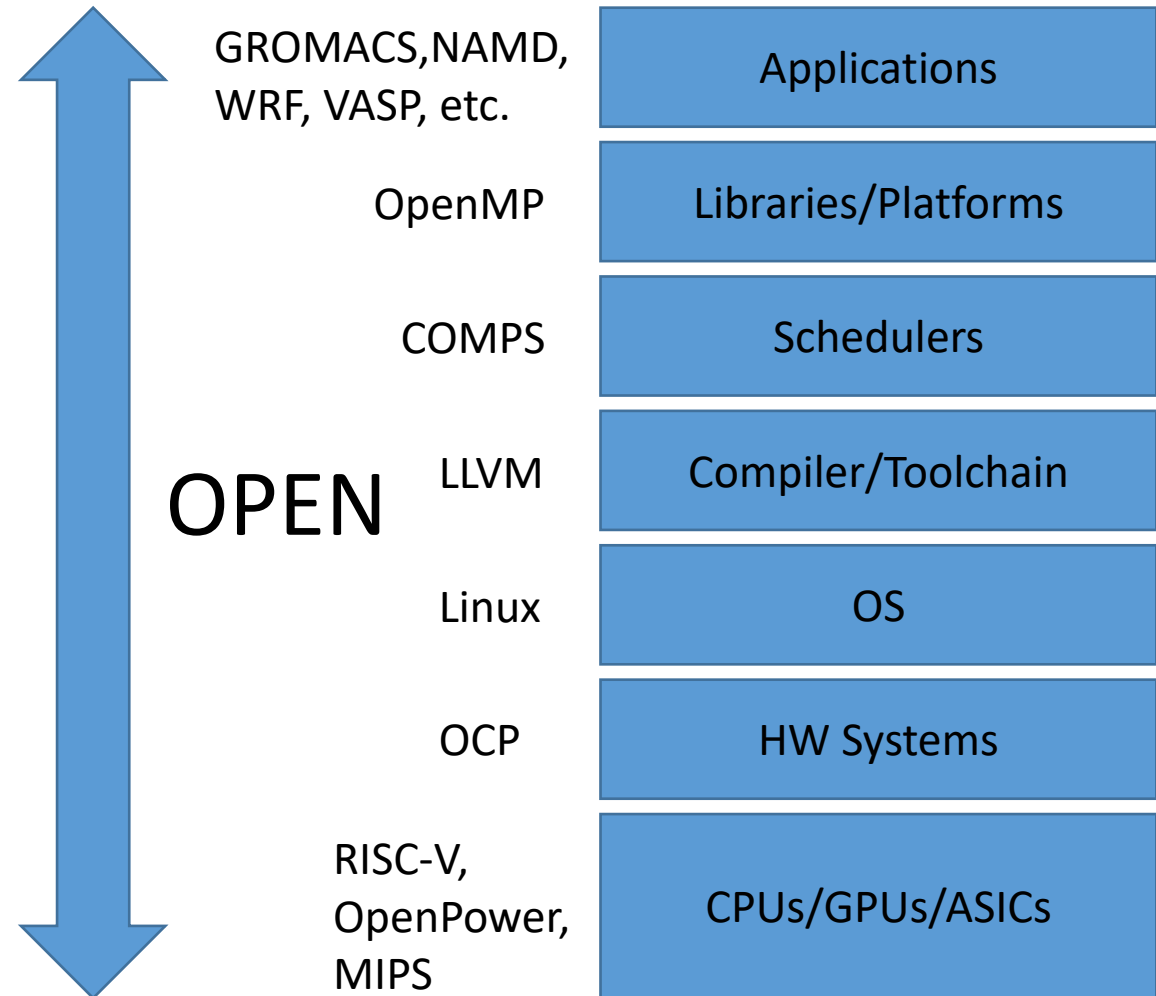
RISC-V is democratizing chip-design



- ➔ 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 Source Beyond 2020

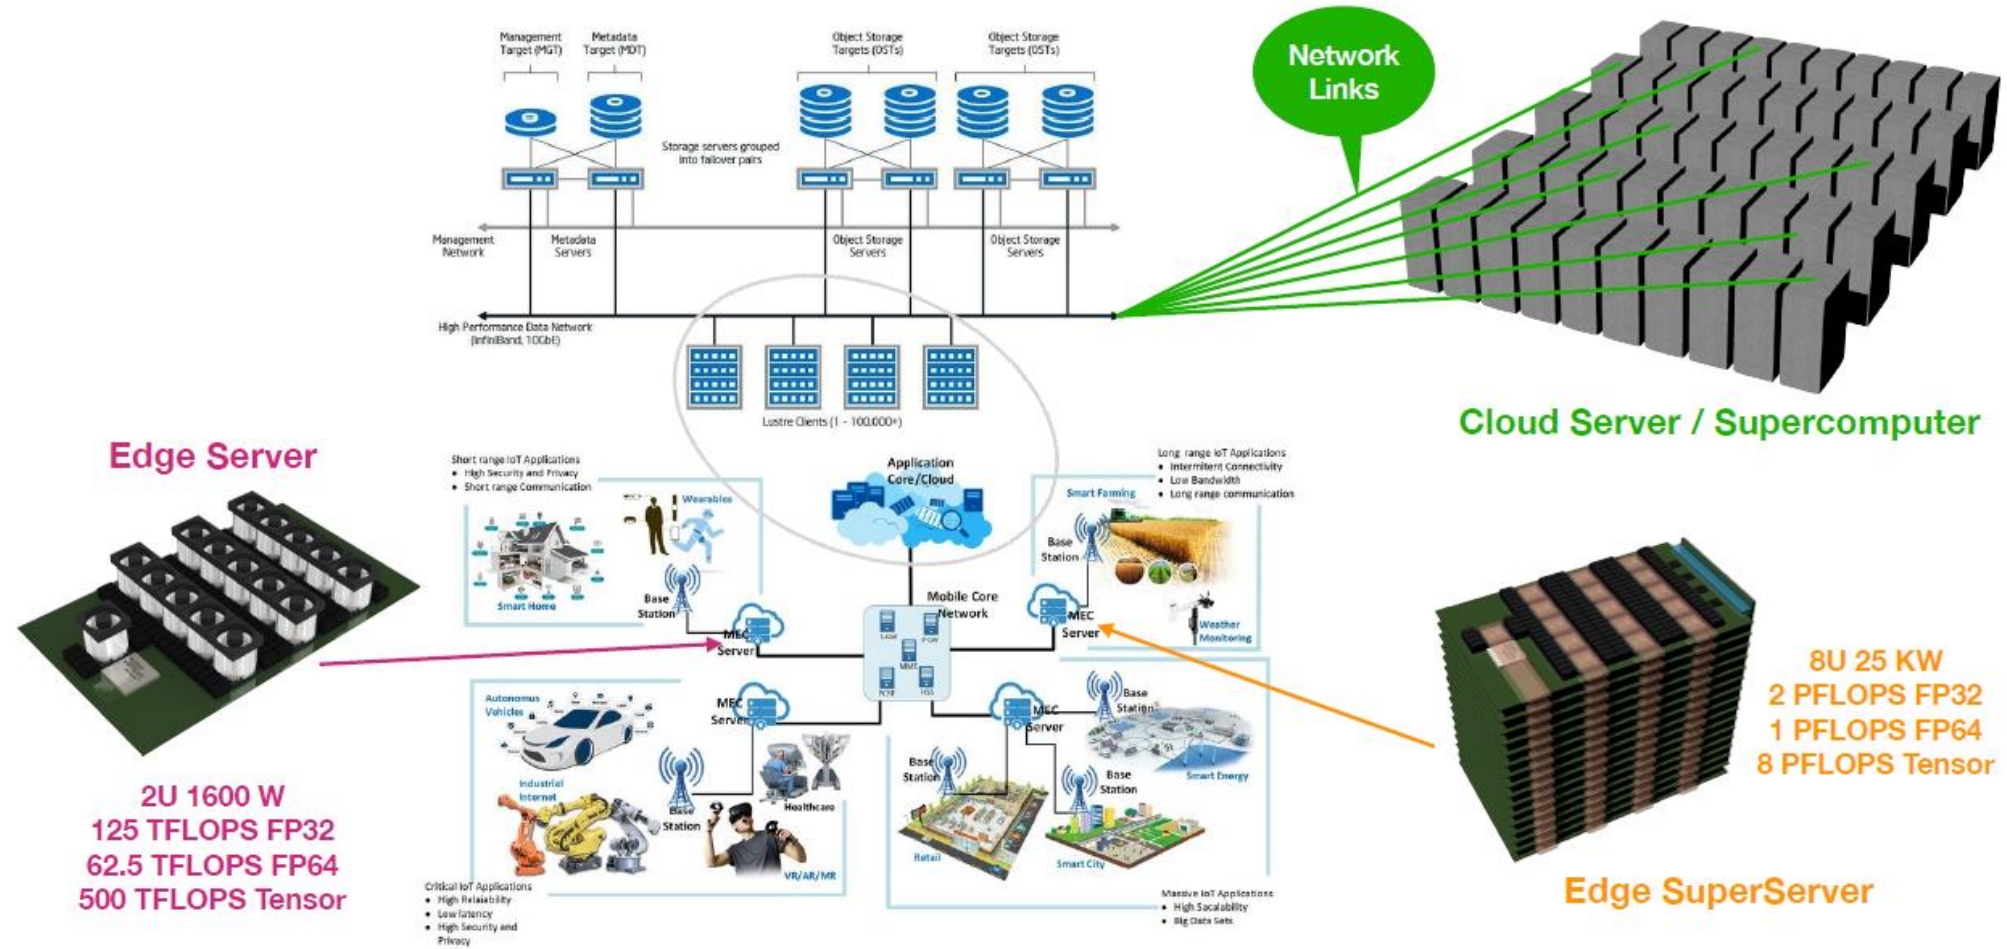


“Open Source has become mainstream across all sectors of the software industry during the past 10 years. To a large extent, open software re-use has proven economically efficient. The level of maturity of Open Source Hardware (OSH) remains far lower than that of Open Source Software (OSS). However, business 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.”

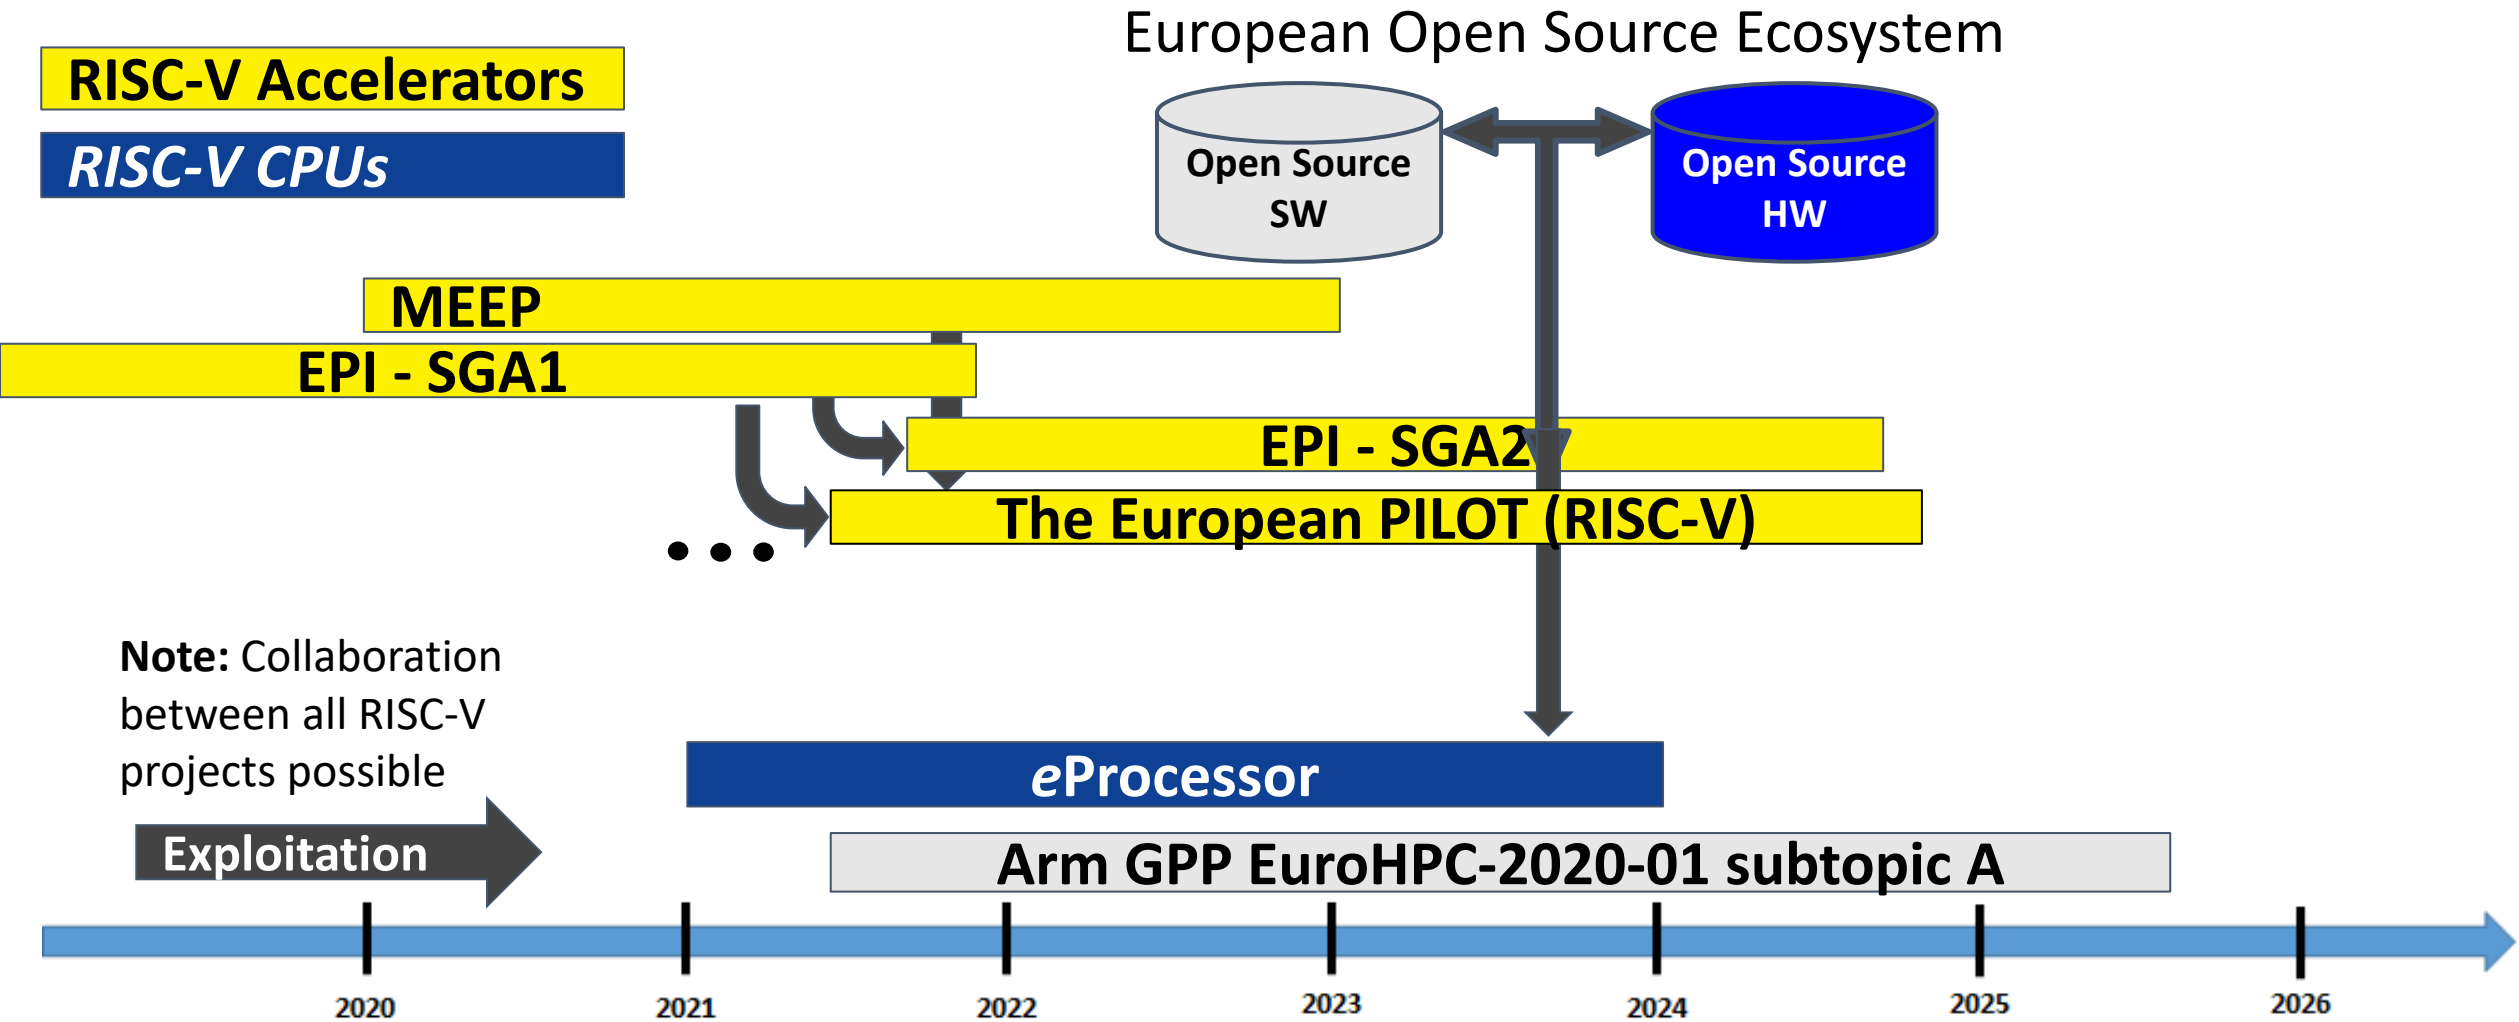
- DG Connect & DG IT Workshop, Brussels, Nov. 14-15, 2019

