



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación



Generalitat de Catalunya
**Departament de Recerca
i Universitats**



GOBIERNO
DE ESPAÑA
MINISTERIO
DE CIENCIA
E INNOVACIÓN



UNIVERSITAT POLITÈCNICA
DE CATALUNYA
BARCELONATECH



UNIÓN EUROPEA
Fondo Europeo de Desarrollo Regional



**Barcelona
Supercomputing
Center**

Centro Nacional de Supercomputación

BSC: AI&HPC, chips, EPI, DARE and Marenosturm6

Prof. Mateo Valero
Director

EPI Forum

9 and 10 October 2024
Barcelona, Spain

Barcelona, October. 9th. 2024



EuroHPC: towards European HPC technologies

#EuroHPC Joint Undertaking

The European High Performance Computing Joint Undertaking (EuroHPC JU) will pool European resources to develop top-of-the range exascale supercomputers for processing big data, based on competitive European technology.

Member countries are Austria, Belgium, Bulgaria, Croatia, Cyprus, Czechia, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Israel, Italy, Latvia, Lithuania, Luxembourg, Malta, Montenegro, Netherlands, North Macedonia, Norway, Poland, Portugal, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden, Türkiye and United Kingdom.



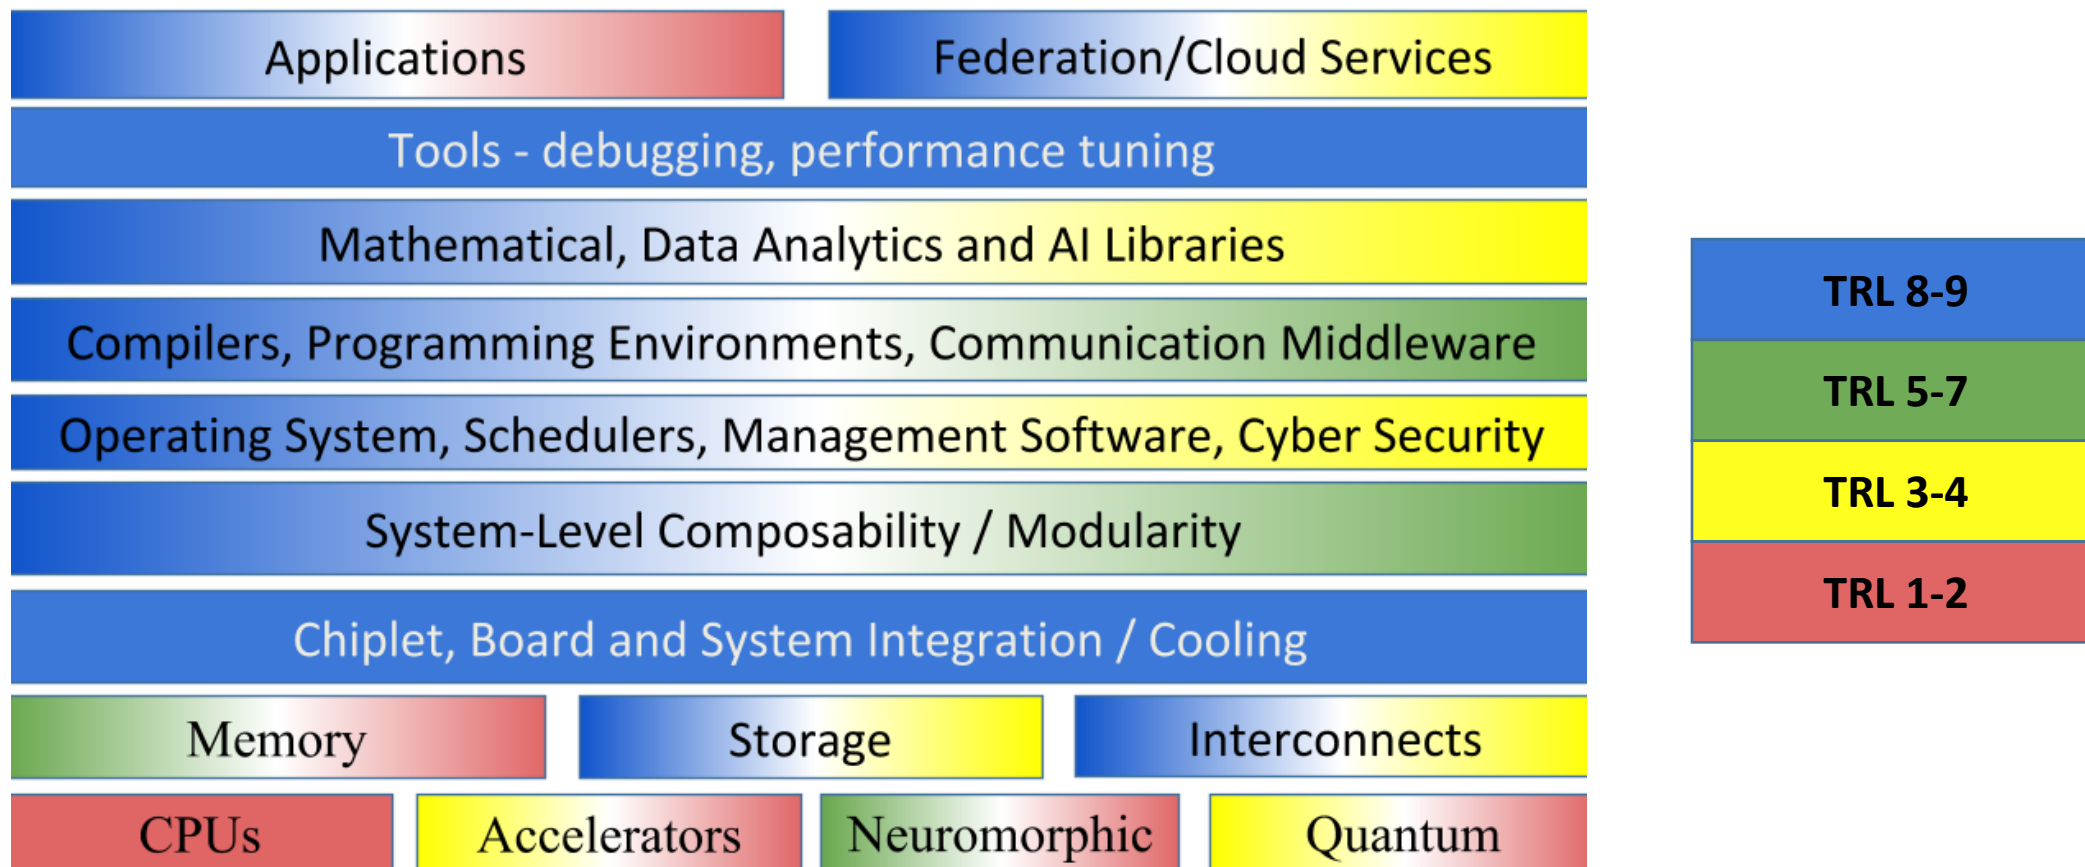
“A new legal and funding structure – the EuroHPC Joint Undertaking – shall acquire, build and deploy across Europe a world-class High-Performance Computing (HPC) infrastructure.

It will also support a research and innovation programme to develop the technologies and machines (hardware) as well as the applications (software) that would run on these supercomputers.”

Top 23, June 2024

Rank	Name	Country	Cores	Accelerators	Rmax [TFlop/s]	Rpeak [TFlop/s]	GFlops/Watts	Computer	Site
1	Frontier	United States	8.699.904	8.138.240	1.206.000	1.714.814	52,93	HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	DOE/SC/Oak Ridge National Laboratory
2	Aurora	United States	9.264.128	8.159.232	1.012.000	1.980.006	26.15	HPE Cray EX - Intel Exascale Compute Blade, Xeon CPU Max 9470 52C 2.4GHz, Intel Data Center GPU Max, Slingshot-11	DOE/SC/Oak Ridge National Laboratory
3	Eagle	United States	2.073.600	1.900.800	561.200	846.835		Microsoft NDv5, Xeon Platinum 8480C 48C 2GHz, NVIDIA H100, NVIDIA Infiniband NDR	Microsoft Azure
4	Fugaku	Japan	7.630.848		442.010	537.212	14,78	Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D	RIKEN Center for Computational Science
5	LUMI	Finland	2.752.704	2.566.080	379.700	531.505	53,43	HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	EuroHPC/CSC
6	Alps	Switzerland	1.305.600	844.800	270.000	353.748	51,98	HPE Cray EX235a, AMD Optimized 3 rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	CSCS
7	Leonardo	Italy	1.824.768	1.714.176	241.200	306.311	32,19	BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, 4x NVIDIA 100	EuroHPC/CINECA
8	MareNostrum 5 ACC	Spain	663.040	591.360	175.300	249.435	42,15	BullSequana XH3000, Xeon Platinum 8460Y+ 40C 2.3GHz, NVIDIA H100 64GB, Infiniband NDR200	EuroHPC/BSC
9	Summit	United States	2.414.592	2.211.840	148.600	200.795	14,72	IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband	DOE/SC/Oak Ridge National Laboratory
10	Eos NVIDIA DGX SuperPOD	United States	485.888	439.296	121.400	188.645		NVIDIA DGX H100, Xeon Platinum 8480C 56C 3.8GHz, NVIDIA H100, Infiniband NDR400	NVIDIA Corporation
11	Venado	United States	481.440	311.520	98.510	130.44	58,29	HPE Cray EX254n, NVIDIA Grace 72C 3.1GHz, NVIDIA GH200 Superchip, Slingshot-11	DOE/NNSA/LANL
12	Sierra	United States	1.572.480	1.382.400	94.640	125.712	12,72	IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband	DOE/NNSA/LLNL
13	Sunway TaihuLight	China	10.649.600		93.015	125.436	6,05	Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway	National Supercomputing Center in Wuxi
14	Perlmutter	United States	888.832	774.144	79.230	112.998	26,90	HPE Cray EX 235n, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-11	DOE/SC/LBNL/NERSC
15	Selene	United States	555.520	483.840	63.460	79.215	23,98	NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband	NVIDIA Corporation
16	Tianhe-2A	China	4.981.760	4.554.752	61.444	100.679	3,32	TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000	National Super Computer Center in Guangzhou
17	CEA-HE	France	389.232	251.856	57.110	112.560	47,32	BullSequana XH3000, Grace Hopper Superchip 72C 3GHz, NVIDIA GH200 Superchip, Quad-Rail BXI v2	Commissariat a l'Energie Atomique (CEA)
18	Explorer-WUS3	United States	445.440	422.400	53.960	86.987		ND96_amsr_MI200_v4, AMD EPYC 7V12 48C 2.45GHz, AMD Instinct MI250X, Infiniband HDR	West US3
19	ISEG	Netherlands	218.880	200.640	46.540	86.792	35,26	Gigabyte G593-SD0, Xeon Platinum 8468 48C 2.1GHz, NVIDIA H100 SXM5 80 GB, Infiniband NDR400	Nebius
20	Adastra	France	319.072	297.440	46.100	61.608	50,03	HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	GENCI-CINES
21	JUWELS Booster Module	Germany	449.280	404.352	44.120	70.980	25,01	Bull Sequana XH2000 , AMD EPYC 7402 24C 2.8GHz, NVIDIA A100, Mellanox HDR InfiniBand/ParTec ParaStation ClusterSuite	Forschungszentrum Juelich (FZJ)
22	MareNostrum 5 GPP	Spain	725.760		40.102	46.371	6,97	ThinkSystem SD650 v3, Xeon Platinum 03H-LC 56C 1.7GHz, Infiniband NDR200	EuroHPC/BSC
23	Shaheen III - CPU	Saudi Arabia	877.824		35.658	39.607	6,73	HPE Cray EX, AMD EPYC 9654 96C 2.4GHz, Slingshot-11	King Abdullah University of Science and Technology

