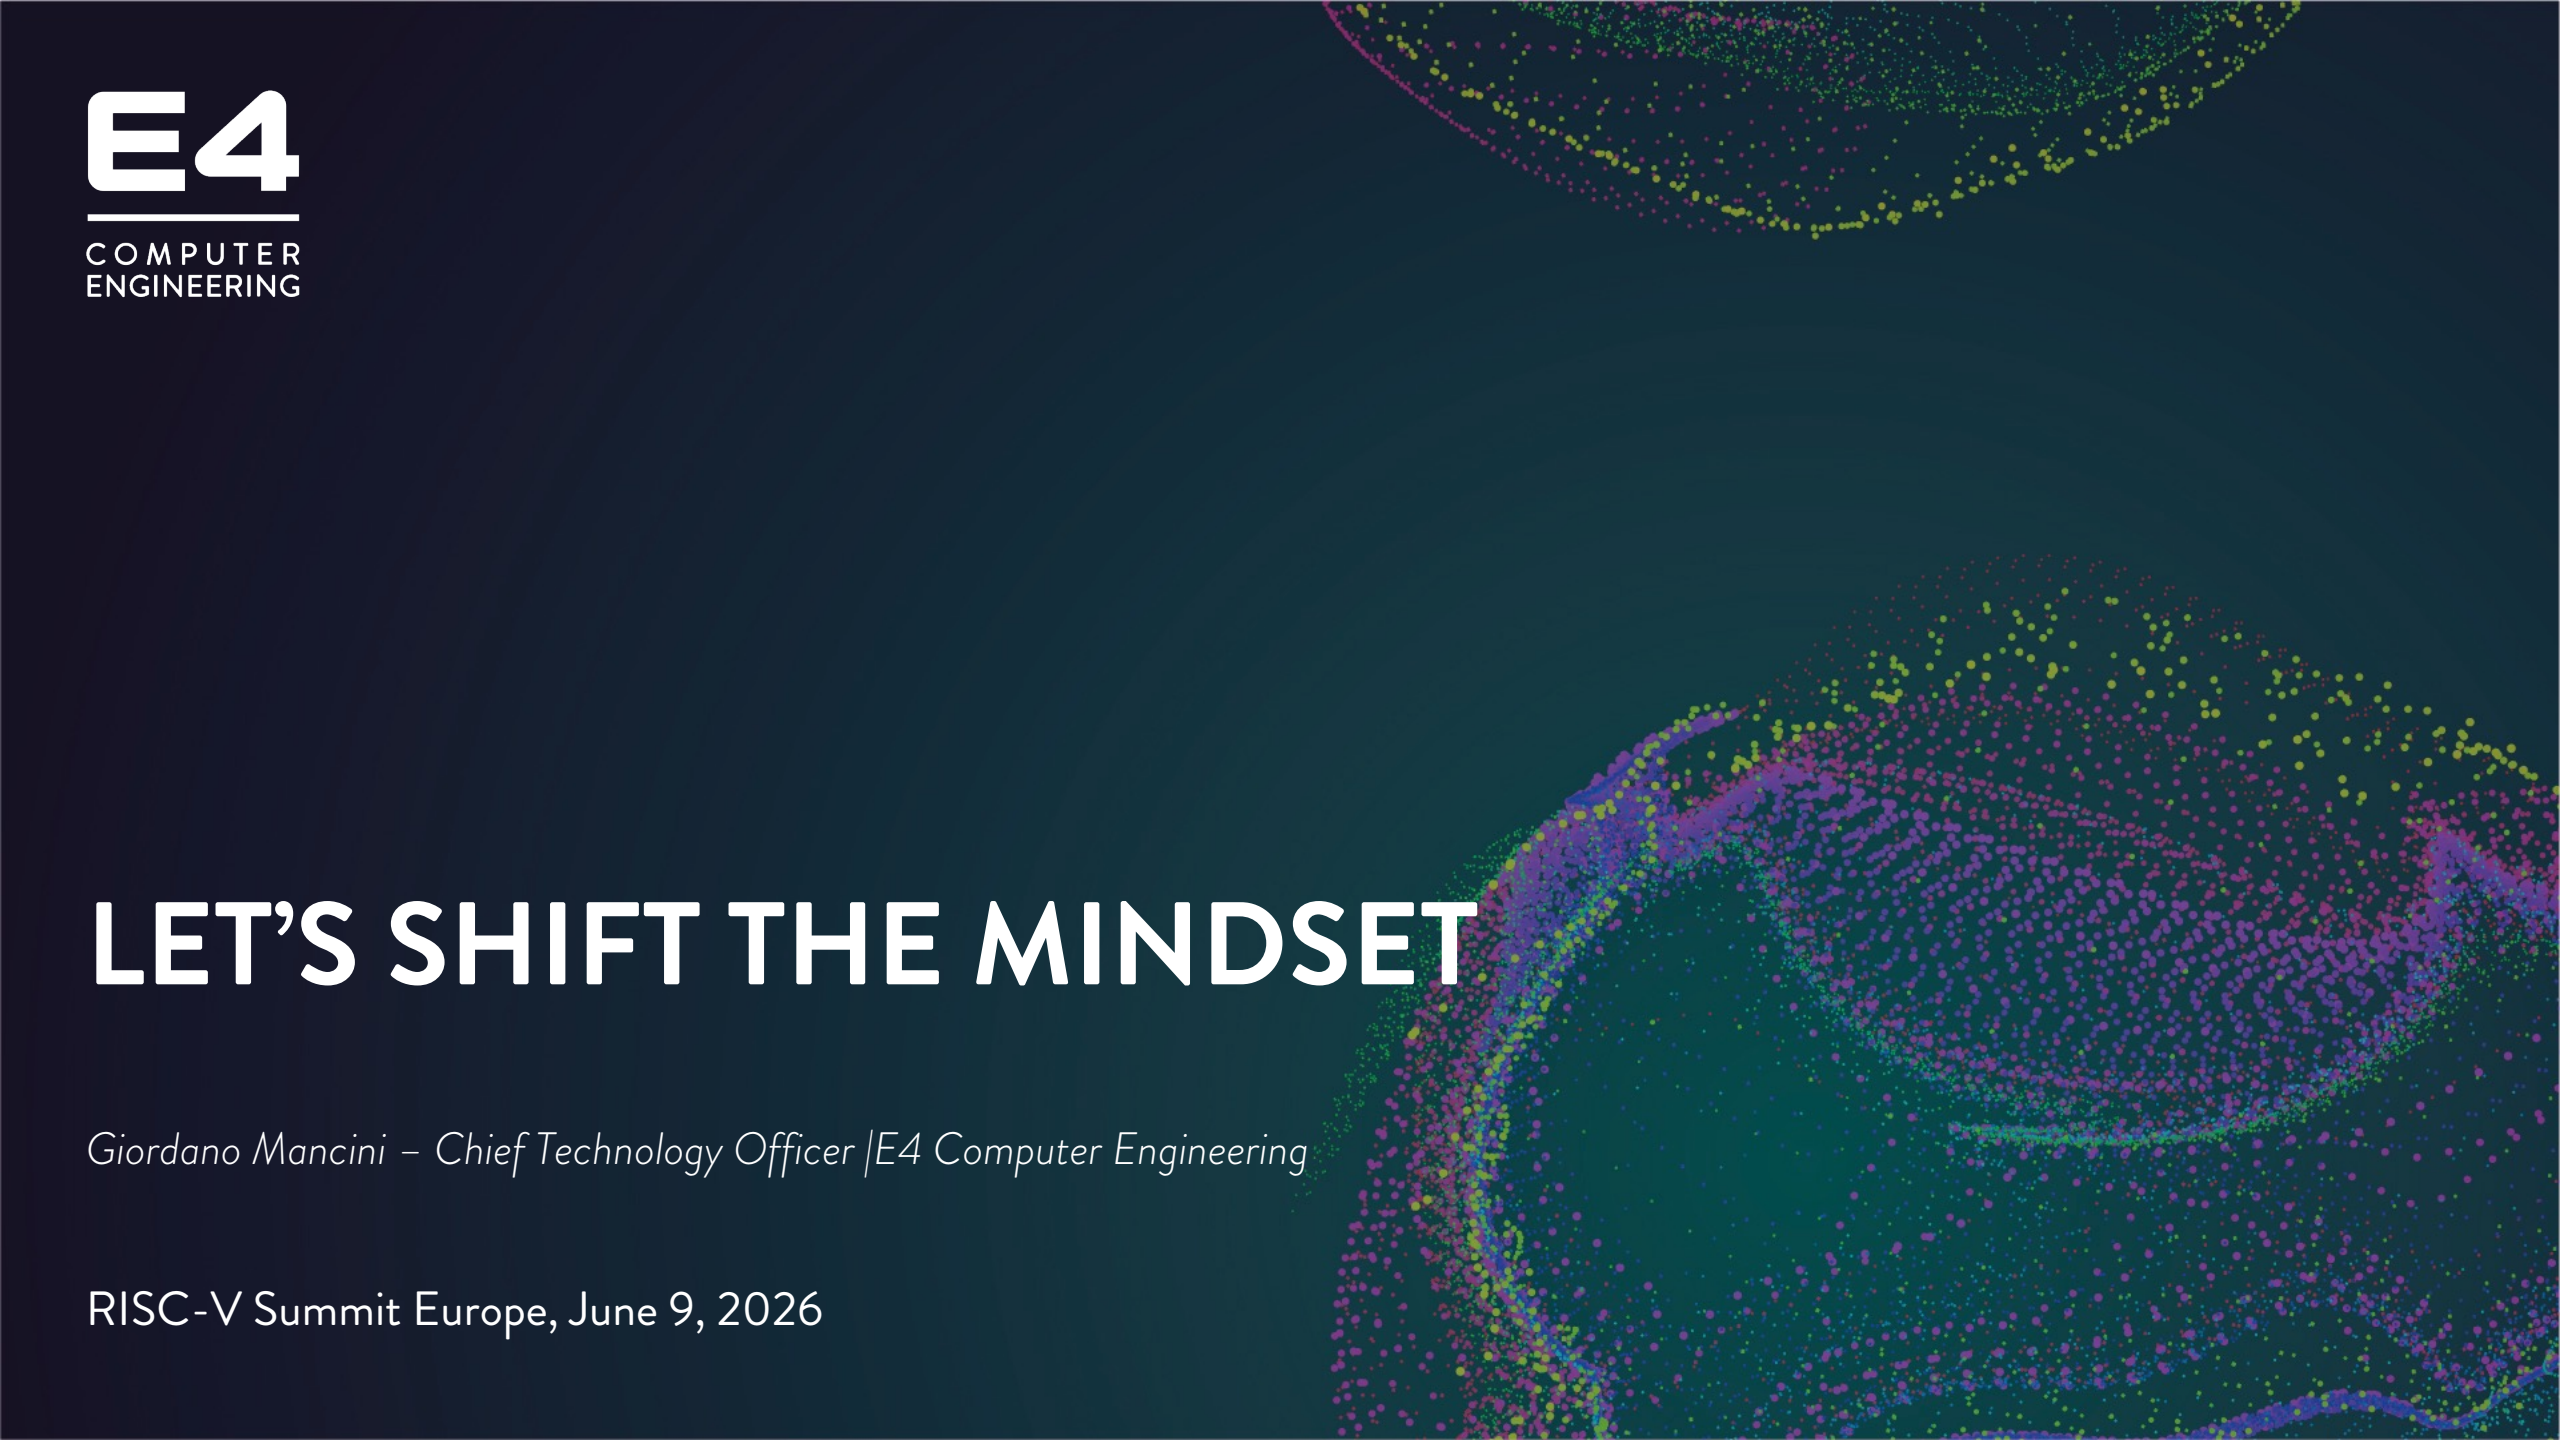

LET'S SHIFT THE MINDSET

Giordano Mancini – Chief Technology Officer |E4 Computer Engineering

RISC-V Summit Europe, June 9, 2026

LET'S SHIFT THE MINDSET –WHY THIS TALK NOW

E4 COMPUTER ENGINEERING

THINK ECOSYSTEM, NOT JUST ISA

Transition community language & approach toward a full ecosystem vision

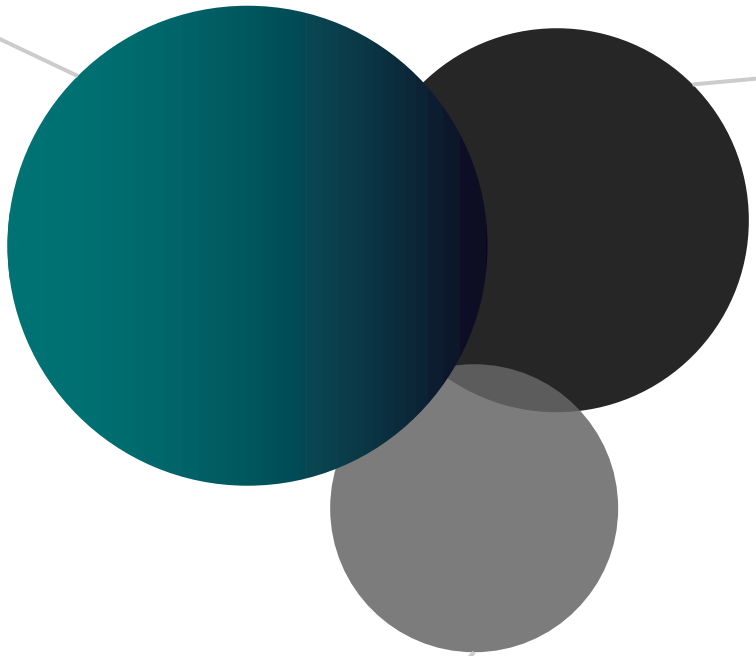

DEVELOPER-FIRST MEETS CUSTOMER-FIRST

Adopt and consistently apply both perspectives together

Actively solicit feedback to identify gaps and missing features — then act on them

LISTEN TO INDUSTRIAL USERS

WHO IS E4?

E4 COMPUTER ENGINEERING

BASED IN ITALY IN THE DATA VALLEY AND ACTIVE SINCE 2002, WE HAVE BEEN DESIGNING AND PROVIDING END-2-END SOLUTIONS FOR **HPC**, **ARTIFICIAL INTELLIGENCE** *AND* **QUANTUM COMPUTING**.
WITH A STRONG PRESENCE IN THE ACADEMIC AND ENTERPRISE MARKETS, WE HAVE BEEN COLLABORATING FOR YEARS WITH THE MAIN RESEARCH CENTERS AT NATIONAL AND INTERNATIONAL LEVEL.

