BSC roadmap for HPC in Europe

Prof. Mateo Valero

Director, Barcelona Supercomputing Center

ISS EUROPE
INDUSTRY STRATEGY SYMPOSIUM

Leveraging European Ecosystems for Digital Innovation

31 MAR – 2 APR 2019
Milan, Italy

semi

Some societal challenges

- Ageing population
- Climate change
- Cybersecurity
- Increasing energy needs
- Intensifying global competition

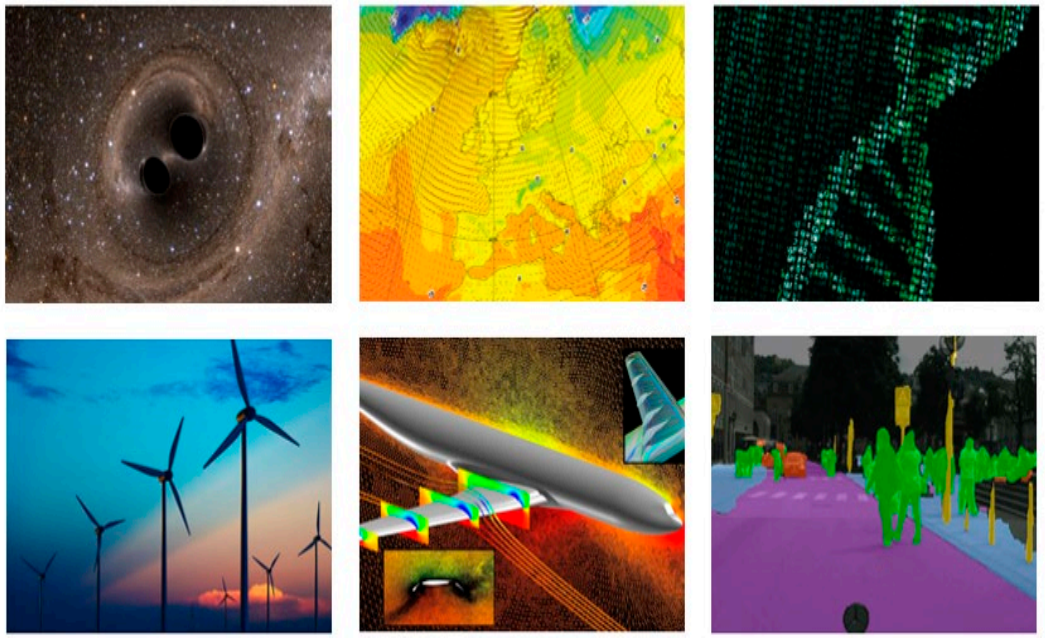

Images courtesy of The PRACE Scientific Steering Committee, “The Scientific Case for Computing in Europe 2018-2026”

Top10, Nov 2018

Rank	Name	Site	Manufacturer	Country	Cores	Accelerators	Rmax [TFlop/s]	Rpeak [TFlop/s]	GFlops /Watts
1	Summit	DOE/SC/Oak Ridge National Laboratory	IBM	US	2,397,824	2,196,480	143,500	200,795	14.67
2	Sierra	DOE/NNSA/Lawrence Livermore National Lab.	IBM/NVIDIA	US	1,572,480	1,382,400	94,640	125,712	12.72
3	Sunway TaihuLight	National Supercomputing Center in Wuxi	NRCPC	China	10,649,600		93,015	125,436	6.05
4	Tianhe-2ª	National Super Computer Center in Guangzhou	NUDT	China	4,981,760	4,554,752	61,445	100,679	3.33
5	Piz Daint	Swiss National Supercomputing Centre	Cray Inc.	Switz	387,872	319,424	21,230	27.154.3	8.90
6	Trinity	DOE/NNSA/LANL/SNL	Cray Inc.	US	979,072		20,158,7	41,461	2.66
7	AI Bridging Cloud Inf.	National Inst. of Adv Industrial Science & Tech.	Fujitsu	Japan	391,680	348,160	19,880	32,577	12.05
8	SuperMUC-NG	Leibniz Rechenzentrum	Lenovo	Germany	305,856		19,477	28,872,86	
9	Titan	DOE/SC/Oak Ridge National Laboratory	Cray Inc.	US	560,640	261,632	17,590	27,113	2.14
10	Sequoia	DOE/NNSA/Lawrence Livermore National Lab.	IBM	US	1,572,864		17,173	20,133	2.18
25	MareNostrum	Barcelona Supercomputing Center	Lenovo	Spain	153,216		6,471	10,296	3.97

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

ISS EUROPE
INDUSTRY STRATEGY SYMPOSIUM

Top10 Europe, Nov 2018

EU	Rank	Name	Computer	Site	Country	Cores & Accel.	Rmax [PFlop/s]	Rpeak [PFlop/s]	GFlops/ Watts
1	5	Piz Daint	Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect , NVIDIA Tesla P100	CSCS	Switzerland	387,872 319,424	21,230	27,154	8,9
2	8	SuperMUC-NG	ThinkSystem SD530, Xeon Platinum 8174 24C 3.1GHz, Intel Omni-Path	Leibniz Rechenzentrum	Germany	305,856	19,476	26,873	
3	15	HPC4	Proliant DL380 Gen10, Xeon Platinum 8160 24C 2.1GHz, Mellanox InfiniBand EDR, NVIDIA Tesla P100	Eni S.p.A.	Italy	253,600 177,520	12,21	18,62	9,25
4	16	Tera-1000-2	Bull Sequana X1000, Intel Xeon Phi 7250 68C 1.4GHz, Bull BXI 1.2	CEA	France	561,408	11,97	23,40	3,77
5	19	Marconi Intel Xeon Phi	CINECA Cluster, Lenovo SD530/S720AP, Intel Xeon Phi 7250 68C 1.4GHz/Platinum 8160, Intel Omni-Path	CINECA	Italy	348,000 204,000	10,384	18,816	
6	23		Cray XC40, Xeon E5-2695v4 18C 2.1GHz, Aries interconnect	UK-MET	United Kingdom	241,920	7,04	8,13	
7	25	MareNostrum	Lenovo SD530, Xeon Platinum 8160 24C 2.1GHz, Intel Omni-Path	BSC-CNS	Spain	153,216	6,47	10,30	3,97
8	26	JUWELS Module 1	Bull Sequana X1000, Xeon Platinum 8168 24C 2.7GHz, Mellanox EDR InfiniBand/ParTec ParaStation ClusterSuite	Juelich	Germany	114,480	6,18	9,89	4,54
9	30	Hazel Hen	Cray XC40, Xeon E5-2680v3 12C 2.5GHz, Aries interconnect	HLRS - Stuttgart	Germany	185,088	5,64	7,40	1,56
10	31	COBRA	Intel Compute Module HNS2600BP, Xeon Gold 6148 20C 2.4GHz, Intel Omni-Path	Max-Planck-Gesellschaft MPI/IPP	Germany	127,520	5,61	9,79	3,43
67	4118	MareNostrum P9 CTE	IBM Power System AC922, IBM POWER9 22C 3.1GHz, Dual-rail Mellanox EDR Infiniband, NVIDIA Tesla V100	BSC-CNS	Spain	19,440 17,280	1,02	1,49	11,87

Why is HPC needed?

- It will save billions by helping us to adapt to climate change
- It will improve human health by enabling personalized medicine
- It will improve fuel efficiency of aircraft & will help design better wind turbines
- It will help us understand how the human brain works

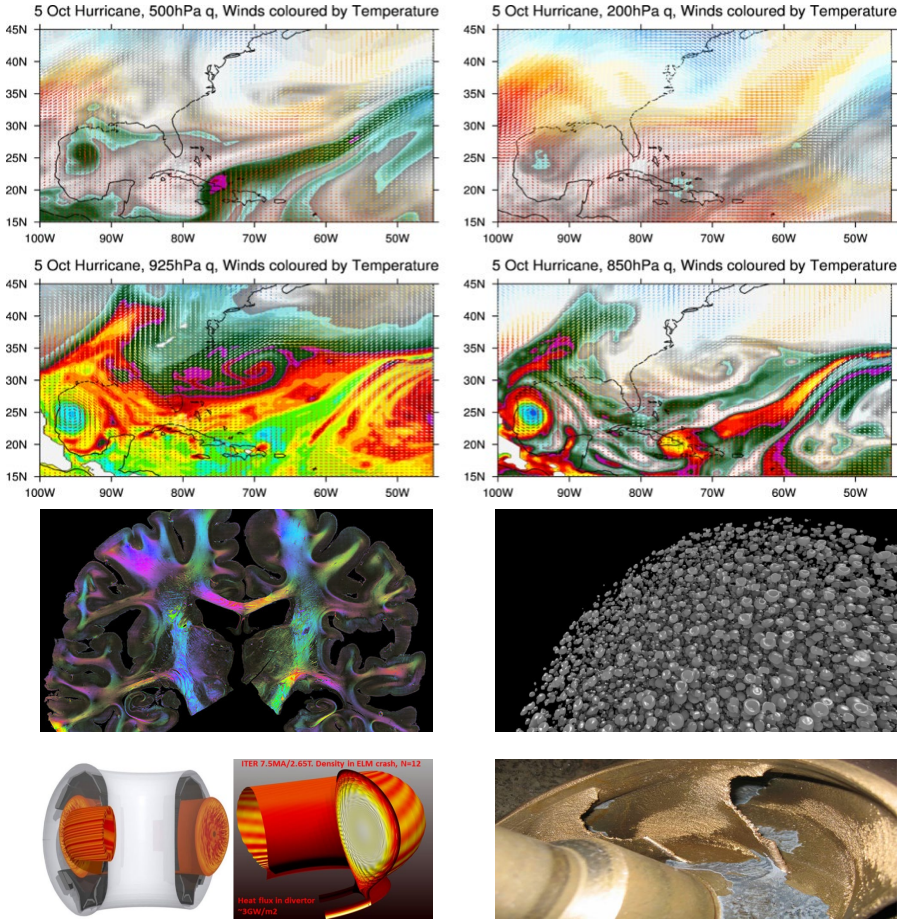

Images courtesy of The PRACE Scientific Steering Committee, "The Scientific Case for Computing in Europe 2018-2026"