RIAG: European Open System Stack Today





The BSC Vision of the Future of European HPC



European
HPC accelerators
& eventually GPP



MontBlanc @ BSC



Fugaku #1 Top500 @ Riken for \$1B over 7 years (CAPEX)

MareNostrum6, ...

European Supercomputers @ Top500

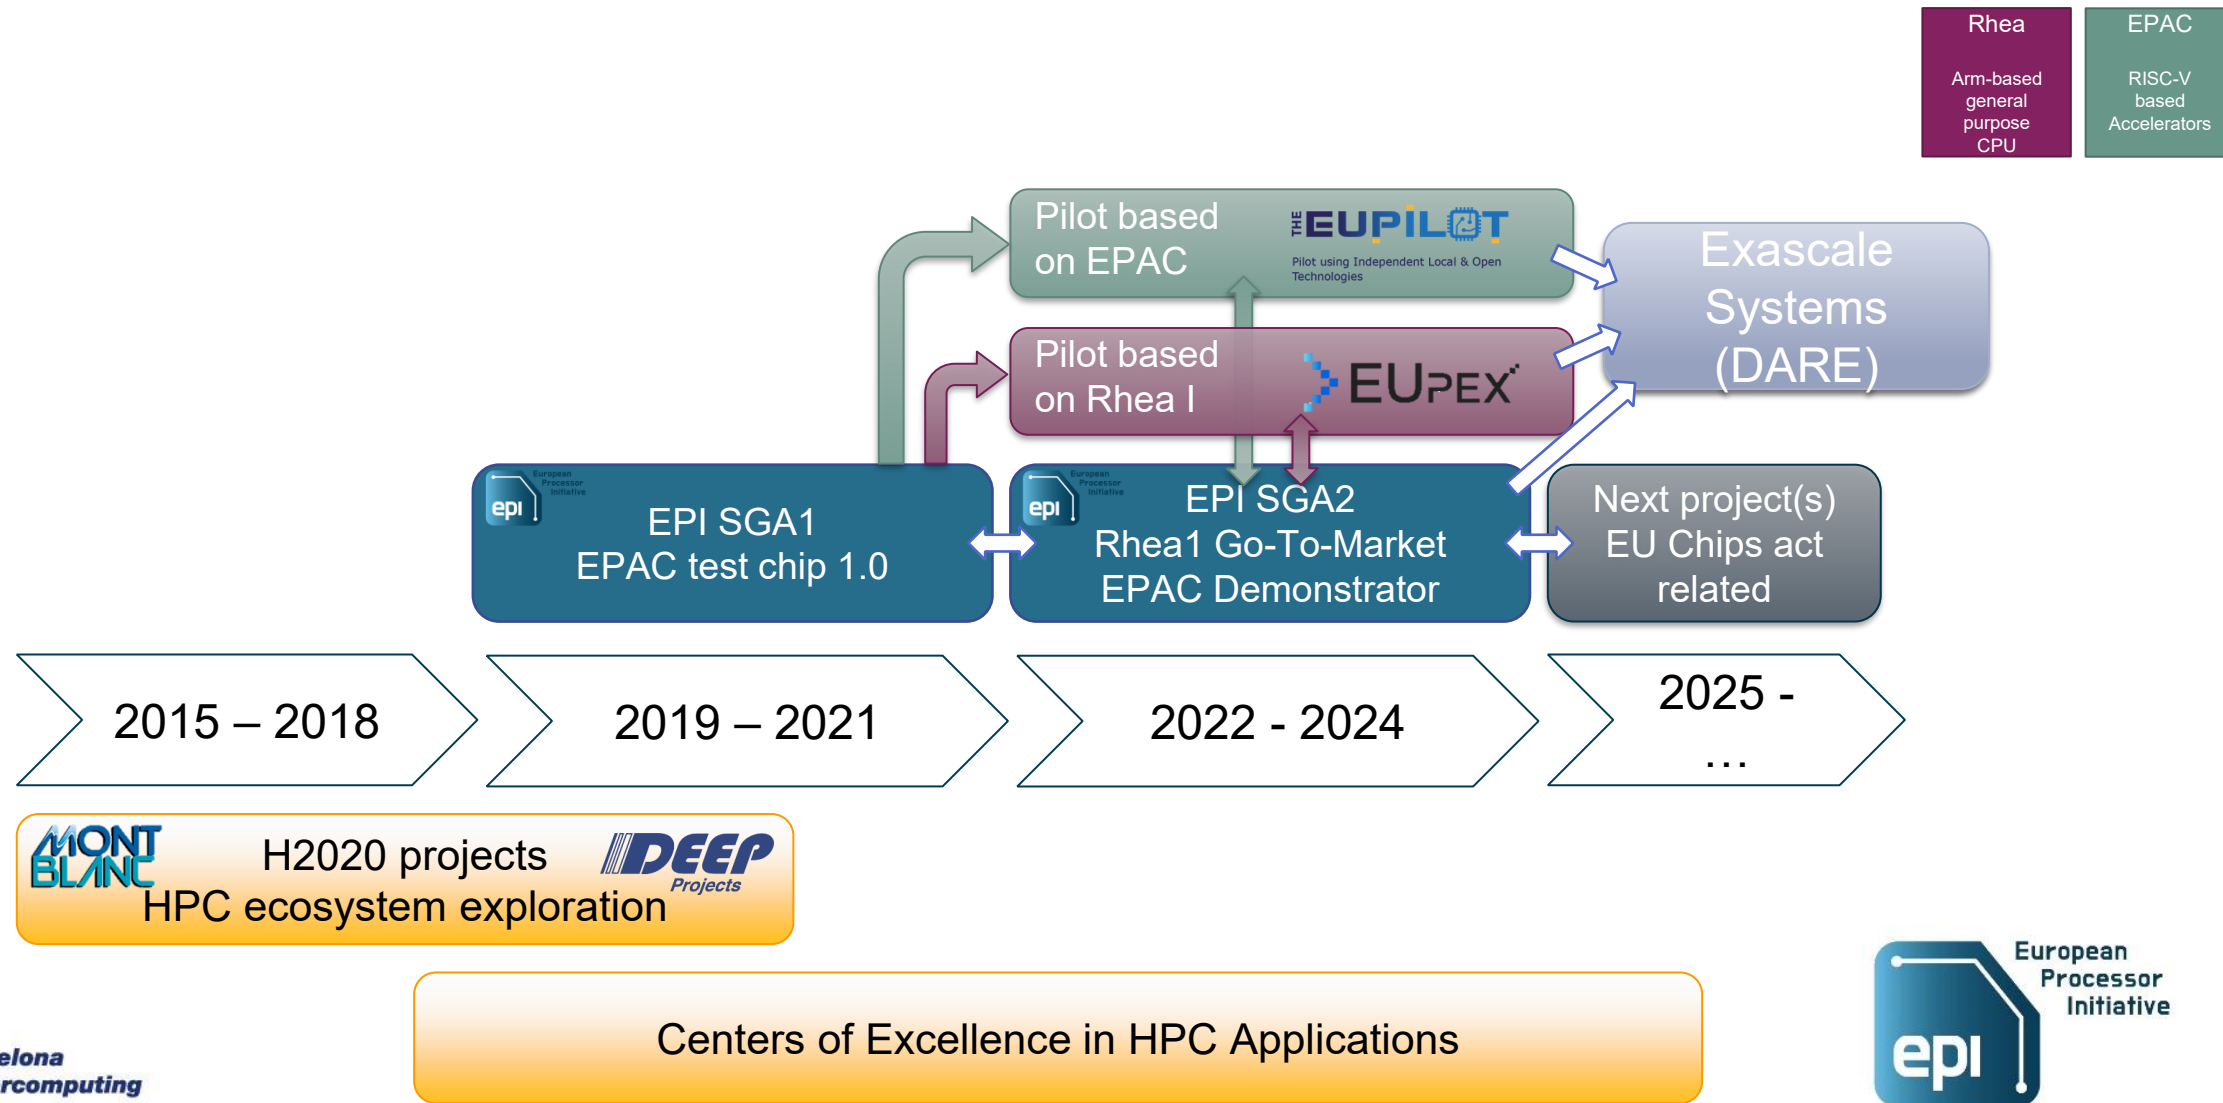
MareNostrum 6



“With MareNostrum5 now successfully installed, our focus has moved to the development of MareNostrum6. Thanks to the unwavering support of the European Commission, the Spanish and Catalan Governments, and the UPC. Plus a strong network of European public and private partners, and the talent of our researchers and collaborators, MareNostrum6 will contain European-designed and owned processors based on RISC-V”

- Mateo Valero


Overall technology roadmap



The European Processor Initiative (EPI)

The European Processor Initiative (EPI) under the SGA1 of the Framework Partnership Agreement (FPA: 800928), to design and implement a roadmap for a new family of low-power European processors for extreme scale computing, high-performance Big-Data and a range of emerging applications.