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

From IoT, Edge Computing, Clouds to Supercomputers

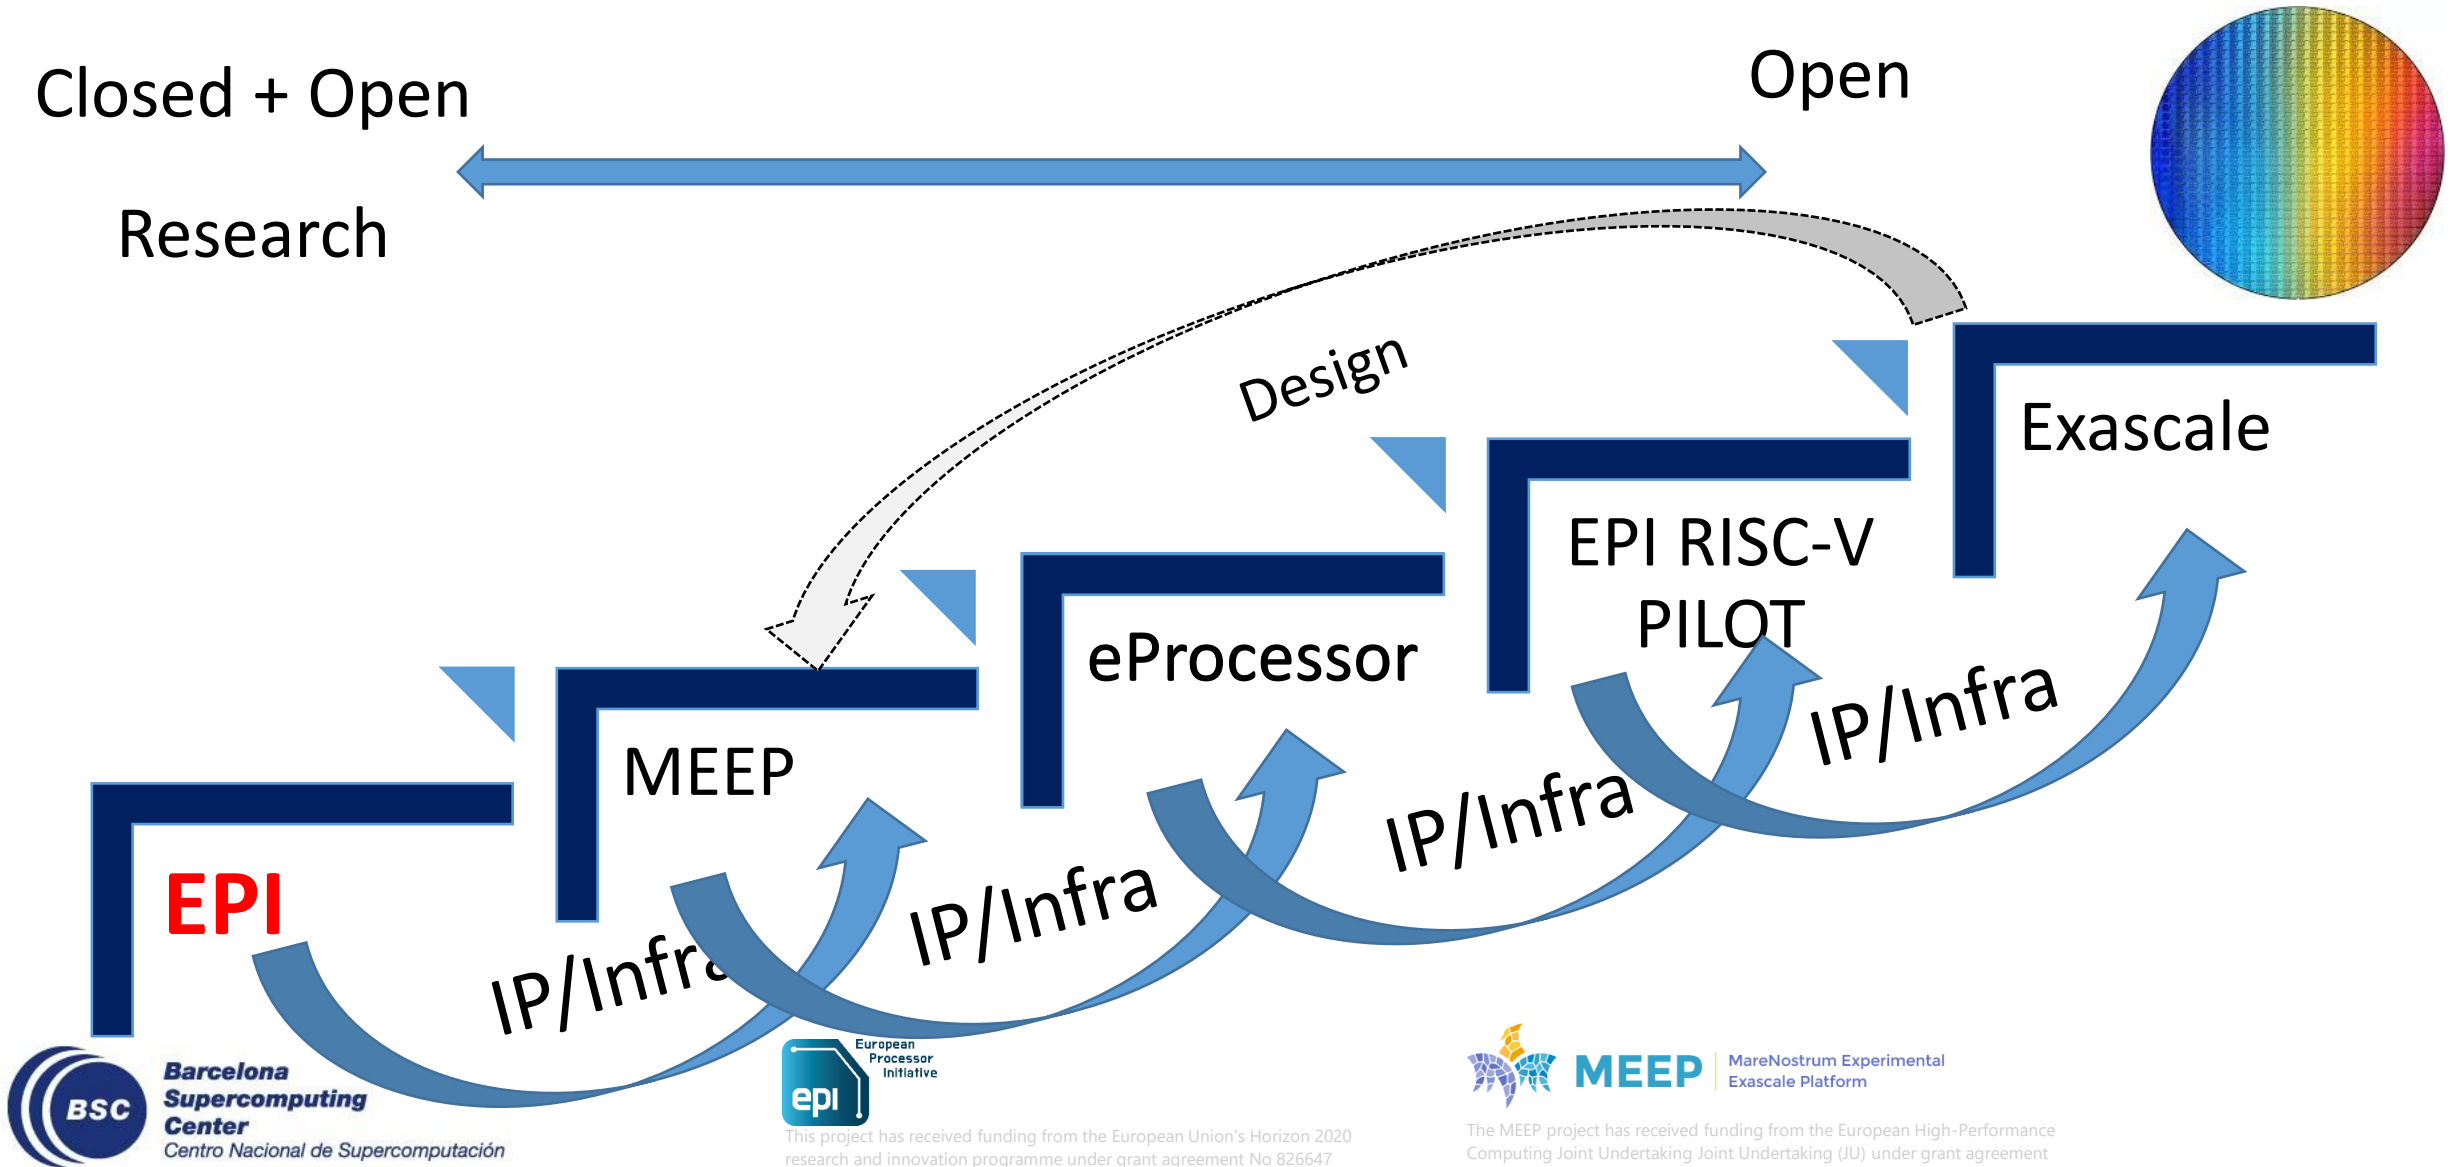


Current BSC RISC-V European Accelerator & CPU Landscape



Rebuilding the European CPU Industry

Production



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 27 partners from leading HPC industrial and academic centres

