EUROPEAN PROCESSOR INITIATIVE (EPI)

A HIGH PERFORMANCE, HIGH EFFICIENCY PROCESSOR FOR HPC

FROM LAST YEAR'S TALK…

- "Is the era of the general-purpose CPU over?"
- "The end of Moore's Law"
- "ARM, a choice for "small cores" ?"

- One year later… an European project to have an accelerated processor with Arm general-purpose cores !

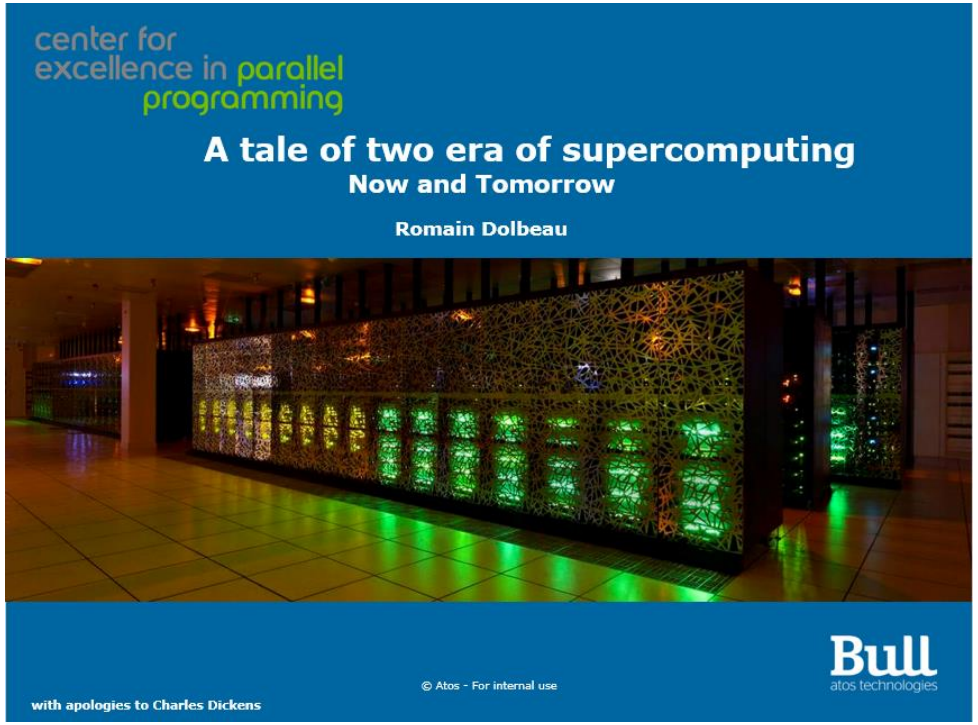

FRAMEWORK PARTNERSHIP AGREEMENT IN EUROPEAN LOW-POWER MICROPROCESSOR TECHNOLOGIES

THIS PROJECT HAS RECEIVED FUNDING FROM THE EUROPEAN UNION'S HORIZON 2020 RESEARCH AND INNOVATION PROGRAMME UNDER GRANT AGREEMENT NO 826647

OFFICIAL TEASER

DRIVERS OF THE EPI PROPOSAL (1)

Societal challenges

- Aging population
- Climate change
- Cybersecurity
- Increasing energy needs
- Intensifying global competition
- Sovereignty (data, economical, embargo)

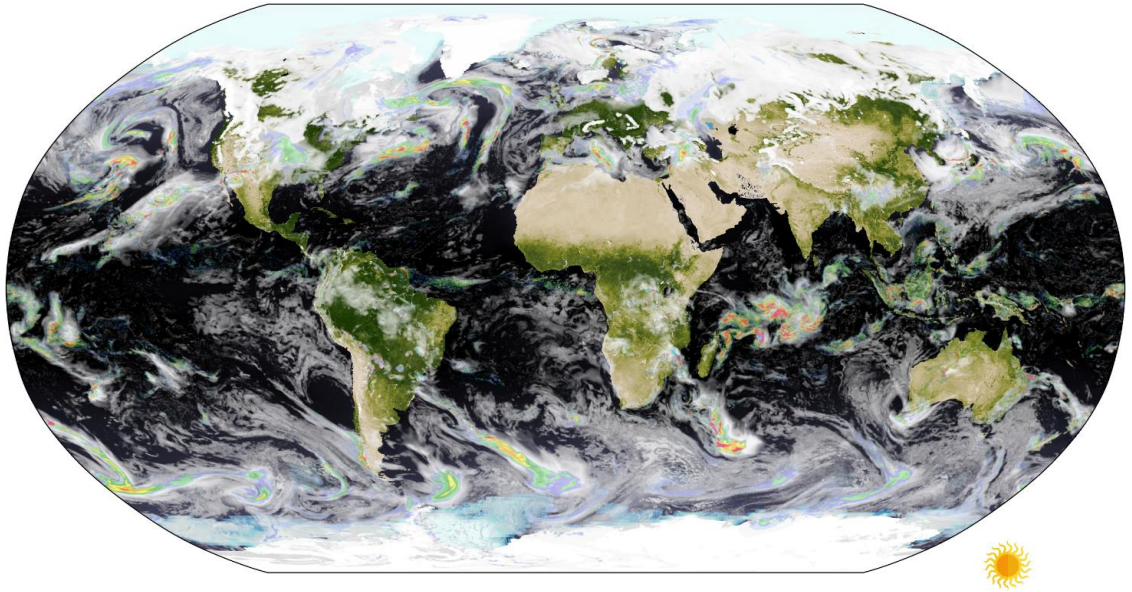

Image/video: courtesy of P.L.Vidale, M.J. Roberts, G.Perez, NCAS, Met Office, University of Reading

DRIVERS OF THE EPI PROPOSAL (1)

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

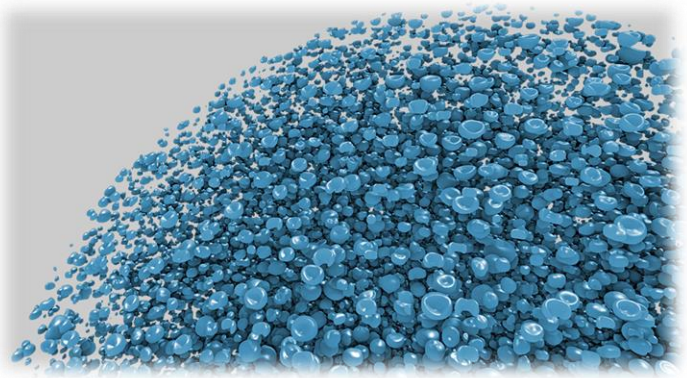

Image courtesy of Petros Koumoutsakos, ETH Zurich

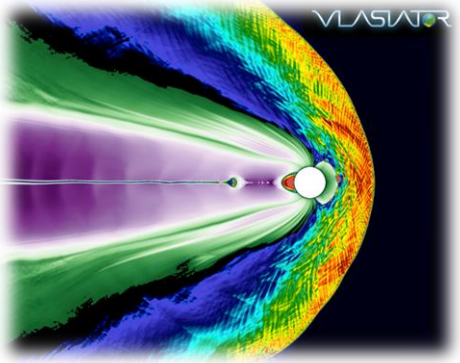

Image courtesy of Minna Palmroth, University of Helsinki

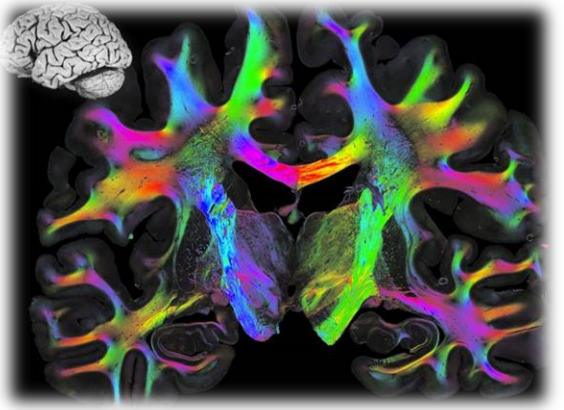

Image courtesy of Axer & Amunts, INM-1, Forschungszentrum Jülich

DRIVERS OF THE EPI PROPOSAL (2)

- Connected mobility & AD *Autonomous Driving computing needs beyond 2023*
- Develop customized processors able to meet the performance needed for autonomous vehicles that would offer:
  - implementation of vehicle perception tasks in real-time in a fail-operational manner
  - increased computing performance, fail-operational, functional safety, cyber-security and real-time behaviour (RT)
  - compute resources with the same characteristics as their “big brothers” in exascale class supercomputers
- Sovereignty (data, economical, embargo)
- EU car manufacturing supremacy

DRIVERS OF THE EPI PROPOSAL (3)