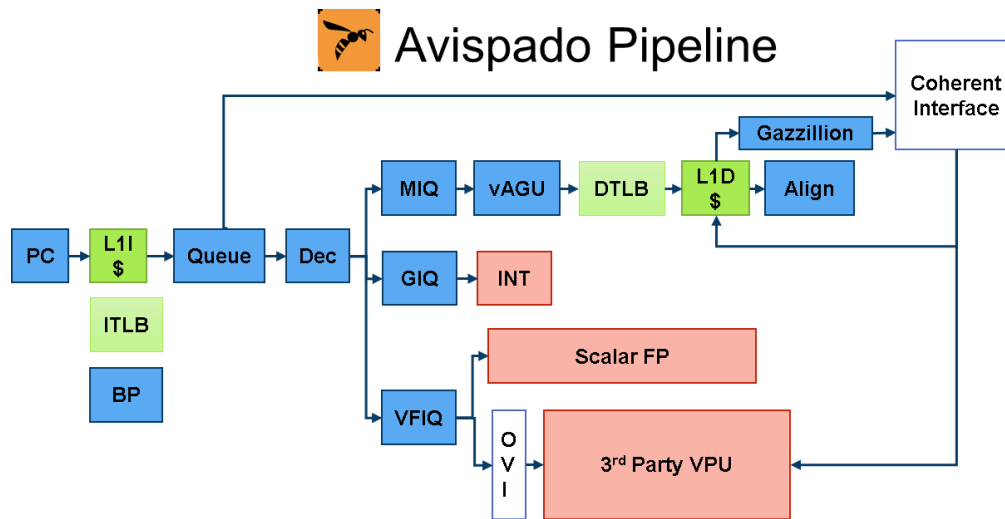
- **History:** Remember MontBlanc? BSC leads the RISC-V HPC accelerator development
- **Consortium (SGA1):** 28 partners from 10 European countries to Coordinate: Bull SAS (France)
- **Budget:** €80M (100% funded)
- **Duration:** 36 months (01/12/2018-31/12/2021)
- **5 Streams** (4 Technical and 1 Management/Exploitation/C&D)



EPI: RISC-V core and VPU

RISC-V core: Avispado

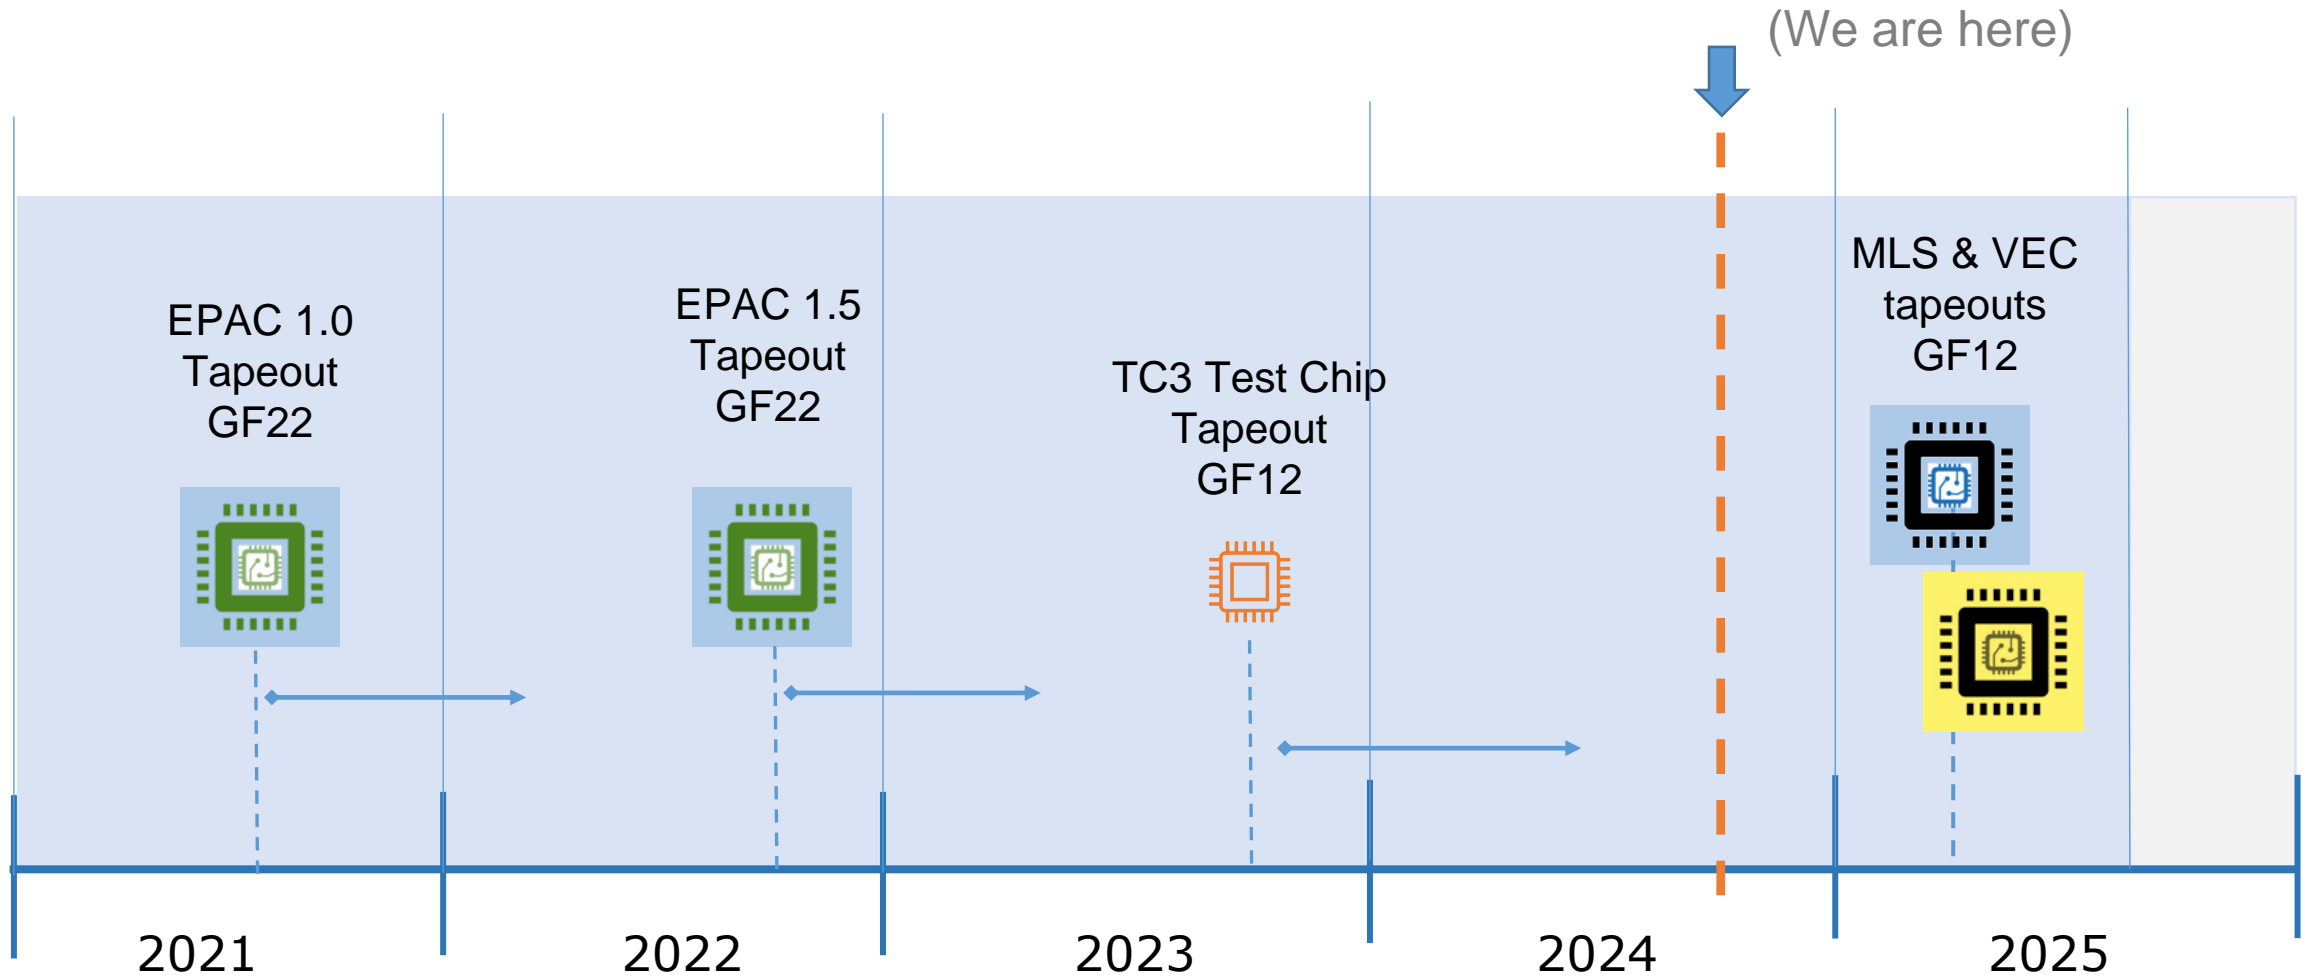
- 2-way in-order core
- Full HW-support for unaligned accesses
- Cache: L1I\$ = 16KB, L1D\$ = 32KB



BSC-VPU: Vitruvius

- Long vectors: 256 DP elements
 - #Functional Units (FUs) \ll Vector Length (VL)
 - 1 vector instruction can take several cycles
- 8 Lanes per core
 - FMA/lane: 2 DP Flop/cycle
- Renaming: 40 physical registers, out-of-order execution
- Chaining from register file
- Decoupled vector architecture
- Area-efficient intra-lane interconnect

EPI / EUPILOT tapeout timeline



EUPILOT chips peak performance projections

VEC chip

- 8-tile area budget of $\sim 49\text{mm}^2$
- Each VPU has 8x FPU
- 2x FLOPs for FMA
- 64-bit ops
- Target at 1.5GHz
- $8 \times 8 \times 2 \times 1.5 = 192$ GFLOPs/chip
- w/RISC-V cores: 216 GFLOPs/chip

~ 1.74 TFLOPs/EAS (64-bit)

