Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Pilots and system-level co-design

The BSC vision

Prof. Mateo Valero

BSC Director

erc

May 2019, Poznan

The Global Race Towards Exascale




2011				K computer. 10 PFlop/s		
2012						
2013		Tianhe-2. 33 PFlop/s				
2014						
2015						
2016		Sunway. 125 PFlop/s				
2017						
2018	Summit 200 PFlop/s.	Sierra 125 PFlop/s.		ABCI 0.55 EFlop/s (DL) 37 Pflop		
2019						
2020		Shuguang 1 EFlop/s	Tianhe-3 1 EFlop/s			
2021	Aurora 1 EFlop/s	Frontier 1. 5 EFlop/s		Post-K computer 1 EFlop/s	>200 PFlop/s	>200 PFlop/s
2022						
2023	El Capitan	ANL system			>1 EFlop/s	>1 EFlop/s

The Exascale Race – The US example

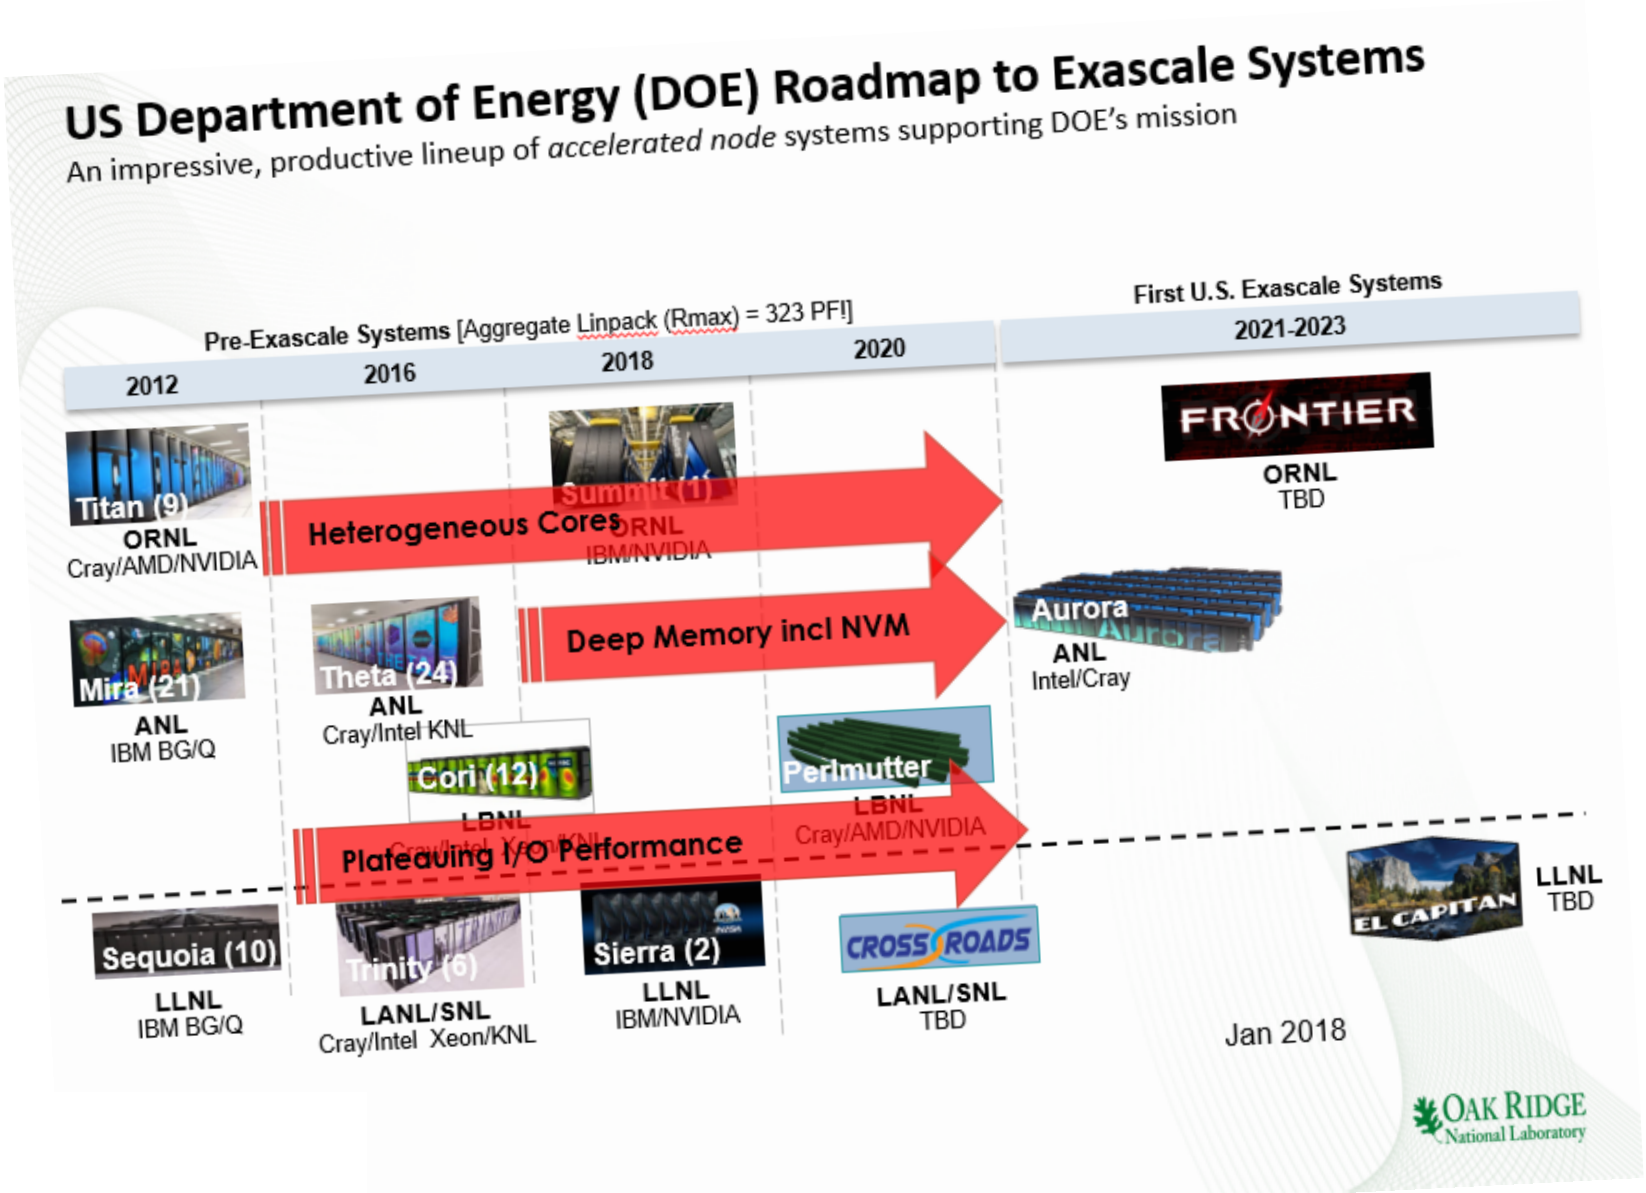

The Exascale Race – The US example

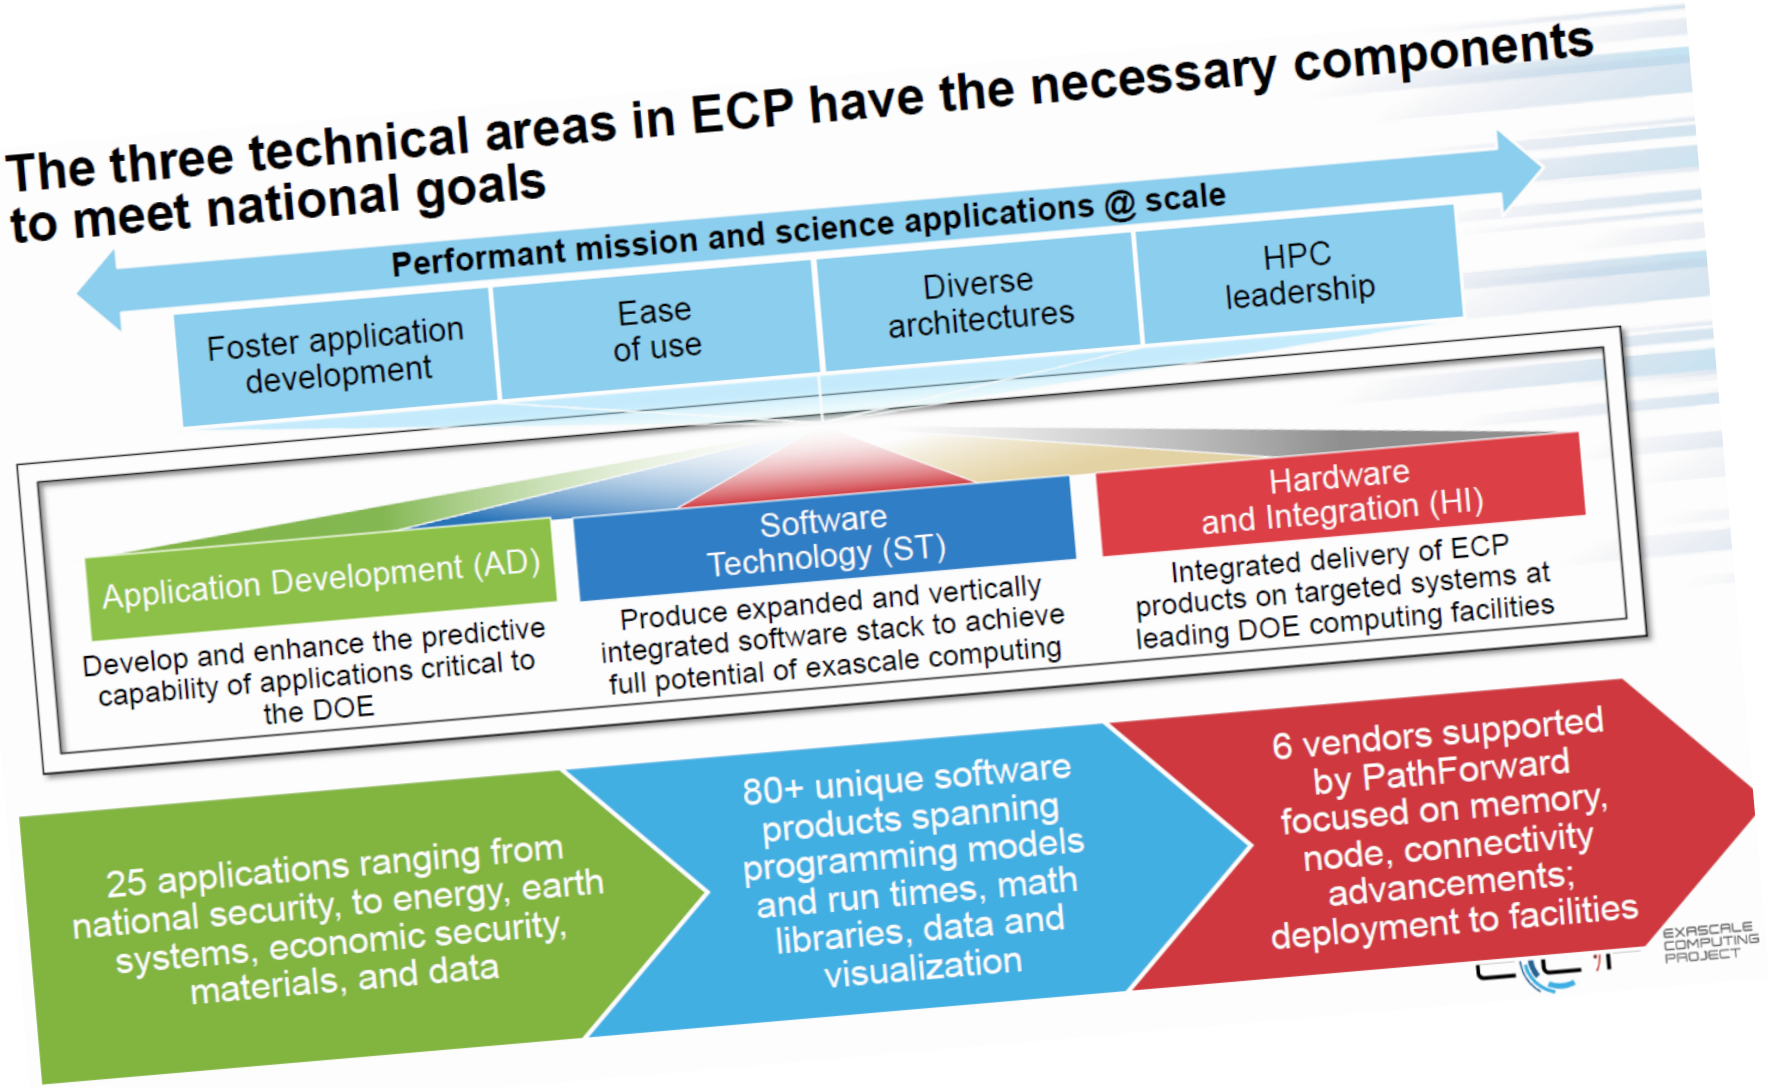

The Exascale Race – The Japanese example

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	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