EPI Partners

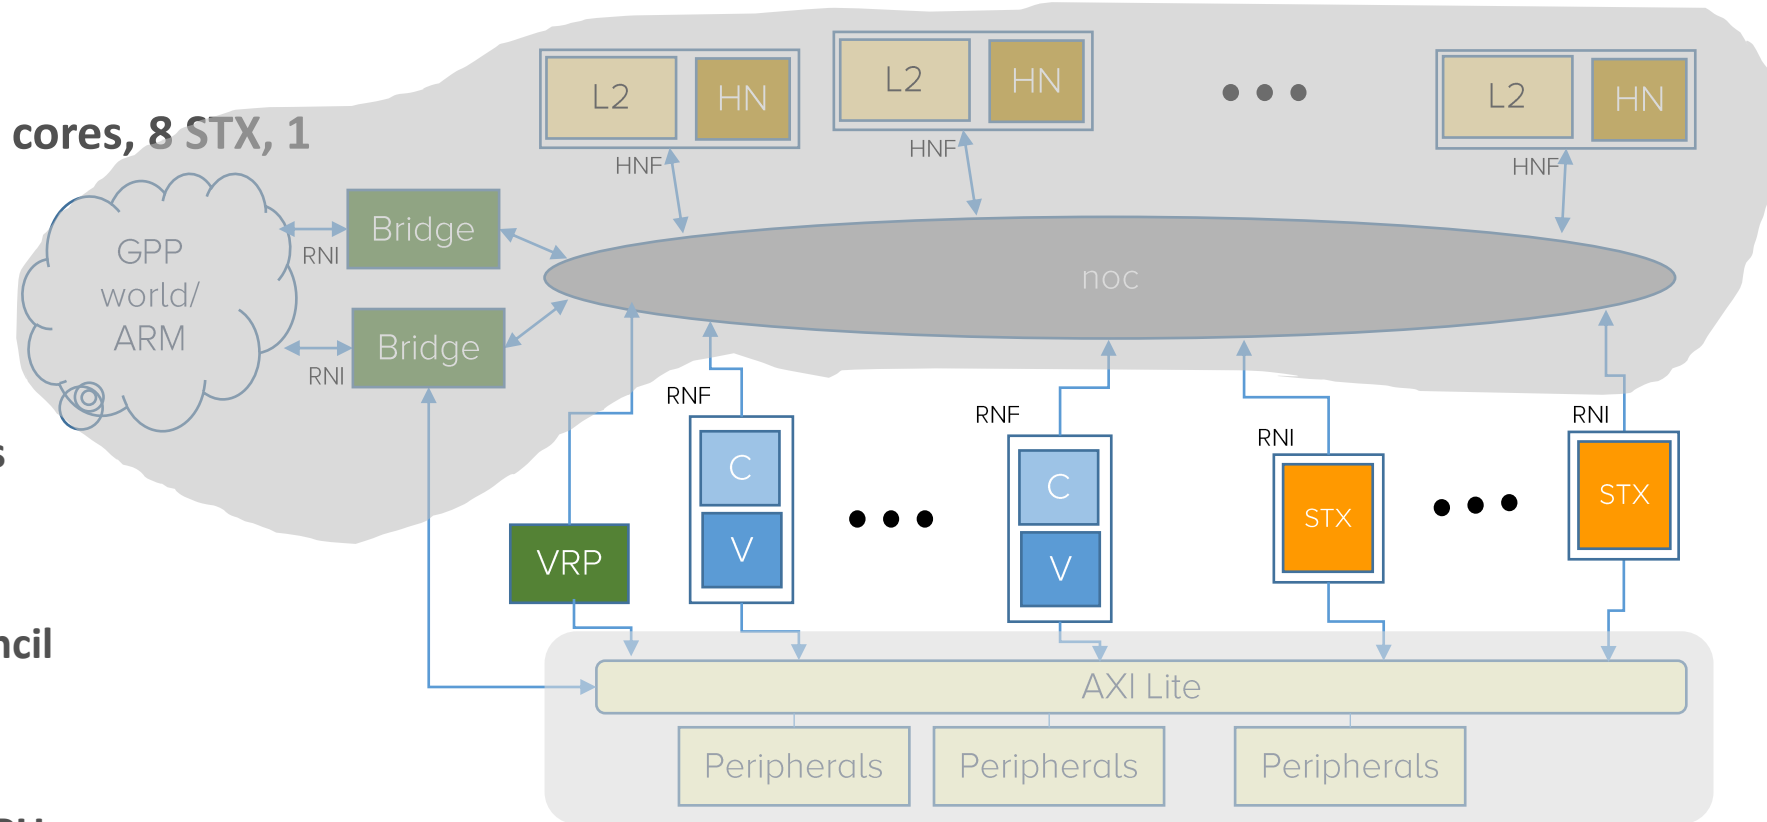


EPAC architecture

- **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**

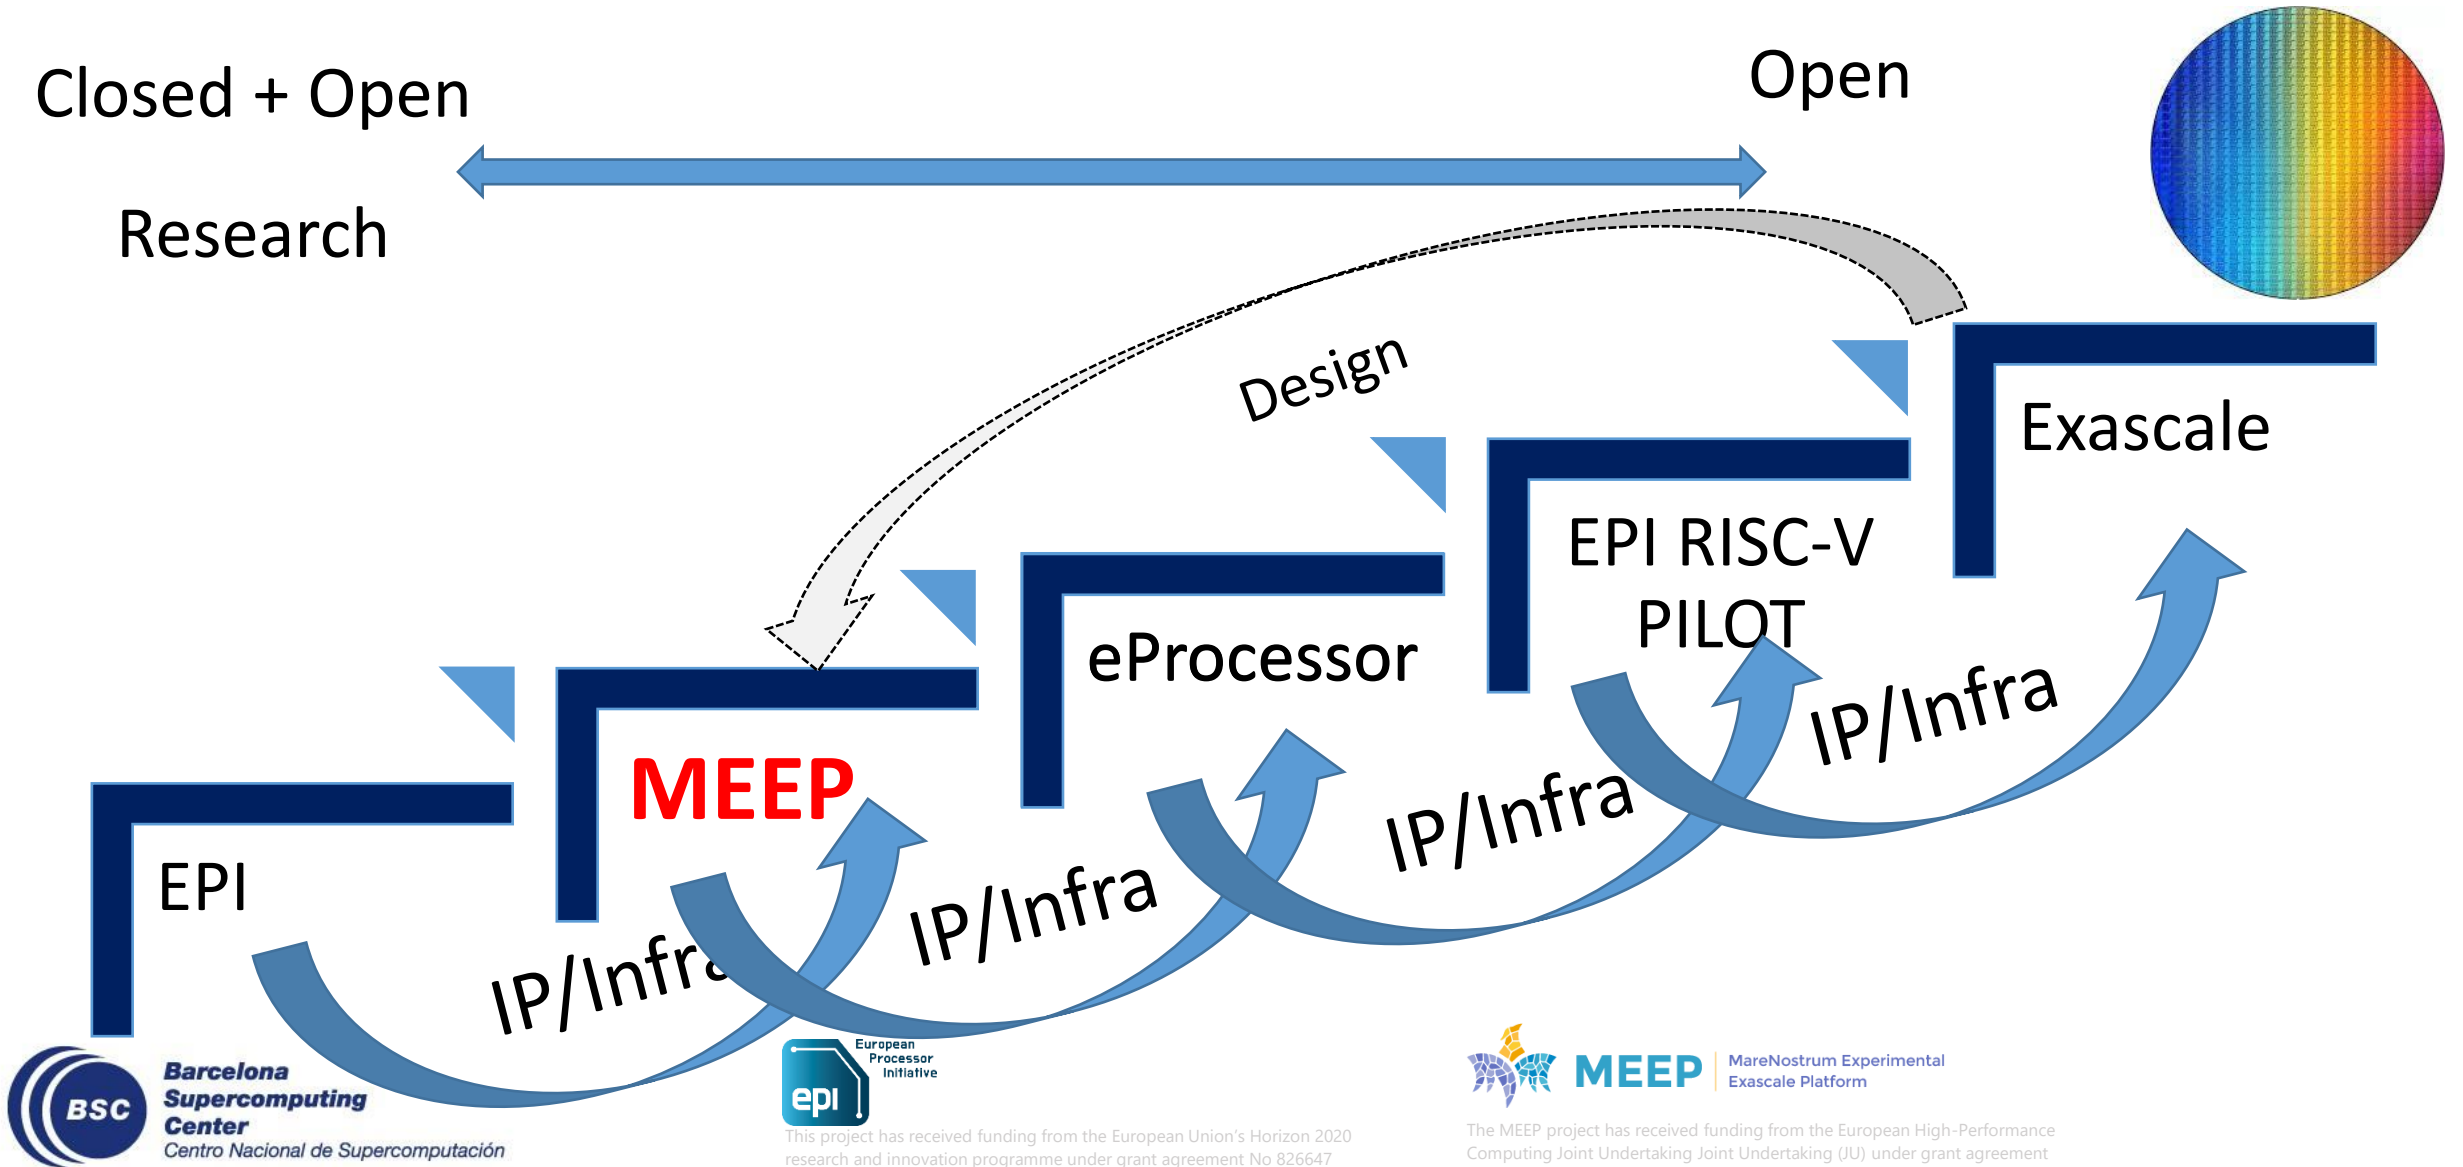
Main Contributors:

- **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
 - RISC-V + Extended precision FPU



Rebuilding the European CPU Industry

Production



MareNostrum 5. A European pre-exascale supercomputer

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

Hosting Consortium:

Spain Portugal Turkey Croatia



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



MEEP: MareNostrum Experiment Exascale Platform

MEEP is a flexible FPGA-based emulation platform that will explore hardware/software co-designs for Exascale Supercomputers and other hardware targets, based on European-developed IP. The project provides two functions:



An evaluation platform of pre-silicon IP and ideas, at speed and scale.



A software development and experimentation platform to enable software readiness for new hardware.