MLS chip

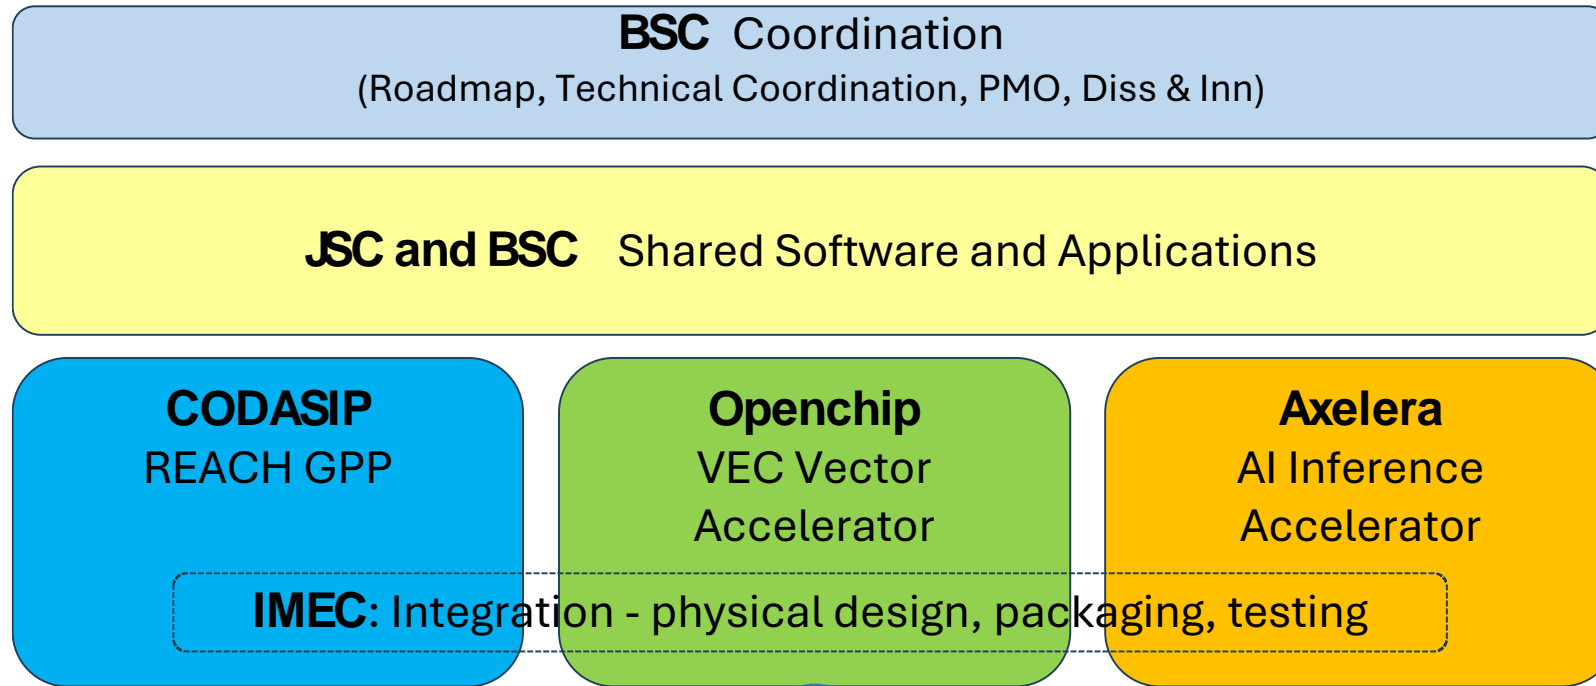
- Mid-area budget of $\sim 35\text{mm}^2$
- 2 Groups, 6 clusters
- Of (8+1) cores = total 108x FPU
- 8-bit FP ops
- Target at 1GHz
- ~ 768 GFLOPs/chip

~ 6.1 TFLOPs/EAS (8-bit)

Openchip

A **fabless** semiconductor company building **RISC-V** based **High Precision**, High Performance **Accelerators** targeting **HPC and adjacent Enterprise** AI/ML/DL real world workloads with dense and **sparse access patterns**.

DARE FPA and SGA1 towards MNx chip (*)



integrating/scaling BSC's VPU into
new chip with focus on AI


+



contributing to **pathfinding** new design for SGA2,
including VPU with IME, memory hierarchy to
support irregular applications, ...

Other partners: EXTOLL, Chalmers, FORTH, UNIZG, ICSC

NVIDIA Blackwell Dual Die



Blackwell GPU

DGEMM	New implementation	>80 TFlops
FP64 FMA		40 TFlops
SGEMM	New implementation	>150 TFlops
FP32 FMA		80 TFlops
HBM Bandwidth	HBM3e	Up to 8 TB/sec
HBM Capacity	HBM3e	Up to 192 GB
FP16 TC*	AI Inference/Training	5 PFlops
FP8 TC*	AI Inference/Training	10 PFlops
FP4 TC*	AI Inference	20 PFlops

*Using sparsity

Simulation capabilities 1.3x – 2x higher than H100



AI SUPERCHIP
208B Transistors



2nd GEN TRANSFORMER ENGINE
FP4/FP6 Tensor Core



5th GENERATION NVLINK
Scales to 576 GPUs



RAS ENGINE
100% In-System
Self-Test



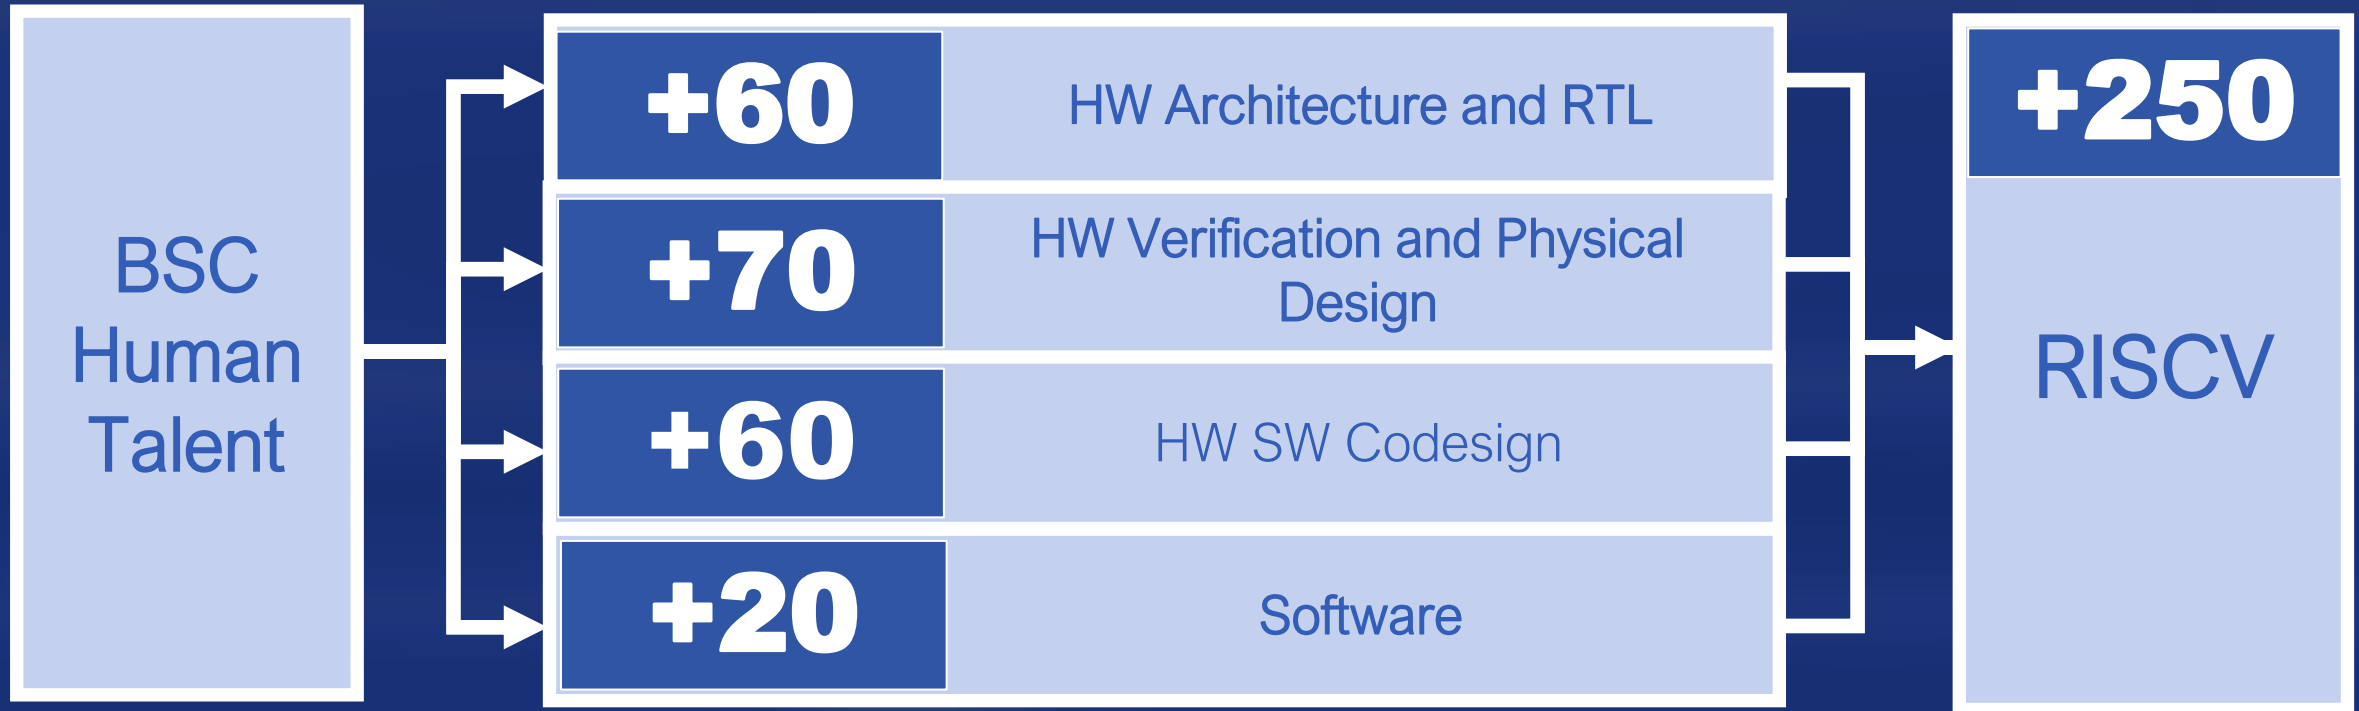
SECURE AI
Full Performance
Encryption & TEE