semi

ISS INDUSTRY STRATEGY SYMPOSIUM EUROPE

Update on the race towards Exascale

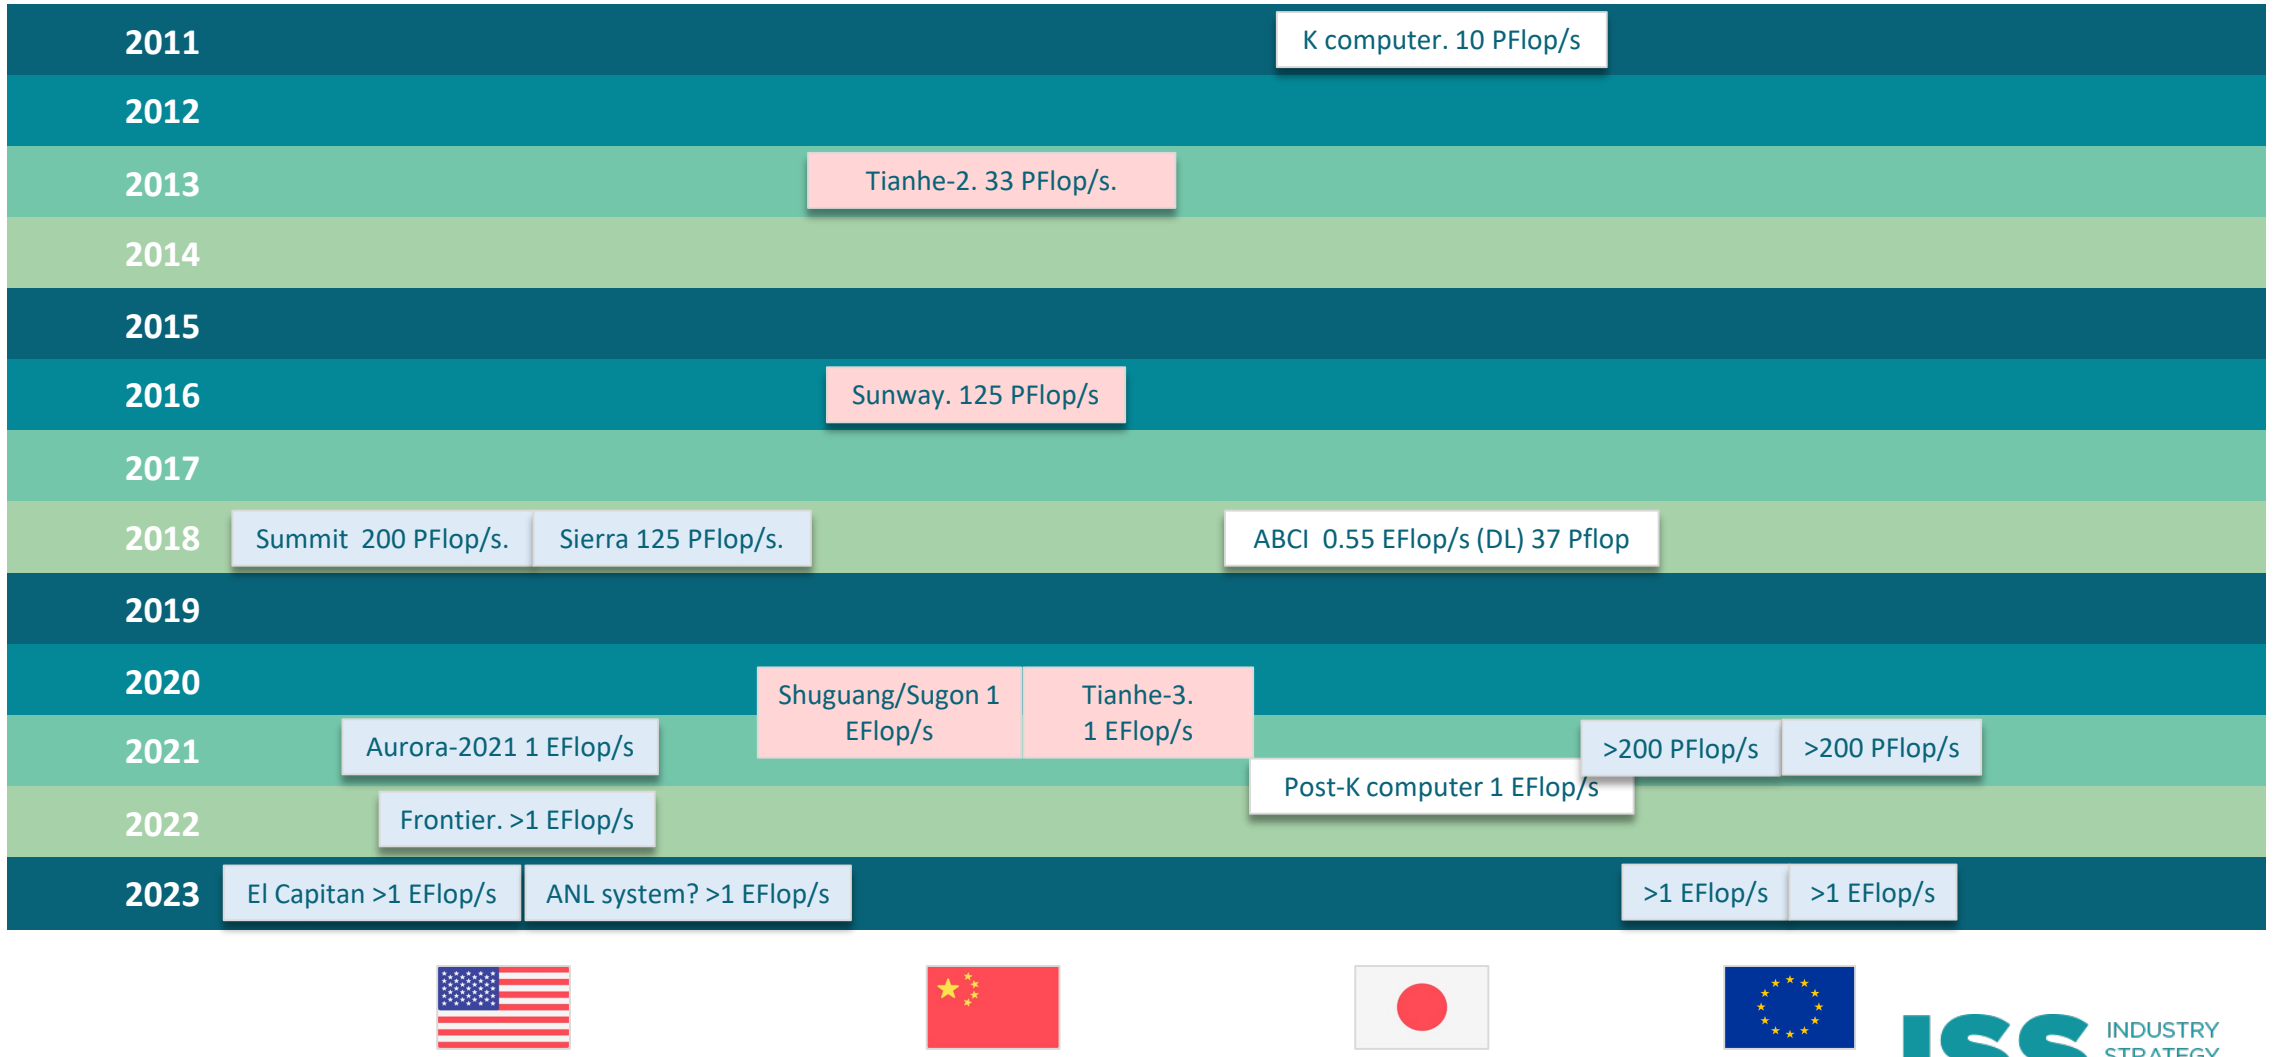

semi

ISS INDUSTRY STRATEGY SYMPOSIUM EUROPE

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

New paradigms for Exascale

- Increasingly complex and heterogeneous hardware, including AI, Neuromorphic, Quantum and other accelerators.
- New memory systems and hierarchies will require a monumental effort to redesign system software and applications for Exascale systems
- AI will be essential to harness the complexity of future Exascale systems

“Killer microprocessors”

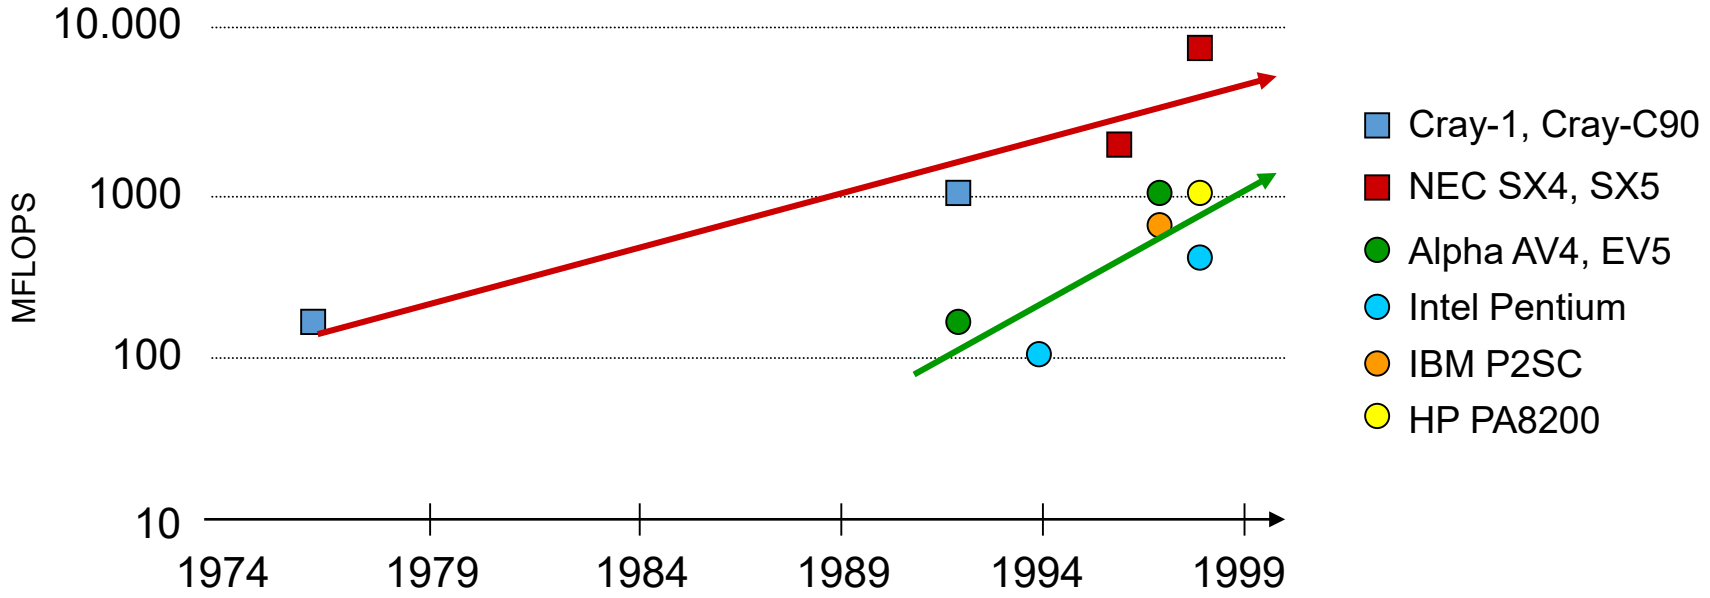

- Microprocessors killed the Vector supercomputers
  - They were not faster ...
  - ... but they were significantly **cheaper** and **greener**
- 10 microprocessors approx. 1 Vector CPU
  - SIMD vs. MIMD programming paradigms

M. Valero. “Vector Architectures: Past, Present and Future”. Keynote talk. ICS-11. IEEE-ACM. Melbourne, 1998

Then, commodity took over special purpose

- ASCI Red, Sandia
  - 1997, 1 Tflops (Linpack),
  - 9298 processors at 200 MHz,
  - 1.2 Terabytes
  - Intel Pentium Pro
    - Upgraded to Pentium II Xeon, 1999, 3.1 Tflops

- ASCI White, Lawrence Livermore Lab.
  - 2001, 7.3 TFLOPS,
  - 8192 proc. RS6000 at 375 MHz,
  - 6 Terabytes,
  - IBM Power 3
  - (3+3) MWats