MEEP

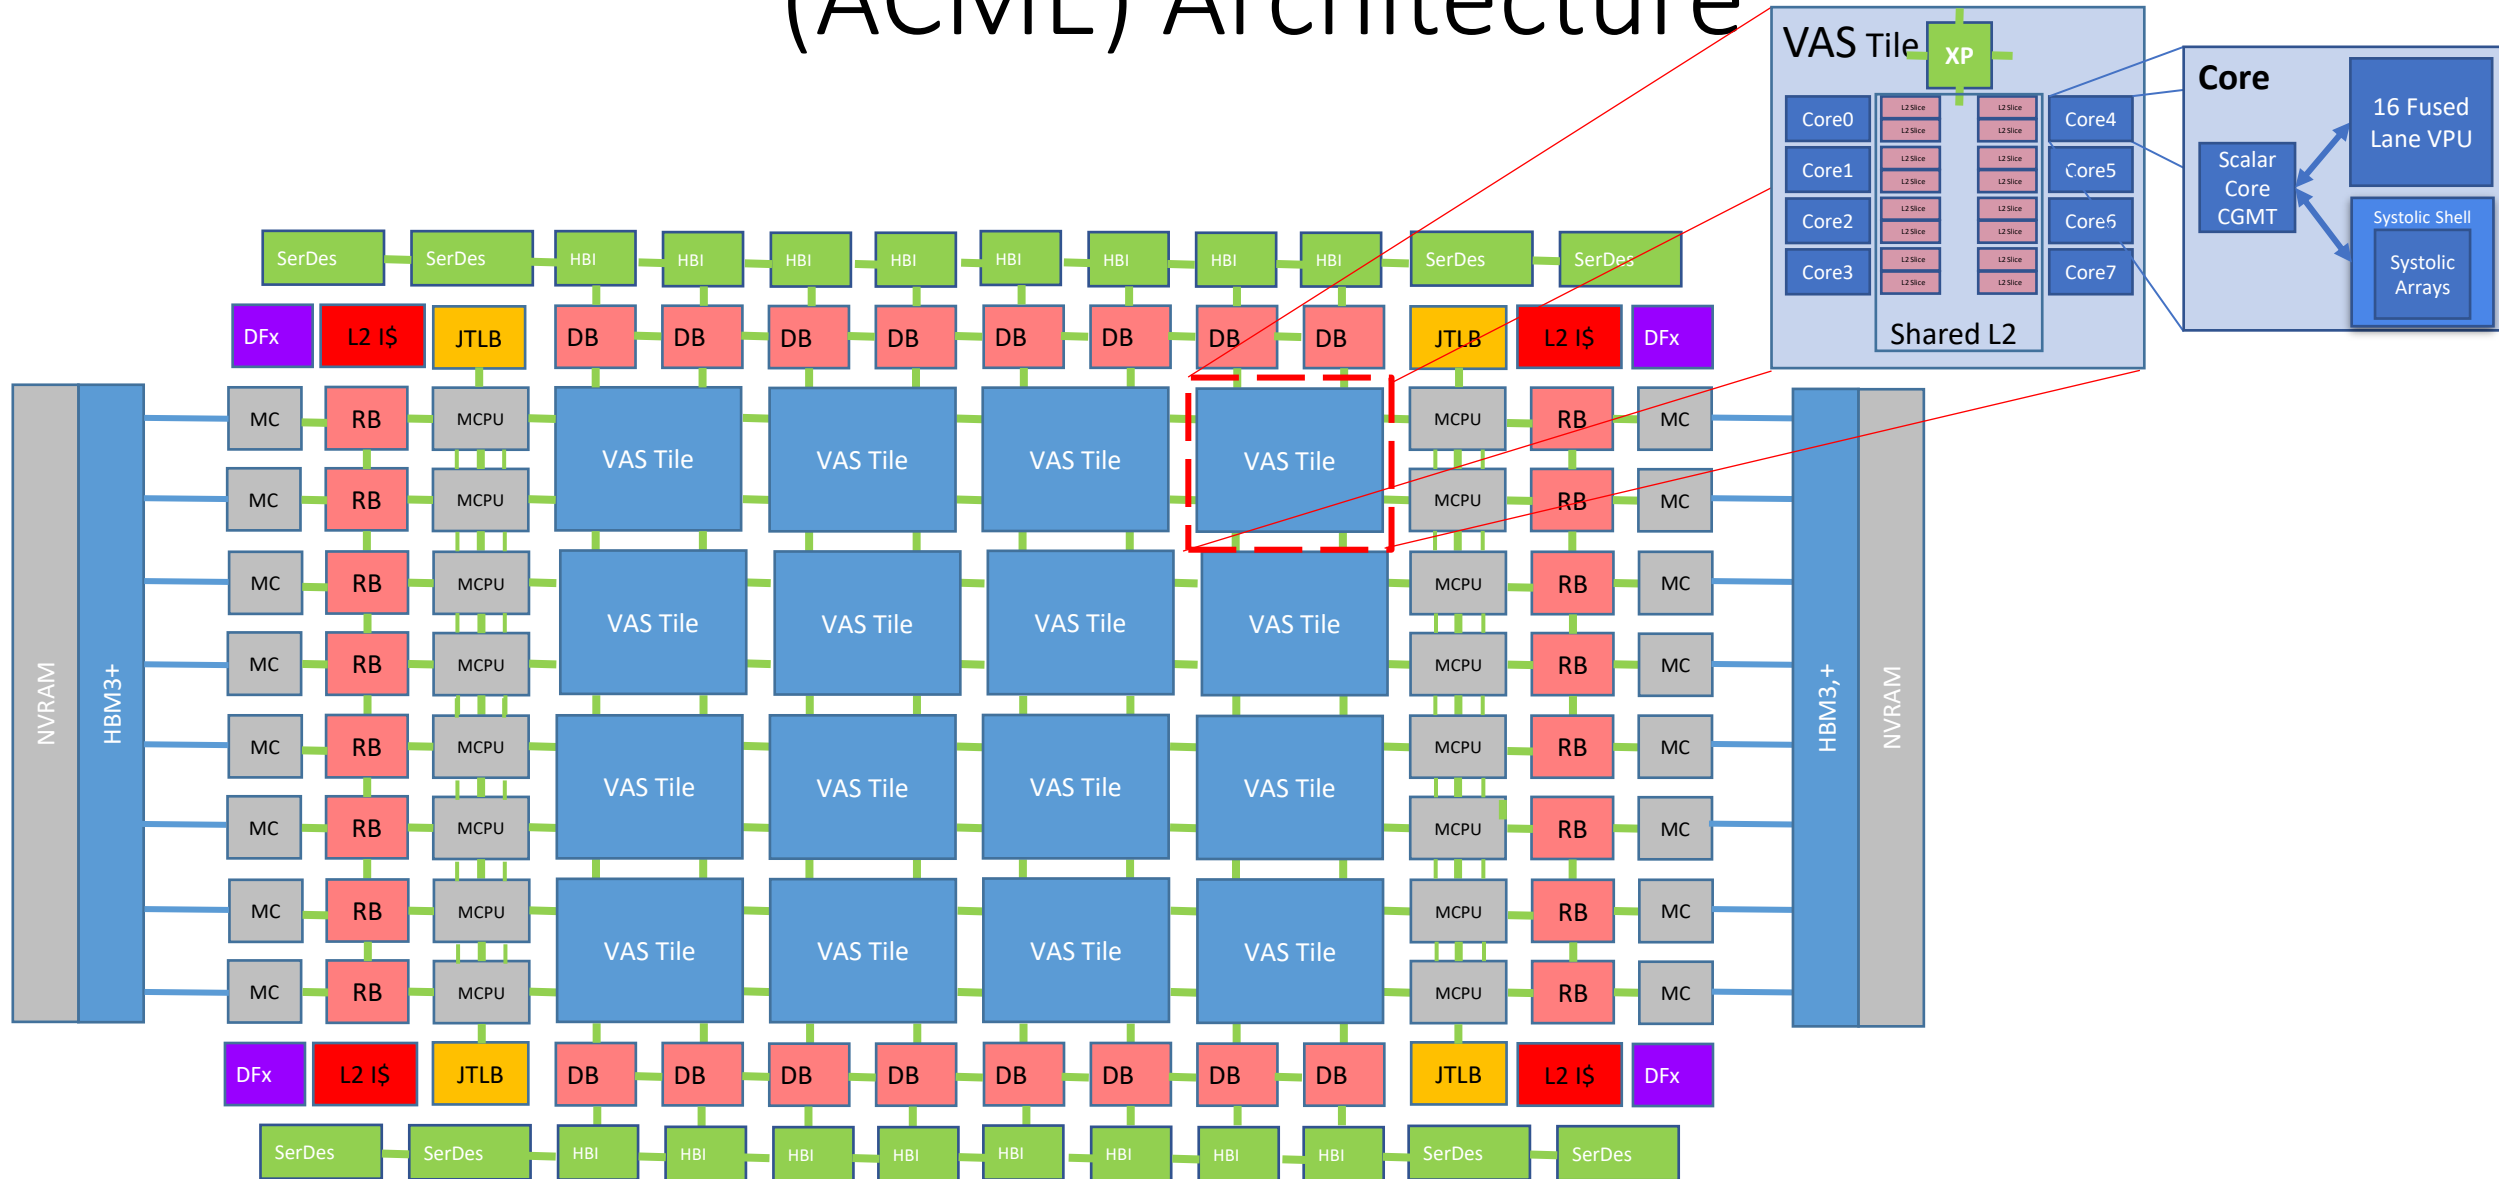
MareNostrum Experimental
Exascale Platform



The MEEP project has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU) under grant agreement No 946002. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Spain, Croatia, Turkey.

MEEP will enable software development, accelerating software maturity, compared to the limitations of software simulation. IP can be tested and validated before moving to silicon, saving time and money.

Accelerated Compute and Memory Engine (ACME) Architecture

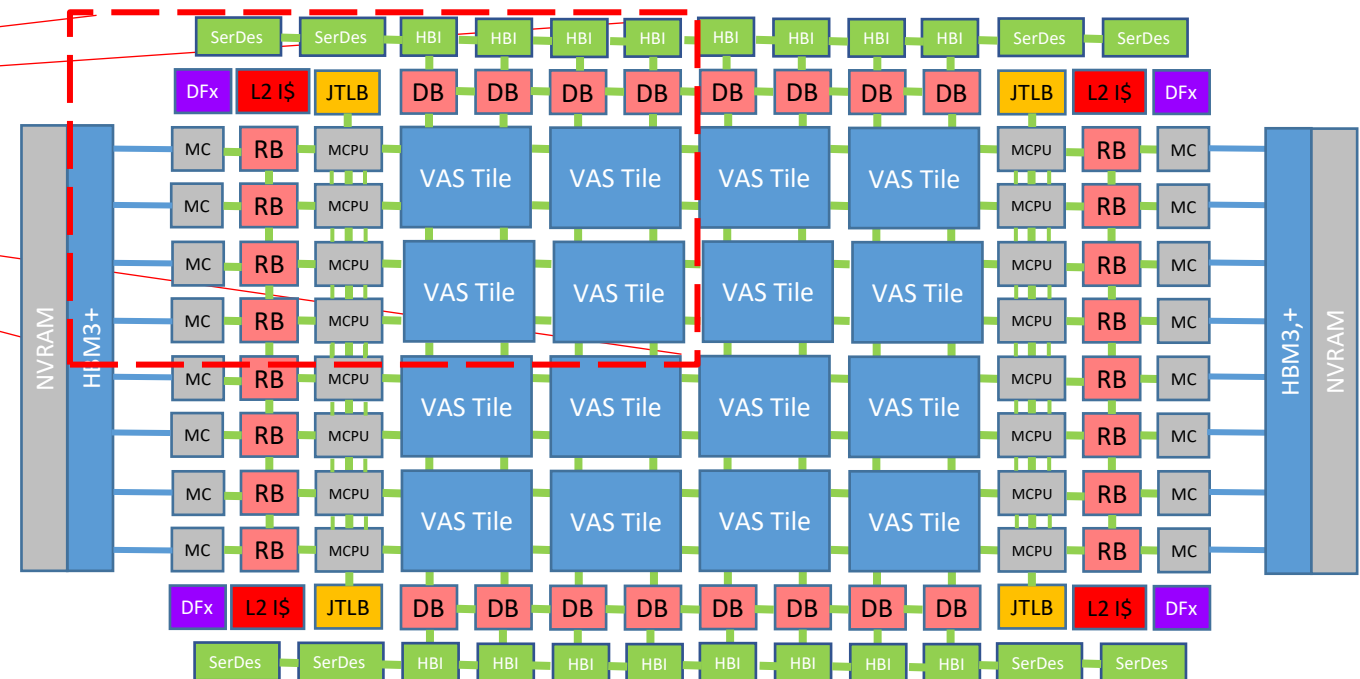
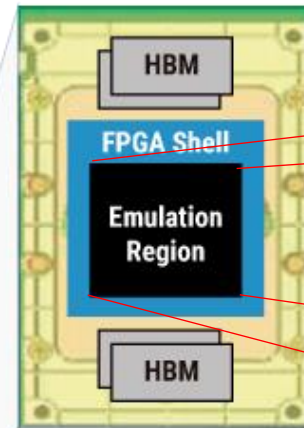
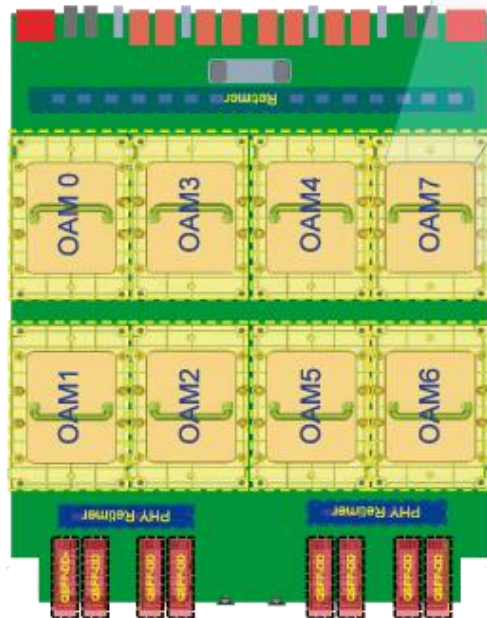