- Servers and Cloud Low Power CPU needs:
  - energy efficiency - lower power consumption
  - new generation of secure and safety-aware virtualization capabilities
- Sovereignty (data, economical, embargo)

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
- Sovereignty (data, economical, embargo)

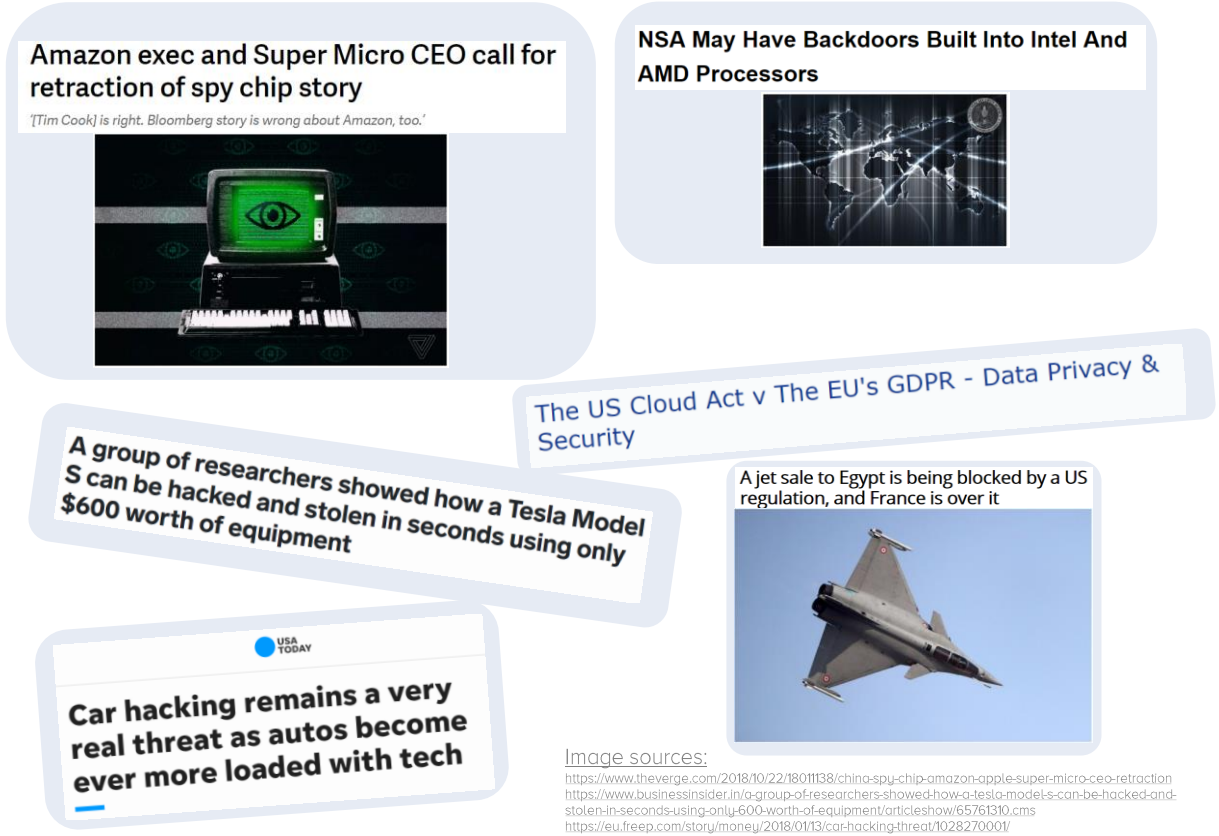

Image sources:

https://www.theverge.com/2018/10/22/18011138/china-spy-chip-amazon-apple-super-micro-ceo-retraction
https://www.businessinsider.in/a-group-of-researchers-showed-how-a-tesla-model-s-can-be-hacked-and-stolen-in-seconds-using-only-600-worth-of-equipment/articleshow/65761310.cms
https://eu.freep.com/story/money/2018/01/13/car-hacking-threat/1028270001/
https://www.eteknix.com/nsa-may-backdoors-built-intel-amd-processors/
https://www.pearse-trust.ie/blog/the-us-cloud-act-v-the-eus-gdpr-data-privacy-security
https://www.defensenews.com/global/europe/2018/08/01/a-jet-sale-to-egypt-is-being-blocked-by-a-us-regulation-and-france-is-over-it/

HOW EUROHPC WILL HELP TO MAKE US STRONGER

- Developing a new European supercomputing ecosystem: HPC systems, network, software, applications, access through the cloud
- Making HPC resources available to public and private users, including SMEs.
- Stimulating a technology supply industry

EUROPEAN PROCESSOR INITIATIVE

- High Performance General Purpose Processor for HPC
- High-performance RISC-V based accelerator
- Computing platform for autonomous cars
- Will also target the AI, Big Data and other markets in order to be economically sustainable

EPI PARTNERS

MISSION

- European independence in High Performance Computing Processor Technologies
- EU Exascale machine based on EU processor by 2023
- Based on solid, long-term economic model, Go beyond HPC market
- Address the needs of European industry (car manufacturing market)
- End-to-end data security

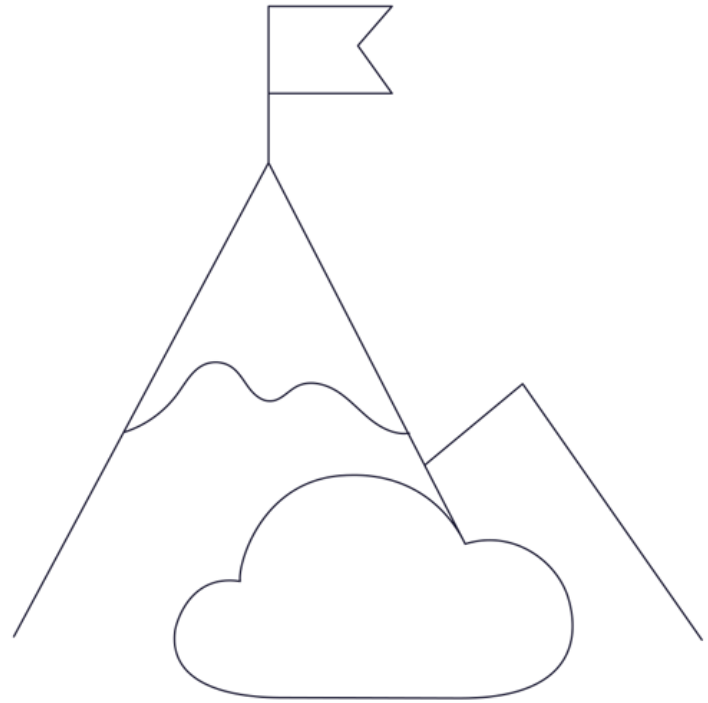

VISION

- High Performance Computing needs for Exascale machines beyond 2022
- Connected mobility & AD Autonomous Driving computing needs beyond 2023
- Low power CPU needs for Servers and Cloud
- Other markets under exploration (Server and Cloud)

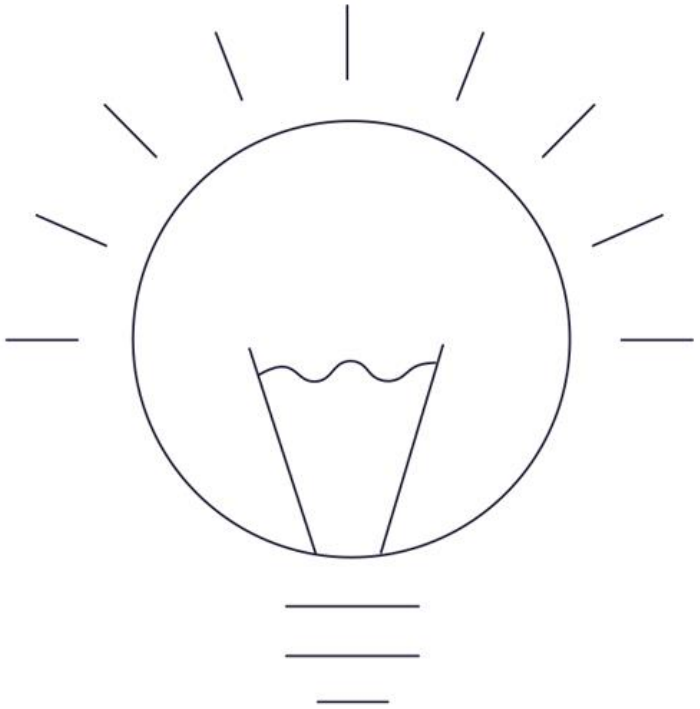

EXPECTED IMPACT

- Strengthening the competitiveness and leadership of European industry and science
- European microprocessor technology with drastically better performance/power ratios
- Tackling important segments of broader and/or emerging HPC and Big-Data markets