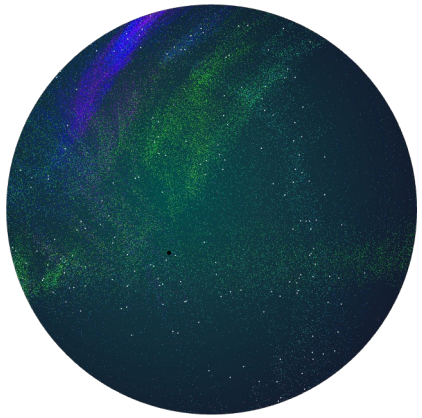

HPC

For more than 20 years, E4 has been innovating and delivering HPC hardware and software solutions for hundreds of customers from the enterprise and the research centers, such as CINECA, CERN, ECMWF, LEONARDO

AI

Building of the infrastructures for the raising AI technologies, helping customers to face and drive the convergence between HPC & AI. Development of vertical GenAI solutions for the secure on-premise scenarios

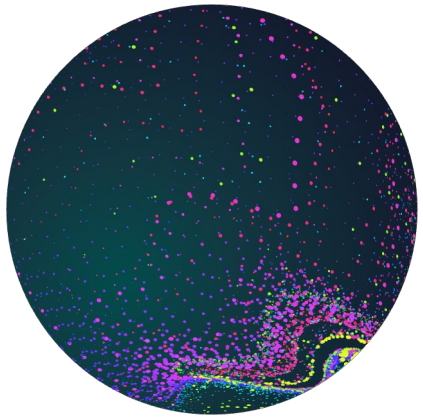

QUANTUM

Pioneering the Quantum Computing technologies, we are involved in several research projects by the European Union, offering our customers expertise for emulation and algorithm optimization

GOING INTERNATIONAL

STRONGER IN OUR REGION,
GLOBAL AMBITION

E4 COMPUTER ENGINEERING
DEUTSCHLAND
MUNICH

E4 COMPUTER ENGINEERING
SCANDIANO HQ
SULMONA
CAGLIARI

E4 COMPUTER ENGINEERING
SPAIN
MADRID

E4 ASSEMBLY LINE

E4 has the facilities, logistics and expertise to ***pre-assemble systems****, up to entire racks, in its own laboratory, where every hardware and software component is* ***checked****,* ***updated*** *and* ***tested*** *before delivery to the customer.*

E4 has the facilities, logistics and expertise to ***pre-assemble systems****, up to entire racks, in its own laboratory, where every hardware and software component is* ***checked****,* ***updated*** *and* ***tested*** *before delivery to the customer.*

The attention paid to every stage of the process is well illustrated by the ***burn-in process****, which consists of a minimum 72-hour test phase during which the system is run at maximum capacity.*

This activity gives back the guarantee to the customer to minimize the failure of components during their lifecycle.

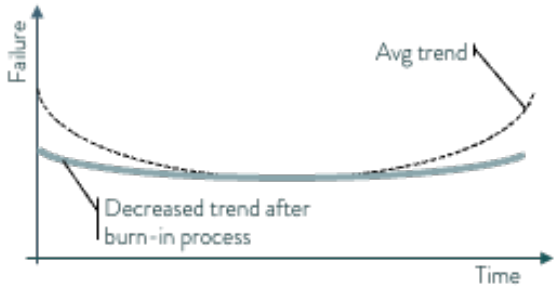

When in production, E4 is able to support the customer in any kind of issue, ***remotely*** *or* ***with on-premise specialistic personnel****. From direct activities on hardware to RMA management and logistics, the team could guarantee the maximum reliability.*

E4 METHODOLOGY

E4 COMPUTER ENGINEERING

Operations

- *Production assembly, verification and testing*
- *Deploy team configures all the software components*
- *Delivery team installs and tests the solution at the customer's premises, until acceptance*
- *Notifies* **Support** *of new completed installations and to* **Tech Factory** *during configuration & installation*

R&D Teams (Tech Factory – AI Factory)

- **Presale** *support sales team analysing opportunities*
- **Presale** *and* **R&D** *team formulate a technical proposal*
- **Presale** *and* **R&D** *team analyse and test new technologies and update E4 solutions and documentation*
- *Provides feedback to* **Operations** *and* **Support** *about opportunities and new technologies*

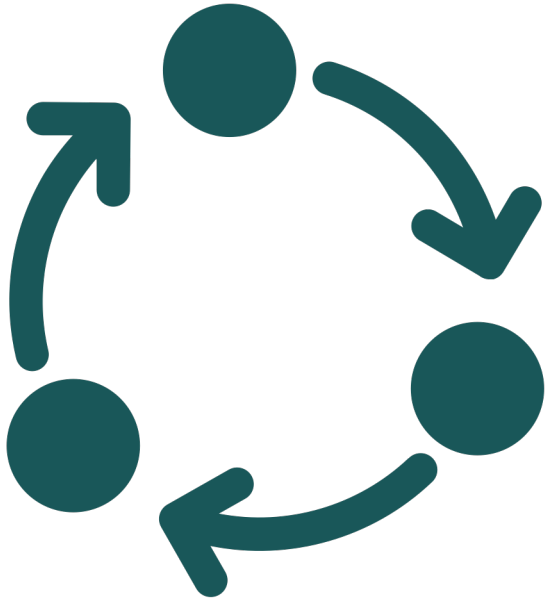

Support

- *All hardware or software issues throughout the life cycle of the solutions*
- *Provides feedback to* **Operations** *and* **Tech Factory** *from customers*