Message-Passing Programming Models

13

The Killer Mobile processors™

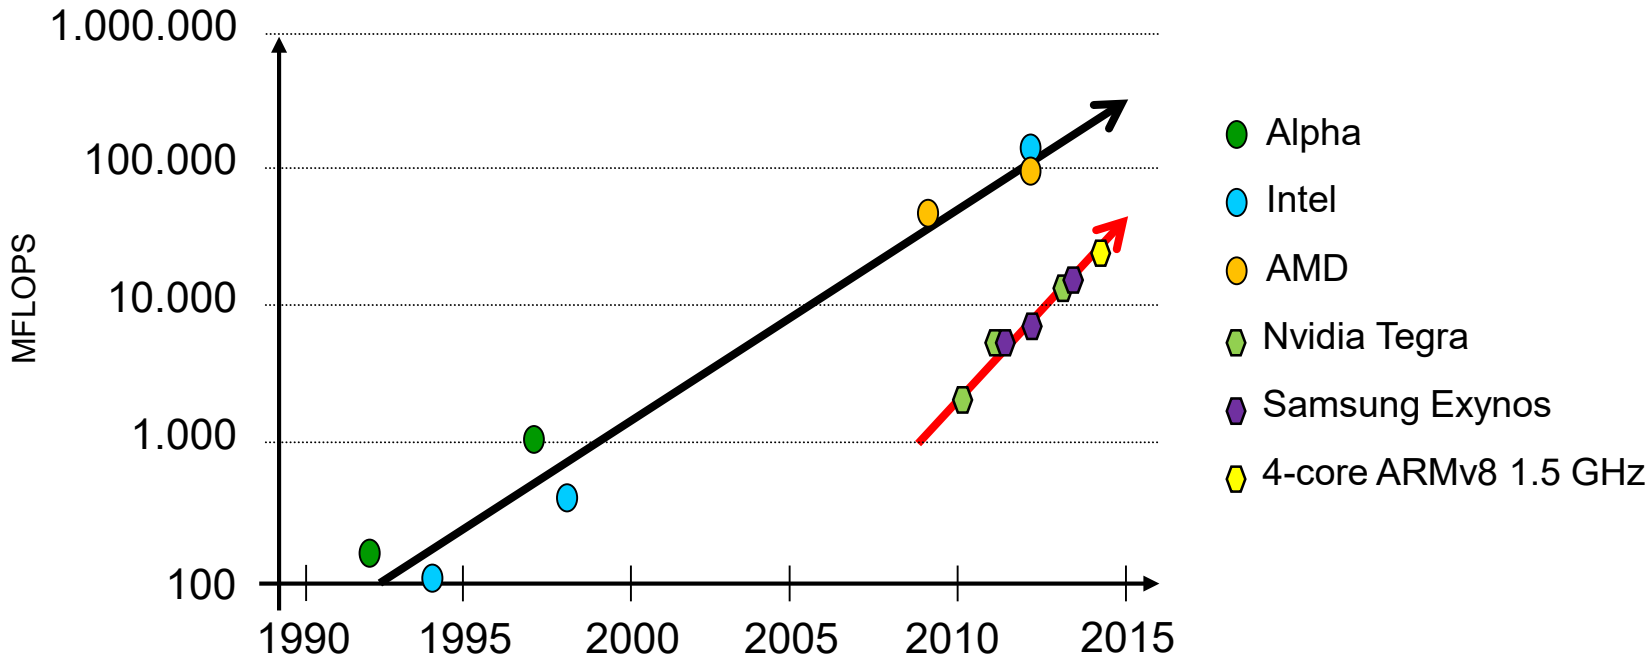

- Microprocessors killed the Vector supercomputers
  - They were not faster ...
  - ... but they were significantly cheaper and greener

- History may be about to repeat itself …
  - Mobile processor are not faster …
  - … but they are significantly cheaper and greener

Steering Committee HiPEAC

Josep Llosa
Spain

Olivier Temam
France

Per Stenström
Sweden

Mike O'Boyle
UK

Coordinator

Mateo Valero
UPC

Theo Ungerer
Germany

Stefanos Kaxiras
Greece

Stamatis Vassiliadis
The Netherlands

Koen De Bosschere
Belgium

Antonio Prete
Italy

semi

ISS INDUSTRY STRATEGY SYMPOSIUM EUROPE

Network of c.2,000 European R+D experts in advanced computing: **high-performance** and **embedded** architecture and compilation

720 members, 449 affiliated members and **871 affiliated PhD students** from **430** institutions in **46** countries.

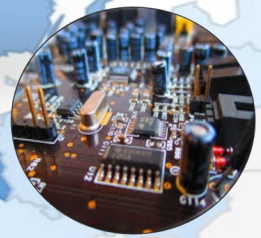

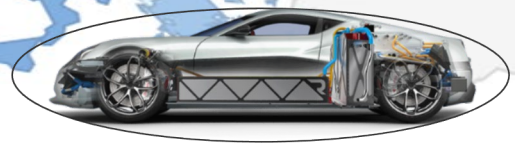

hipeac.net

HiPEAC has received funding from the European Union's Horizon2020 research and innovation programme under grant agreement number 779656.

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 increasing protectionism
- **A competitive EU supply chain** for HPC technologies will create jobs and growth in Europe

ARM-based prototypes at BSC

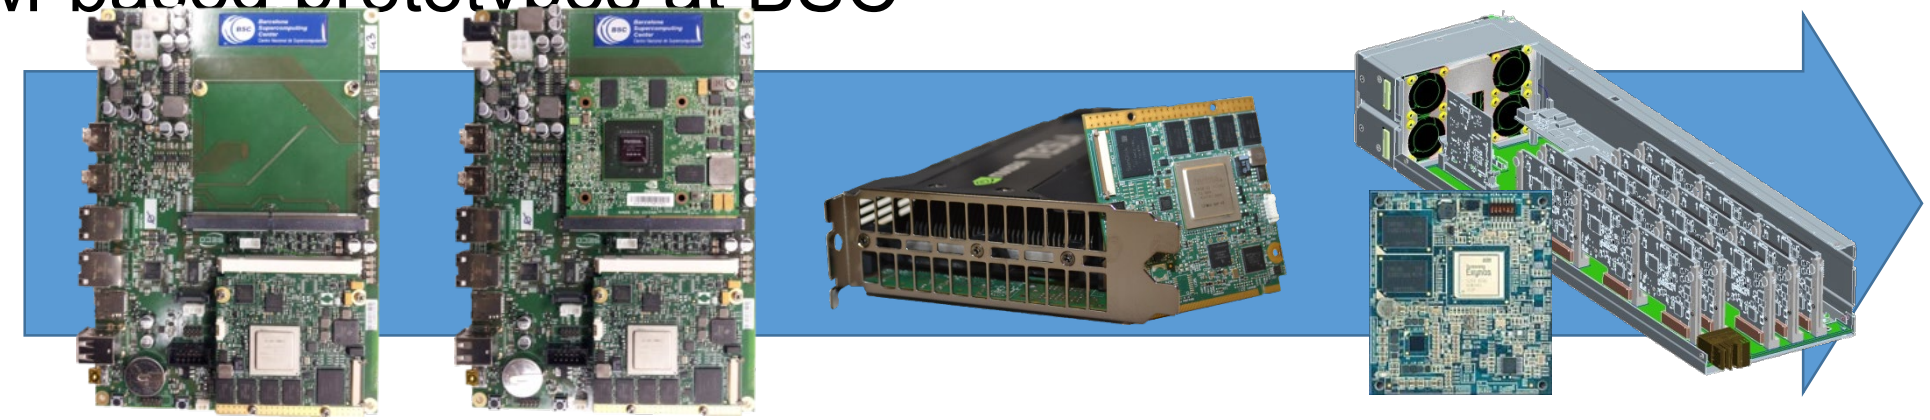

2011
Tibidabo

ARM multicore

2012
KAYLA

ARM + GPU
CUDA on ARM

2013
Pedraforca

ARM + GPU
Inifinband
RDMA

2014
Mont-Blanc

Single chip ARM+GPU
OpenCL on ARM GPU

semi
ISS INDUSTRY STRATEGY SYMPOSIUM EUROPE

Tibidabo: The first ARM HPC multicore cluster

Q7 Tegra 2
2 x Cortex-A9 @ 1GHz
2 GFLOPS
5 Watts (?)
0.4 GFLOPS / W

Q7 carrier board
2 x Cortex-A9
2 GFLOPS
1 GbE + 100 MbE
7 Watts
0.3 GFLOPS / W

1U Rackable blade
8 nodes
16 GFLOPS
65 Watts
0.25 GFLOPS / W

2 Racks
32 blade containers
256 nodes
512 cores
9x 48-port 1GbE switch

512 GFLOPS
3.4 Kwatt
0.15 GFLOPS / W

- Proof of concept
  - It is possible to deploy a cluster of smartphone processors
- Enable software stack development

semi
ISS INDUSTRY STRATEGY SYMPOSIUM EUROPE

Mont-Blanc HPC Stack for ARM

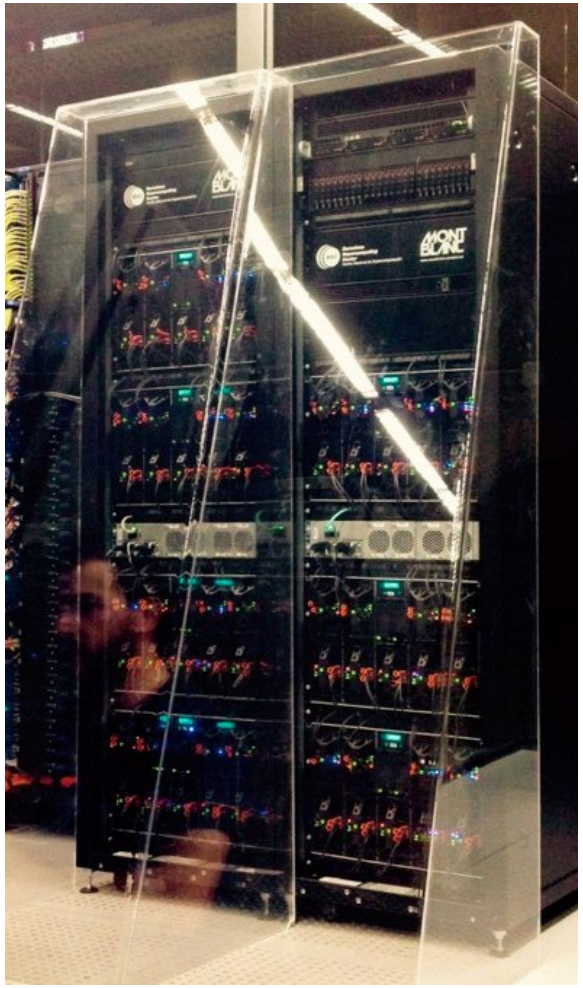

Industrial applications

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, Barcelona Supercomputing Center, JÜLICH FORSCHUNGSZENTRUM, Bull atos technologies

Hardware

ARM, Barcelona Supercomputing Center, Bull atos technologies, cea, UC UNIVERSIDAD DE CANTABRIA

BullSequana compute blade: X1310
Marvell ThunderX2™ (ARMv8) processor

MONT-BLANC
EUROPEAN APPROACH TOWARDS ENERGY EFFICIENT HIGH PERFORMANCE

BSC

Bull
atos technologies

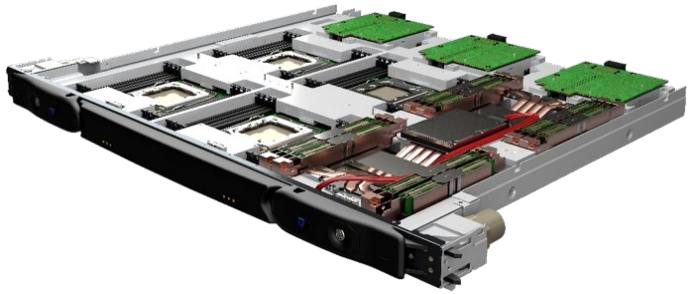

- **Atos and ARMv8**
  - Atos is the industrial pivot of MontBlanc3
  - ARM is one of the Atos strategic directions for the next years
  - Europe is leading in ARM development
  - Montblanc project is proceeding as expected

- **BullSequana X1310 blade**
  - Up to 288 nodes in one BullSequana X1000 and up to 96 nodes in one BullSequana XH2000 with:
  - 3 compute nodes with 2 Marvell ThunderX2 (ARMv8) processors
  - Up to 1024 GB of Memory per node – DDR4 @2666 MT/s (w/64GB DIMMs)
  - High-speed Ethernet, InfiniBand EDR, HDR or HDR100 on the mezzanine interconnect
  - Up to 192 cores per blade

311 TFops

cea

92 BullSequana X1310 blades, three compute nodes per blade, dual Marvell ThunderX2 processors of 32 cores @ 2.2 GHz, based on the Armv8-A instruction set, with 256 GB per node and Infiniband® EDR interconnect.