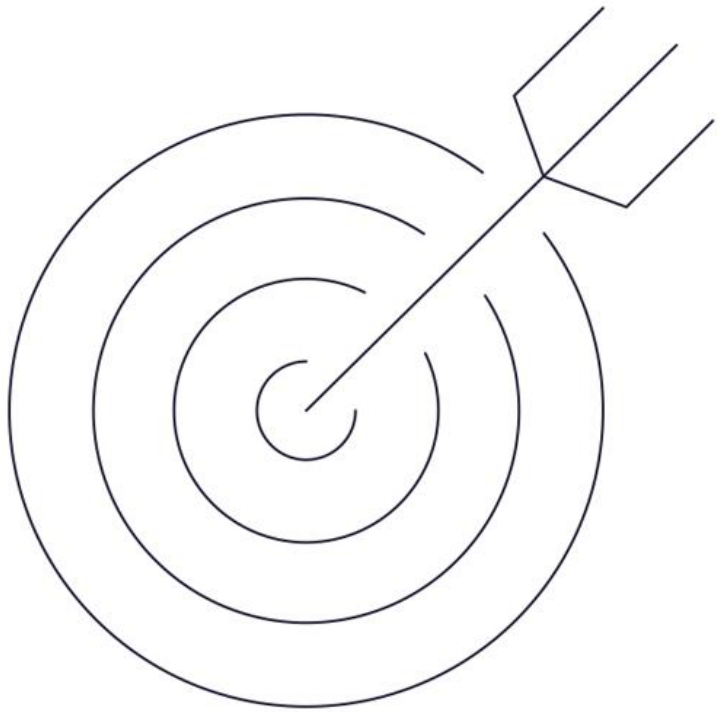

ROADMAP

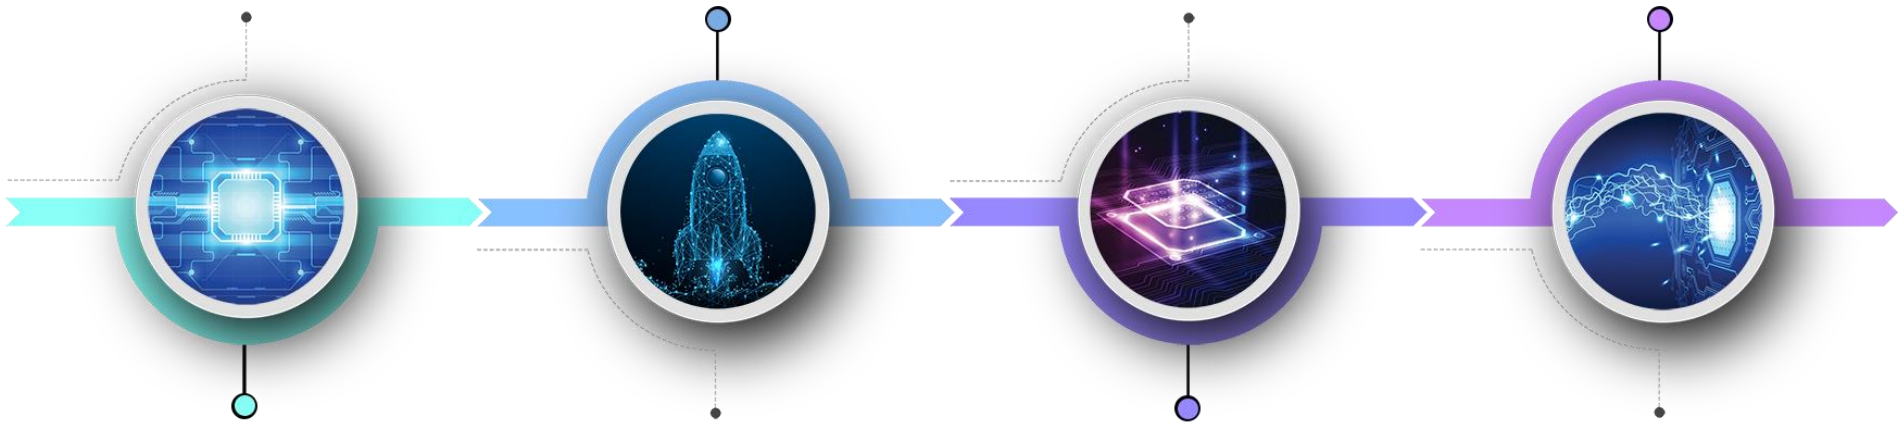

2021

Rhea Family - Gen1 GPP

EPI Common Platform
ARM & RISC-V
External IPs
•
HPC System PreExascale
Automotive PoC

2021-2022

EPI IP's Launch Pad
&
Pan European Research Platform for HPC and AI

2022-2023

Cronos Family - Gen2 GPP

EPI Common Platform
ARM & RISC-V
•
HPC System Exascale
Automotive CPU

Gen3 GPP Family

2024-...