TWO DECADES INTEGRATING NON-MAINSTREAM ARCHITECTURES

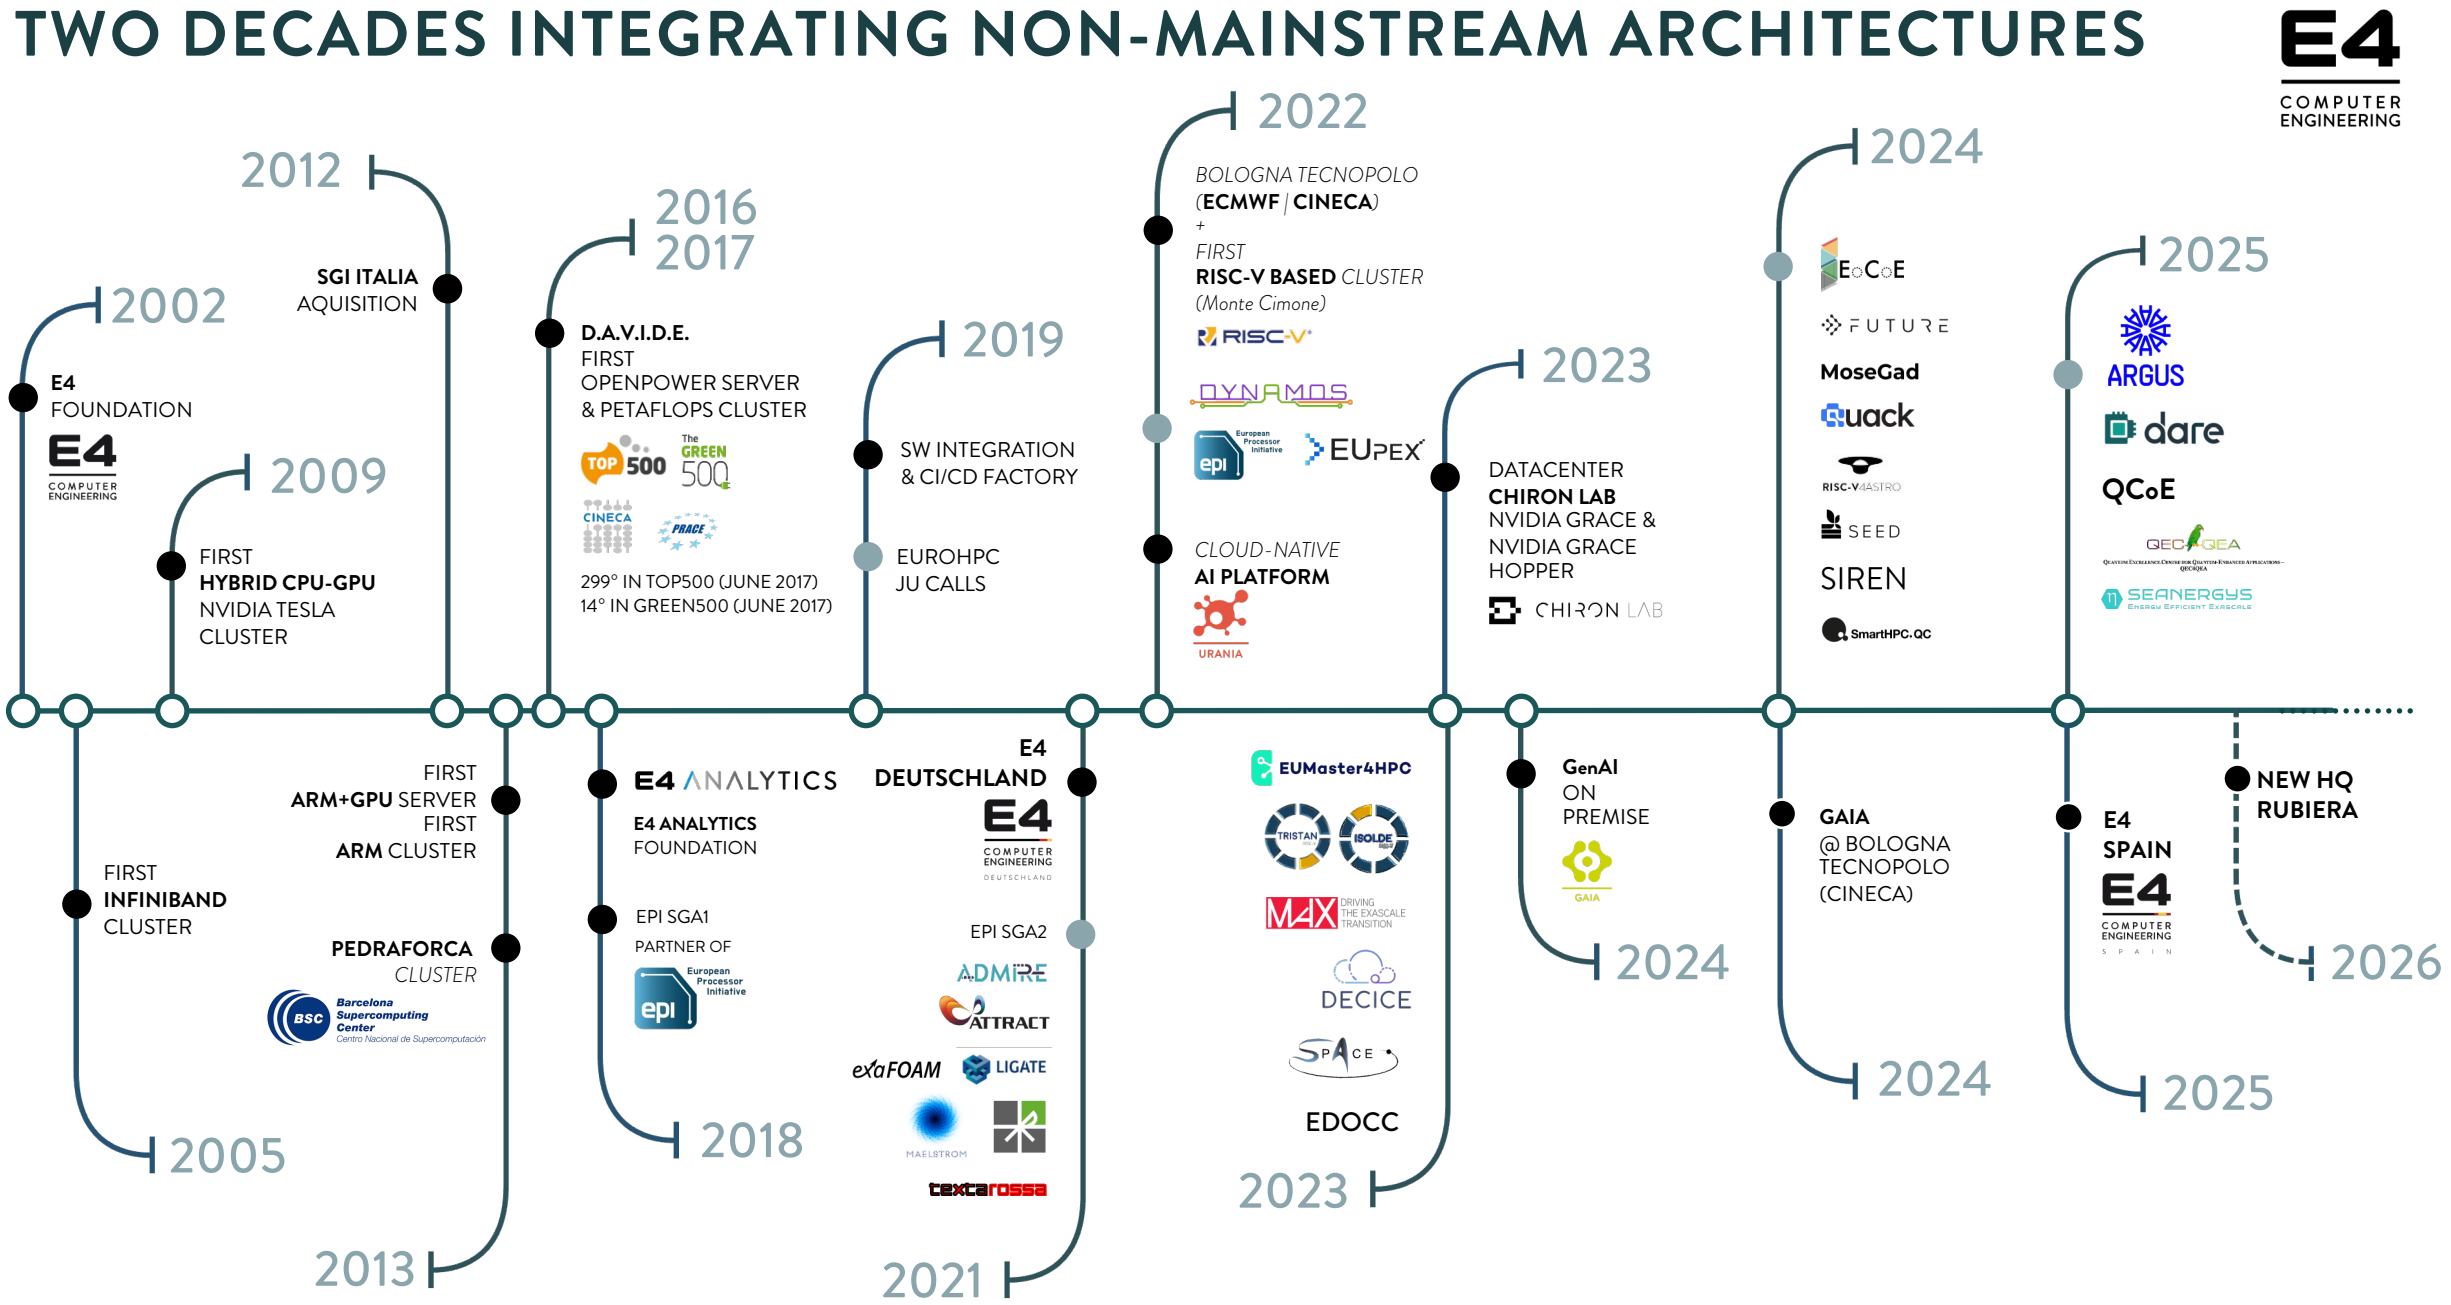

RISC-V & MONTE CIMONE TIMELINE

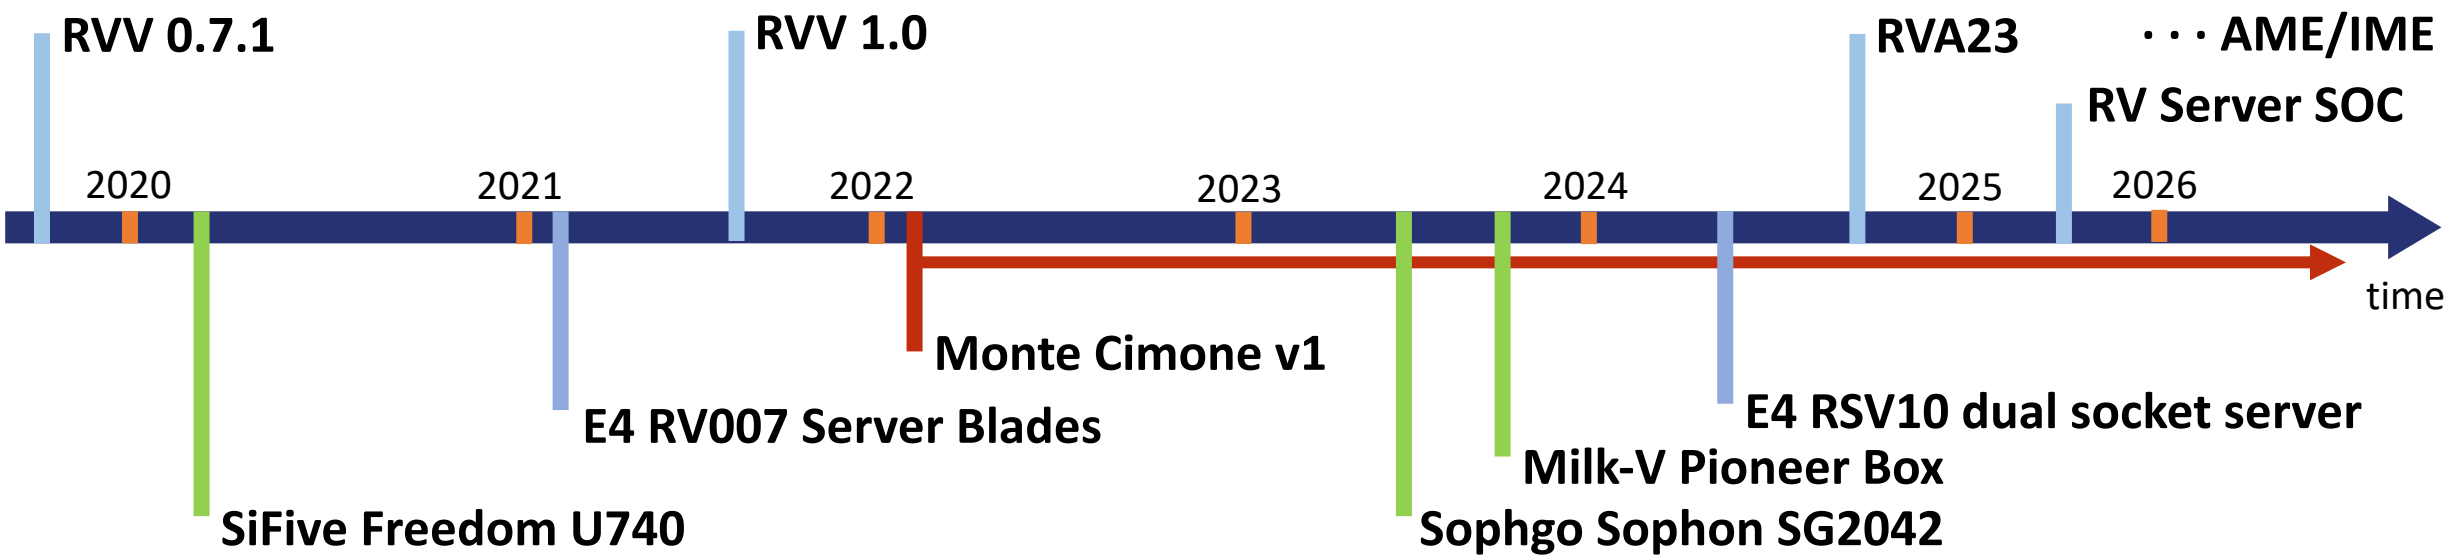

E4 RSV10 dual socket server

- Dual socket SG2042
- 256 GB RAM

MONTE CIMONE 1

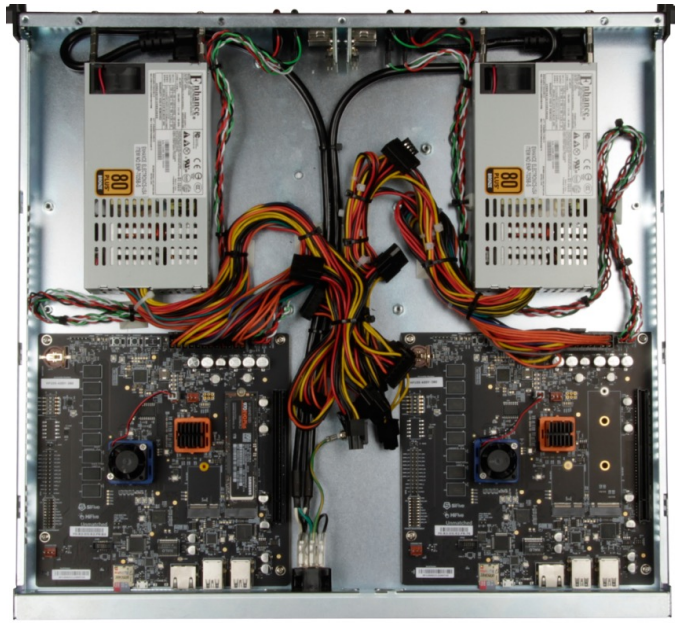

SIFIVE HIFIVE UNMATCHED BOARD

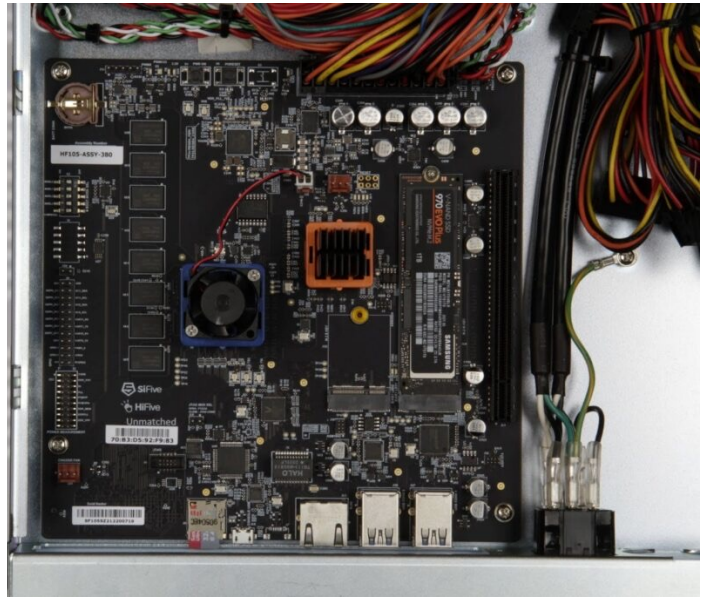

4X E4 RV007 1U CUSTOM SERVER BLADES

- *2x SiFive U740 SoC with 4x U74 RV64GCB cores*
- *16GB of DDR4*
- *1TB node-local NVME storage*
- *PCIe expansion card w/InfiniBand HCAs*
- *Ethernet + IB parallel networks*

SIFIVE U740 SOC W. 7 SEPARATED POWER RAILS

- *Core complex, IOs, PLLs, DDR subsystem and PCIe one*
- *Board implements distinct shunt resistors*