DECOMPRESSION ENGINE
800 GB/s

RISCV Hardware and Software Researchers at BSC

The four key technical inputs to RISCV developments



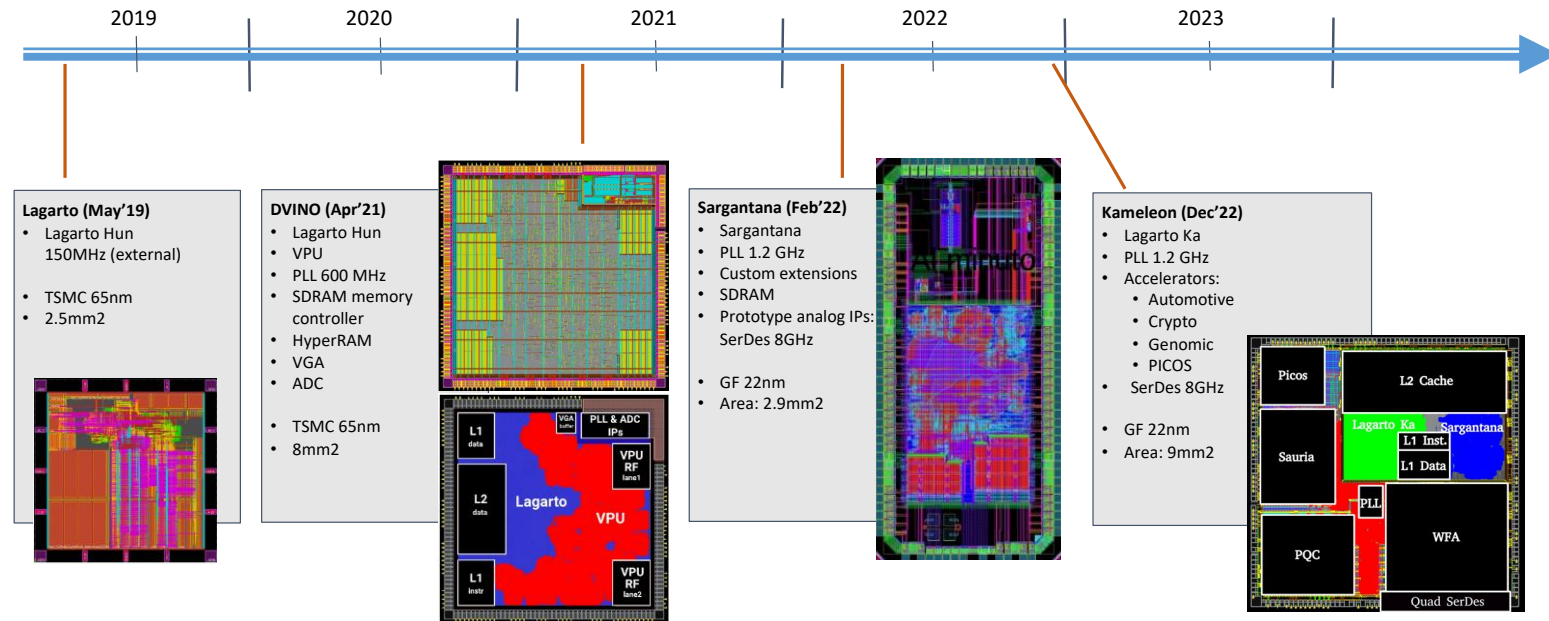
RISC-V@BSC: Research



RISC-V cores with vector/matrix extensions

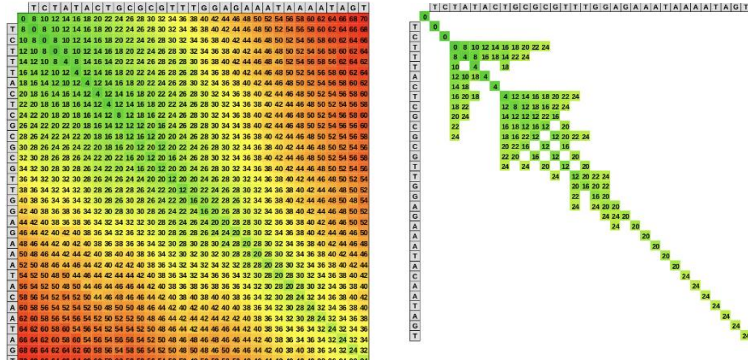
Design of RISC-V cores

- **In-order** Lagarto (5-stage) and Sargantana (7-stage)
- **Out-of-order** Lagarto Ka (11-stage) and Lagarto Ox (12-stage)
- **VPU supporting RVV vector extensions** (long vectors with 256 DP elements)
- **Matrix extensions** (contribution to RISC-V Integrated Matrix Extension IME Task Group)
- Design of functional units supporting **data formats and optimizations for Artificial Intelligence (AI)**
- Cores for **near-memory** and **CXL acceleration**



Acceleration in domain-specific architectures

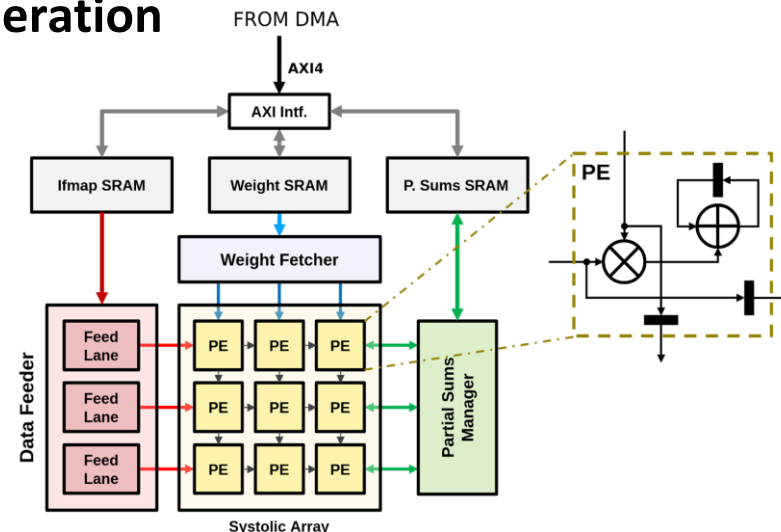
High performance algorithms and HW accelerators for genomic analytics



Other accelerators:



- Post Quantum cryptography
- Database management systems (DBMS)
- Graph processing
- ...

SAURIA: Systolic Array tensor Unit for aRtificial Intelligence Acceleration



Hardware acceleration of parallel runtimes


Managing tasks and dependences generated from the execution of parallel programs



Contributing to EU chips


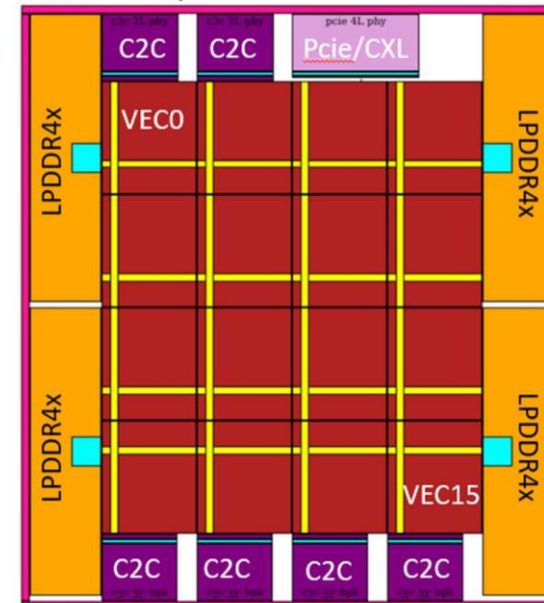


Vitruvius VPU in the EPI EPAC
(EuropeanProcessor Accelerator)