GPP AND COMMON ARCHITECTURE

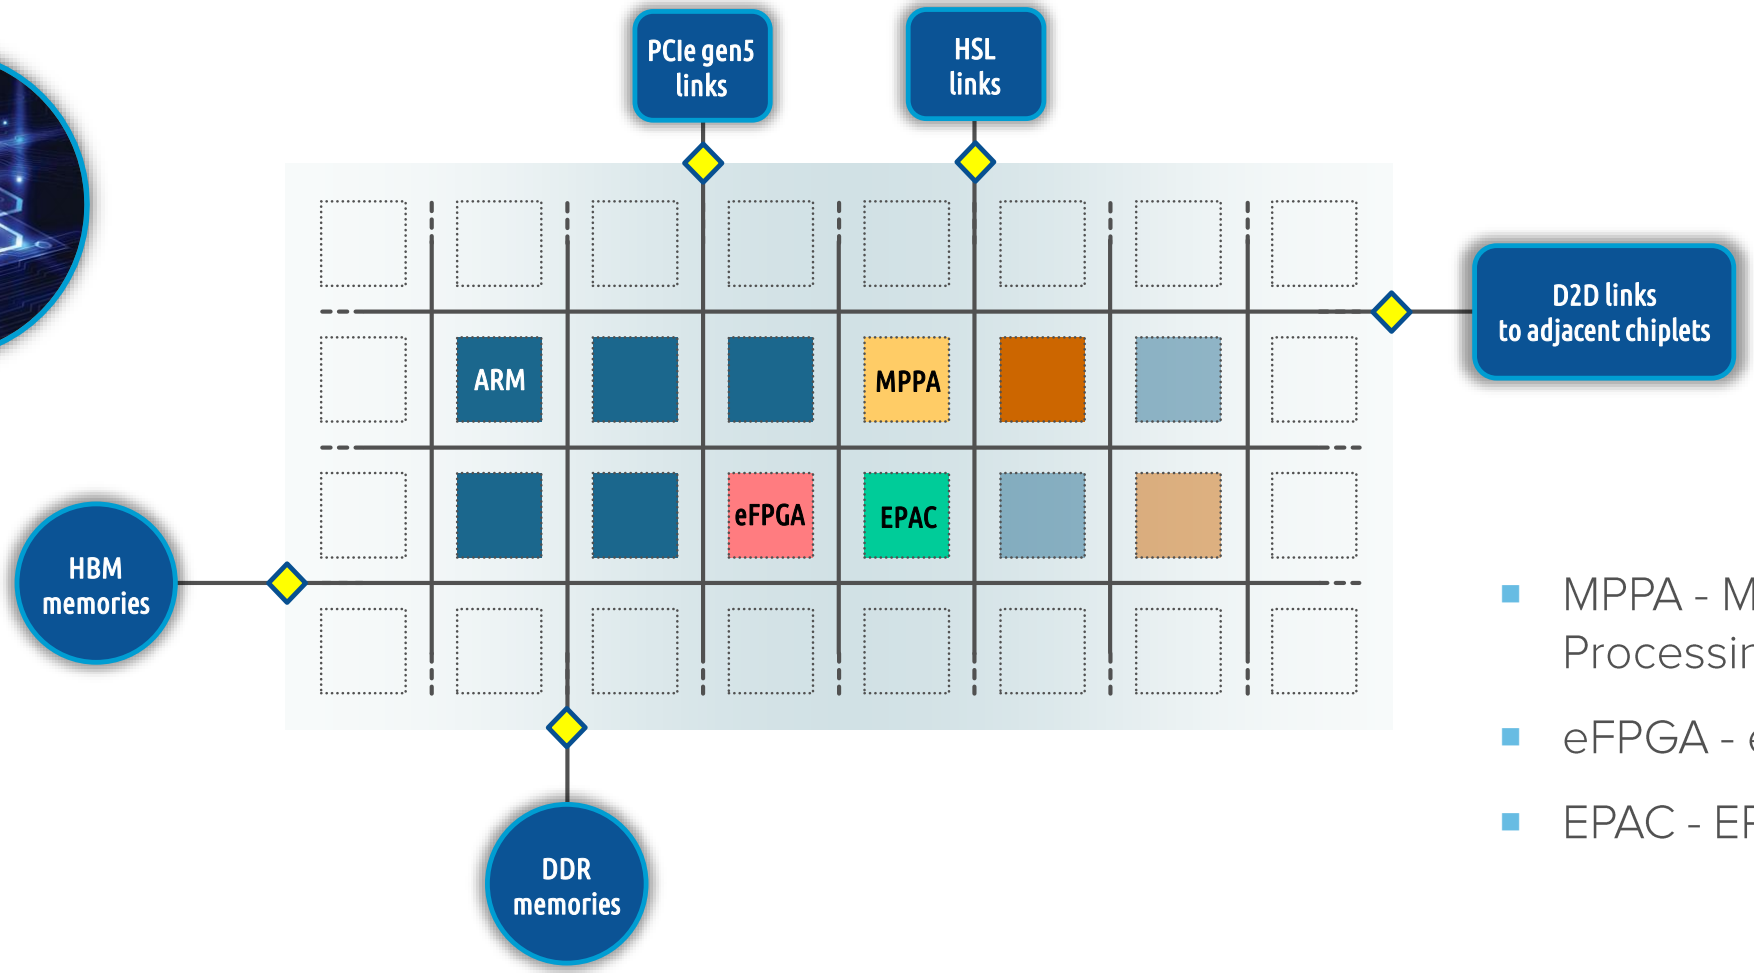

- MPPA - Multi-Purpose Processing Array
- eFPGA - embedded FPGA
- EPAC - EPI Accelerator

EPAC – RISC-V ACCELERATOR

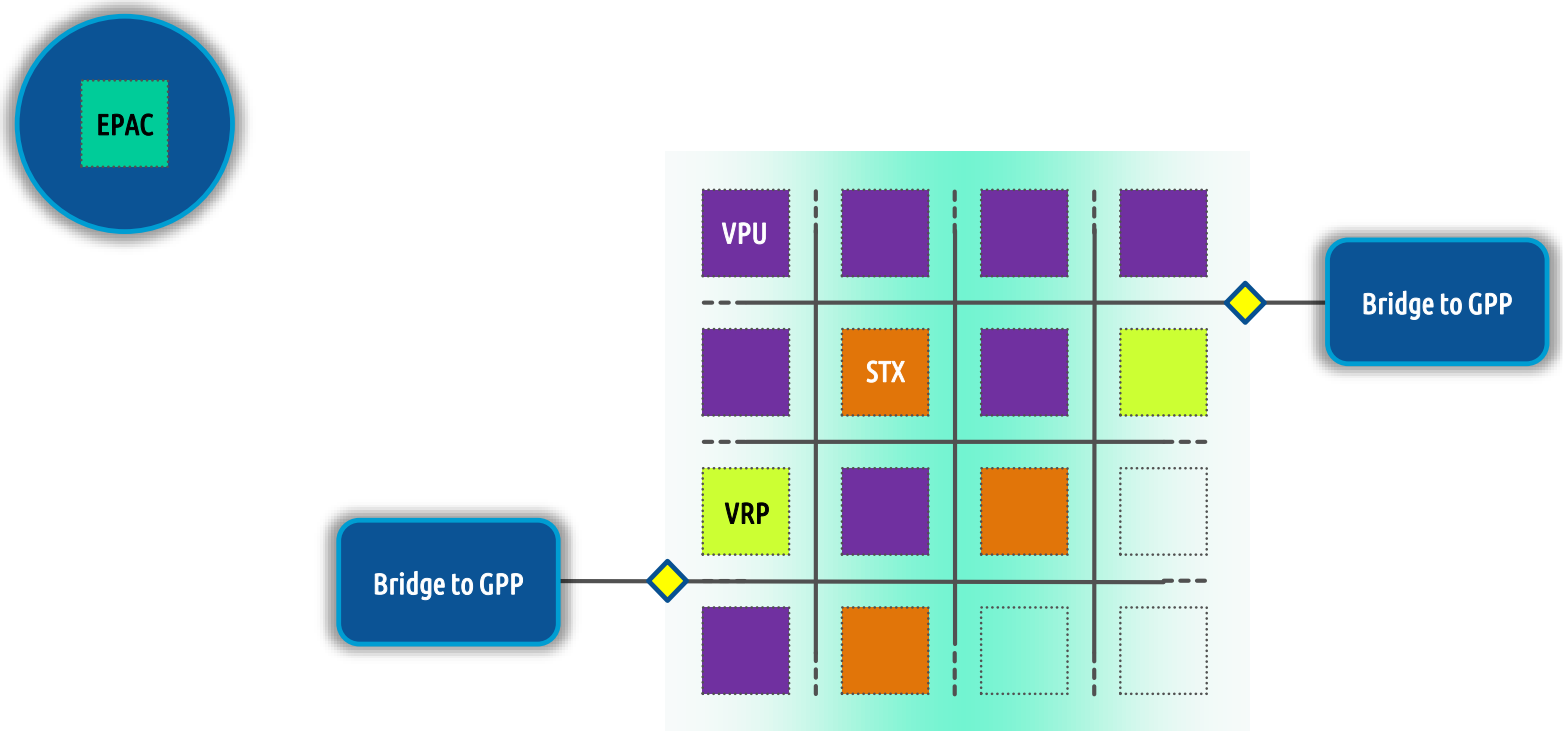

- EPAC - EPI Accelerator
- VPU – Vector Processing Unit
- STX – Stencil/Tensor accelerator
- VRP - VaRiable Precision co-processor

EPI AUTOMOTIVE

- Autonomous driving systems
- Connected mobility
- EPI: A powerful data fusion platform – the automotive embedded HPC platform
- EPI heterogeneous multicore architecture can provide enough performance and low power consumption in parallel