ARM processor – a credible alternative to X86 processor clusters

Fujitsu Processor Development

FUJITSU

Persistent Evolution > 60 years

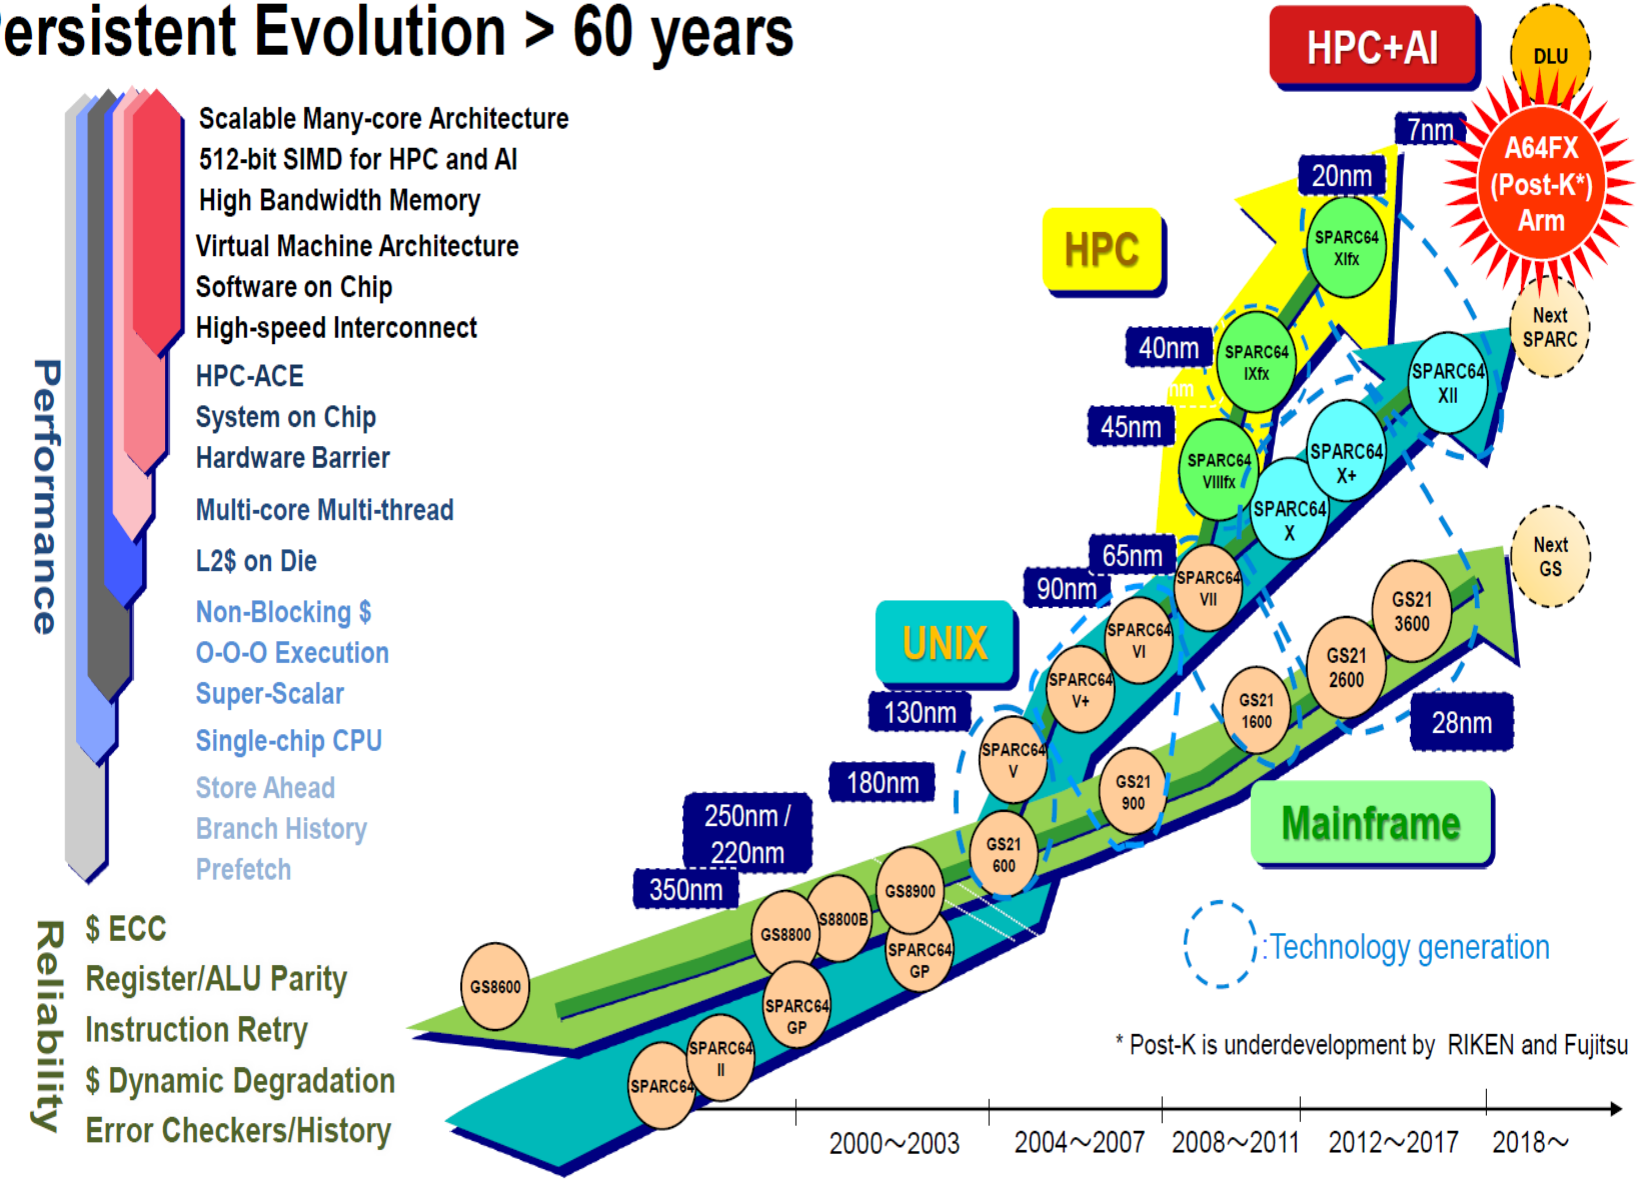

All Rights Reserved. Copyright © FUJITSU LIMITED 2018

sem

INDUSTRY STRATEGY SYMPOSIUM EUROPE

USA

Sandia Labs:

- HPE, Astra Supercomputer
- 2592 nodes, 28 core dual
- 2.3 petaflops/peak, 1.529, Linpack)
- # 203, Top500 (Nov. 2018)
- #36, HPCG (Nov. 2018)

Others: (smaller systems)

- Nercs Labs: Cray, 1080 cores
- Argonne Labs: HPE, Comanche system

Some context

- I was kindly invited by Roberto Viola to give a talk at ICT in Lisbon in October 2015
- I tried to emphasise the importance of HPC and of producing European technology
- A lot has happened since then…

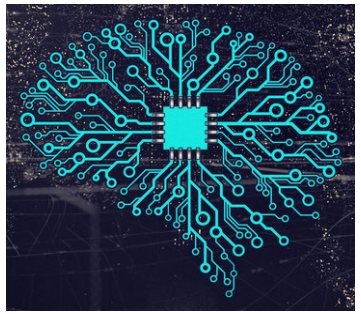

The Open-Source Hardware Opportunity

- In 2015 I said I believed a European Supercomputer based on ARM was possible (Mont-Blanc).
- Even though ARM is no longer European, it can take part in the short-term solution
- The fastest-growing development in computer architecture is Open-Source and is called RISC-V
- The future is Open and RISC-V is making HPC chip-design more affordable.

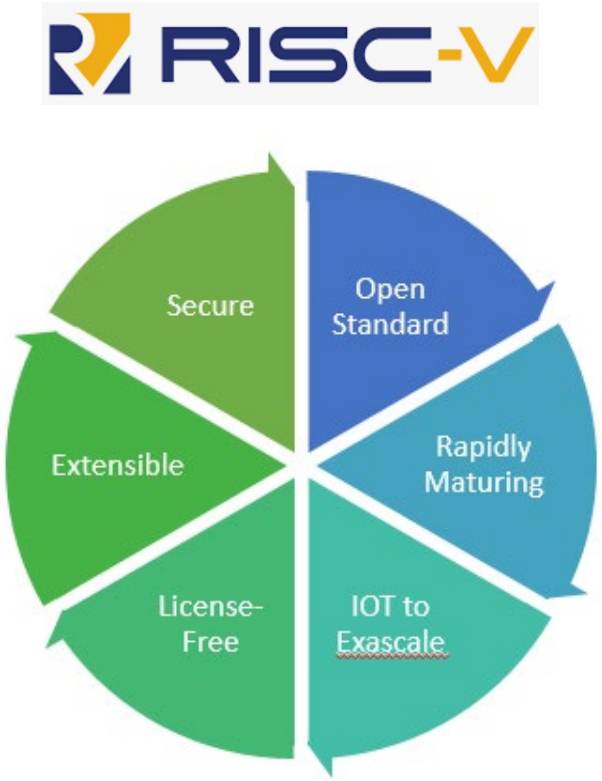

BSC and the EC

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

ISS INDUSTRY STRATEGY SYMPOSIUM EUROPE

HPC is a global competition

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

European Commission President
Jean-Claude Juncker (27 October 2015)

"With the EuroHPC initiative we want to give European researchers and companies world-leading supercomputer capacity".