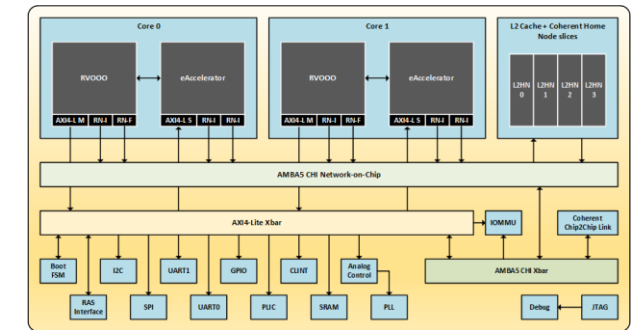


New VPU and OVI in EUPilot chip with
8 cores and multichip support

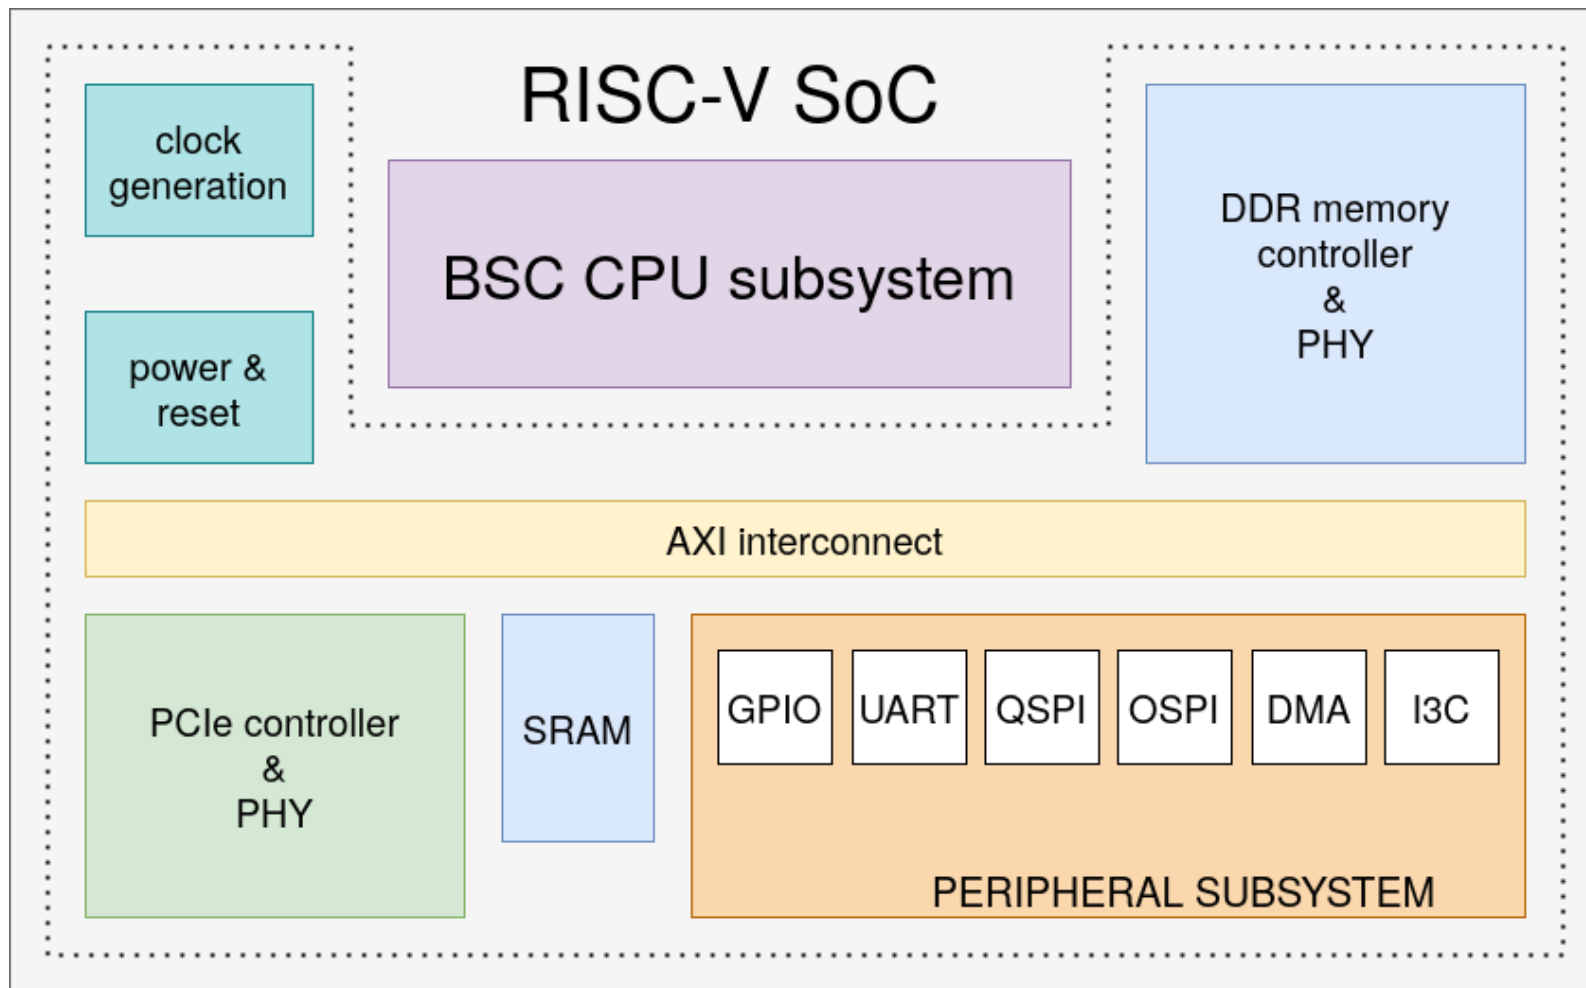
Area = ~45 sqmm



VPU in eProcessor eAccelerator



BZL Test Chip 1 (TC1) Cinco Ranch



BARCELONA ZETTASCALE LAB

Contributing to RISC-V ecosystem



Degree in Informatics Engineering
Degree in Artificial Intelligence
Master in Informatics Engineering




Spanish Open Hardware Alliance

BSC Advanced Training Center (BETC)



European Laboratory for Open Computer Architecture

<https://github.com/bsc-loc>



- Co-chair of the Vector C Intrinsic technical group
- Vice-chair High Performance Computing (SIG-HPC)
- Vice-chair SIG-Safety
- Integrated Matrix Extension IME task group
- RVV vector extensions task group

**Compute architectures
are changing forever...**




**... there will be pervasive use
of AI across all applications**

18 AI chip makers by category

Vendor	Category	Selected AI chip*	Date of Announcement
NVIDIA	Leading producer	Blackwell	March 2024
AMD	Leading producer	MI350	June 2024
Intel	Leading producer	Gaudi 3	April 2024
IBM	Public cloud & chip producer	NorthPole	October 2023
Alphabet	Public cloud & chip producer	Trillium	May 2024
AWS	Public cloud & chip producer	Trainum2	November 2023
Alibaba	Public cloud & chip producer	ACCEL**	November 2023
Apple	Upcoming producer	M4	May 2024
Meta	Upcoming producer	Artemis	April 2024
Microsoft Azure	Upcoming producer	Maia 100	November 2023
SambaNova Systems	AI startup	SN40L	September 2023
Cerebras Systems	AI startup	WFE-3	March 2024
Groq	AI startup	LPU Inference Engine	November 2023
Preferred Networks	AI startup	MN-Core 2	August 2024
Rebellions	Upcoming producer	Rebel	January 2024
Graphcore	Other producers	Bow IPU	March 2022
Sifive	RISC V start-up	XM-series	July 2024
Tenstorrent	RISC V startup	Blackhole	May 2023
Ventana	RISC V startup	Veyron 2	January 2023

Veyron V2: Momentum to Mainstream with Complete Platform


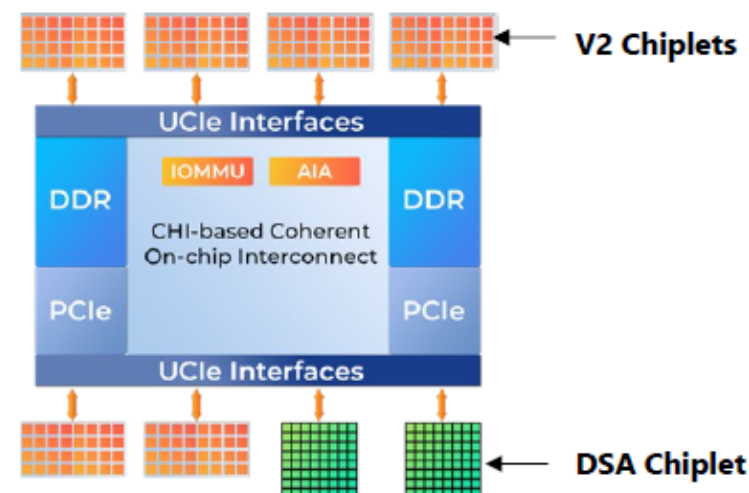
Highest Performance RISC-V Processor for Data Center, Automotive, Gen AI and Client



Highlights

- +40% performance, 32 cores per cluster, 4nm
- UCle chiplet
- RISC-V Vector Extension support
- Ventana AI Matrix Extensions
- Server-class IOMMU
- RISE support
- Domain Specific Acceleration

Available as UCle-Compatible Chiplets or IP



AI Silicon Roadmap



tenstorrent

2021

High Perf AI ASIC

Grayskull

AI Processor



- 120 Tensix Cores
- 12nm
- 276 TOPS (FP8)
- 100 GB/s LPDDR4
- Gen4x16

GEN 1

2022

Scalability

Wormhole

Networked AI Processor



- 80 Tensix+ Cores
- 12nm
- 328 TOPS (FP8)
- 336 GB/s GDDR6
- Gen4x16
- 16x100 Gbps Ethernet

GEN 1

2023

Heterogeny

Blackhole

Standalone AI Computer



- 140 Tensix++ Cores
- 6nm
- 745 TOPS (FP8)
- 512 GB/s GDDR6
- Gen5x16
- 10x400 Gbps Ethernet
- 16 RISC-V CPU cores