END2END SECURITY – FROM THE AUTOMOTIVE SYSTEM TO THE CLOUD

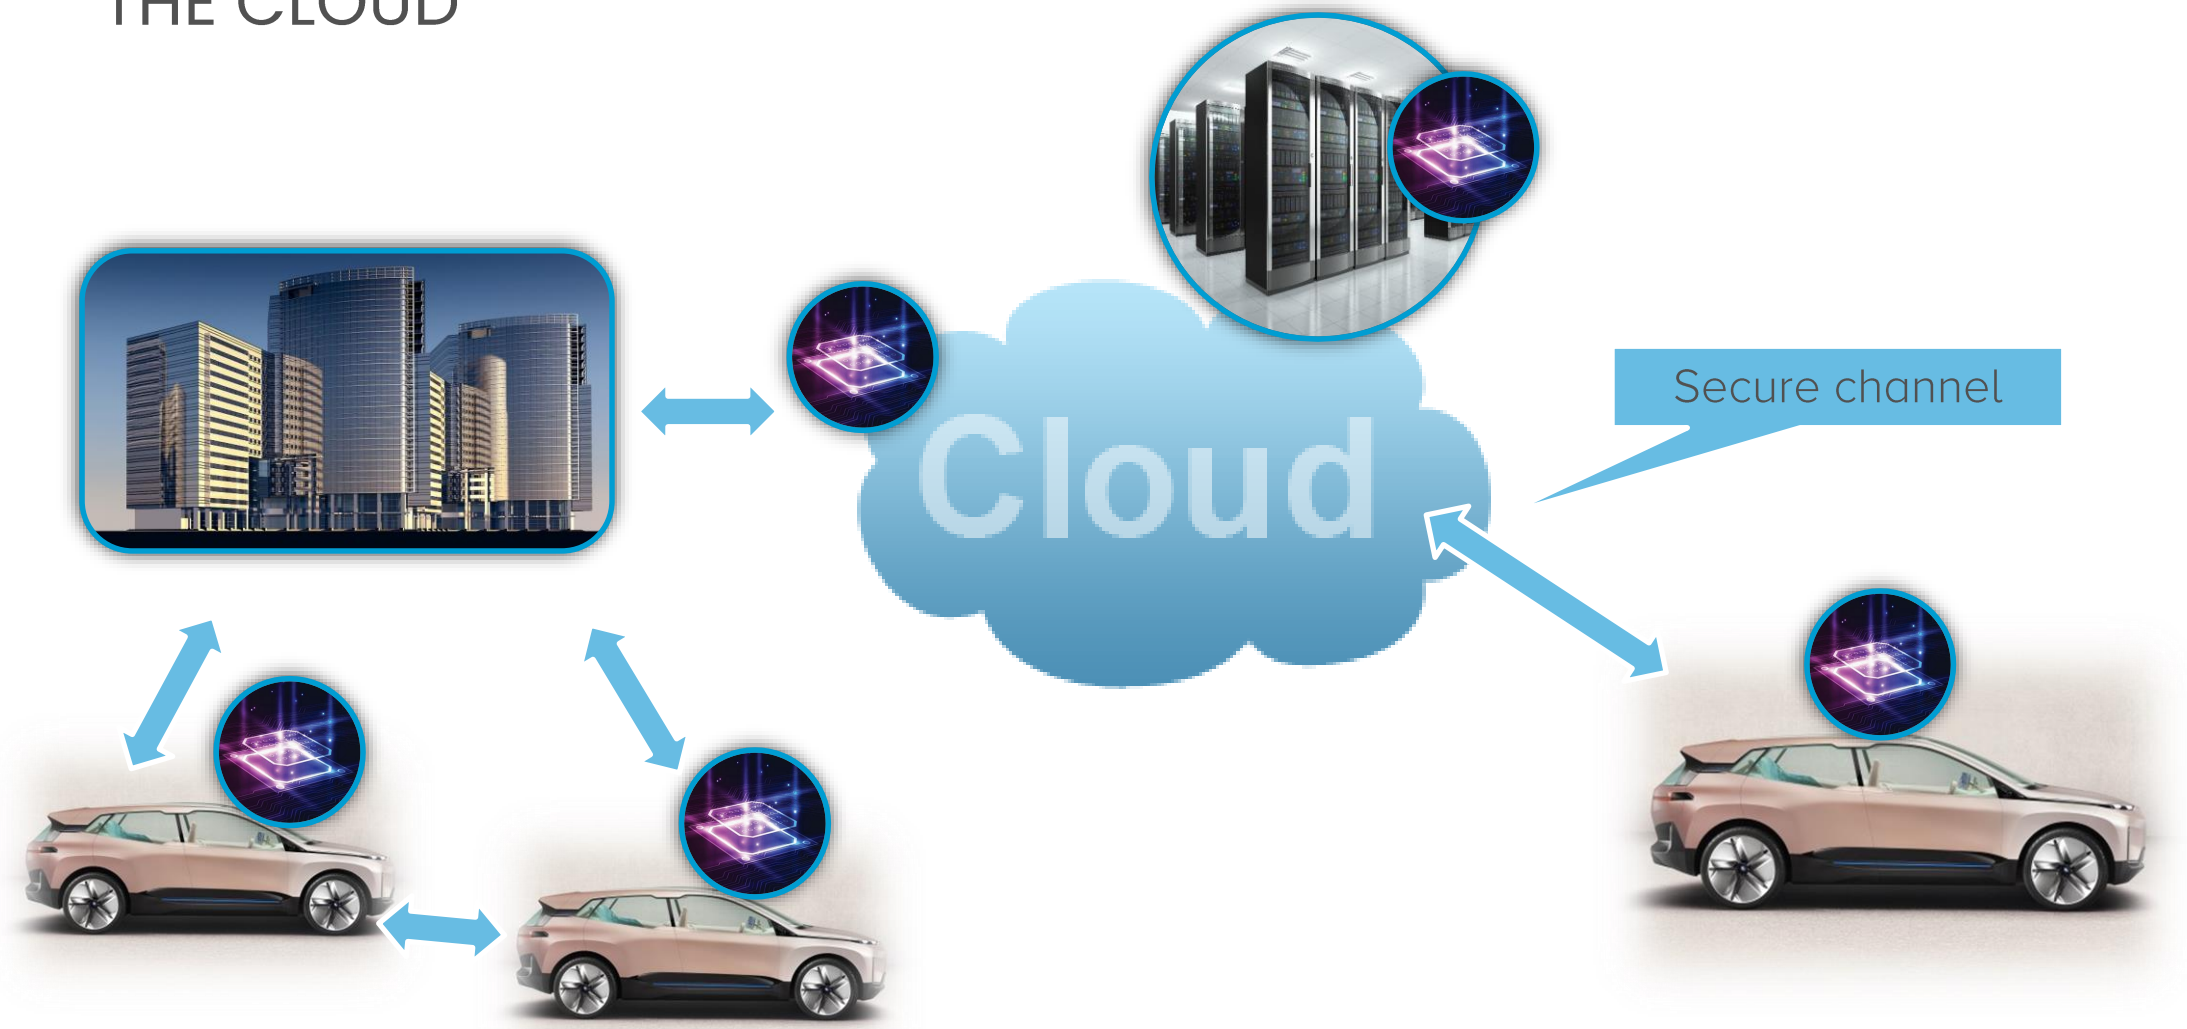

EPI FABLESS COMPANY

- EPI's Fabless company
  - licence of IPs from the partners
  - develop own IPs around it
  - licence the missing components from the market

  - generate revenue from both the HPC, IA, server and eHPC markets
  - integrate, market, support & sales the chip
  - work on the next generations