"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

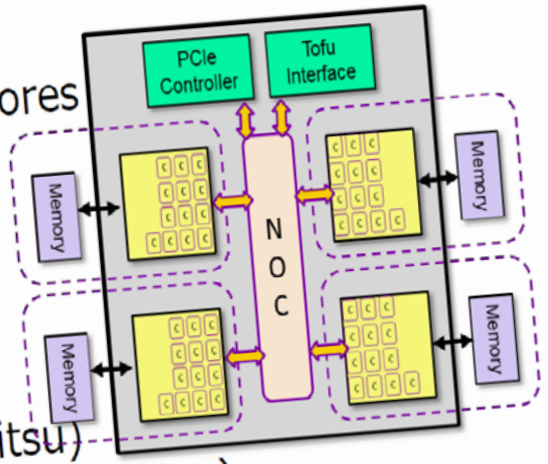

- **…but also a GPU-like processor**
  - SVE 512 bit vector extensions (ARM & Fujitsu)
    - Integer (1, 2, 4, 8 bytes) + Float (16, 32, 64 bytes)
  - Cache + access localization (sector cache) – similar to scratchpad
  - HBM2 OPM – Massive Mem BW (1TByte/s, Bytes/DPF ~0.4 same as K)
    - Streaming 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…**

11

Why Europe needs its own Processors

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

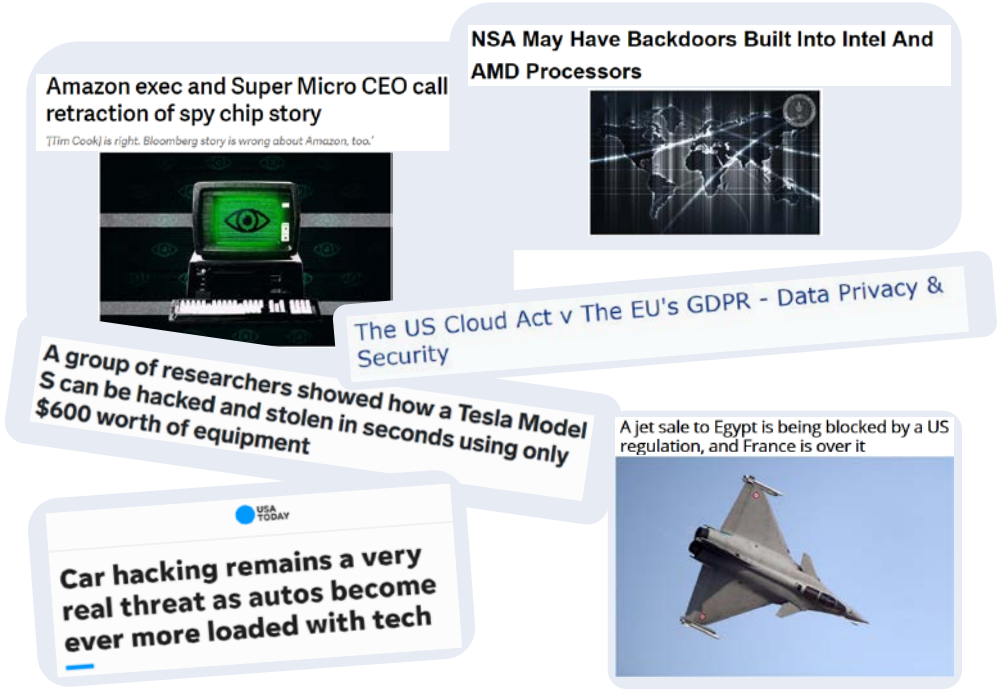

Images courtesy of European Processor Initiative

The Open-Source Hardware Opportunity

- The fastest-growing movement in computing at the moment is Open-Source and is called RISC-V
- The future is open and RISC-V is democratising chip-design