European Commission Vice President
Andrus Ansip (11 January 2018)

The new European framework

Horizon Europe | The EU's next research and innovation programme (2021 – 2027)

EuroHPC Joint Undertaking

EUROPEAN CLOUD INITIATIVE
UNLOCKING THE POWER OF BIG DATA FOR OPEN SCIENCE

European HPC roadmap

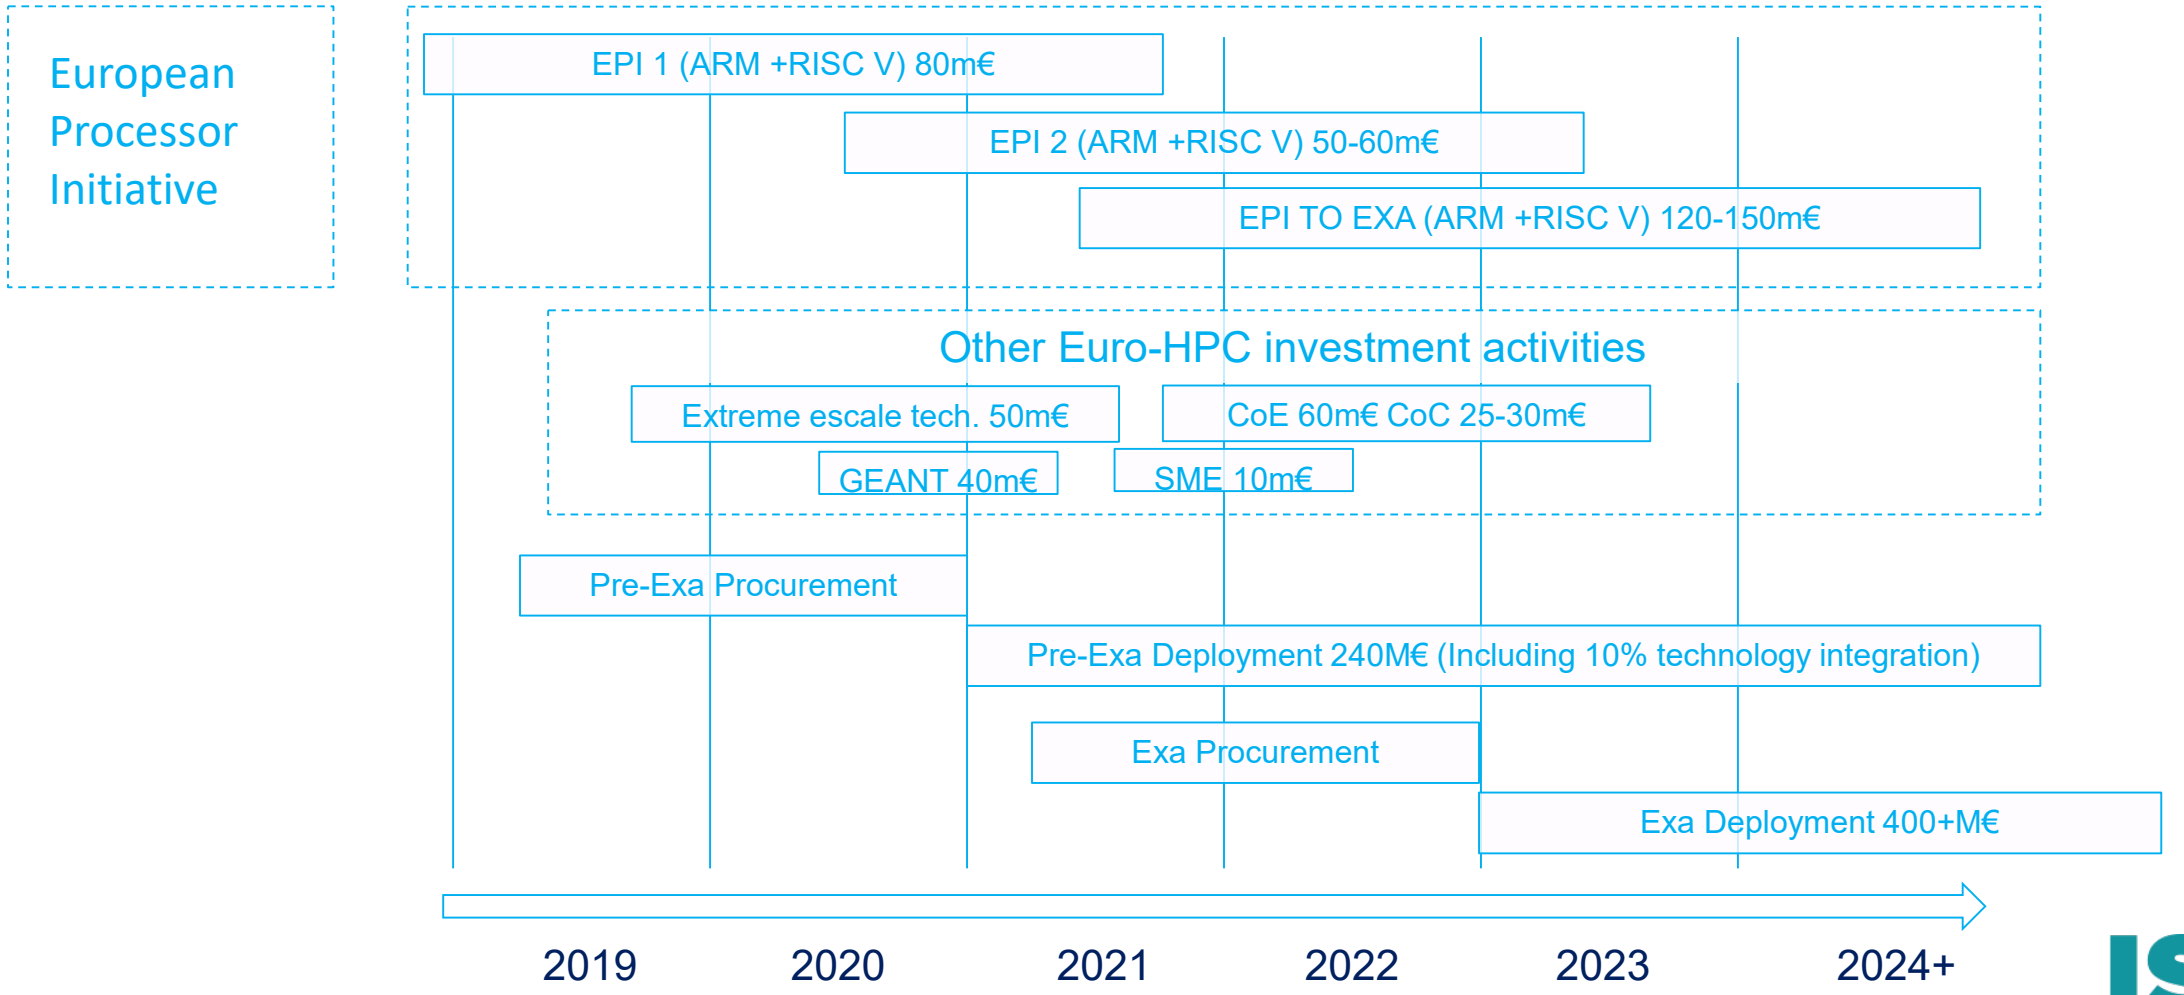

semi
ISS INDUSTRY STRATEGY SYMPOSIUM EUROPE

from consortium to EU high tech fabless

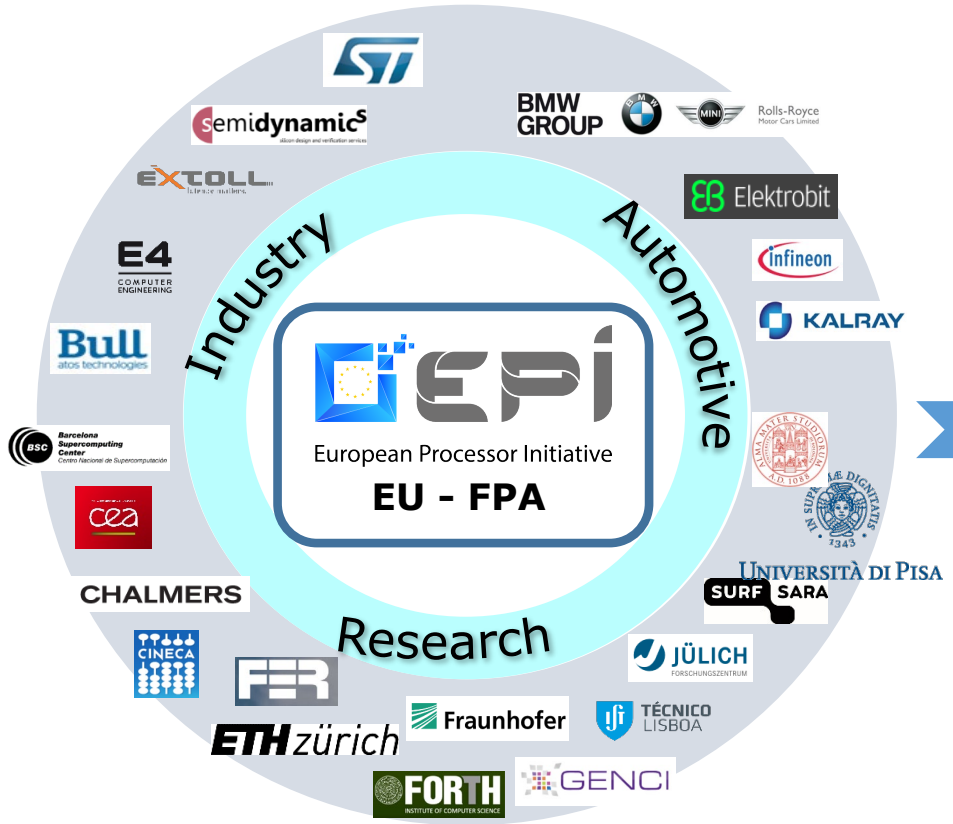

Semiconductor

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

EPI Common Platform

1st EPI production

EPI ROADMAP

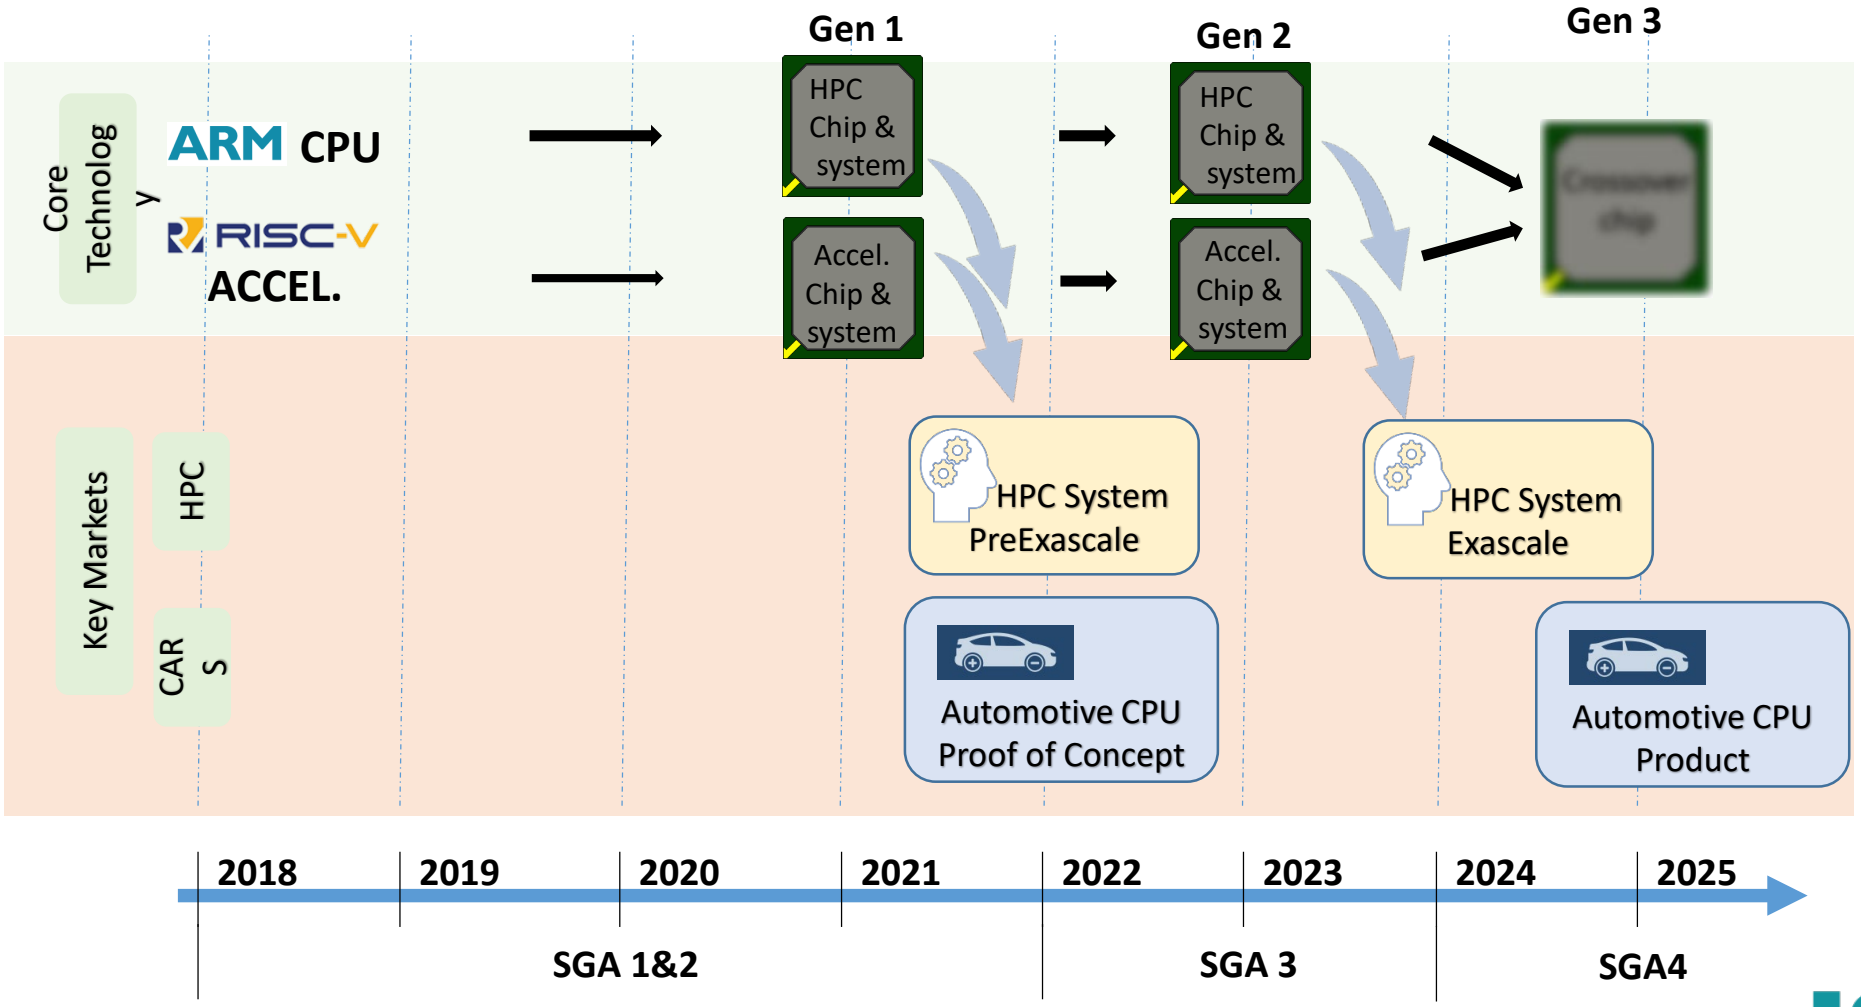

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

European Processor Initiative

ISS EUROPE
INDUSTRY STRATEGY SYMPOSIUM

BSC HPC roadmap

- FIRST PHASE: DEVELOP FOUNDATION IP
  - EPI first phase
  - Small RISC-V core with vector unit ready for NOC integration
- SECOND PHASE: DEVELOP SW