GEN 2

Blackhole - A Standalone AI Computer

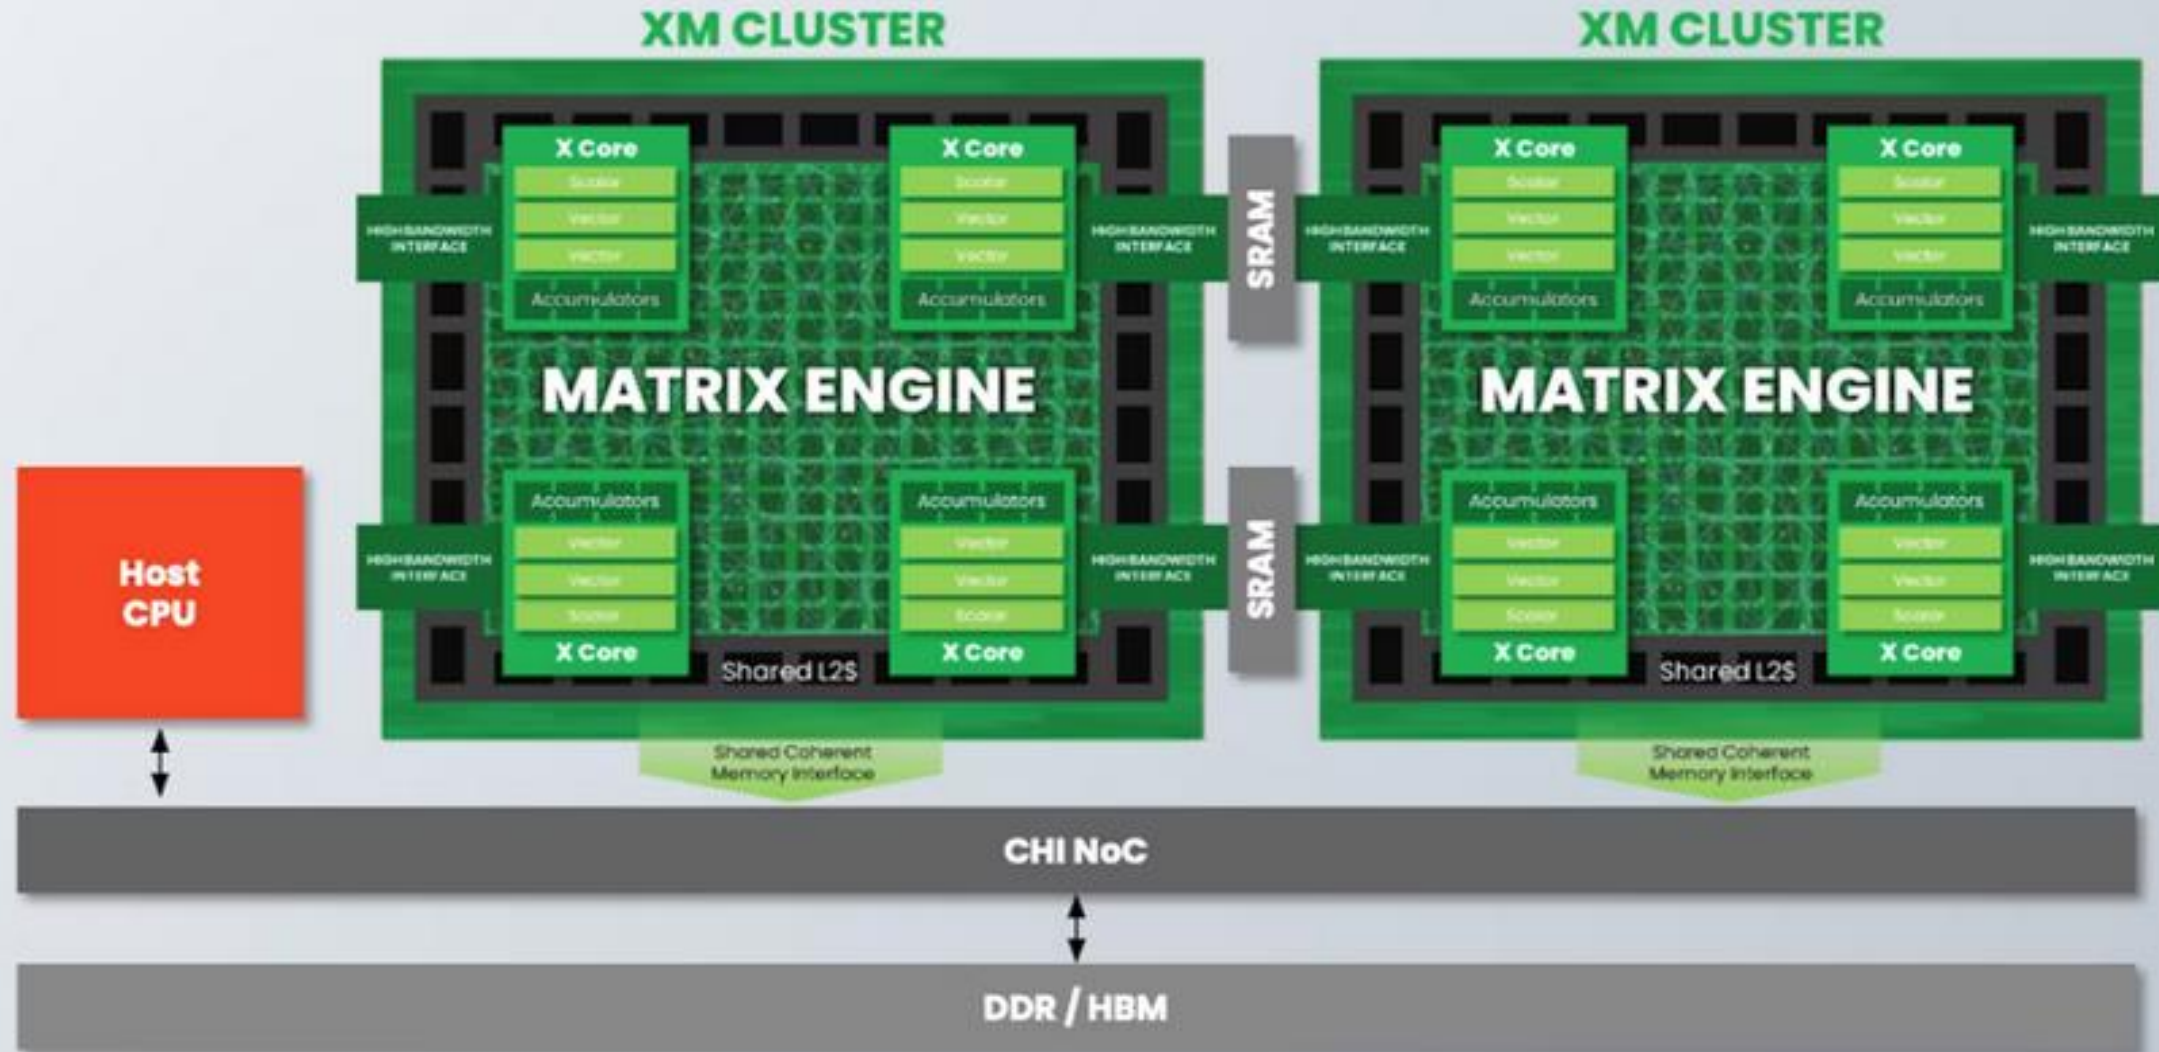


tenstorrent

Feature	Spec
Tensix	745 TFLOPs (8-bit) 372 TFLOPs (16-bit)
SRAM	241 MBs
Ethernet	10x 400 Gbps
DRAM	512 GB/s BW 32 GBs capacity
Baby RISC-Vs	752
Big RISC-Vs	16
PCIe	Gen5x16, 64 GB/s
NoC	2 NOCs 2D Torus 256 B per core



XM enables high scalability



XM specifications



- SiFive Matrix Engine
 - Fat Outer Product design
 - Tightly integrated with 4 X-Cores
 - Deep fusion with vector units
- 4 X-Cores per cluster
 - Each with dual vector units
 - Executes all other layers e.g. activation functions
 - New exponential acceleration instructions
- New matrix instructions
 - Fetched by scalar unit
 - Source data comes from vector registers
 - Destination to each matrix accumulator
- 1 Cluster = 16 TOPS (INT8), 8 TFLOPS (BF16) per GHz
- 1TB/s sustained bandwidth per XM Series cluster
- XM clusters connect to memory in 2 ways:
 - CHI port for coherent memory access
 - High bandwidth port connected to SRAM for model data
- Host CPU can be RISC-V, x86 or Arm (or not present)
- System can scale across multiple dies using CHI

The road to the highest AI compute in a single chip

- Build each GPU to reticle limit as intra-GPU communication provides:
 - Highest communication density
 - Lowest latency
 - Optimal energy efficiency




Volta

>21 billion transistors
815mm²
TSMC 12nm FFN




Ampere

>54 billion transistors
826 mm²
TSMC N7



Hopper


>80 billion transistors
814 mm²
TSMC 4N



Blackwell

>208 billion transistors
>1600 mm²
TSMC 4NP

GB200 NVL 72 & NVL 36



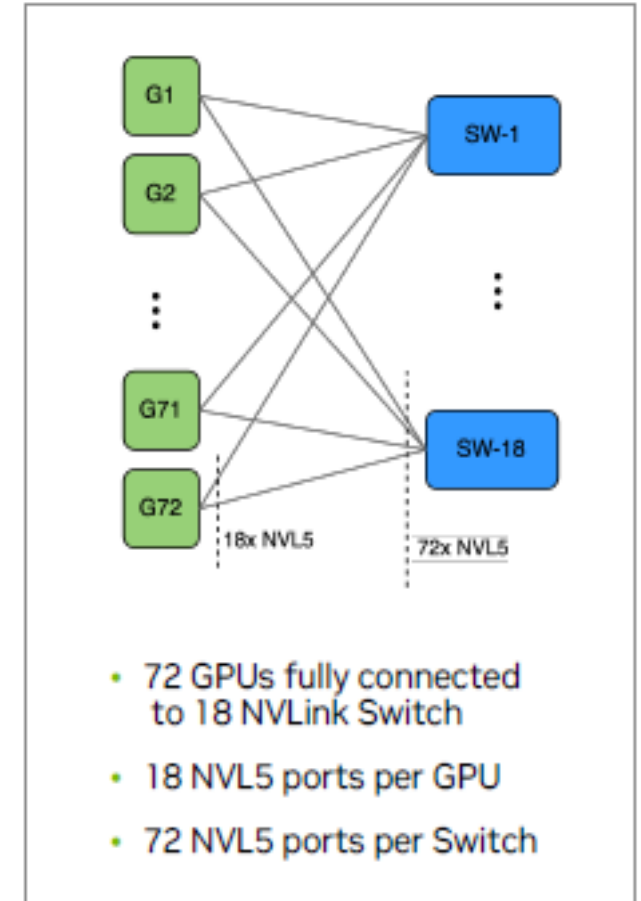
GB200 NVL72

36 Grace CPUs
72 Blackwell GPUs
Fully connected NVLink Switch rack

Training	720 PFLOPs
Inference	1,440 PFLOPs
NVL Model Size	27 Trillion params
Multi-Node Bandwidth	130 TB/s
Multi-Node All-Reduce	260 TB/s

GB200 NVL36

18 Grace CPUs
36 Blackwell GPUs
Fully connected NVLink Switch rack



Von der Leyen gives nod to €100 billion 'CERN for AI' proposal

Re-elected European Commission President Ursula von der Leyen's political guidelines addressed calls for huge artificial intelligence (AI) research investments under the "CERN for AI" banner, but proponents and critics say the plan is lacking crucial details.



“ With the launch of AI Factories, we are leveraging one of Europe’s biggest assets: our world-class supercomputers. AI Factories will serve as a one-stop shop for Europe's AI start-ups, helping them develop the most advanced AI models and industrial applications. This will make Europe the best place in the world for trustworthy AI. ”

Thierry Breton, Commissioner for Internal Market

EU boosts European AI developers with the AI Factories call for proposals

PAGE CONTENTS

[Top](#)

[Quote\(s\)](#)

[Related topics](#)

[Print friendly pdf](#)

[Contacts for media](#)

Today, the Commission has launched a call for setting up AI Factories to boost European leadership in trustworthy artificial intelligence (AI). AI Factories will be created around the EU's world-class network of [European High-Performance Computing \(HPC\)](#) supercomputers and will be available to a range of European users, such as startups, industry and researchers.