Mapping ACME Architecture to an FPGA

IoT to HPC, AI/ML/DL

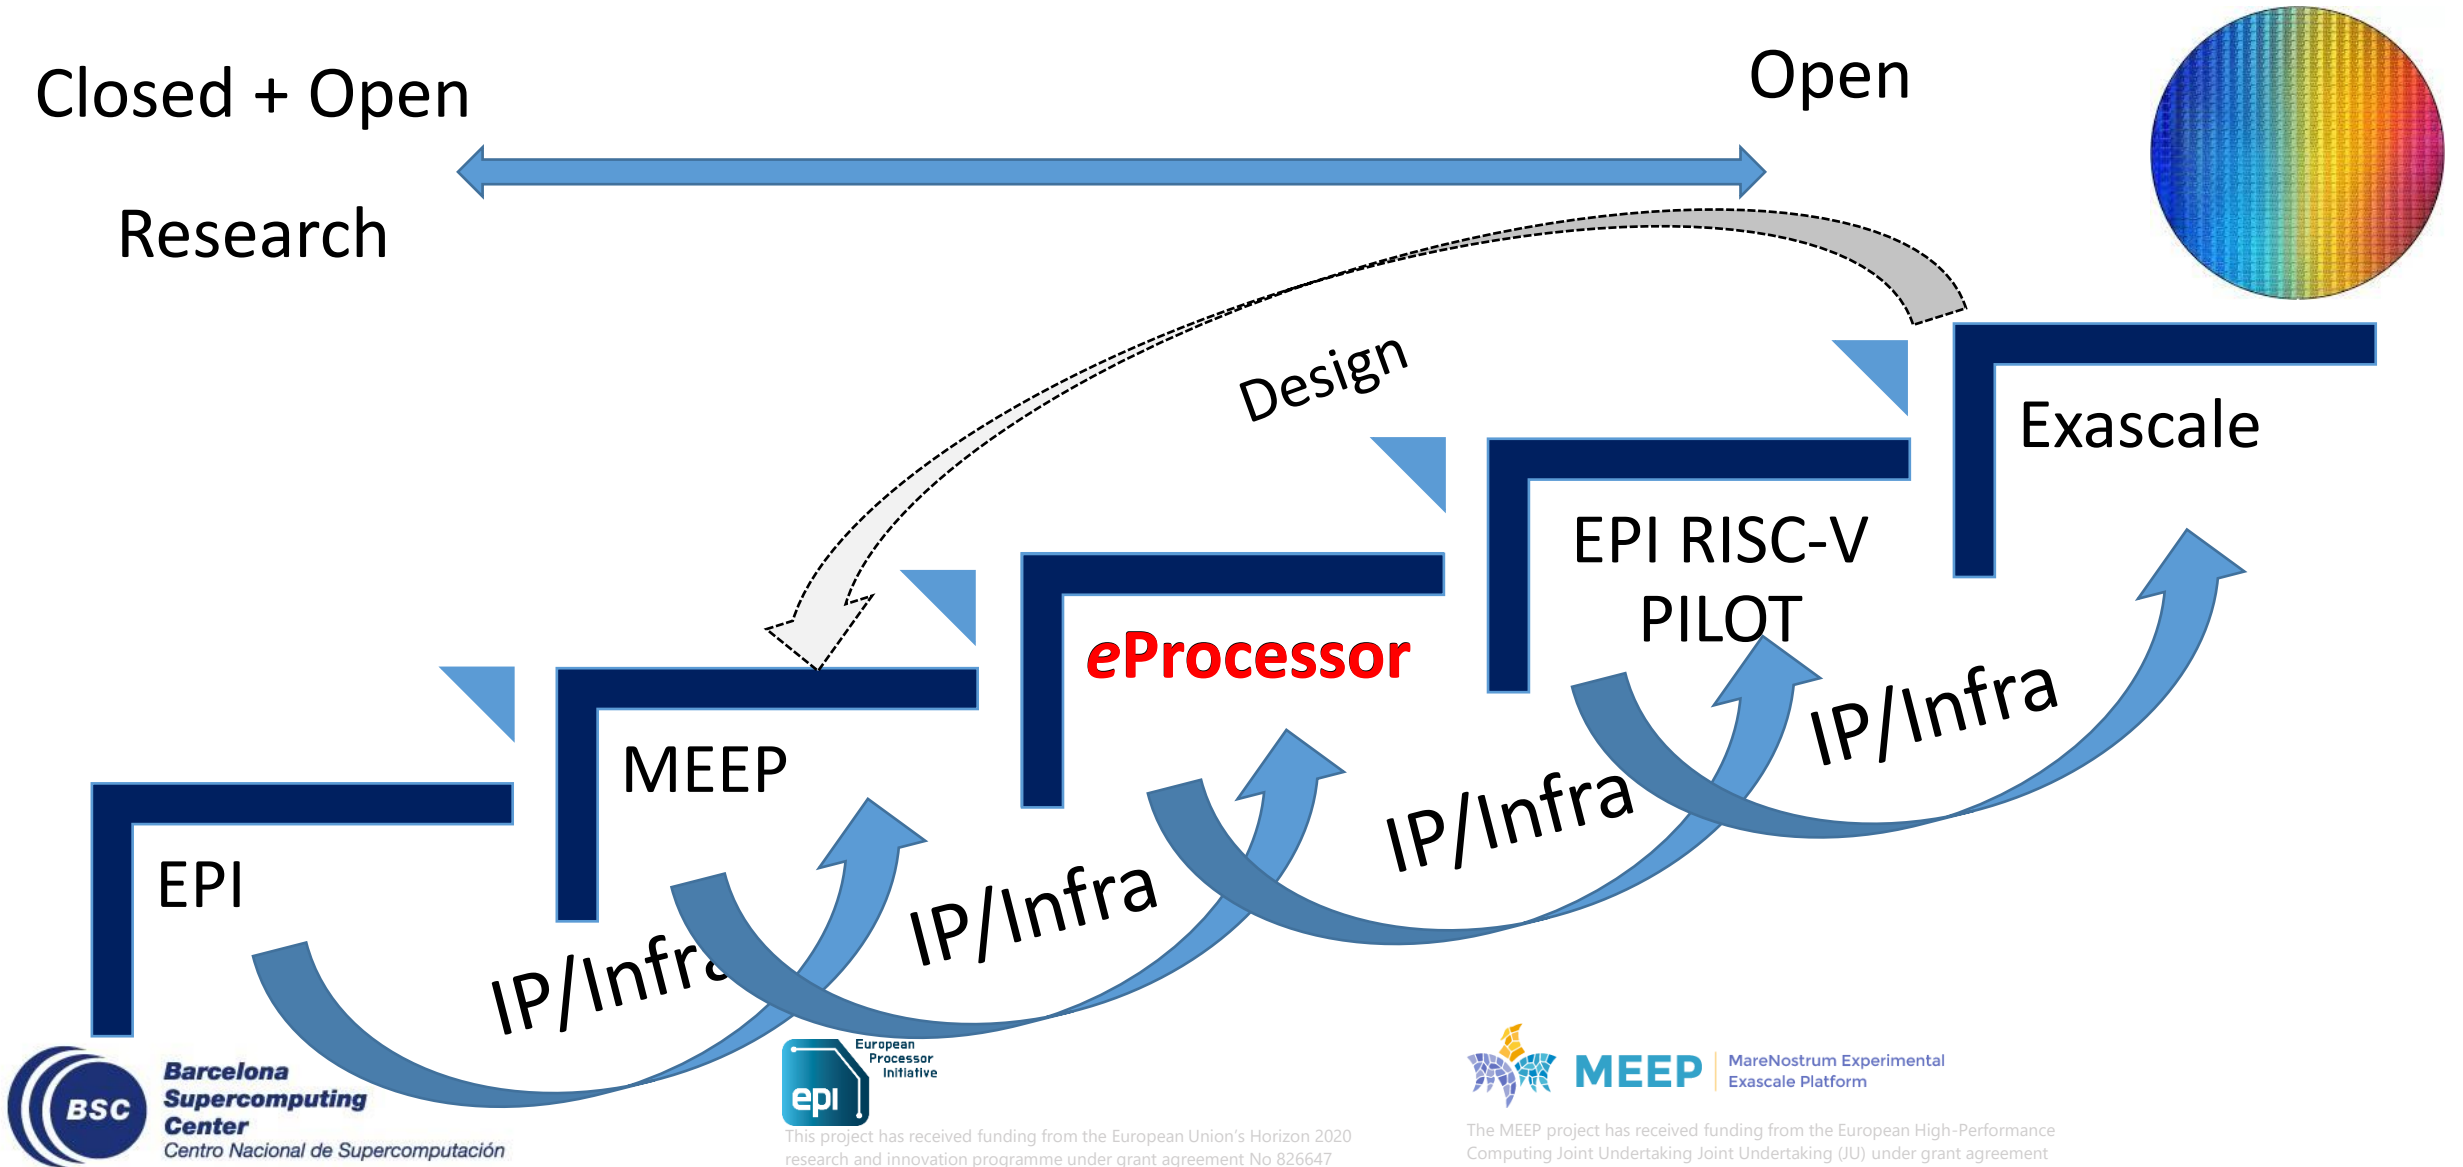


LINUX

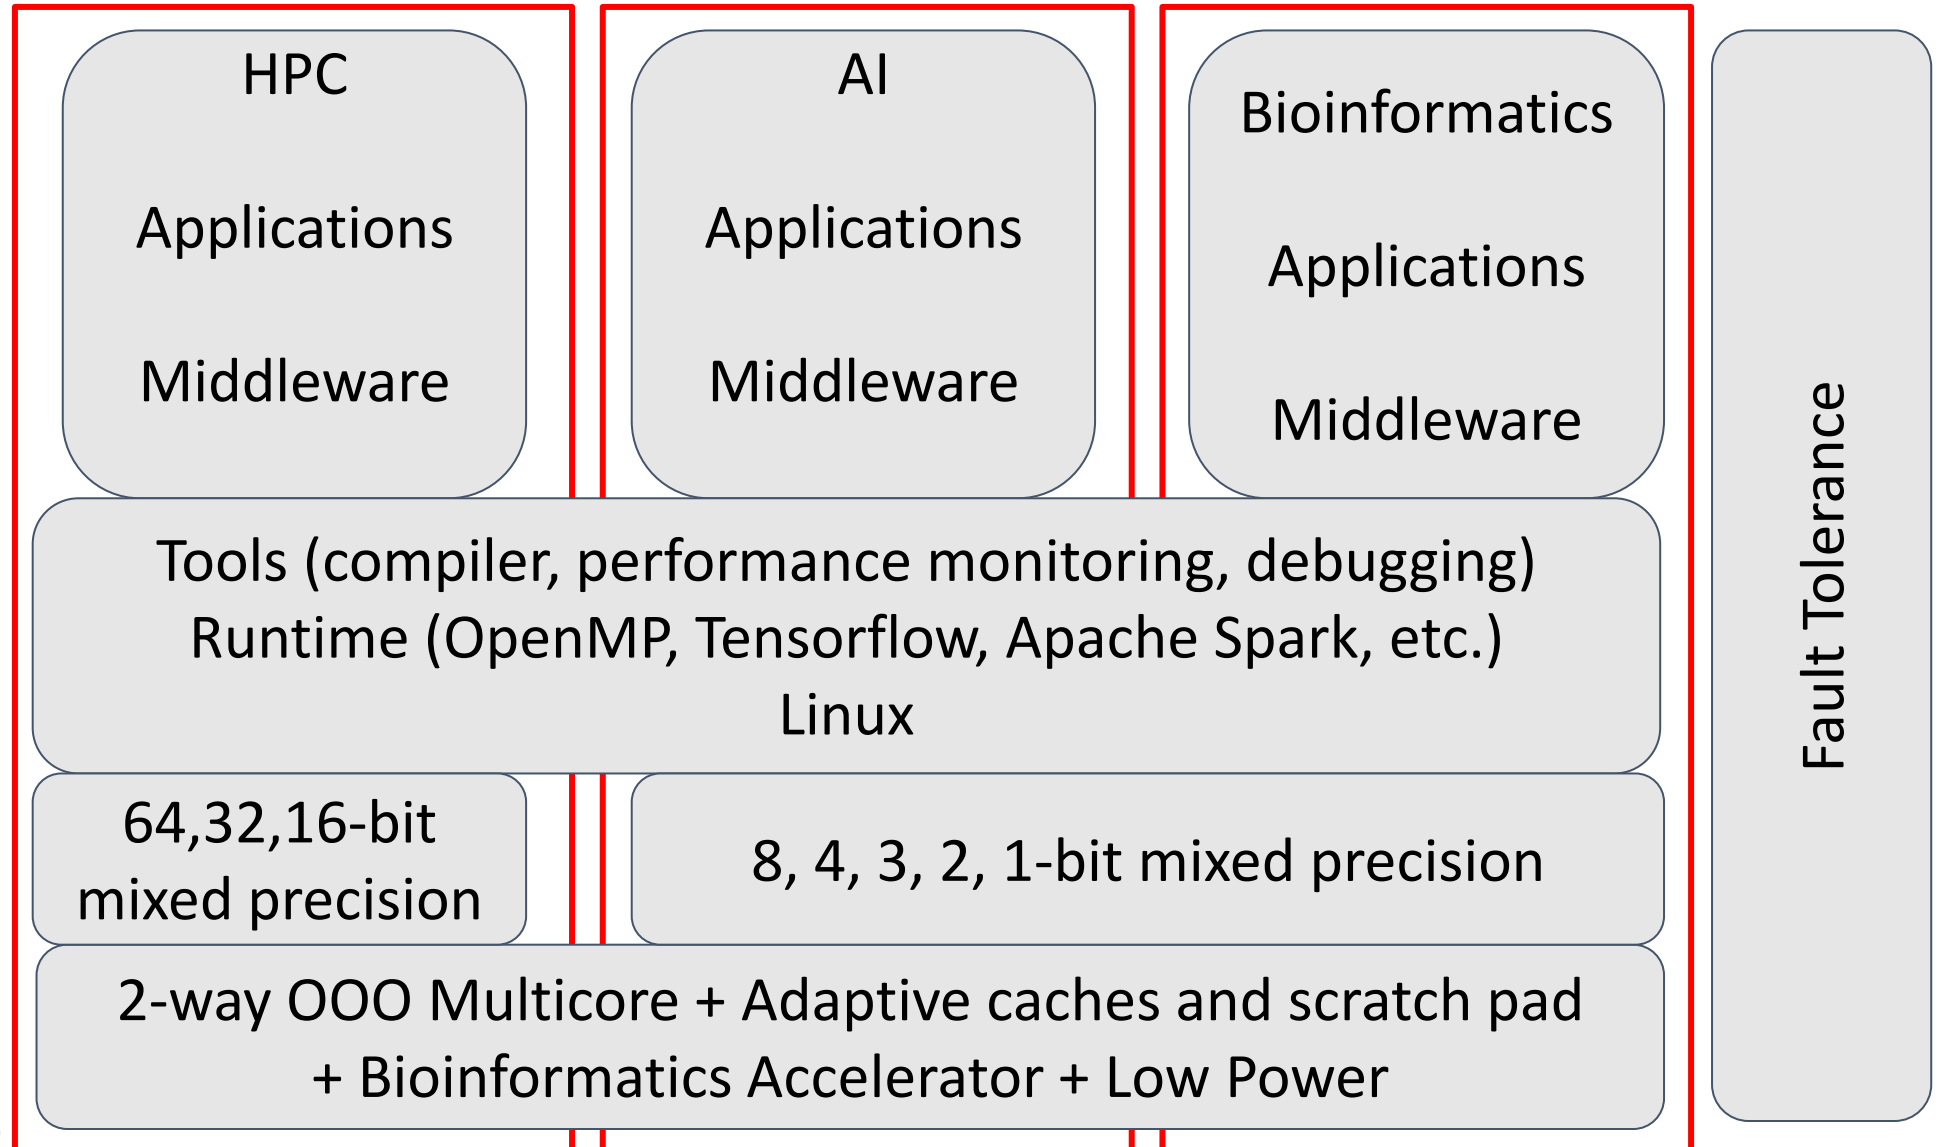


Rebuilding the European CPU Industry

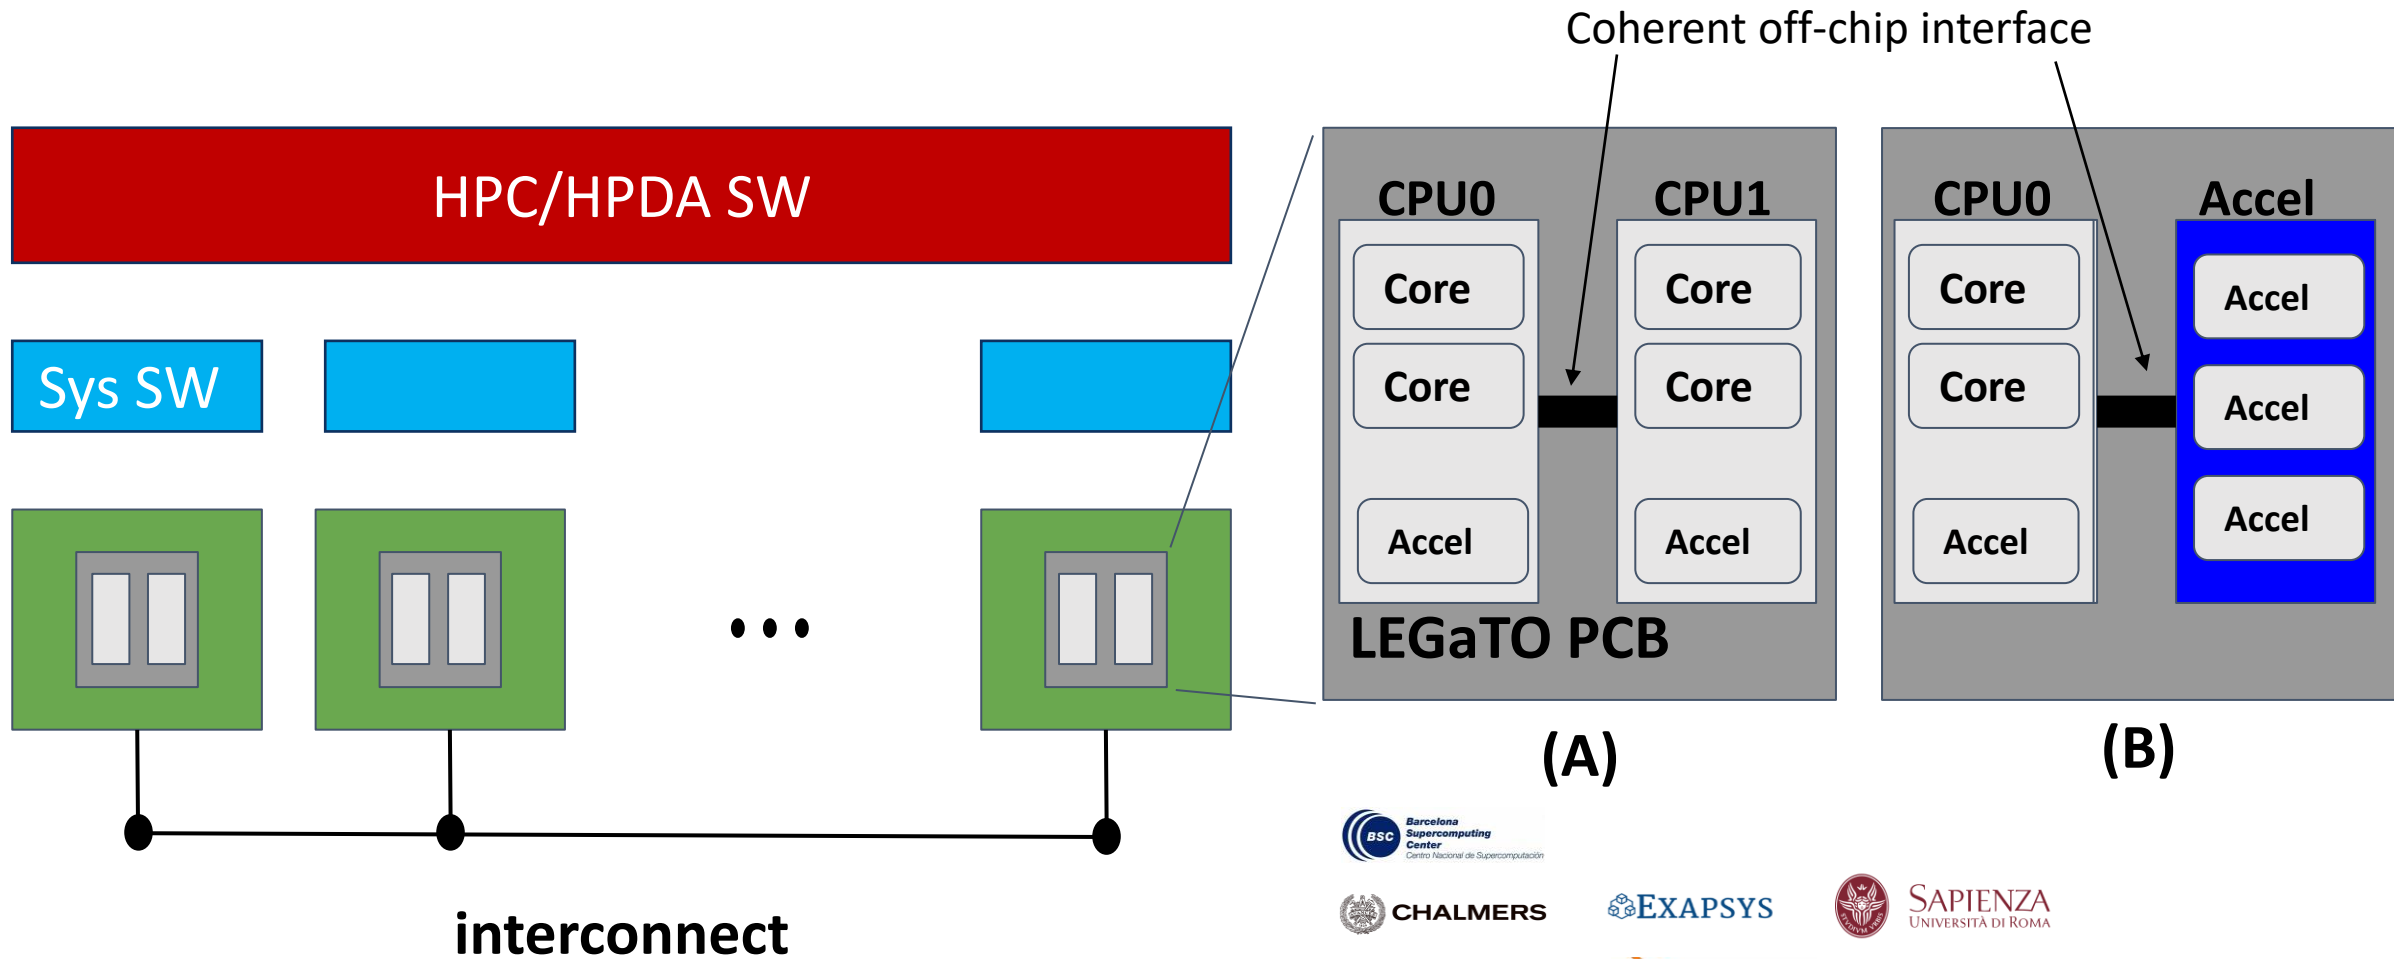
Production