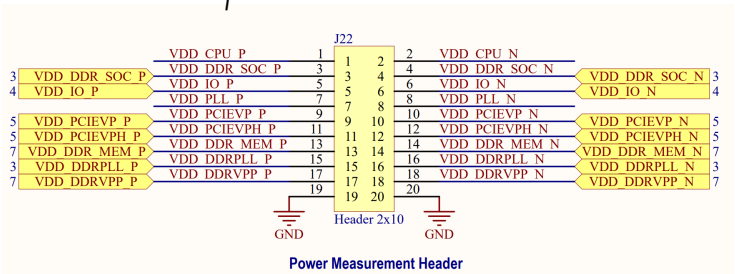

POWER CHARACTERIZATION

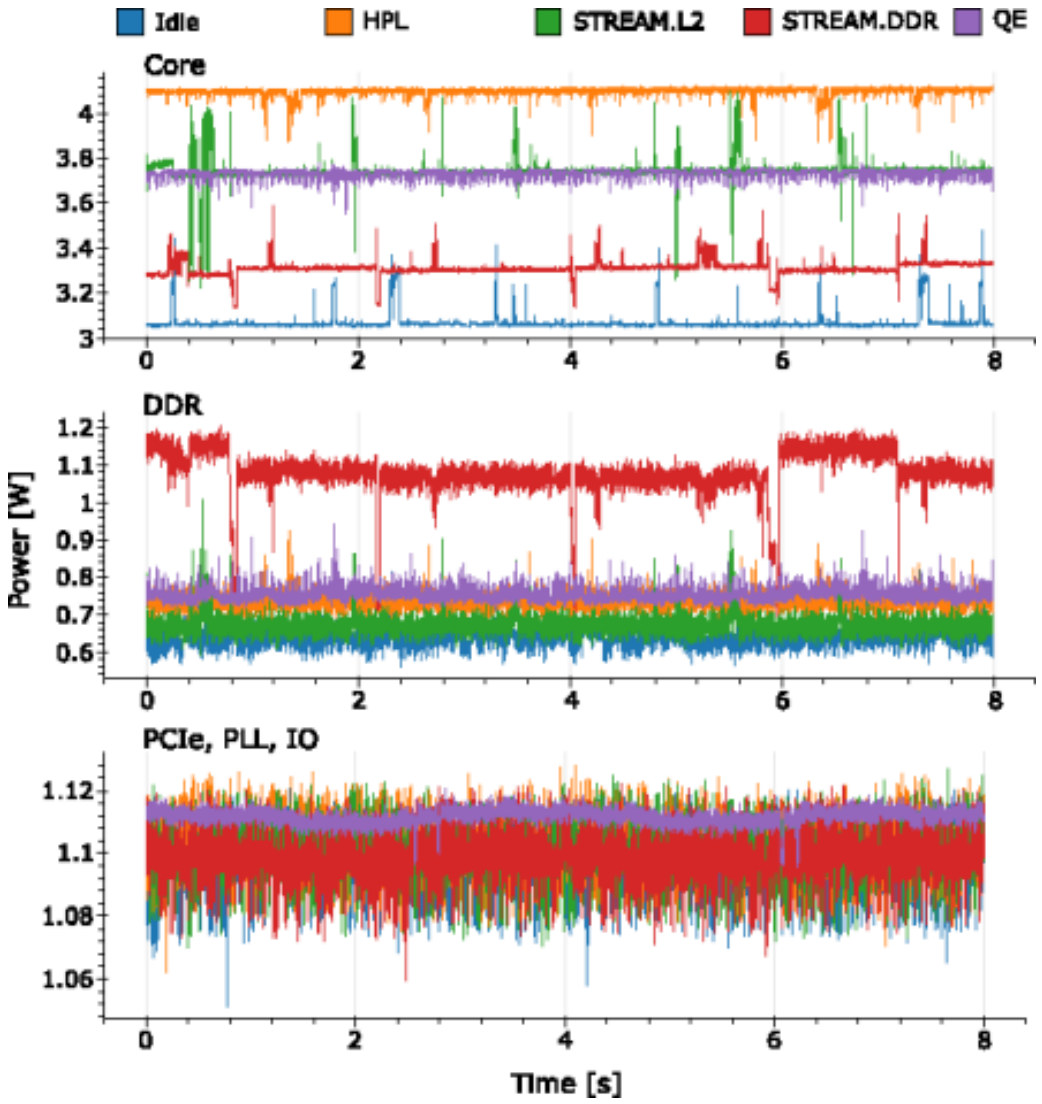

Line	Idle		HPL		STREAM.L2		STREAM.DDR		QE	
	[mW]	[%]	[mW]	[%]	[mW]	[%]	[mW]	[%]	[mW]	[%]
core	3075	64	4097	69	3714	68	3287	62	3825	67
ddr_soc	139	3	177	3	170	3	232	4	176	3
io	20	0	20	0	20	0	20	0	20	0
pll	1	0	1	0	1	0	1	0	1	0
pcievp	521	11	527	9	524	10	522	10	530	9
pcievph	555	12	554	9	554	10	555	10	561	10
ddr_mem	404	8	440	7	401	7	592	11	434	8
ddr_pll	28	1	28	1	28	1	28	1	28	1
ddr_vpp	67	1	90	2	73	1	98	2	95	2
Total	4810	100	5935	100	5486	100	5336	100	5670	100

- **IDLE:**
  4.81W (64% of core power, 13% related to DDR and 23% of related to PCI subsystem)
- **HPL:**
  5.935W (69% of core power, 14% related to DDR and 18% related to PCI subsystem)

FROM MONTE CIMONE TO MONTE CIMONE 2

E4 COMPUTER ENGINEERING

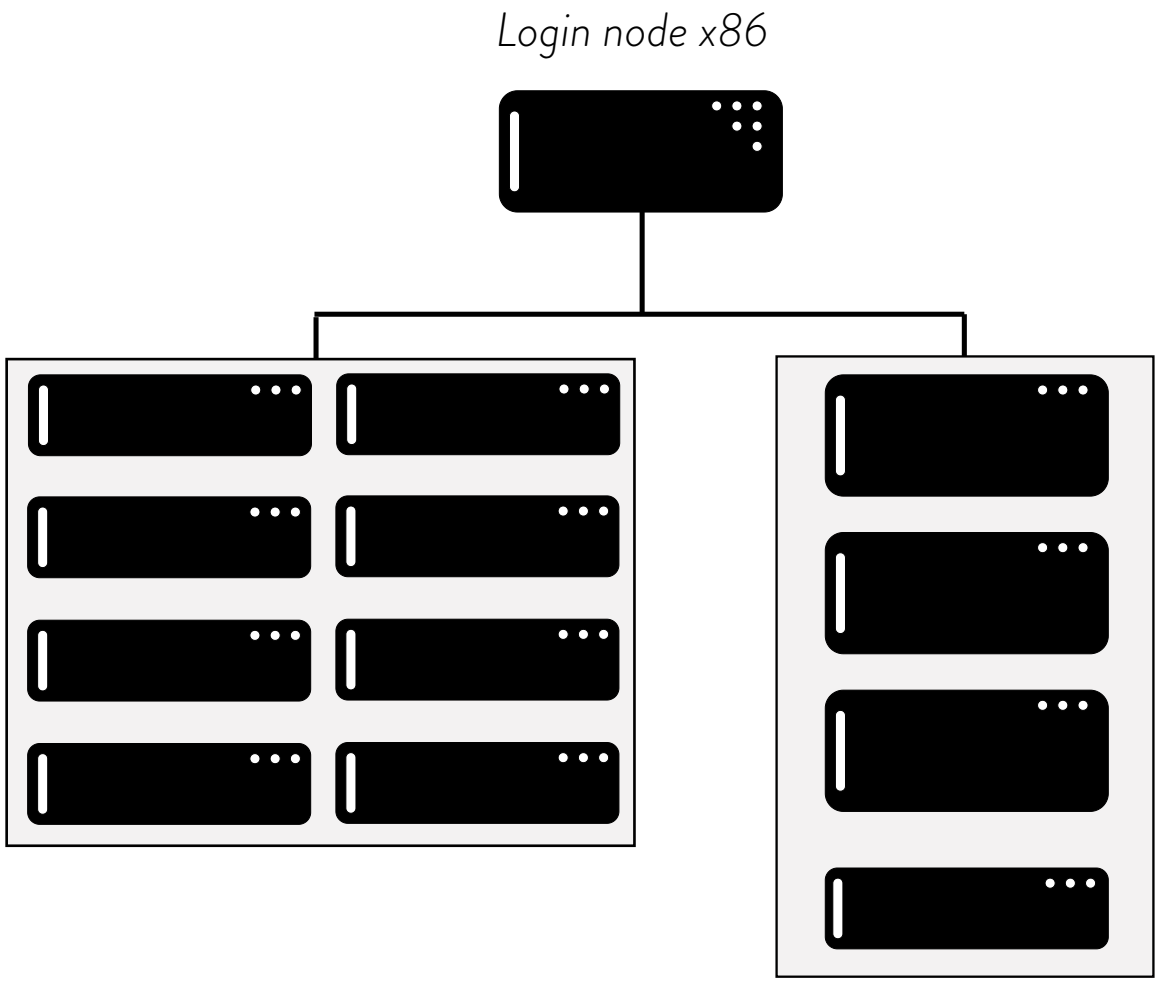

Partition 1:
8 SiFive

Partition 2:
3 Milk-V + 1 E4
RSV10

FROM MONTE CIMONE 2 TO MONTE CIMONE 3

Cluster environment:

- SLURM
- Examon
- Shared NFS
- Spack + tools
  - Added compilers:
    - Xuantie GNU Toolchain (vendor GCC w. RVV 0.7.1 support)
    - GCC 14 (first version w. xtheadvector)
  - Added BLAS libraries:
    - OpenBLAS (vanilla no RVV, RVV1, vendor RVV 0.7.1)
    - BLIS RVV1, RVV 0.7.1

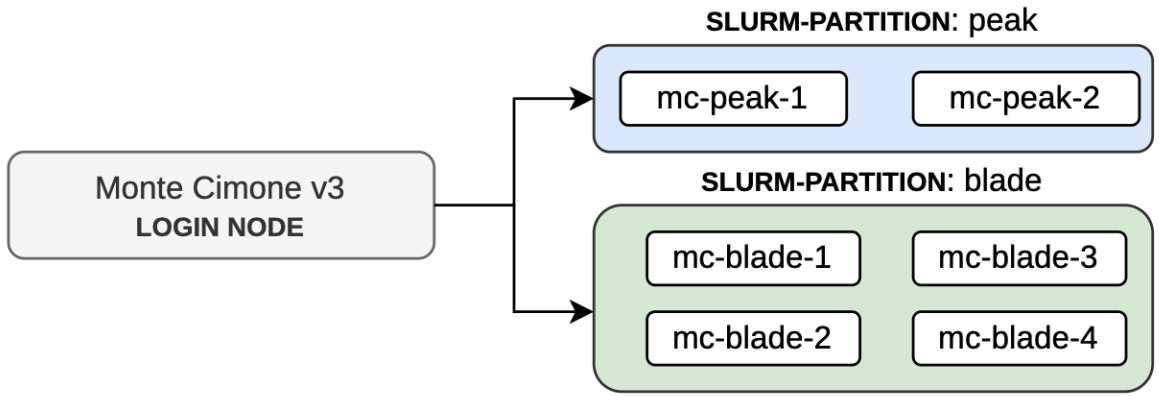

EXAMON - A DATA LAKEHOUSE FOR HPC

E4 COMPUTER ENGINEERING

Every signal from the infrastructure flows into a single, schema-less, scalable store and immediately available to operators, researchers, and AI agents.

Collection

Plugin SDK
Declarative SDK, built-in resilience

Publishers
IPMI, Redfish, GPU, Slurm, ...

MQTT/Kafka
Pub/sub transport, M2M native

CMDB/DCIM
Site inventory friendly, for automated onboarding

Any source, any protocol, zero schema required

Storage & Query

Schema-Less time series
KairosDB, TDengine, ...

Job accounting & metadata
Apache Cassandra

Unified federated SQL
Trino

Cold Tier for long-term History
Apache Iceberg - S3

One SQL interface, all data, all time ranges, local

Intelligence & Insight

Grafana
User friendly dashboards

Apache Superset
SQL-based BI

Jupyter notebook
Advanced analytics/ML

ZenML / Urania –E 4
MLOps

Custom Tools/CLI
Anomaly detection, RCA

AI Agents
NL interface & agentic loops

Any user, any question, any skill level

BENCHMARKING MONTE CIMONE - HPL

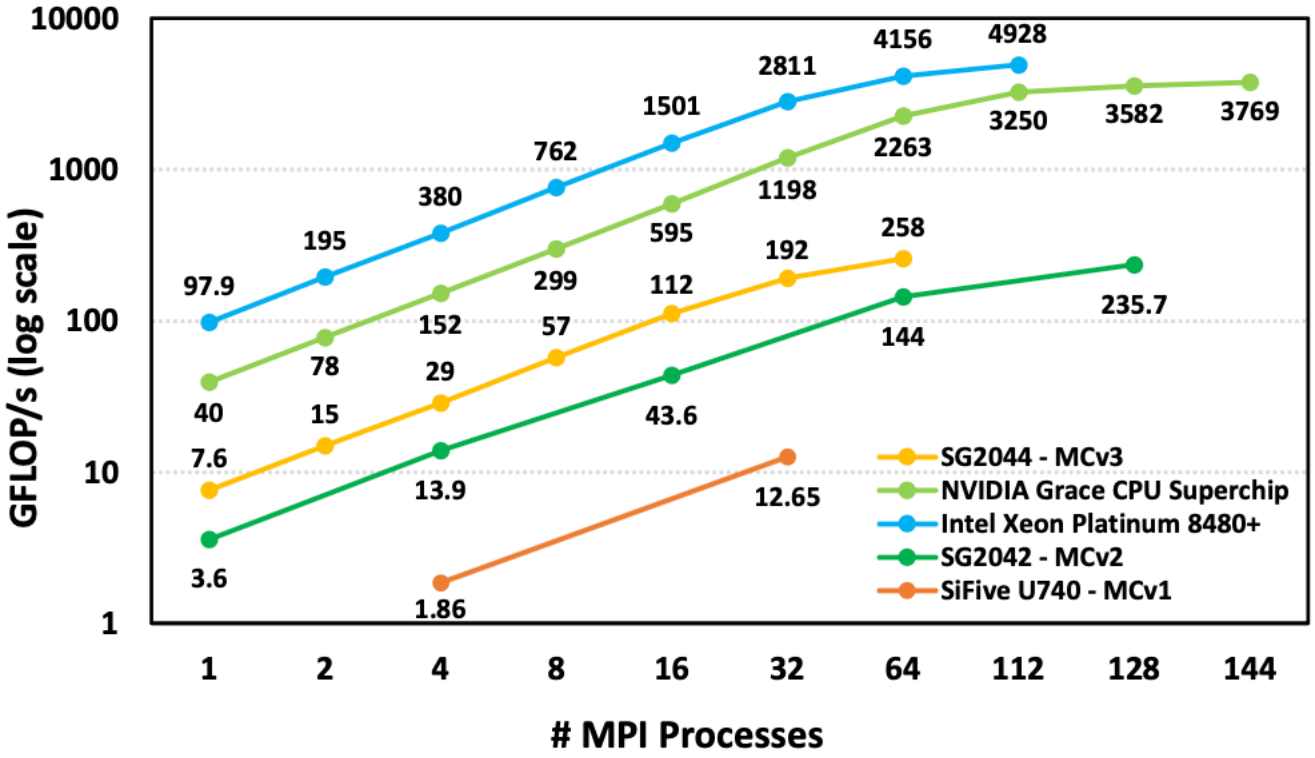

	SG2044	Intel SR	NVIDIA GS
ISA	RISC-V	x86-64	Armv9
Cores / Node	64	112 (2×56)	144 (2×72)
Vector ISA	RVV 1.0	AVX-512	SVE2
Vector length	128 bits	2×512 bits	4×128 bits
Frequency [GHz]	2.6	2.0	3.1

Energy Efficiency

	MCv1	MCv3	NVIDIA GS	Intel SR
Avg Power [W]	5.9	83.9	828	1208
HPL [GFLOPs]	1.86	258.0	3769	4928
GFLOPs/W	0.31	3.08	4.55	4.08