President of the Commission, Ursula **von der Leyen** said: “*Europe is already leading the way with the EU AI Act, ensuring AI is safer and more trustworthy. Earlier this year, we fulfilled our promise by opening our high-performance computers to European AI start-ups. Now, Europe must also become a global leader in AI innovation. AI Factories will help secure our position at the forefront of this transformative technology.*”

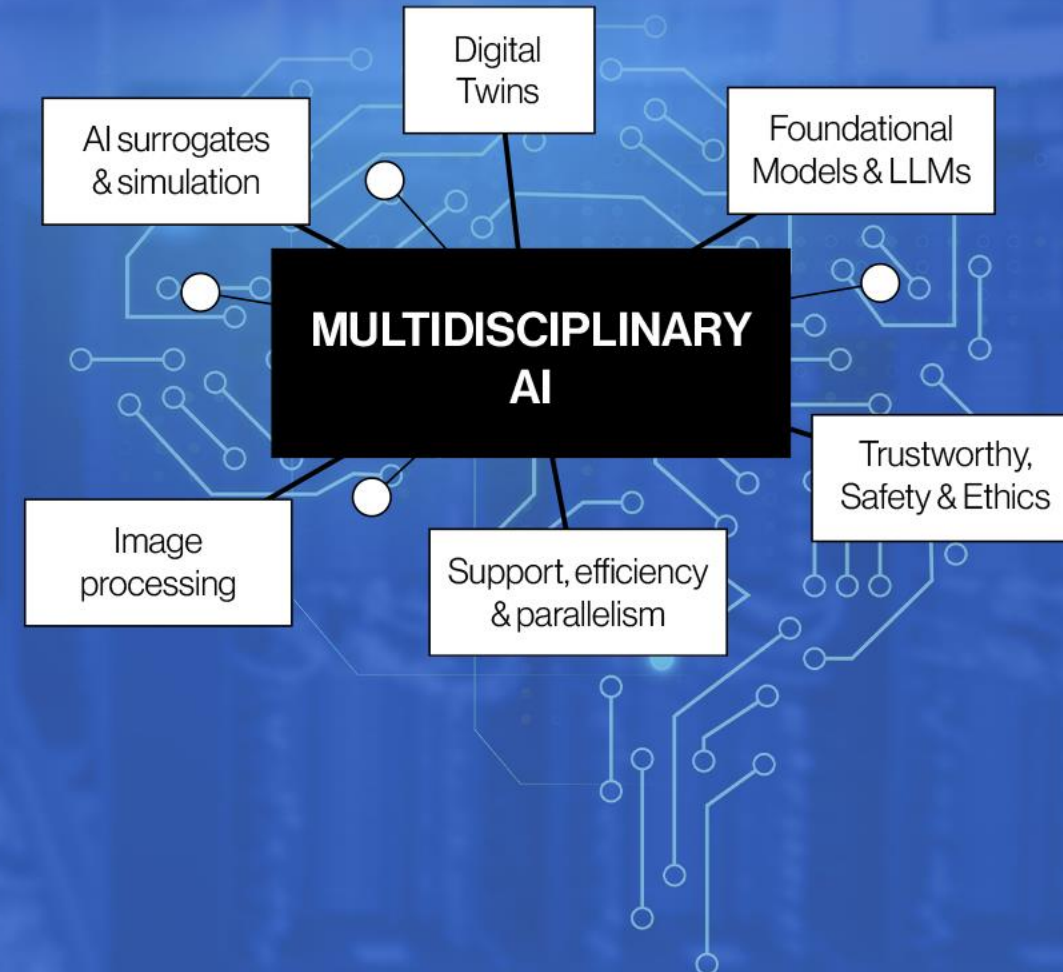


“ Europe is already leading the way with the EU AI Act, ensuring AI is safer and more trustworthy. Earlier this year, we fulfilled a promise from my State of the Union address by opening our high-performance computers to European AI start-ups. Now, our next goal is clear: Europe must become a global leader in AI innovation. AI Factories will help secure position at the forefront of this transformative technology, outlined in my political guidelines for the next Commission.

ula von der Leyen, President of the European Commission

BSC capabilities in Artificial Intelligence (AI)

+80	AI researchers	
+40	AI open Models	(+1.7M Downloads)
+30	Datasets released	
+100	AI Papers last 3 years	(+1,500 Citations)
32	European AI-related competitive projects	

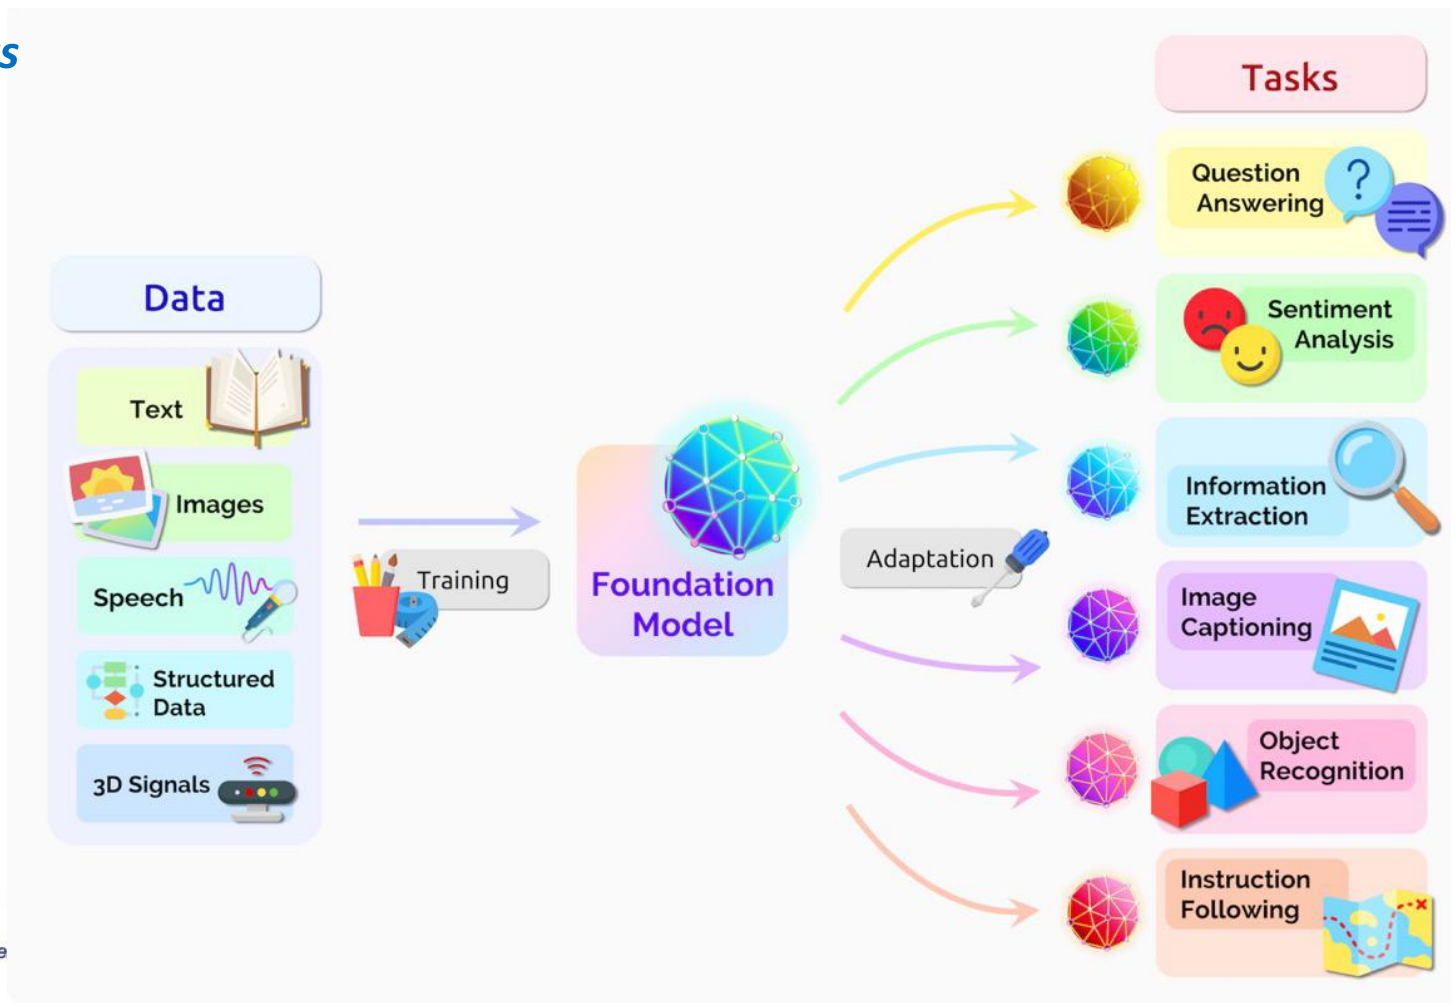


Trillion Parameters Consortium (TPC)

Leveraging Community Efforts to Build Foundation Models for Science

Scientific & Engineering Datasets

Mathematics
Biology
Materials
Chemistry
Particle Physics
Nuclear Physics
Computer Science
Climate
Medicine
Cosmology
Fusion Energy
Accelerators
Reactors



Exemplar Tasks

Scientific Discovery

Digital Twins

Inverse Design

Code Optimization

Accelerated Simulations

Autonomous Experiments

Secure Data Infrastructure

Co-Design

Example Focus Areas (1)

Model Architecture and Performance Evaluation

Architectures for LLMs are continuously evolving. Variants of transformers, their mixture-of-experts-based extensions, and state-space models are being proposed on a weekly basis. Frameworks such as Megatron-LM and DeepSpeed, and their various forks each cover a different subset of architectures, distributed parallelism, and compute/memory/IO optimizations.

Determining the optimal architecture for training a trillion parameter model on scientific data, and the best framework to accomplish this on today's Exascale platforms is critical for the creation of a new class of AI models for a broad range of scientific applications.

Prasanna Balaprakash (ORNL)
Rio Yokota (TiTech)
Irina Rish (UdeM/MILA)

Data, Training Workflows, LLM strategies

In the era of exponential data growth, this session addresses critical challenges and innovative strategies in harnessing vast datasets needed for training large language models (LLMs) in domains such as chemistry/materials, biology, climate science, and high energy physics. This session will discuss the complexities of developing a data-focused infrastructure, including streamlined data curation pipelines, the refinement of data curation practices, and the application of pre-training methodologies. It will explore how the incorporation of domain-specific knowledge into these processes can significantly enhance the performance and applicability of LLMs in scientific research, emphasizing the critical role of targeted data selection and preparation in advancing AI capabilities.

Neeraj Kumar (PNNL)
Ian Foster (Argonne)

Skills, Safety, and Trust Evaluation

One of the main thrusts behind the rapid evolution of LLMs is the availability of benchmarks that assess the skills and trustworthiness of LLMs. Not only do they enable a rigorous evaluation of LLMs skills and trustworthiness from accepted metrics, but they also generate competition between LLM developers. While several frameworks/benchmarks have emerged as de facto standards for the evaluation of general-purpose LLMs (Eleuther AI Harness and HELM for skills, DecodingTrust for trustworthiness), only very few of them specifically are related to science.

Franck Cappello (Argonne)
Bo Li (UChicago)
Javier Aula-Blasco (BSC)



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación

THANK YOU!

mateo.valero@bsc.es