eProcessor Stack

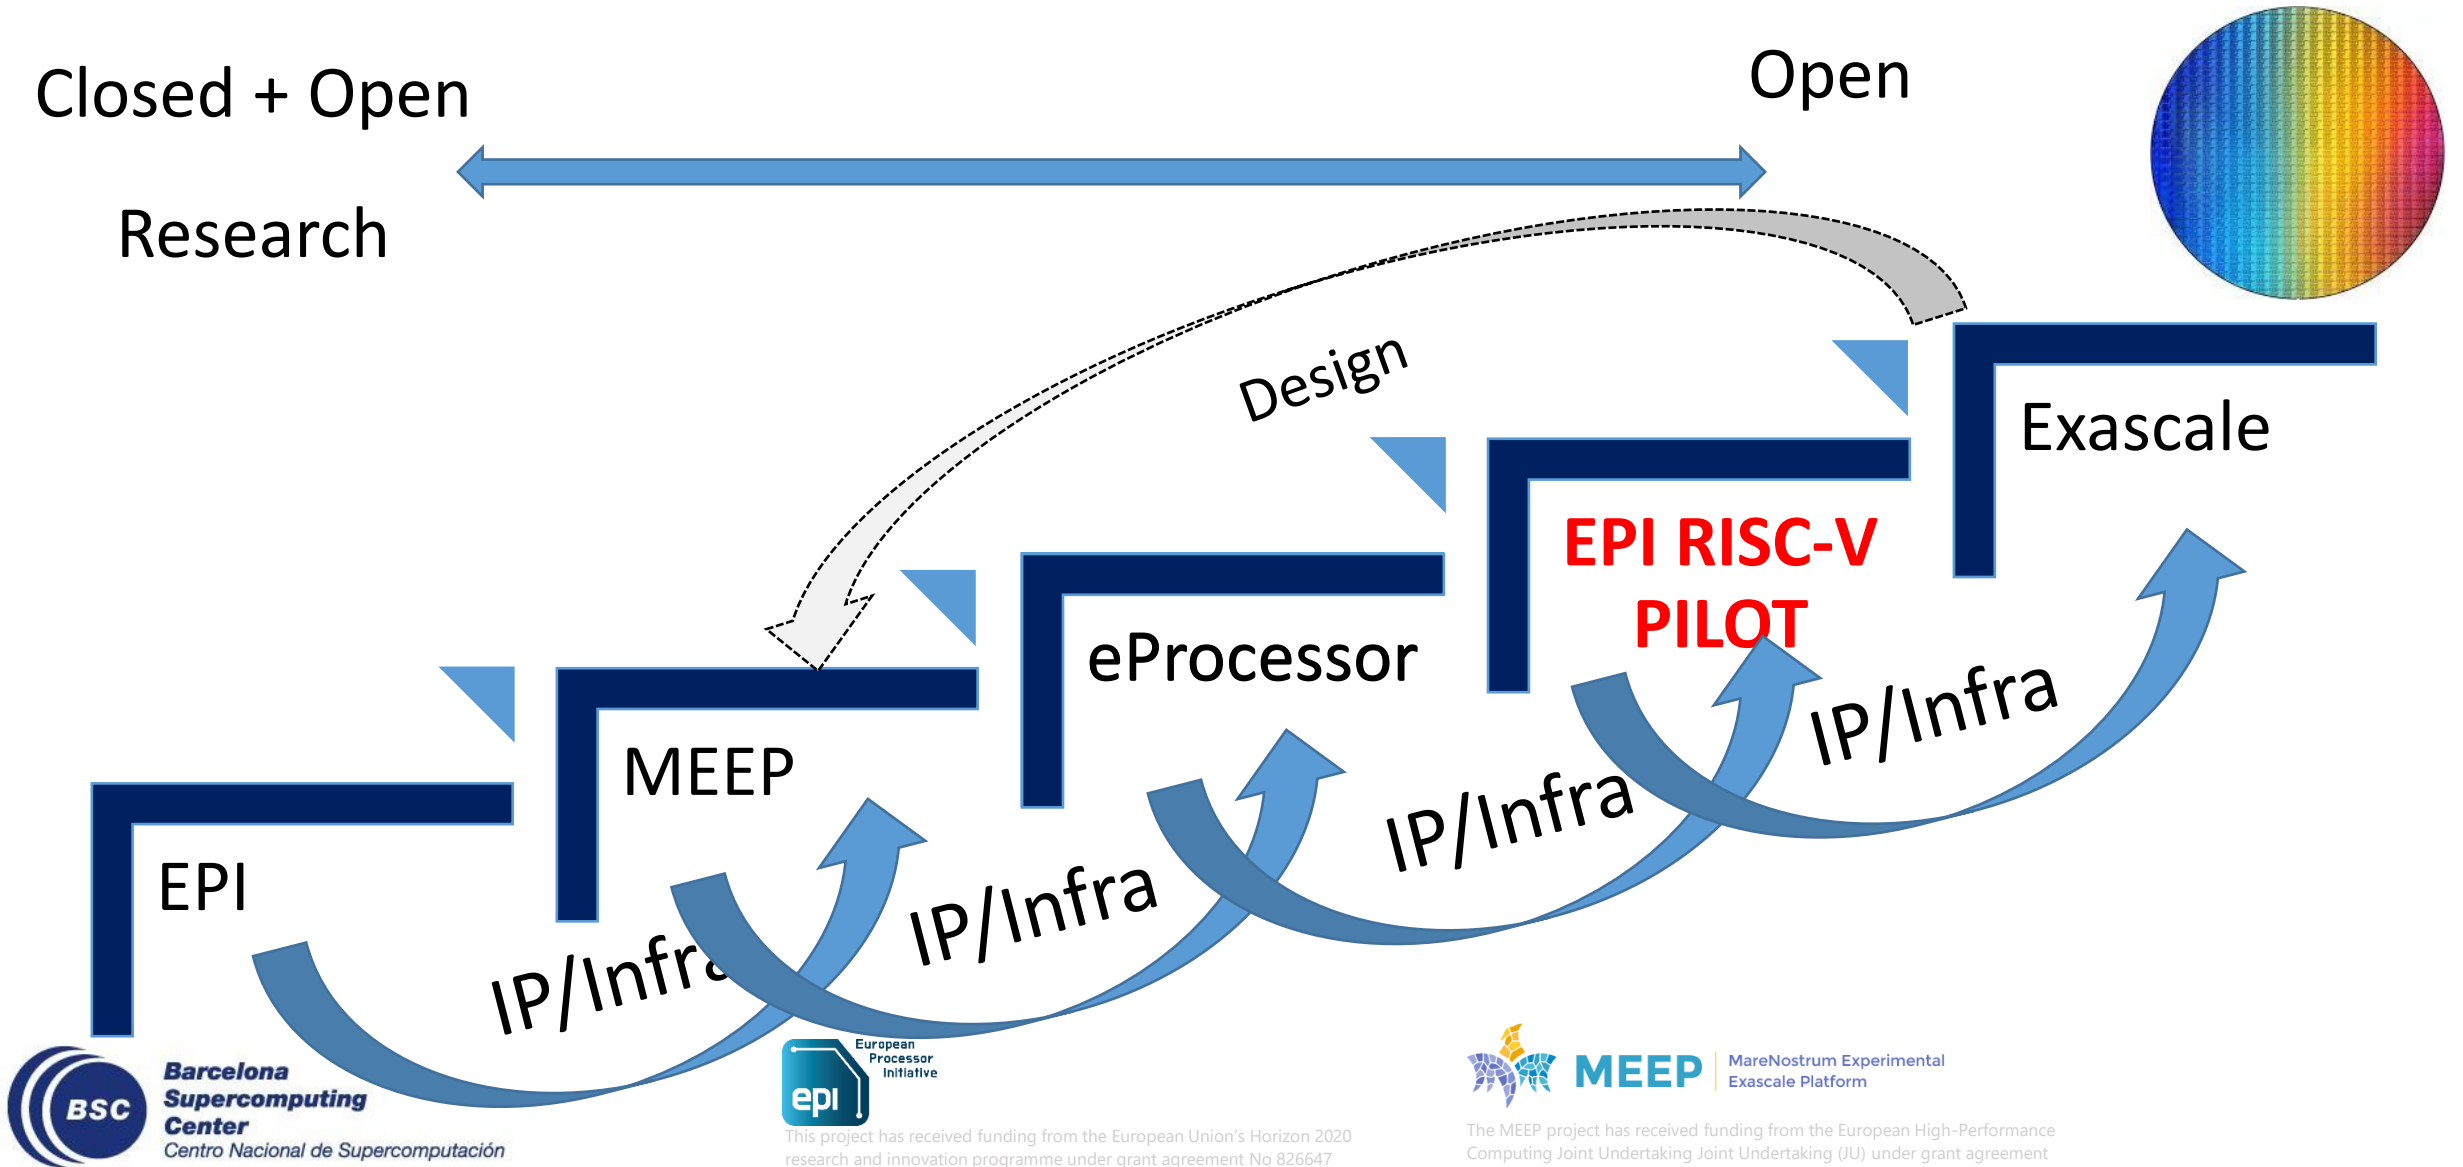


eProcessor System diagram



Rebuilding the European CPU Industry

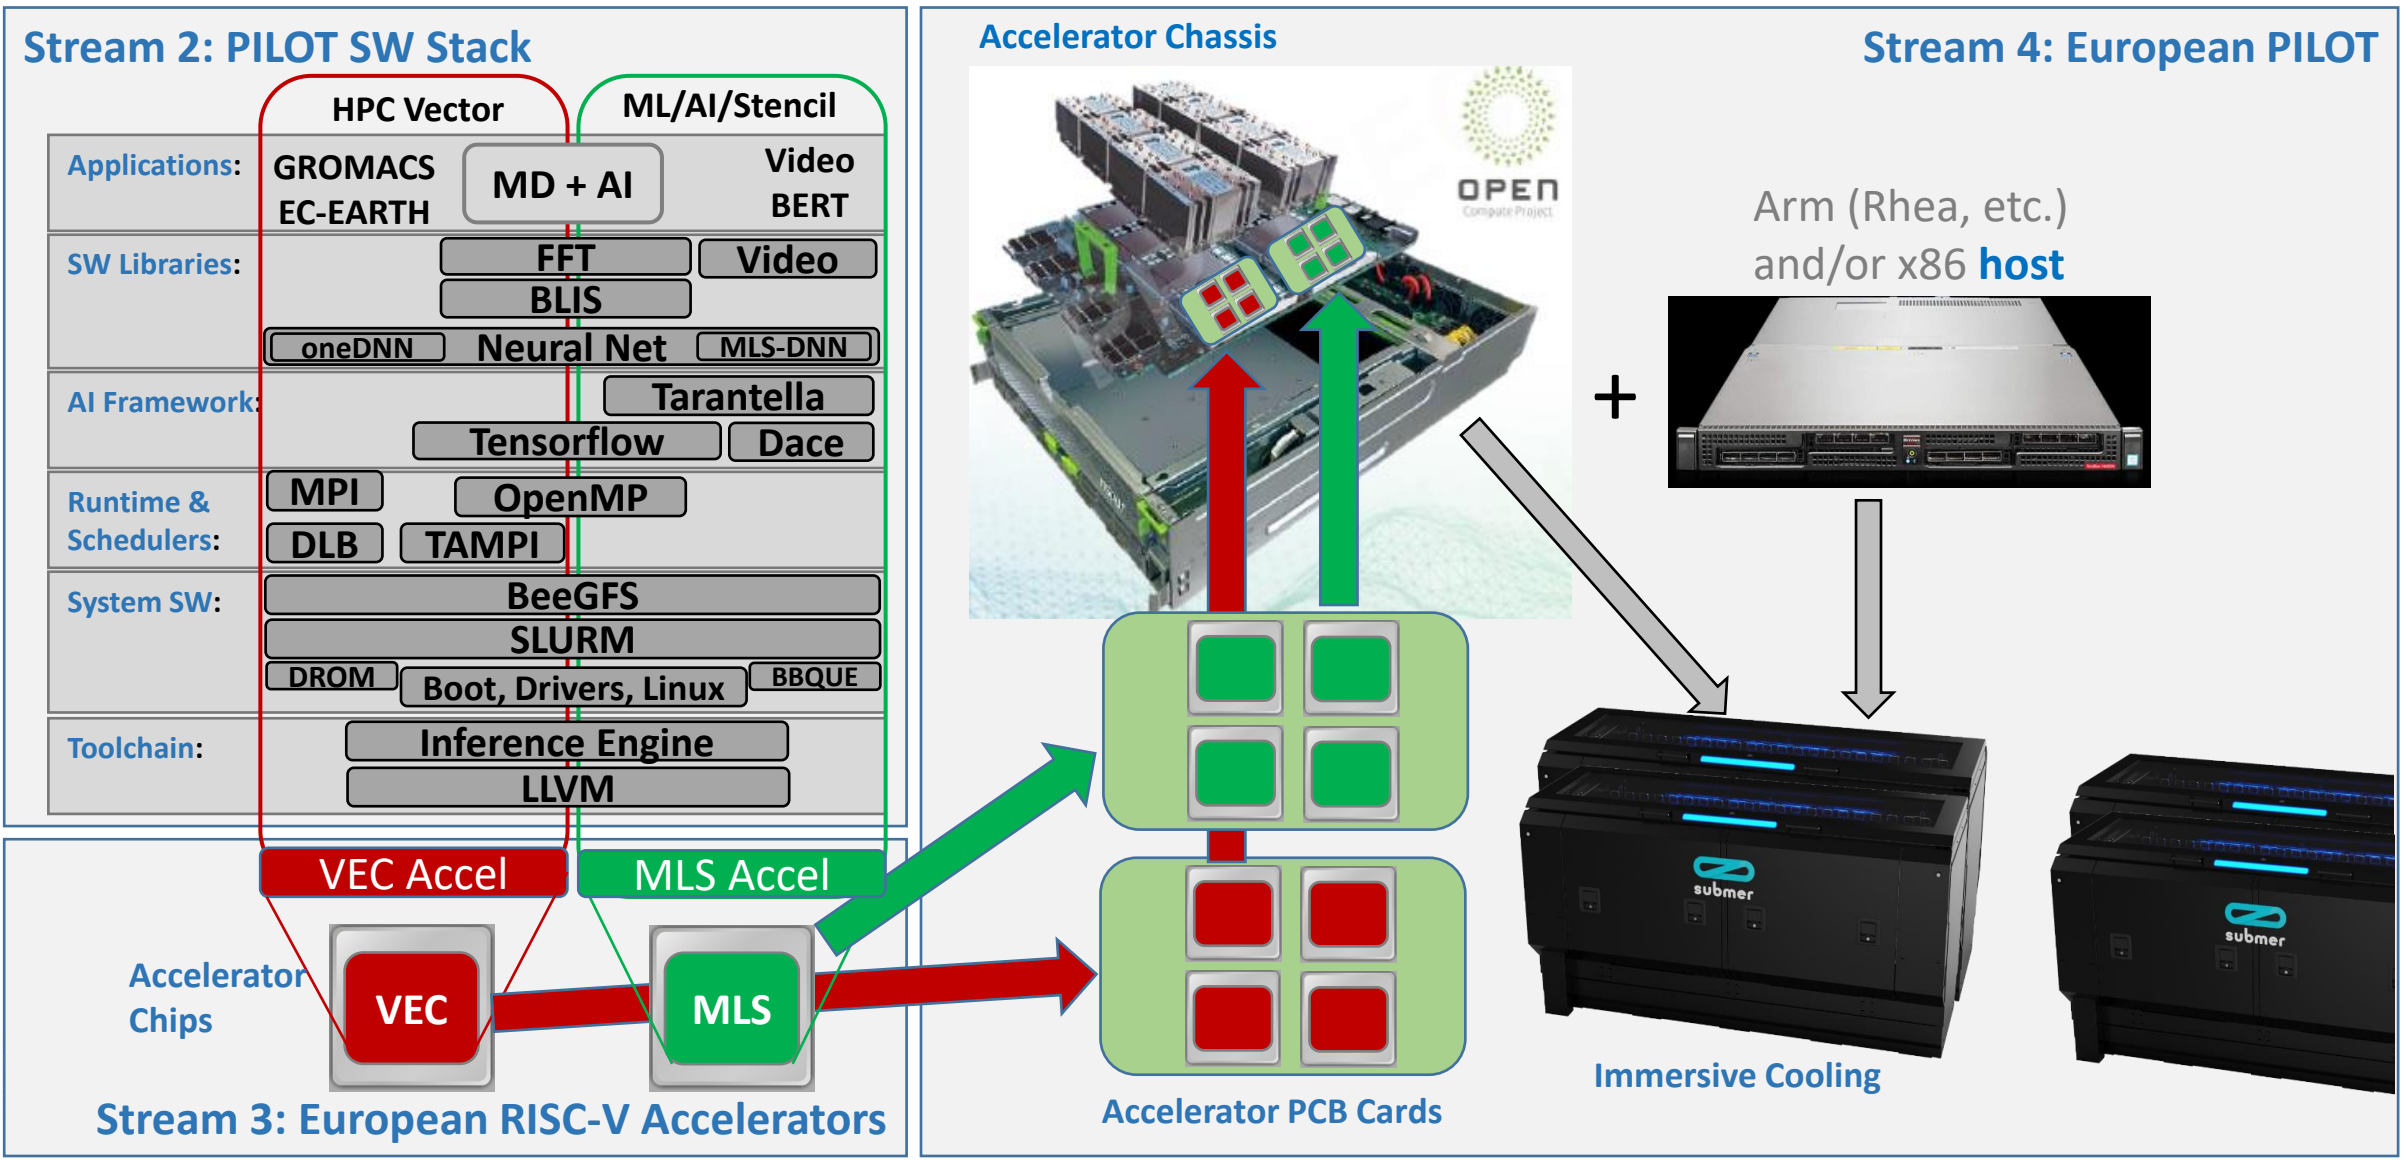
Production



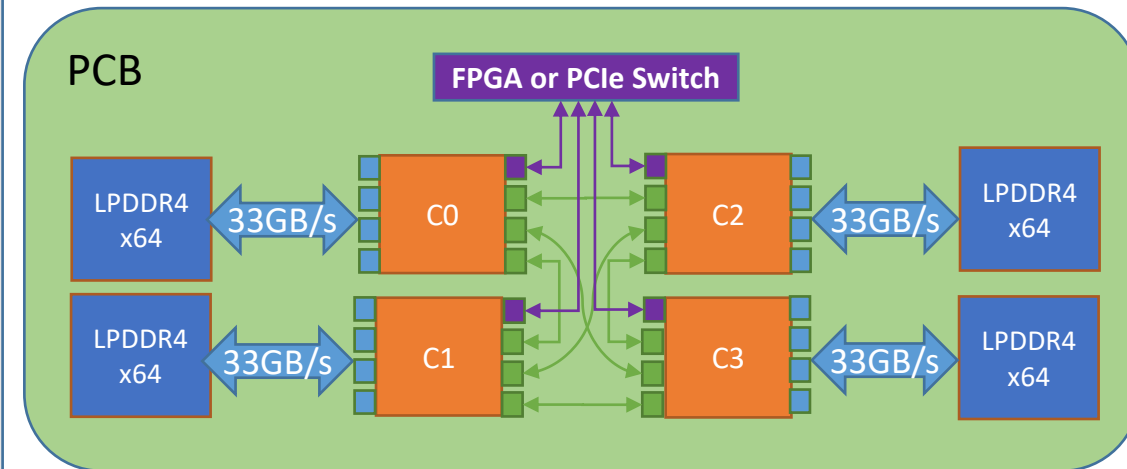
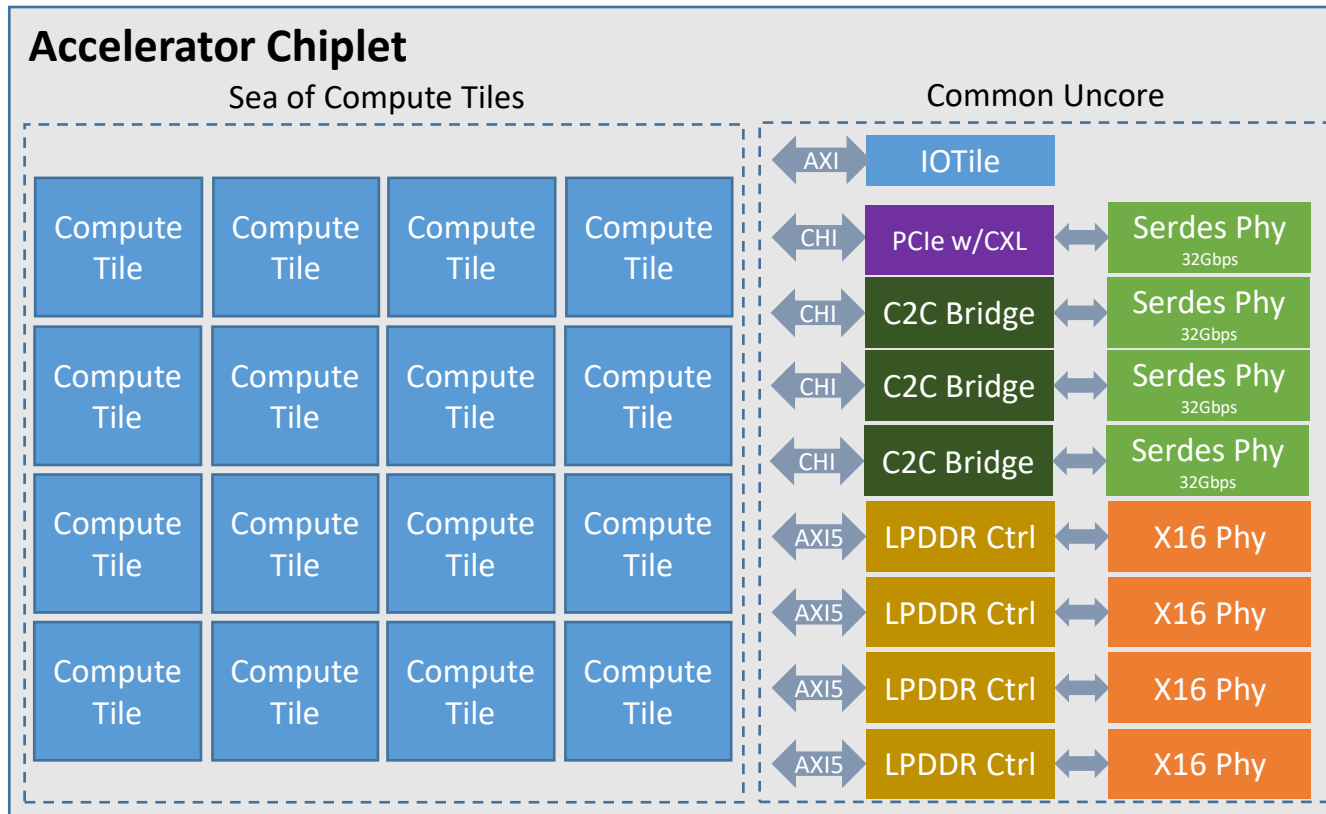
This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826647

The MEEP project has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU) under grant agreement No 946002. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Spain, Croatia, Turkey.