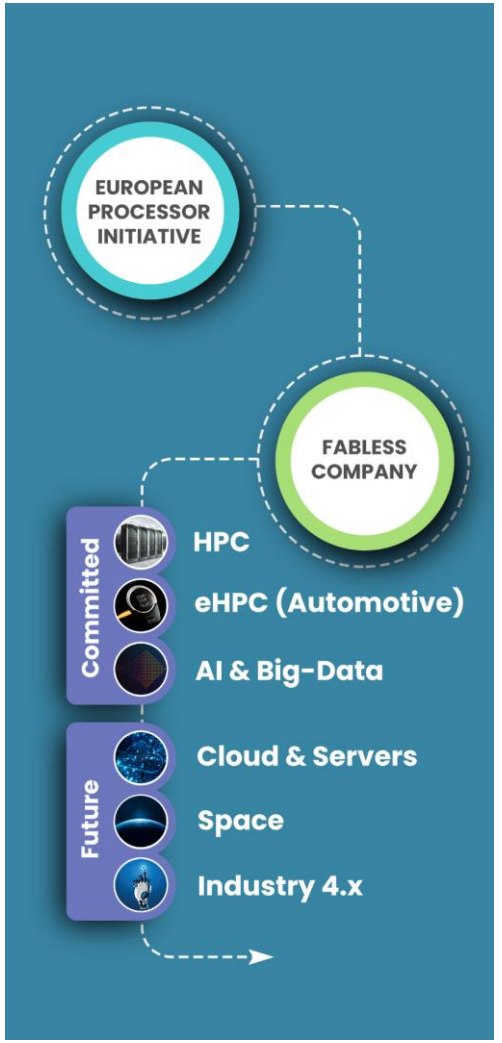

SCALABILITY ALLOWS WIDE MARKET POTENTIAL COVERAGE

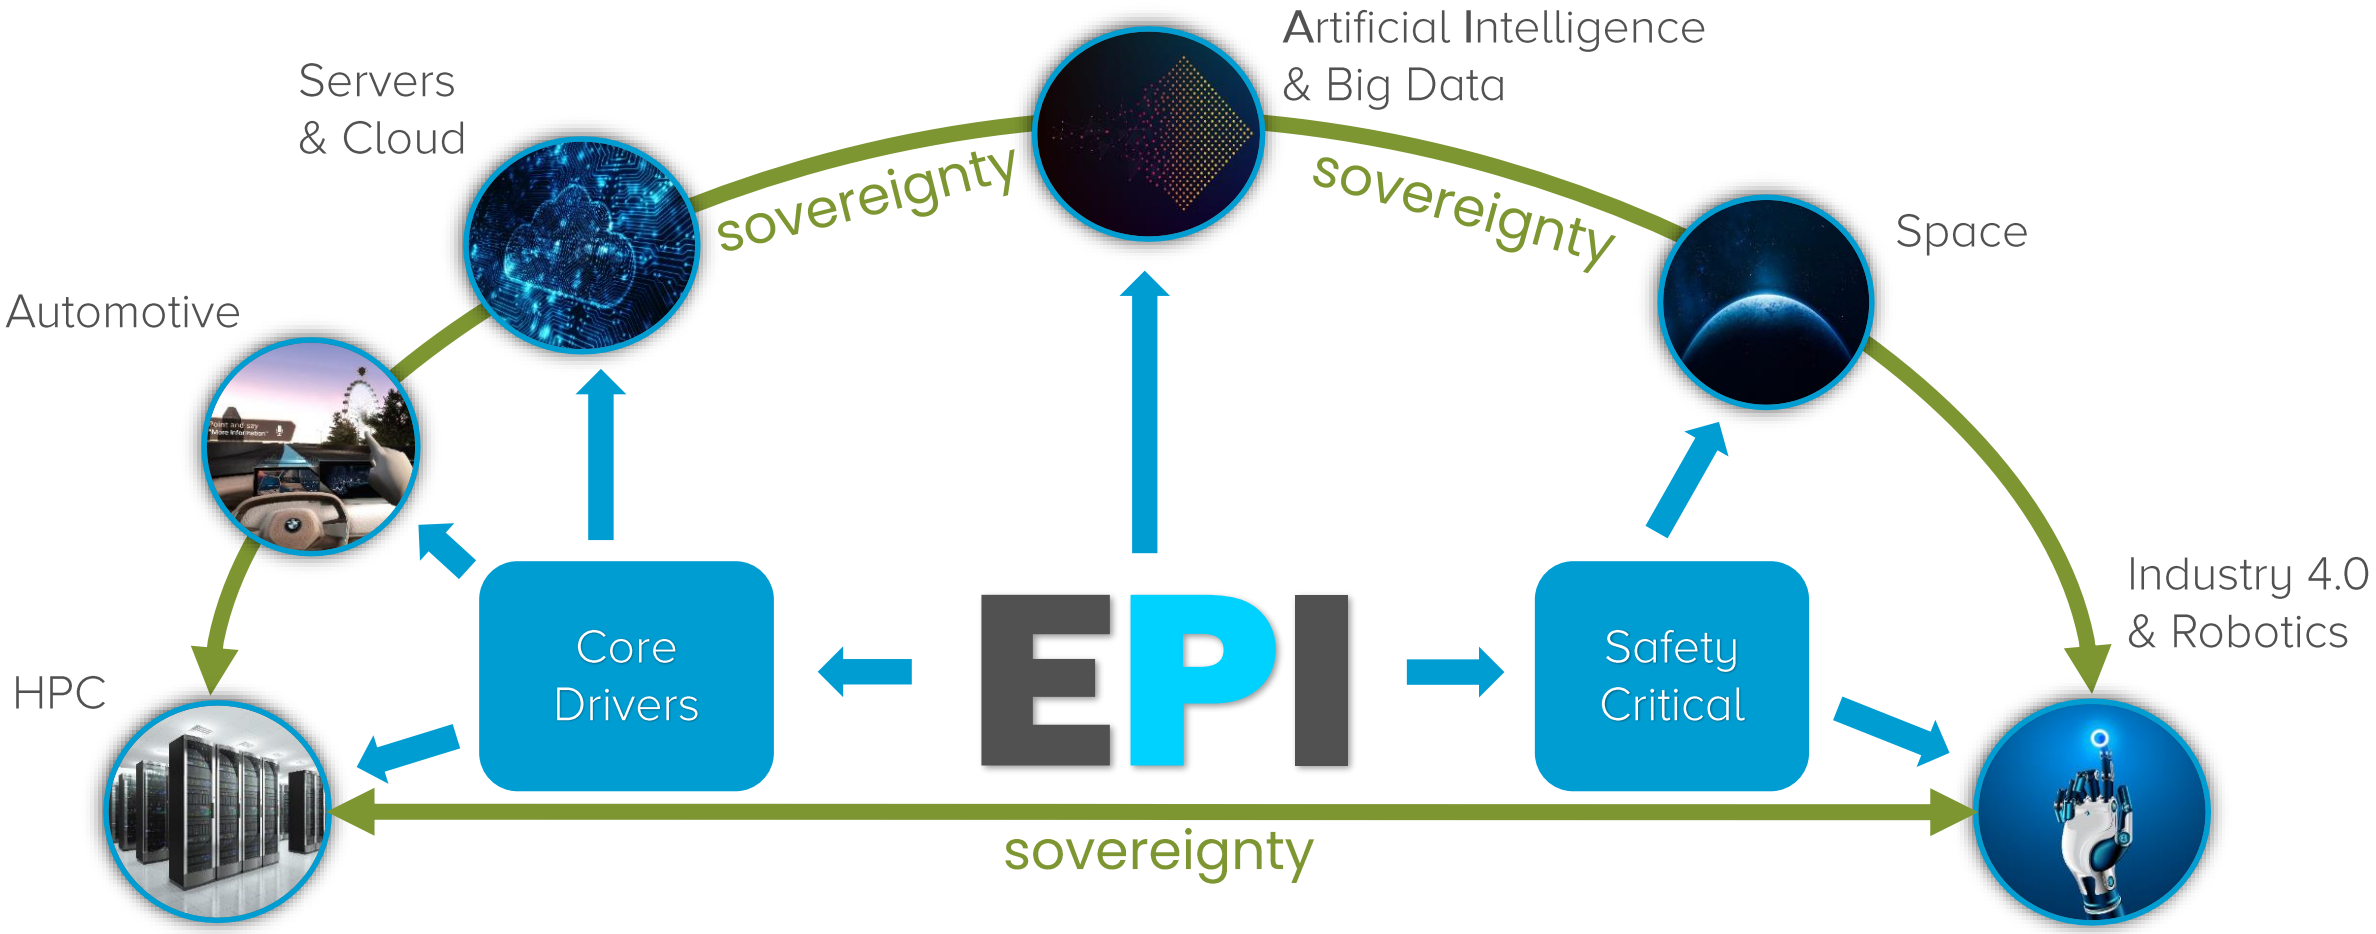

CONCLUSION

- 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 long-term
- The chip design effort must continue for the EU's security and competitiveness, and should create a processor ecosystem covering IoT, servers, cloud, autonomous connected vehicles and HPC