  - Use commercial RISC-V cores
  - Integrate using Exascale demonstrator program
  - Provide a usable platform to develop RISC-V SW
- THIRD PHASE: INDUSTRIAL PARTNERSHIP
  - Uses foundation IP from Phase 1
  - Uses SW from Phase 2
  - Partner with experienced SOC integrators
  - Partner with Exascale system vendors
  - Deliver Accelerators based on European IP

EPI Accelerator roadmap

- Core + VPU IP delivered to JV
- SOC tapeout
- Board ready
- RISC-V SW delivered to JV

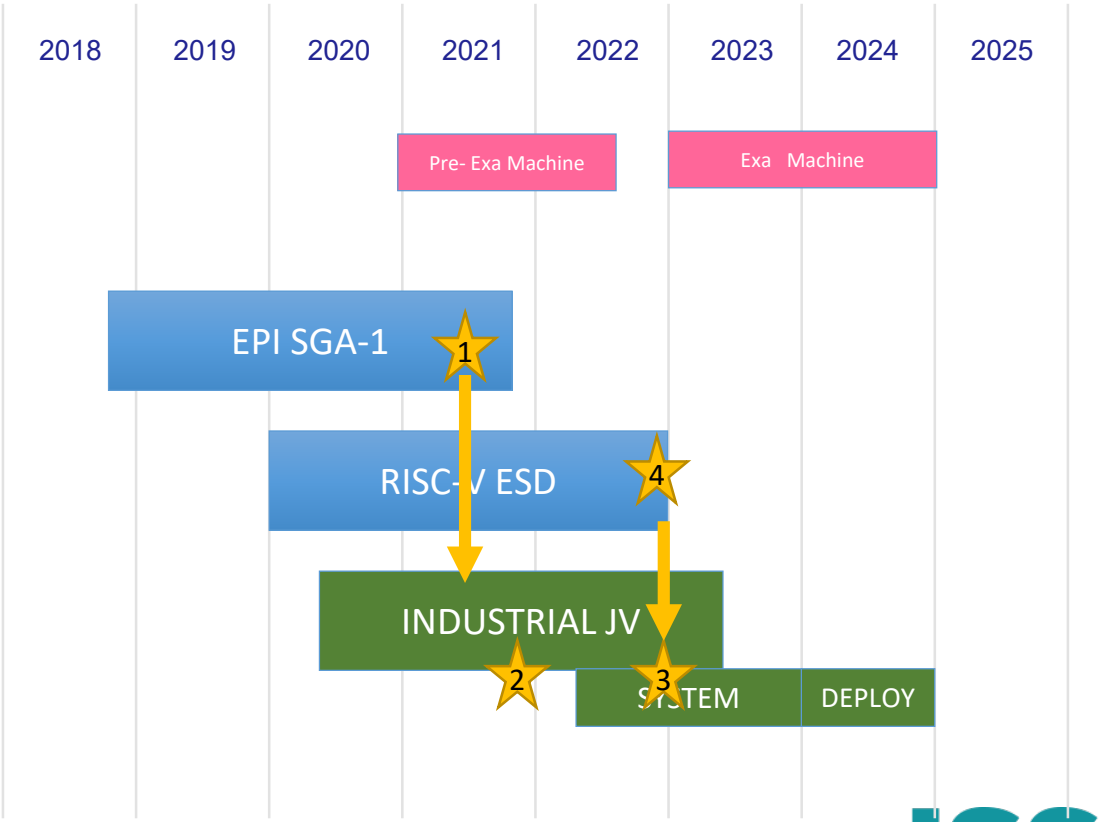

EPI GPP development

- REQUIRES MORE TIME & EFFORT THAN GPU
  - Large OOO, high IPC core more complex
  - Multiple generations needed to achieve parity
- START FOUNDATION IP DEVELOPMENT 2020
  - Use FET-HPC funding starting 2020
  - Pool academia and EPI resources around EU
- THIRD PHASE: INDUSTRIAL PARTNERSHIP
  - Same scheme as accelerator
  - Same JV partners
  - Be ready for 2025