The European PILOT

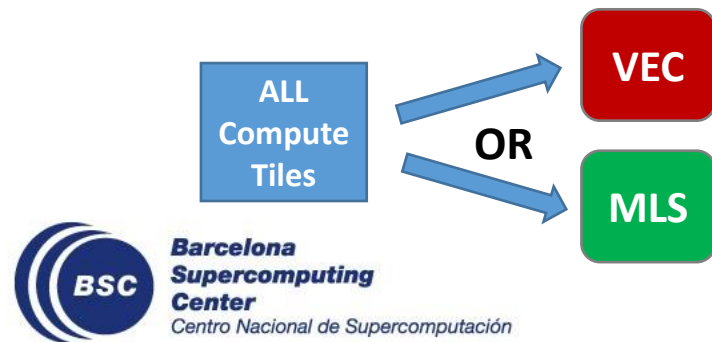


The European PILOT Accelerator Architecture



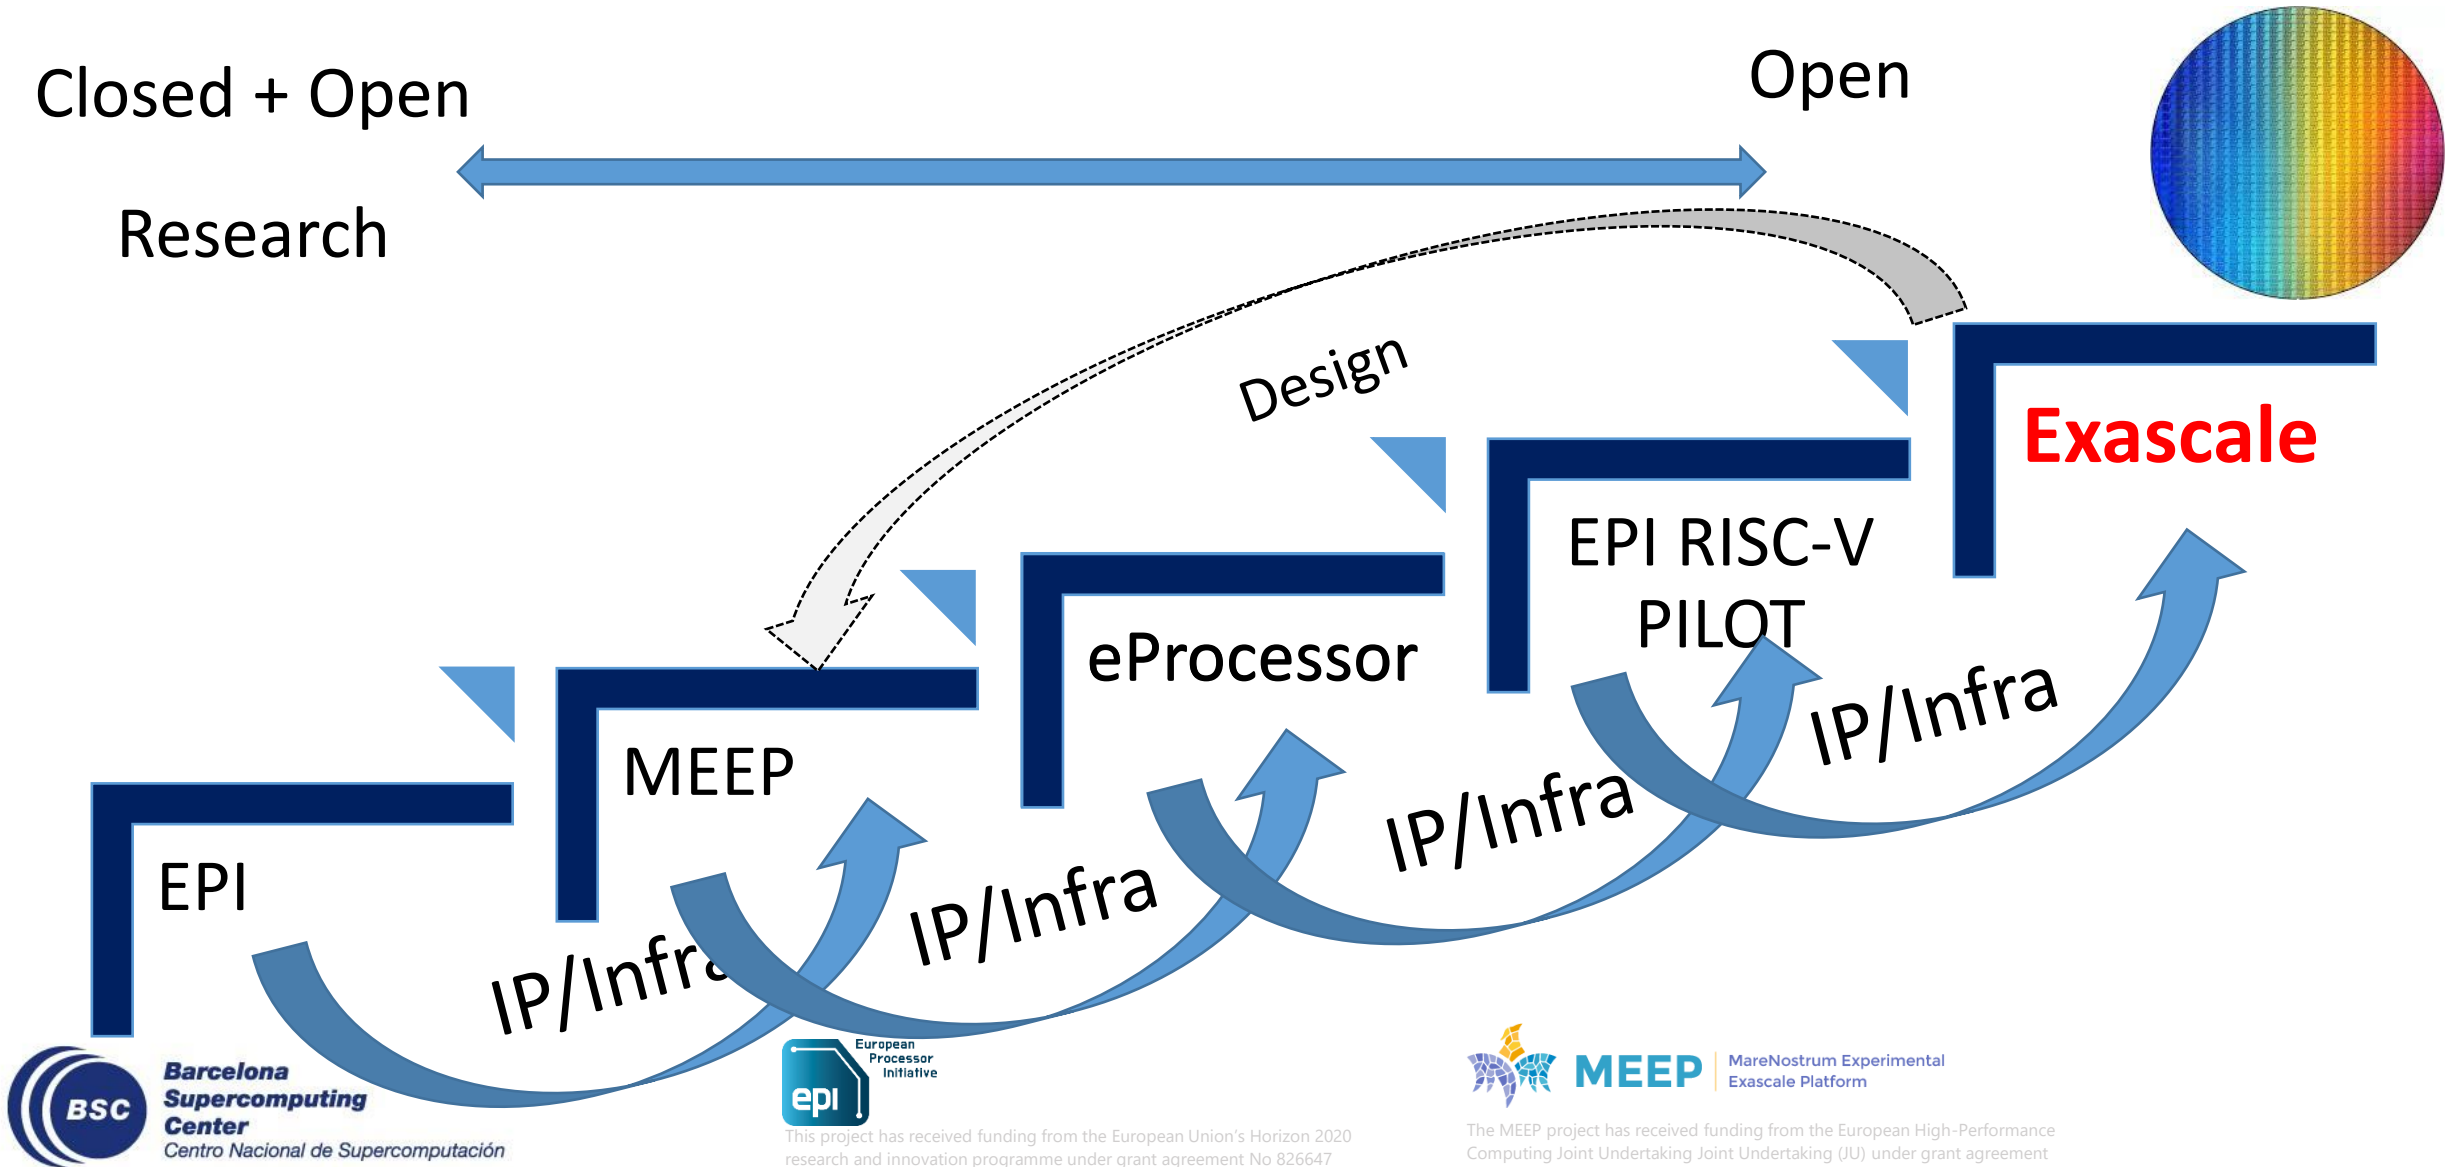
Accelerator Card

- 3 C2C interfaces per chiplet, 8 lanes each
- 1 PCIe 6.0 w/ CXL 3.0 interface per chiplet, 4 lanes
- 4 x16 LPDDR controllers per chiplet
- Homogeneous or Heterogeneous PCB
- Optional FPGA/PCIe Switch



Rebuilding the European CPU Industry

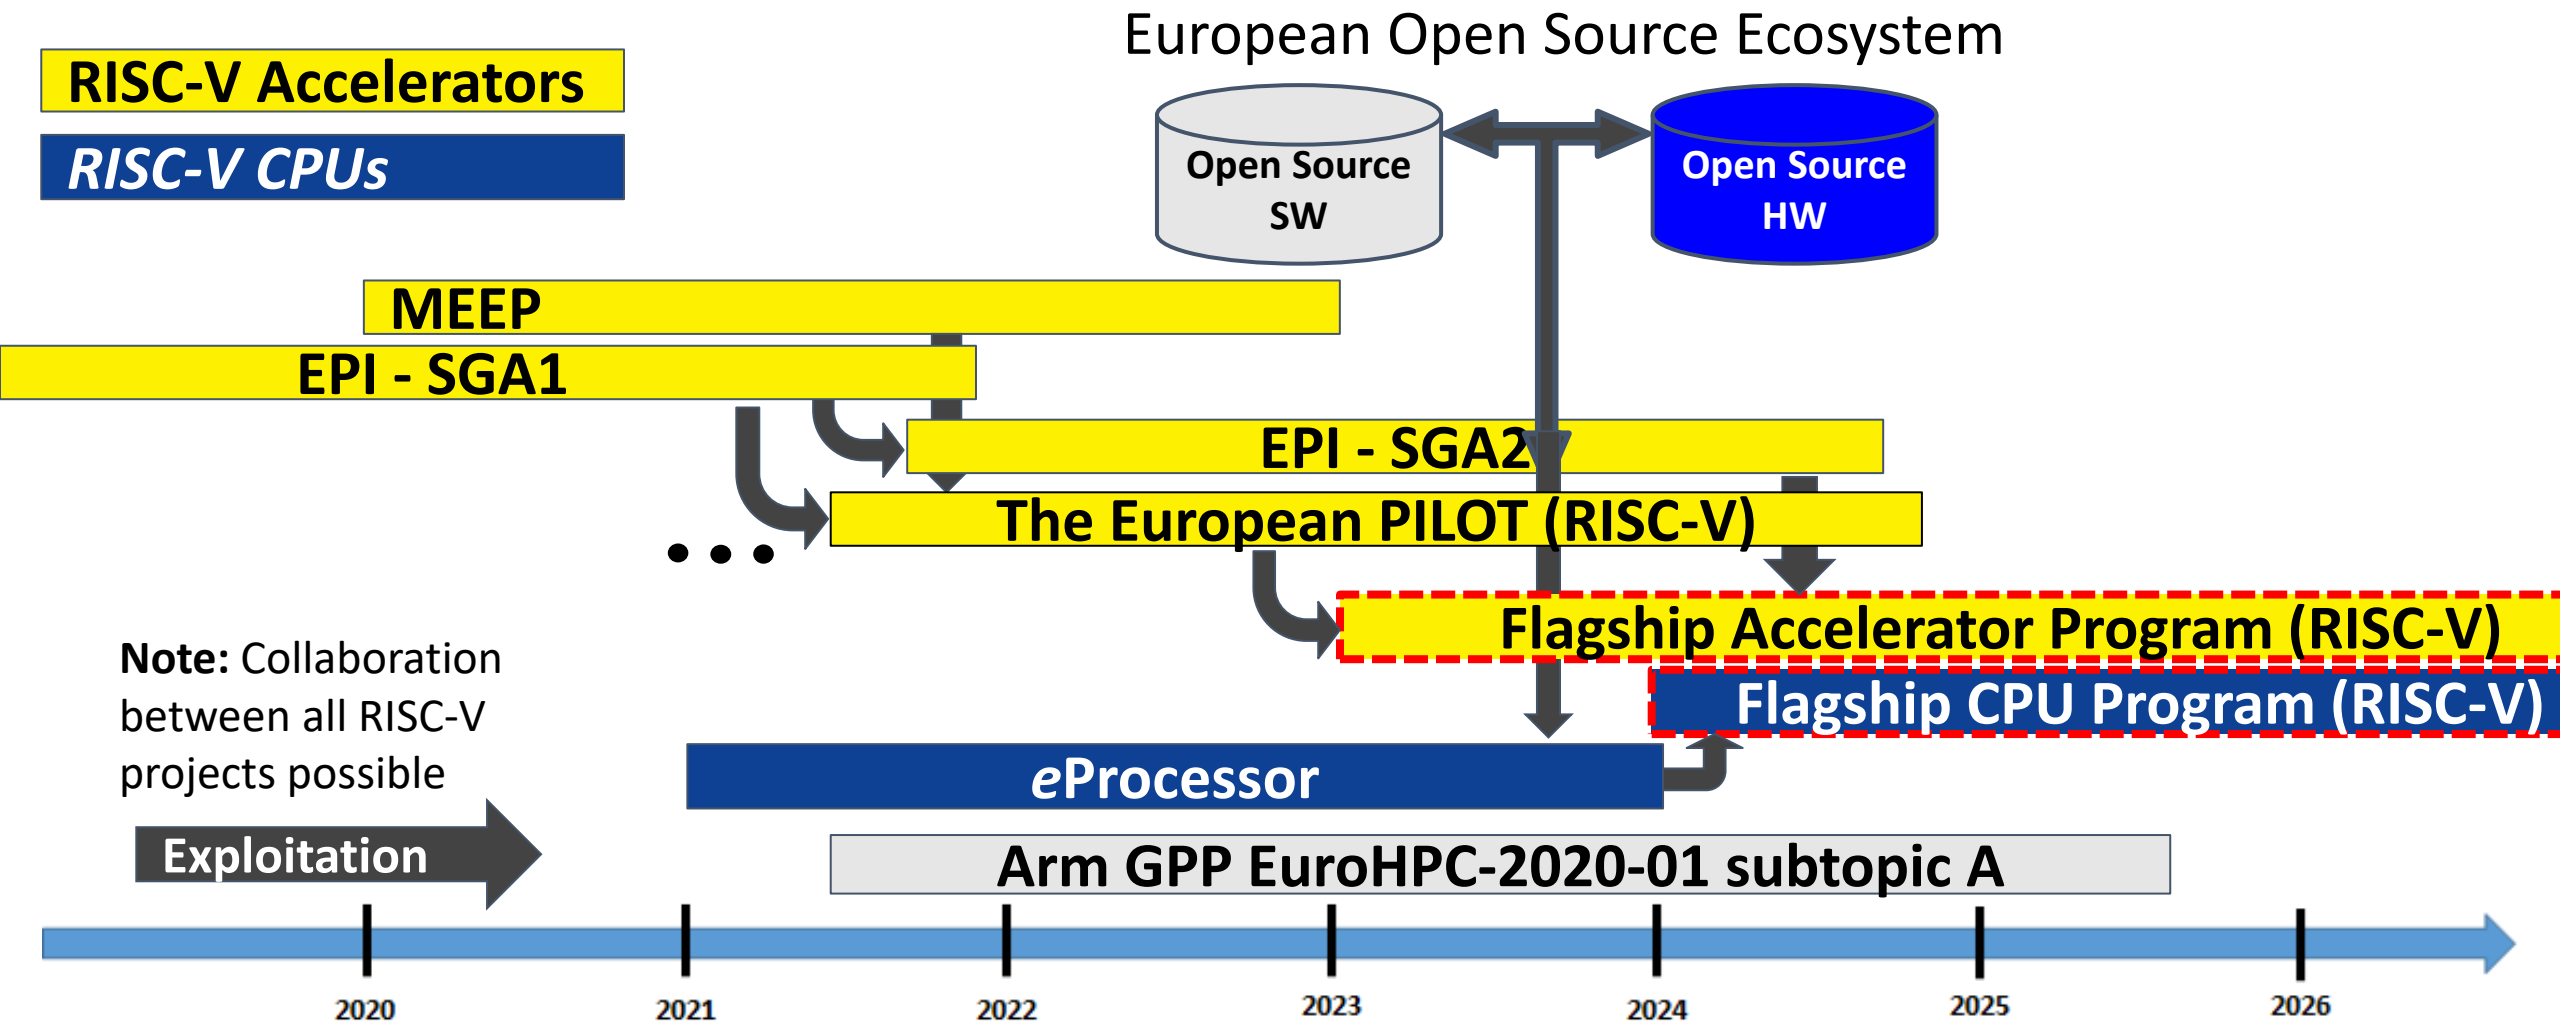
Production



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826647

The MEEP project has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU) under grant agreement No 946002. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Spain, Croatia, Turkey.

The Future: Flagship RISC-V Exascale Accelerator



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 goal: 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**




LOCA @ BSC



European Laboratory for Open
Computer Architecture

BSC full stack



European Laboratory for Open
Computer Architecture



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
 - Provide proven/usable Open Source H/W
 - Intersection of academia and industry
 - Open European IP repository → rapid implementation
 - Catalyst to reinvigorate European ICT industry
 - **Global** collaboration and training center
 - Incubator for European IP



European Collaboration & Education



Casteller

(human tower)



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

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

LOCA is Bigger than Barcelona!



Casteller
(human tower)



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



Thank you

mateo.valero@bsc.es