www.european-processor-initiative.eu

@EuProcessor

European Processor Initiative

European Processor Initiative

SOME MORE TECHNICAL DETAILS

ON SOME SLIDES FROM THE OFFICIAL TEASER

GPP AND COMMON ARCHITECTURE

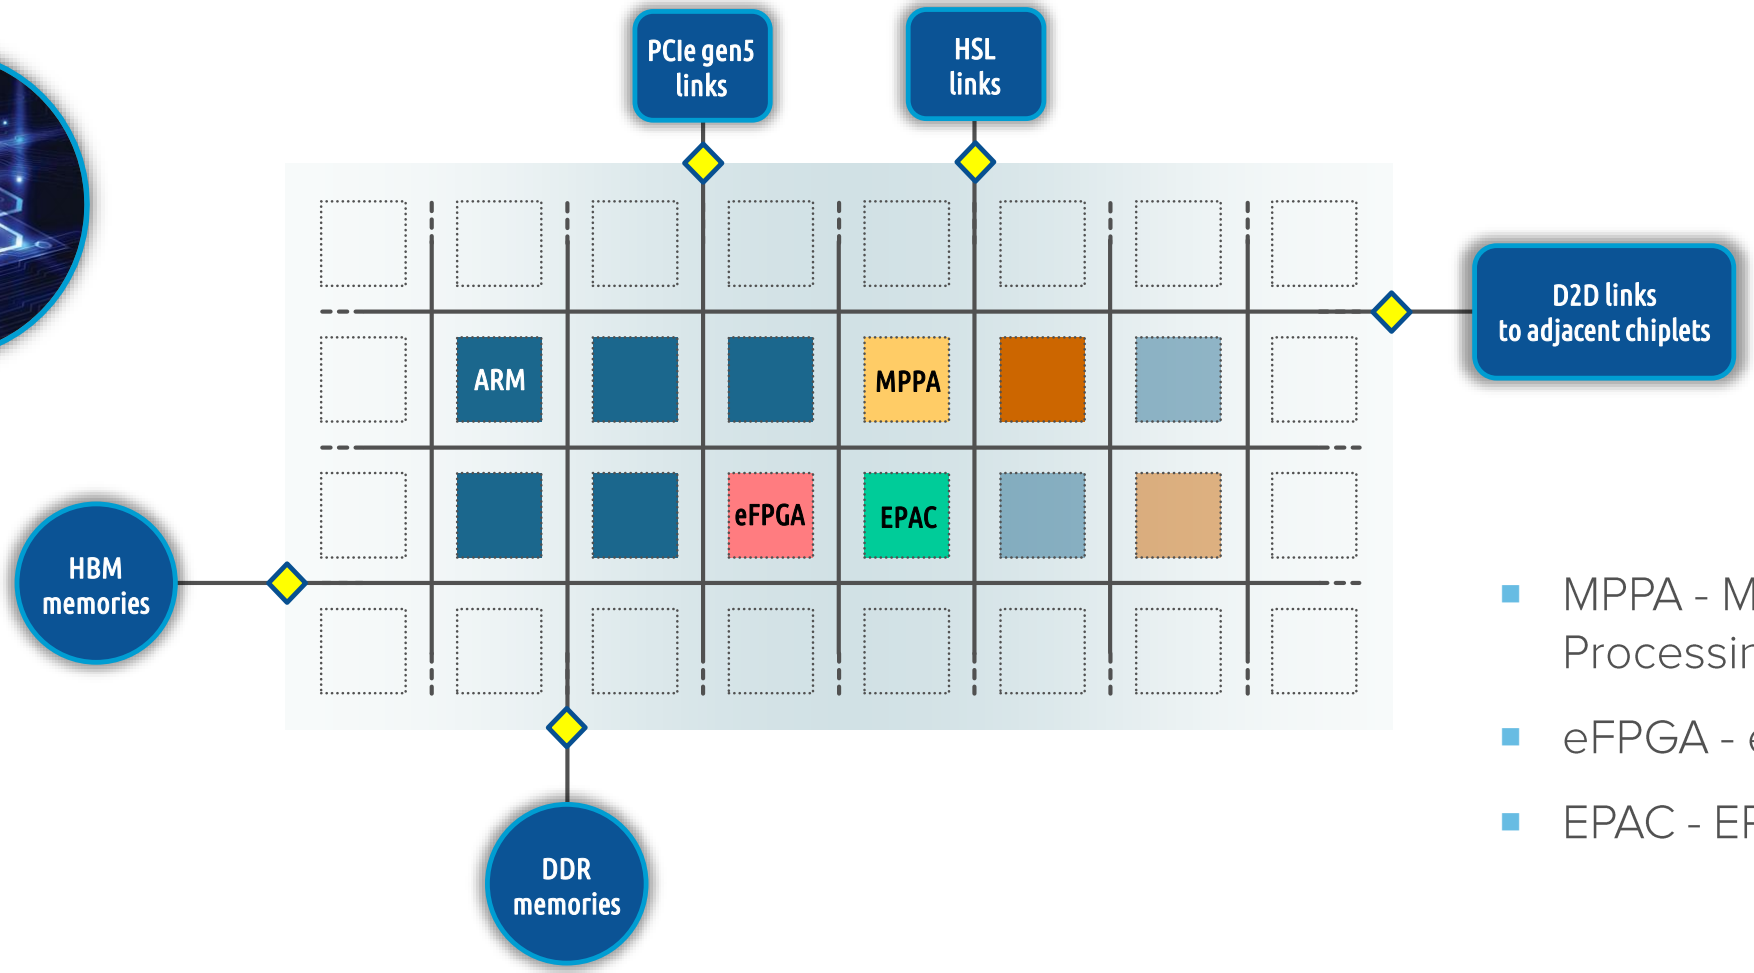

- MPPA - Multi-Purpose Processing Array
- eFPGA - embedded FPGA
- EPAC - EPI Accelerator

ARM & COMMON PLATFORM

- High-performance general-purpose cores based on the Arm v8.2 (or later) architecture
- Using the Scalable Vector Extension for faster computations
  - Replaces NEON as the FPU of choice
  - Implementation-defined vector width, from 128 to 2048 bits
    - Theoretically… 256 or 512 bits in practice
    - Because of alignment (avoid any non-power-of-two), cache line size (practical upper bound), …
- Handle the general-purpose aspect of the Common Platform
  - Operating system, control-driven codes, …
- Repetitive computations offloaded to dedicated accelerators
  - “Dark Silicon” to preserve power & thermal characteristics
- MPPA as a time-predictable accelerator
  - (Soft) Real Time, automotive market
- eFPGA for highly custom functionalities
- EPAC (next slides) for HPC-style computations
- Dedicated subsystems for power management & security management
- Everything fed by multiple stacks of HBM + DDR5
  - Targeting > 1 TB/s of bandwidth in a socket

EPAC – RISC-V ACCELERATOR

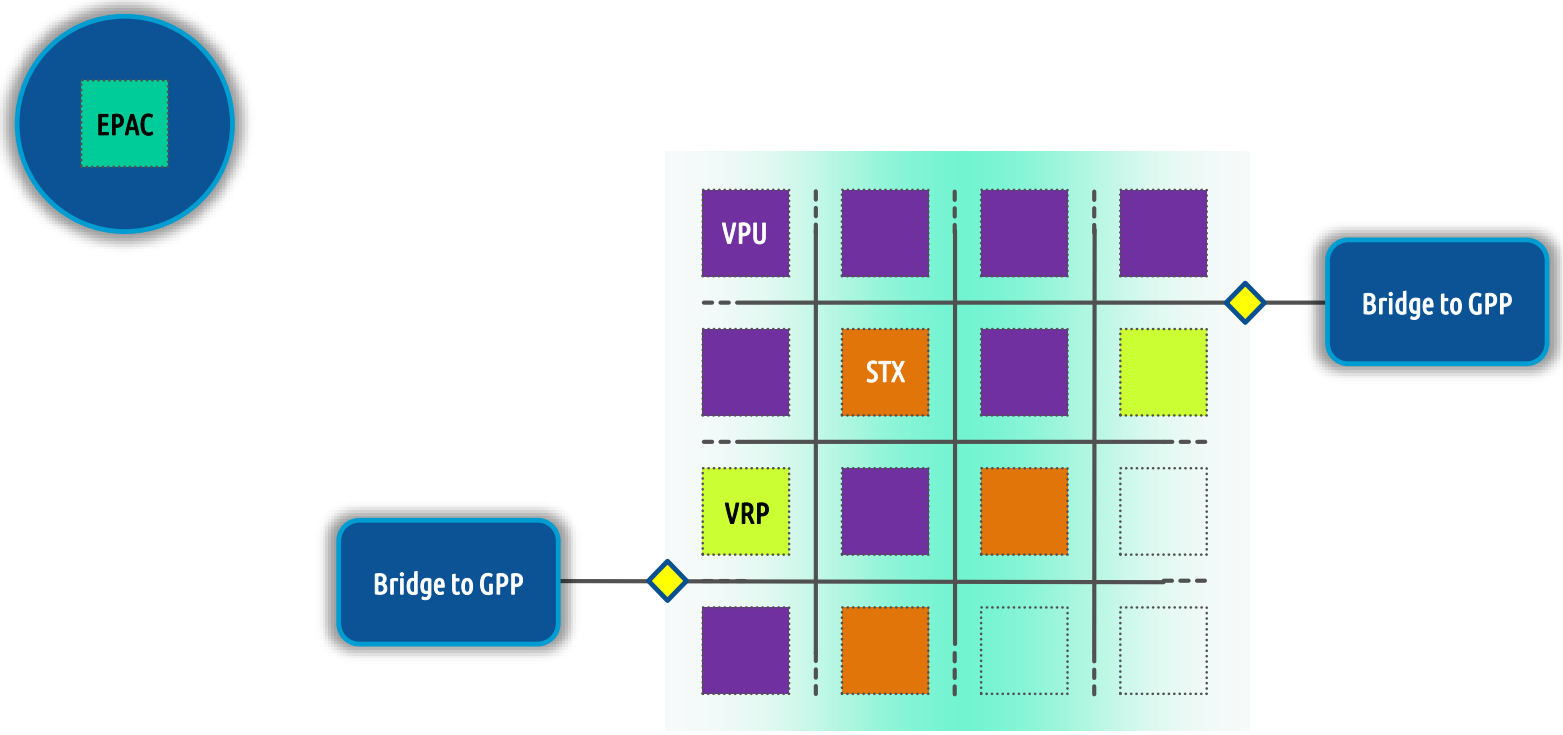

- EPAC - EPI Accelerator
- VPU – Vector Processing Unit
- STX – Stencil/Tensor accelerator
- VRP - VaRiable Precision co-processor

EPAC

- Multiple functionalities embedded in an accelerator “tile” in the design
- Based on the RISC-V open architecture
  - Leveraging the existing work in processor core, “uncore”, compilers, …
- VPU leverages the RISC-V “V” (Vector) extension
  - Much “smaller” core than the GPP cores
  - Large vector for high flops
  - Specific design points (memory hierarchy, …) to sustain the VPU throughput
  - Almost “general-purpose” – HPC-oriented
- STX is for stencil/tensor accelerator
  - For highly specific workload
  - Some HPC workloads (for stencil)
  - Neural network-style workloads (for tensor)
- VRP for variable-precision
  - When accuracy issues forces to switch to more accuracy (double instead of single, multi-precision instead of double)
  - For some solver-style workloads
- Technology proof of concepts