EPI GPP roadmap

- T- GPP TAPEOUT

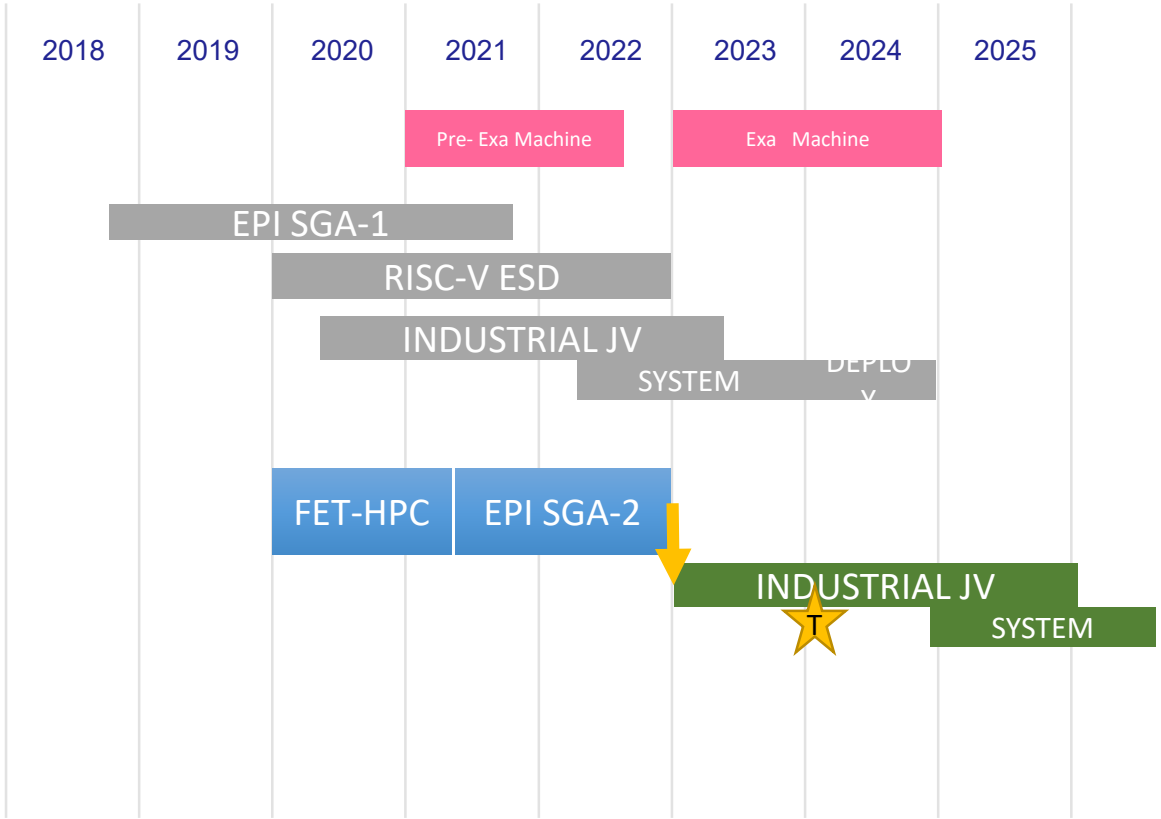

Full system roadmap

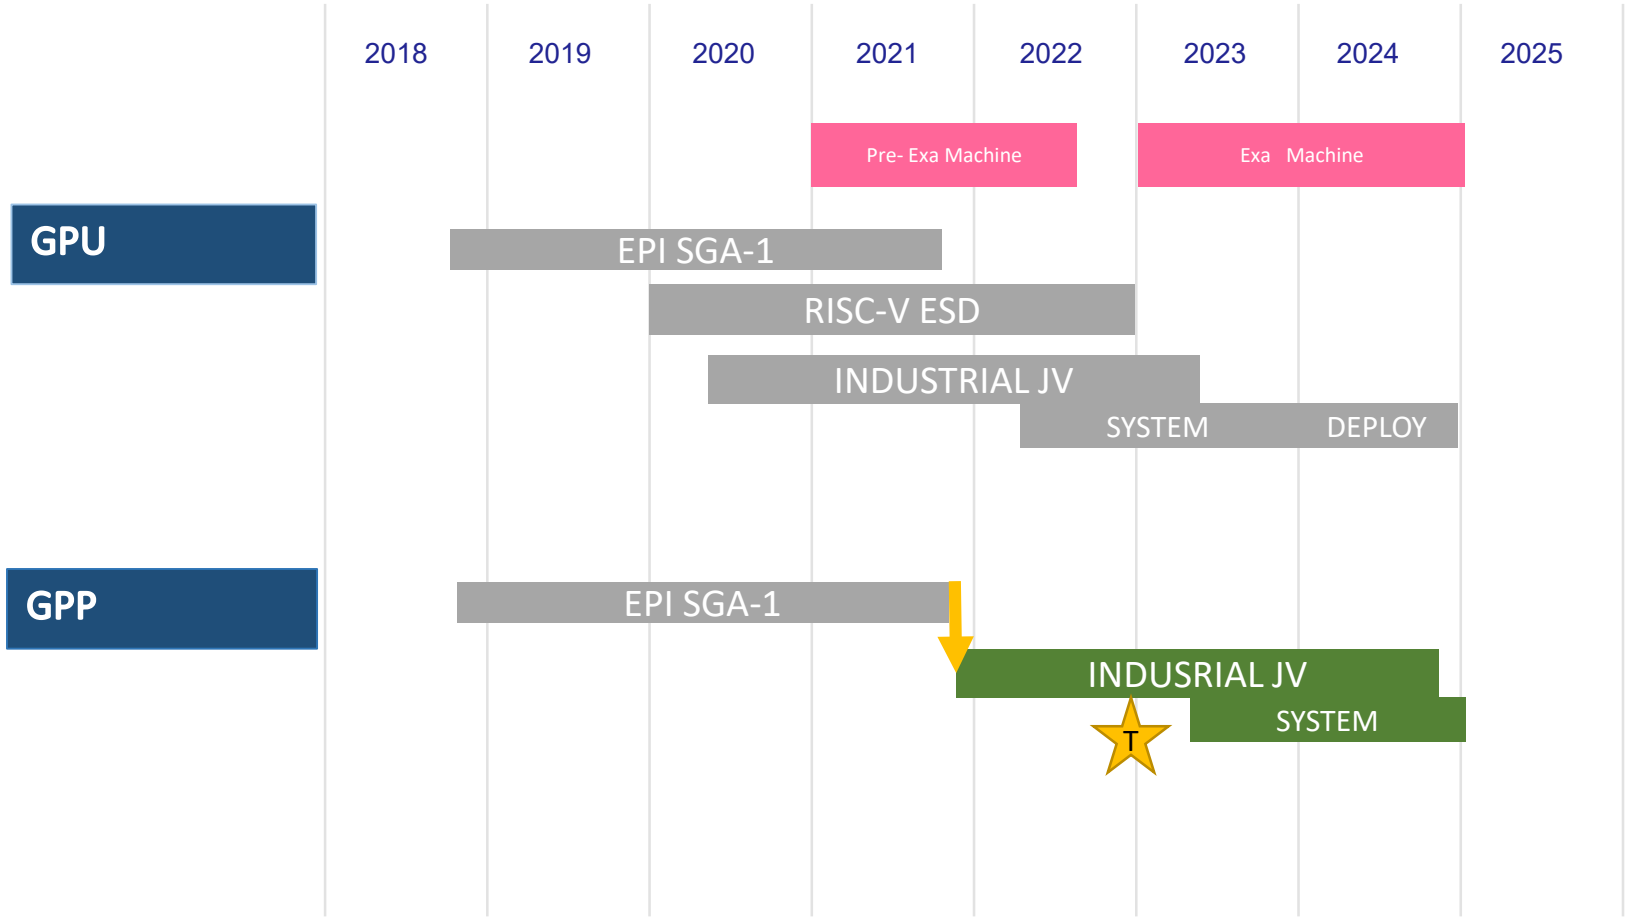

Copyright European Processor Initiative 2019

Scalability allows wide market potential coverage

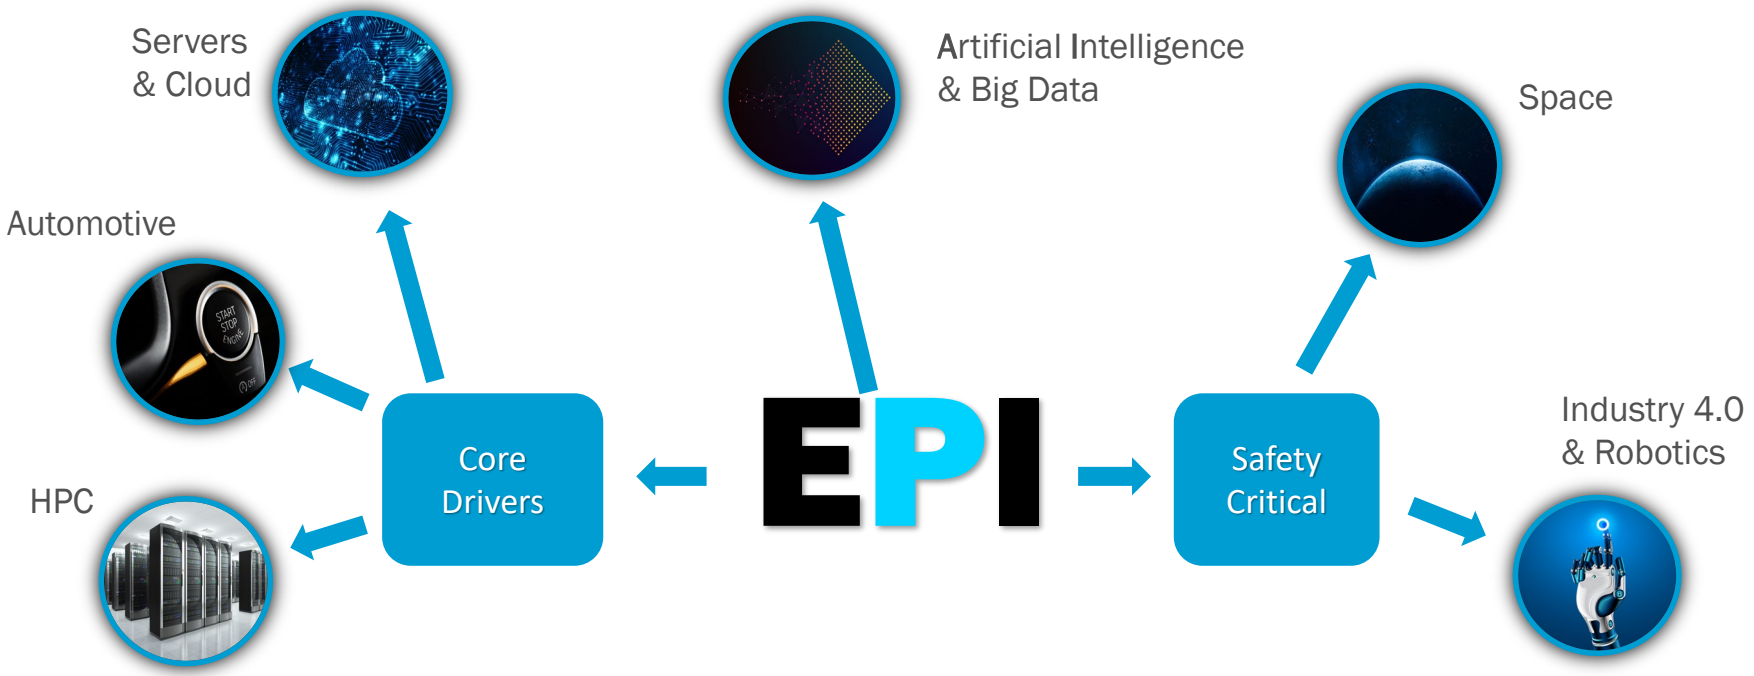

JV partners

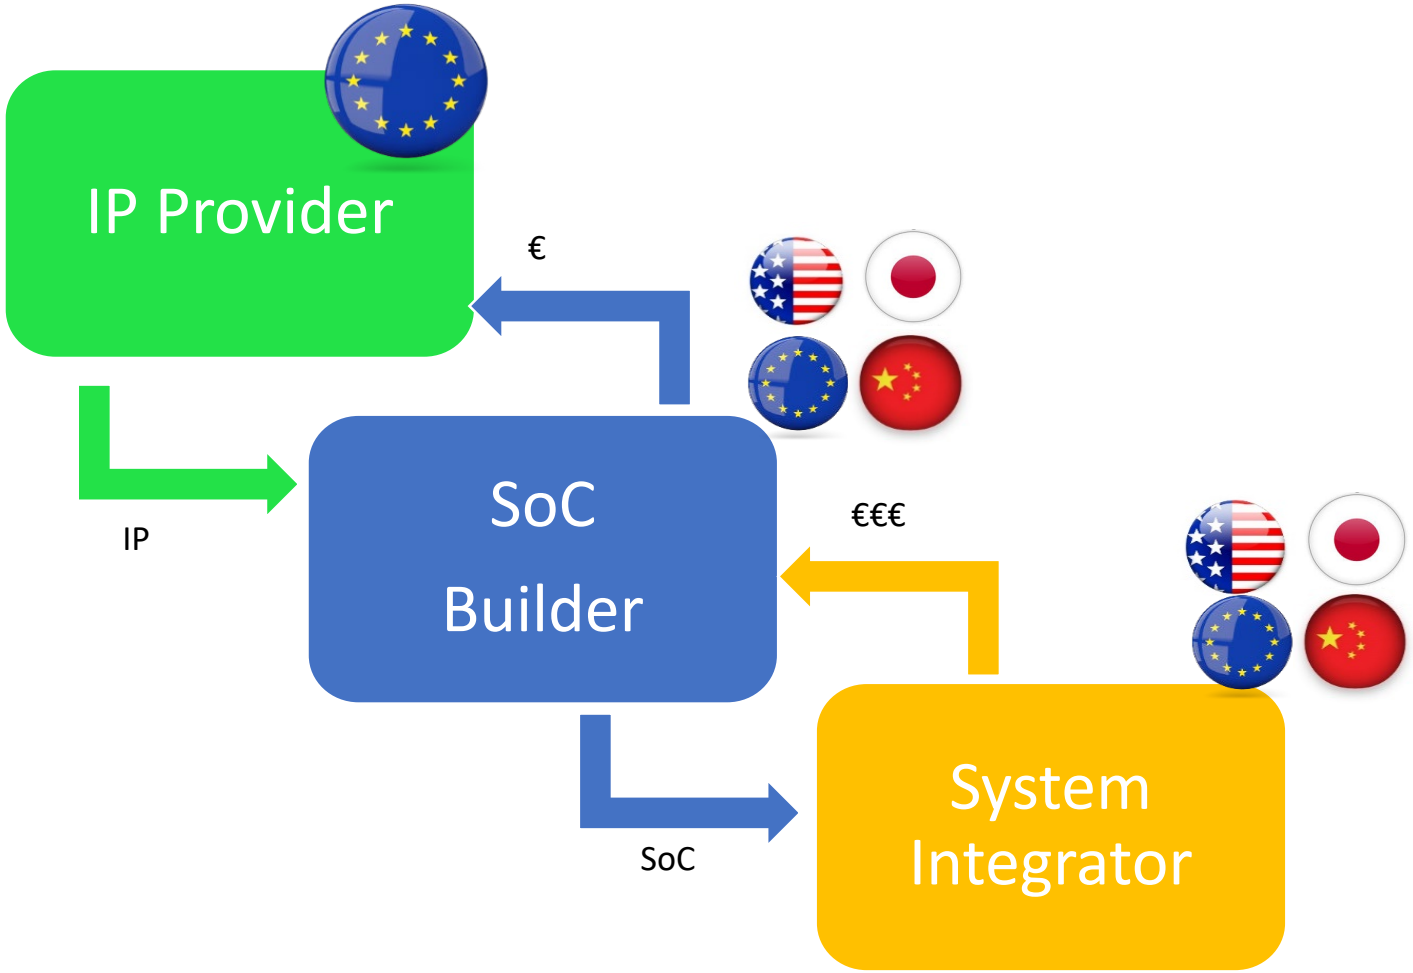

RISC-V ecosystem

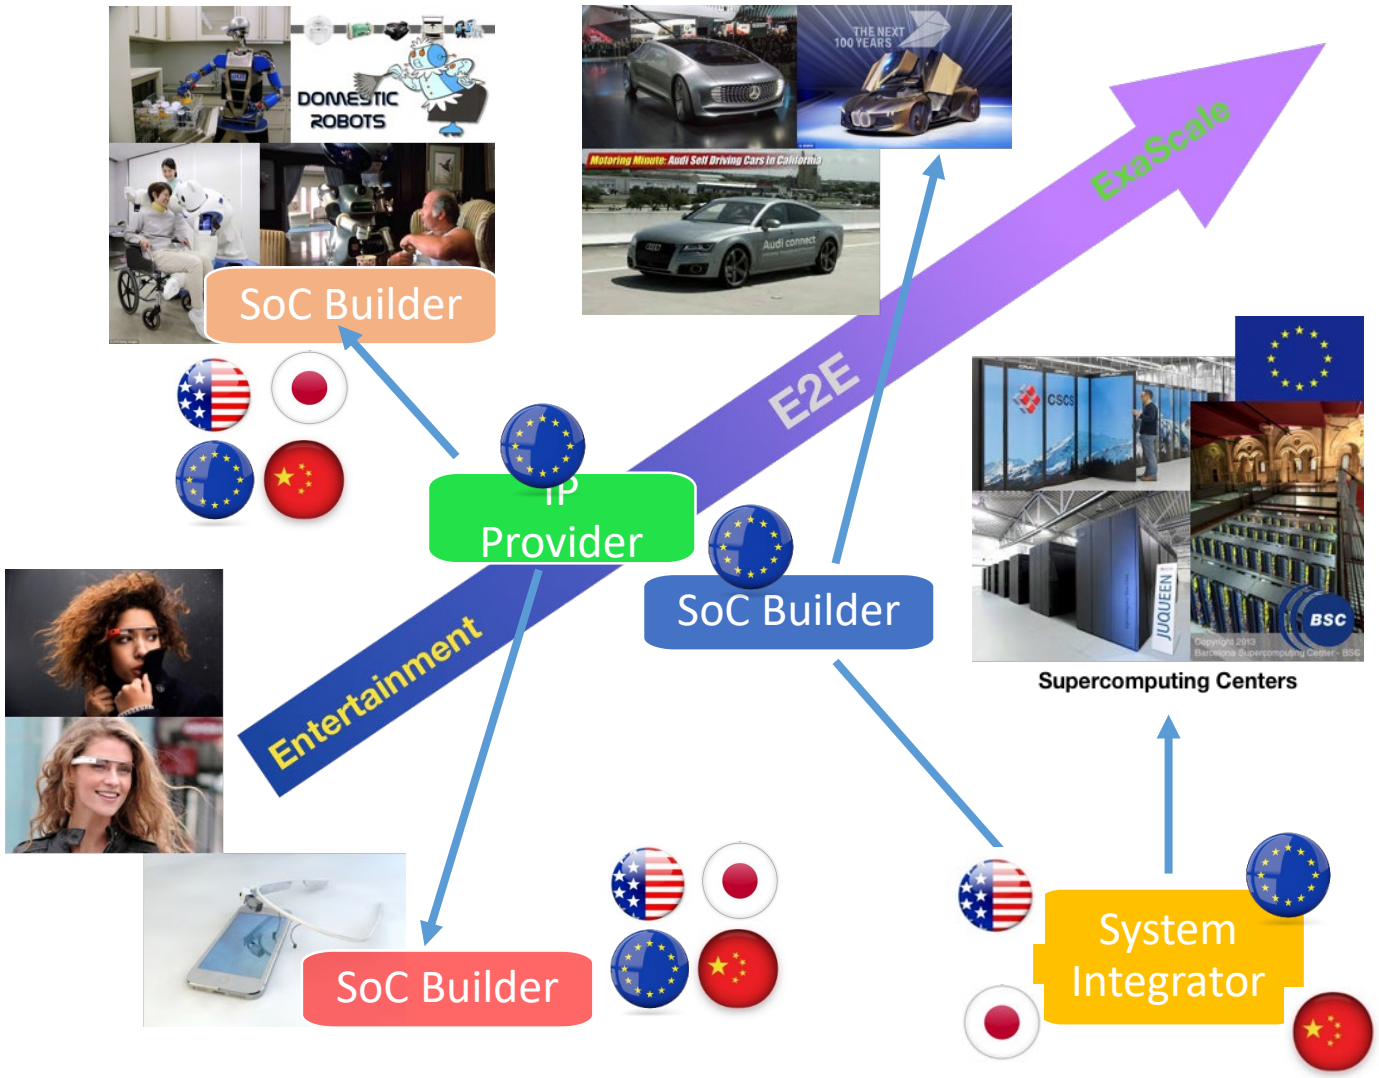

Conclusions

- HPC is crucial to resolve societal challenges and preserve European competitiveness.
- Europe is going in the right direction with EuroHPC. This must be sustained in the longer term.