EPI. From research to industry

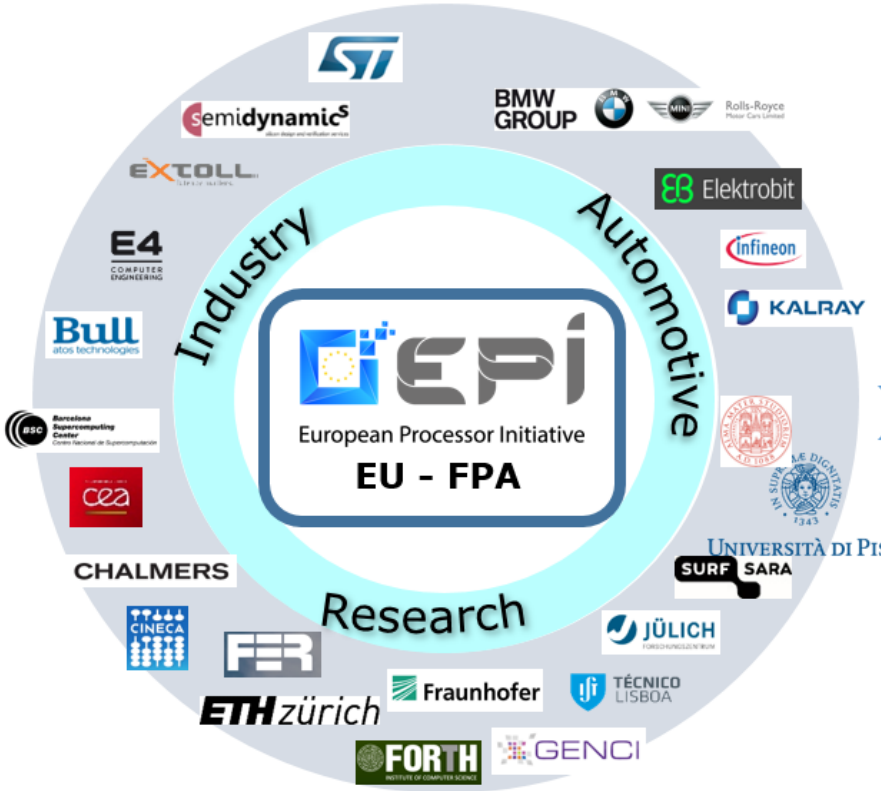

Semiconductor

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

EPI Common Platform

1st EPI production

The European HPC Roadmap

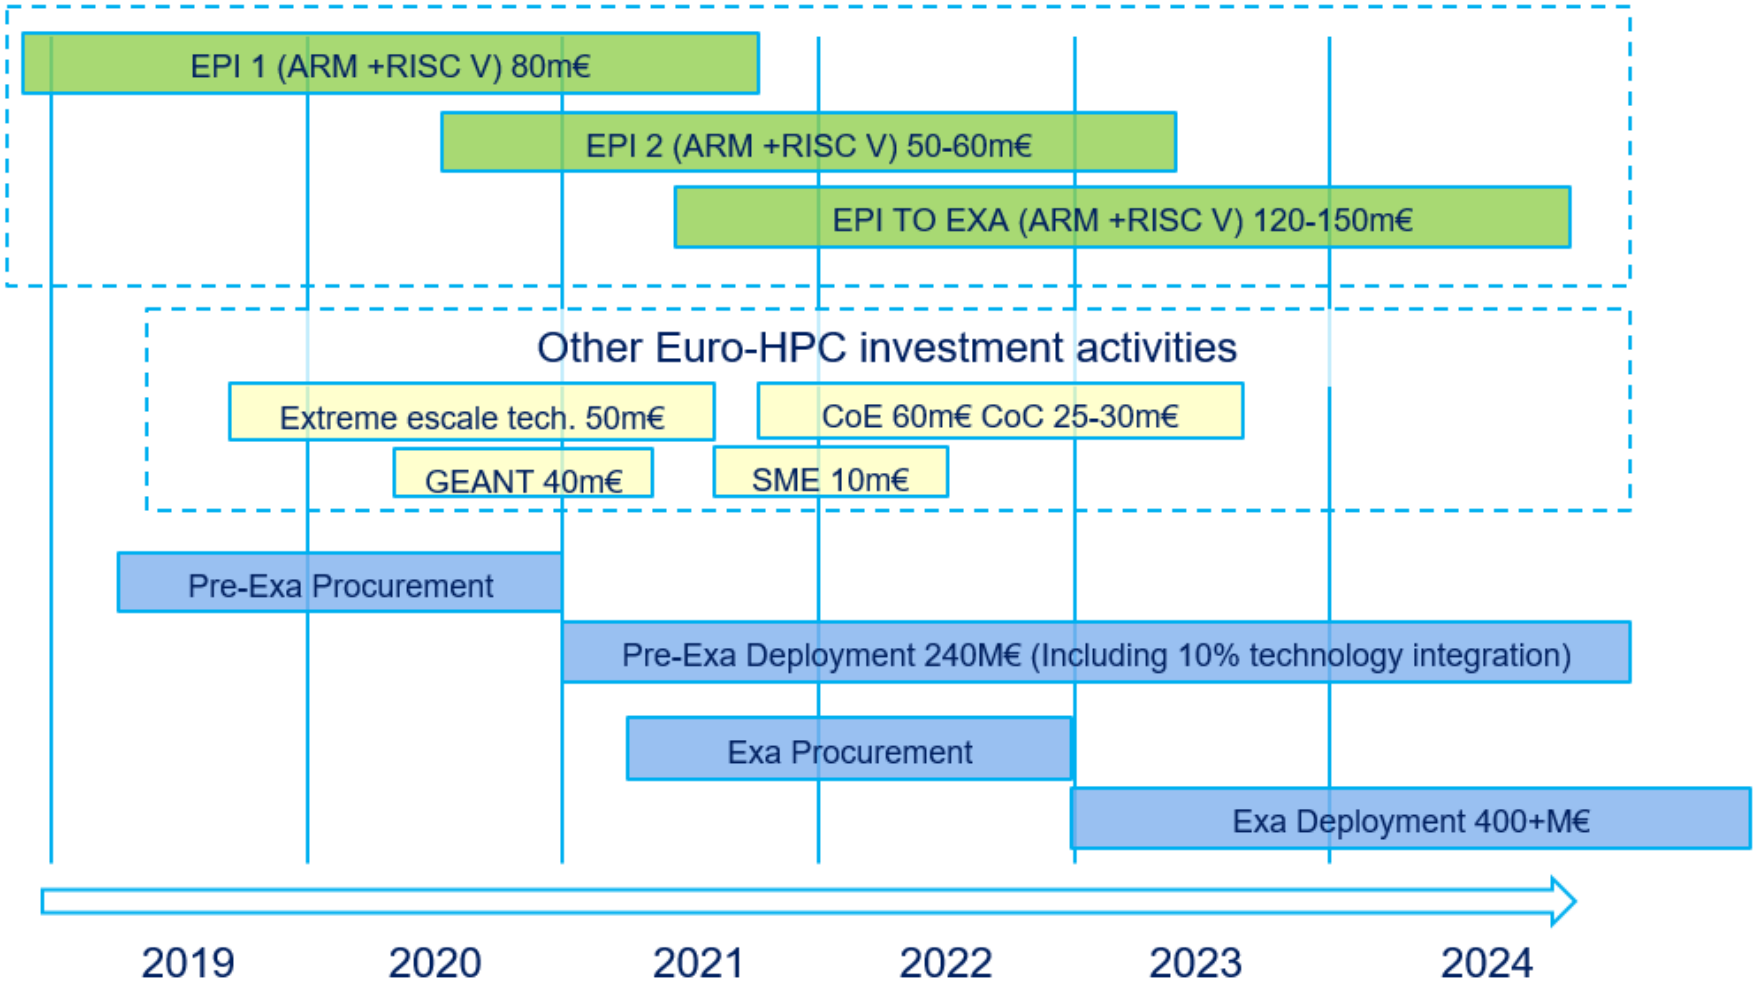

The Killer Mobile processors™

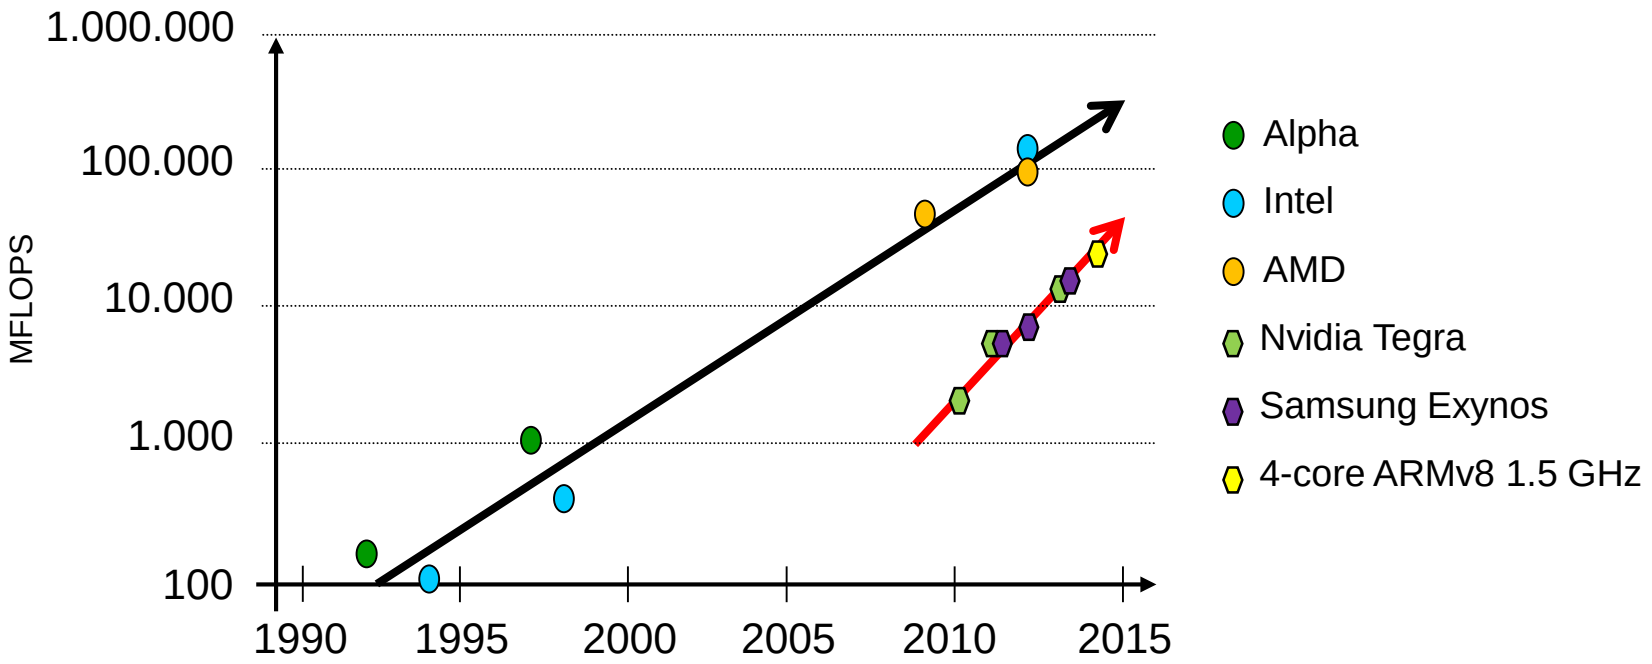

- Microprocessors killed the Vector supercomputers
  - They were not faster ...
  - ... but they were significantly cheaper and greener
- History may be about to repeat itself ...
  - Mobile processor are not faster ...
  - ... but they are significantly cheaper and greener

Mont-Blanc ARM-based Prototypes

2011	2012	2013	2014
Tibidabo	KAYLA	Pedraforca	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

The 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

The EPI Pilot as an Opportunity

- To co-develop an HPC/AI software ecosystem for RISC-V
- Using the hardware coming from EPI (ARM+RISC-V), other EU projects and from any other relevant national and international actors
- Gathering together a wide variety of stakeholders
- BSC would like to collaborate with everyone who can add value

The Tortoise and the Hare

- Is Aesop's Fable relevant here?
- Are the hares so confident they will win the race, that they are resting?

The Tortoise and the Hare

- We are now turbo-charging our tortoise!

www.bsc.es

Thank you !!!

For further information please contact
mateo.valero@bsc.es