ROADMAP

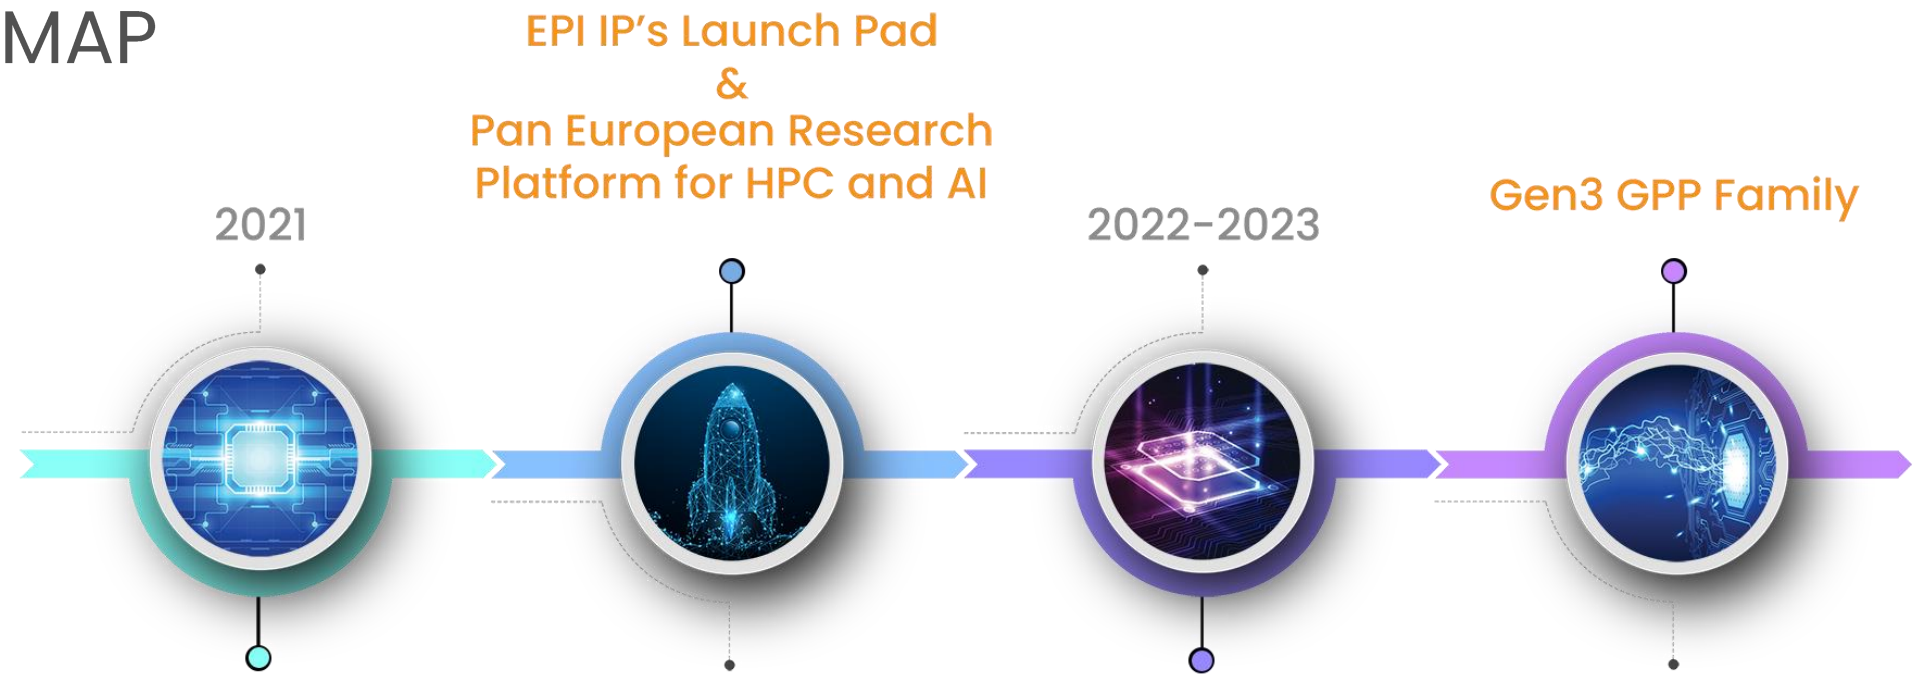

EPI VIEW OF EXASCALE PROCESSORS

- As an ExaScale processor
  - Specialization is the only way toward energy efficiency
  - Bytes/FLOP has to be improved for new HPC workloads
- As a consequence for the processor implementation in EPI:
  - Use/Design specialized computing units (ARM/SVE + EPAC + MPPA + ..)
  - Ease heterogeneous integration of above computing units thanks to a common design platform at SoC level and package level.
  - Put as much as possible large amount of memory close to the processing units (HBM)
  - Adapt the NoC and Die-2-Die BW requirements to the use of HBM with heterogeneous processing units

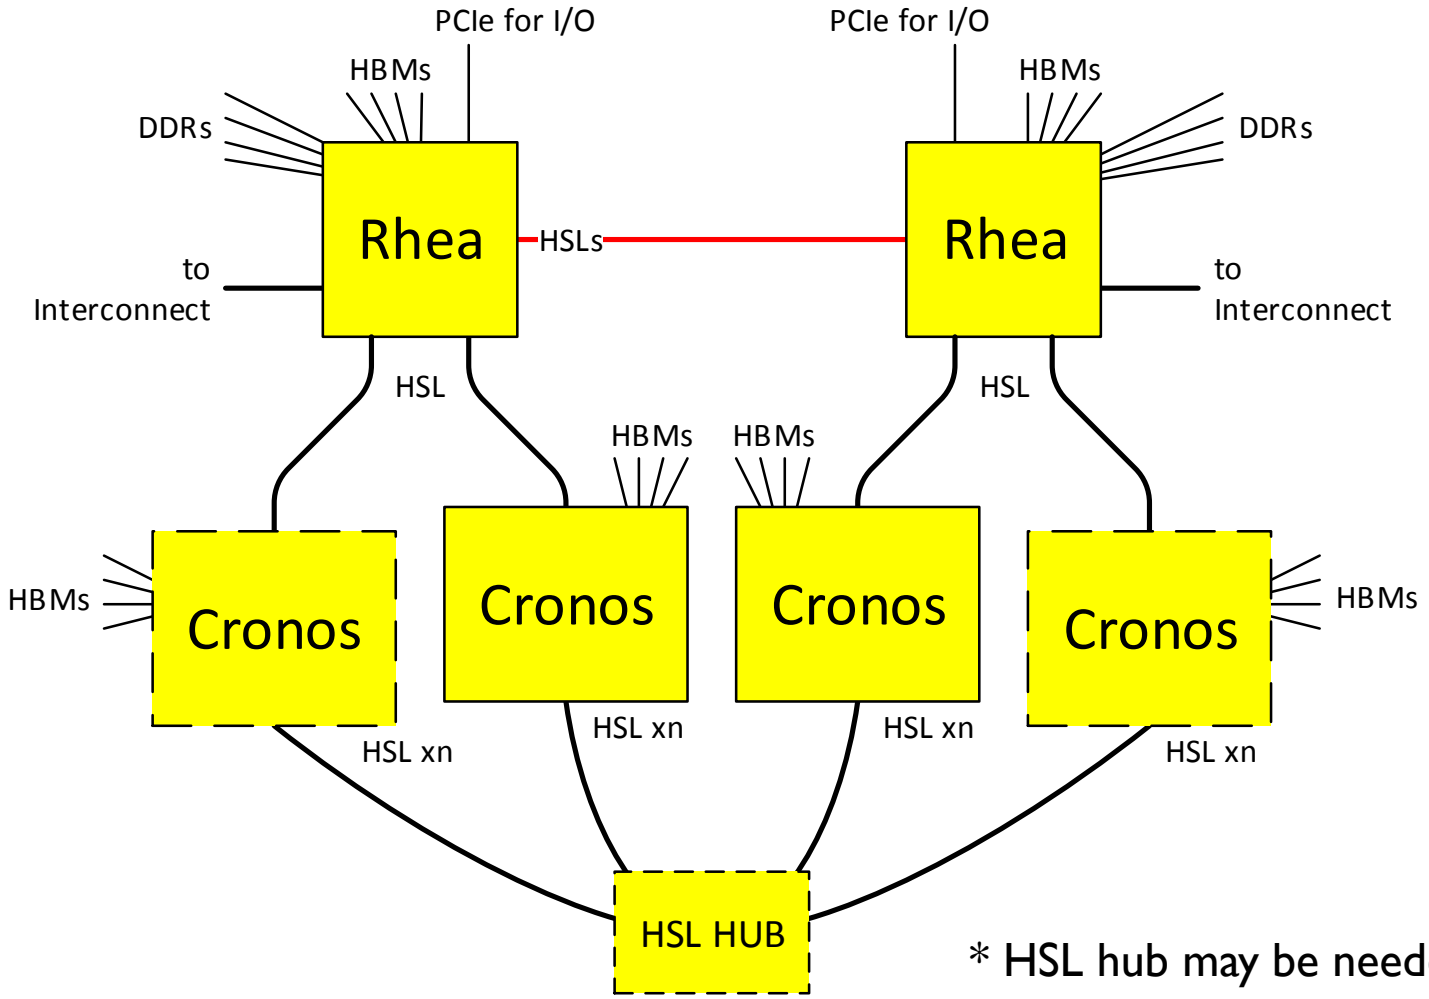

* HSL hub may be needed for more than 2 chips

THANK YOU

ROMAIN.DOLBEAU@ATOS.NET