- The chip design effort must continue for the sake of EU's security and competitiveness, and should create an entire ecosystem from IoT to HPC.
- BSC is fully behind the EuroHPC initiative with major investments and ongoing intense activities for the design and development of the first EPI HPC accelerator, entirely based on RISC-V.

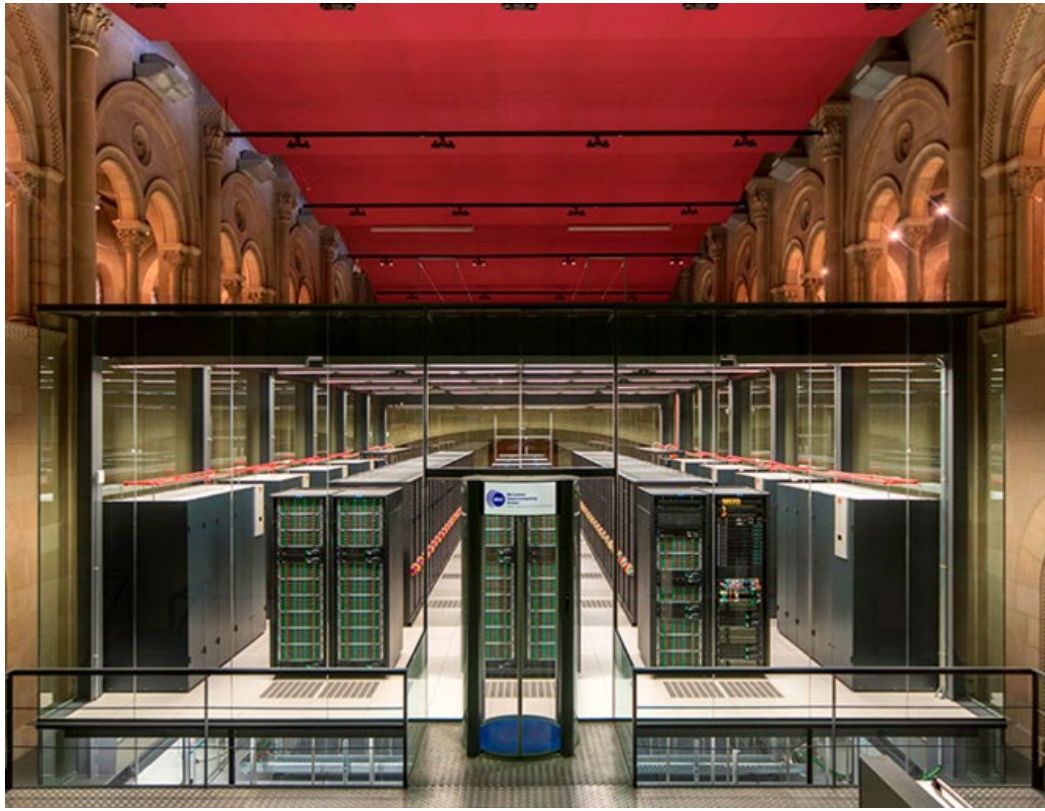

EuroHPC
Joint Undertaking

EuroHPC opens a window of opportunity to create the Airbus/Galileo of HPC

semi

ISS INDUSTRY STRATEGY SYMPOSIUM EUROPE

Thank you!

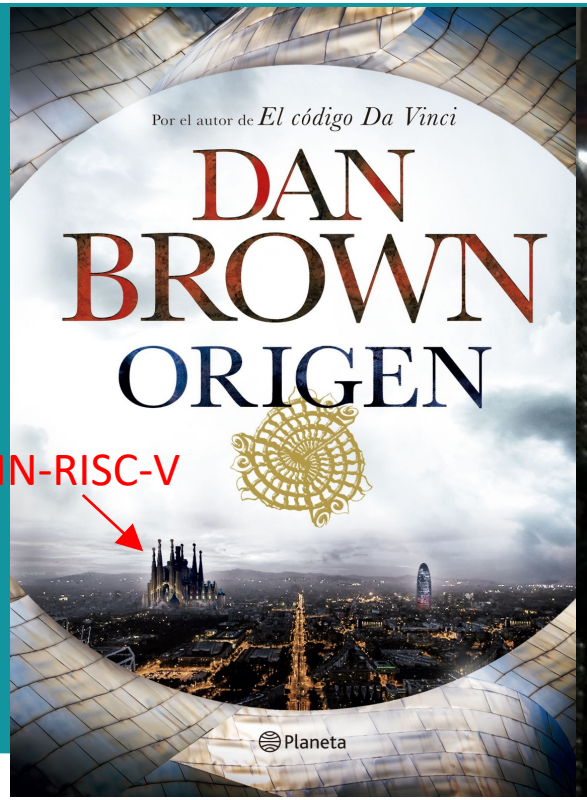
Por el autor de El código Da Vinci
DAN
BROWN
ORIGEN
MN-RISC-V
Planeta