MONTE CIMONE IN PRODUCTION

E4 COMPUTER ENGINEERING

Preliminary results of OpenFOAM on RISC-V with Vector Extension 1.0

Elisabetta Boella, Fabrizio Magugliani

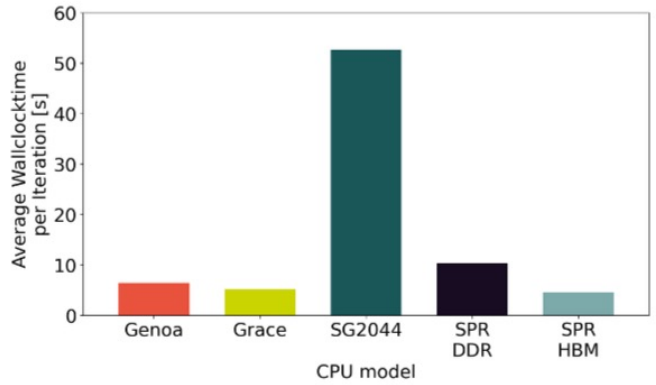

Processor	Architecture	Cores/node
AMD Genoa	x86_64	128
NVIDIA Grace	ARM	144
INTEL SPR HBM	x86_64	96
INTEL SPR DDR	x86_64	96
SOPHON SG2044	RISC-V	64

Validation and Performance Assessment of a Massively Parallel Astrophysical Application on RISC-V CPUs

Jenny Lynn Almerol[1[x]], Nitin Shukla[2[xx]], Geray Karademir[3[xxx]], Federico Ficarelli[1[xxx]], Andrea Bartolini[4[xxx]], Emanuele Venieri[4[xxx]], Giacomo Madella[4[xxx]], and Elisabetta Boella[3[xxx]]

[1] SuperComputing Applications and Innovation Department, CINECA, Via Magnanelli 6/3, Bologna, 40033, Italy
[2] Universitäts-Sternwarte, Fakultät für Physik, Ludwig-Maximilians-Universität München, Munich, Germany
[3] ABC Institute, Rupert-Karls-University Heidelberg, Heidelberg, Germany
[4] University of Bologna Bologna, Italy

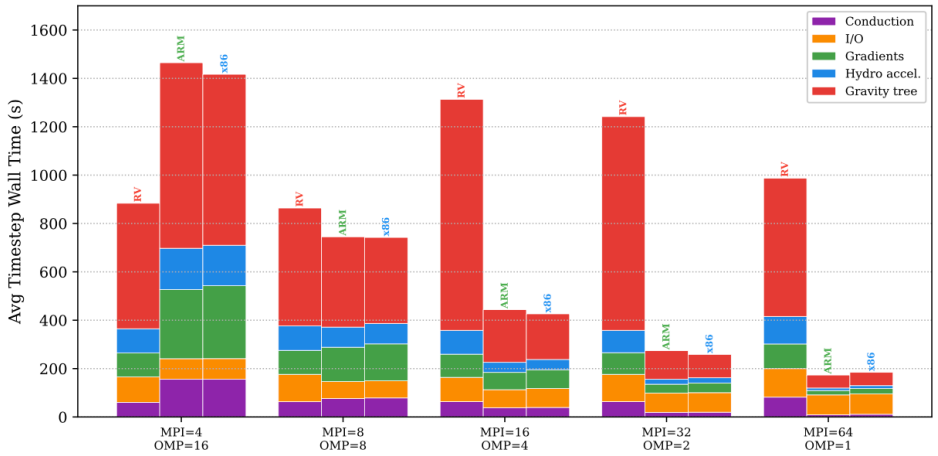

OpenGadget3

MONTE CIMONE: CONCLUSIONS

E4 COMPUTER ENGINEERING

MCv1 vs. MCv2: rapidly evolving RISC-V ecosystem

In two years a performance improvement obtained in average in eight years (Top500)

MCv3 vs. Intel SF & NVIDIA GS

Memory bandwidth gap: less than 2x (vs Sapphire Rapids) / 4x (vs NVIDIA Grace)

FLOPS gap: less than 2× / less than 3×

Energy efficiency gap: less than 1.5×

What's next? MC AI

Fully RV AI machine

```
[evenieri.local@mc-peak-2 ~]$ sudo lscpu | head -n 10
Architecture:                         riscv64
Byte Order:                           Little Endian
CPU(s):                               64
On-line CPU(s) list:                  0-63
Vendor ID:                            0x5b7
BIOS Vendor ID:                       SOPHGO
Model name:                           -
BIOS Model name:                      SG2044 Not Set CPU @ 2.6GHz
BIOS CPU family:                      513
CPU family:                           0x80000000090c0d00
[evenieri.local@mc-peak-2 ~]$ lspci | grep accelerators
0000:01:00.0 Processing accelerators: Tenstorrent Inc Blackhole
0006:14:00.0 Processing accelerators: Axelera AI Metis AIPU (rev 02)
```

EUROPEAN PARTITION IN IT4LIA AI FACTORY

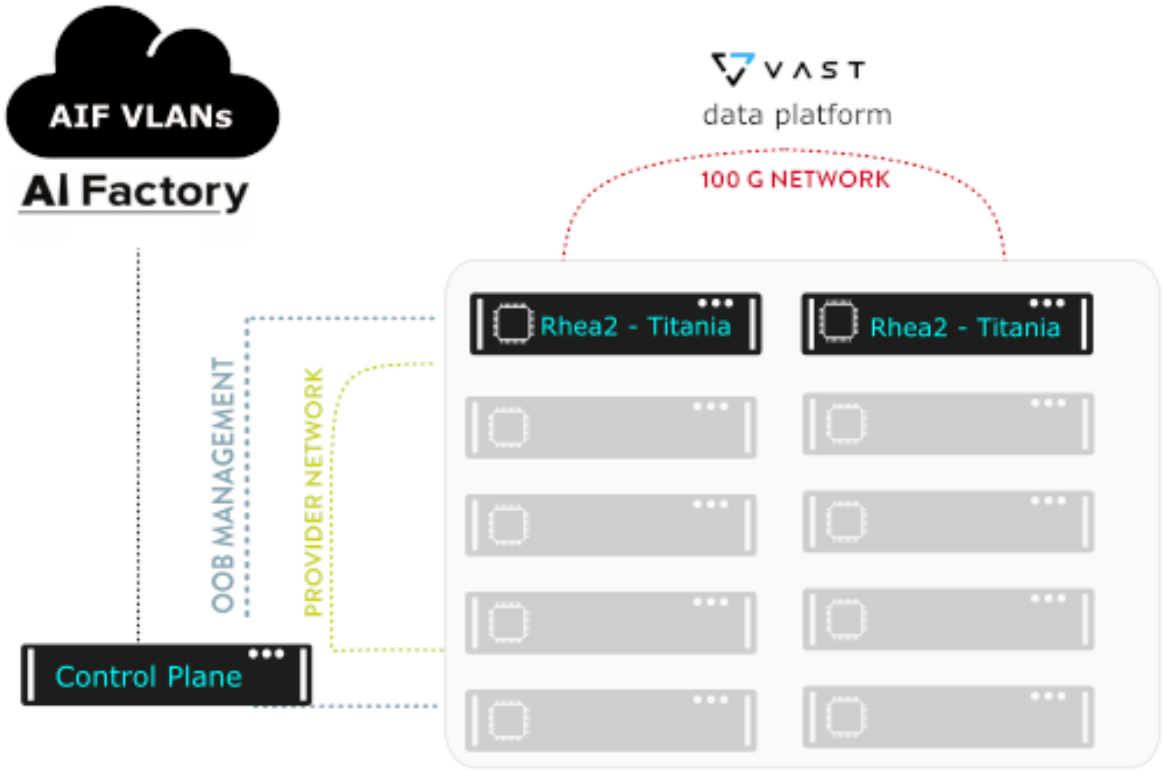

- 10 Next-Gen-EU nodes inside the IT4LIA system
- Per node: 2× SiPearl Rhea2 + 8× Axelera Titania AIPUs
- 1 TB accelerator memory per node, 200 PFlop/s FP8 partition peak
- Data fabric: 200 GbE Ethernet, fat-tree, shared VAST 100 PB storage
- RISC-V enters production not as a replacement, as a **specialised European path**

THIS IS THE TIME FOR TRANSITIONING TO INDUSTRY–GRADE RISC-V!

E4 COMPUTER ENGINEERING

ISA MINDSET *(starting point)*

- *ISA / non-ISA specs*
- *Profiles*

ECOSYSTEM MINDSET *(target)*

TECHNOLOGY

- *High-performance, industry-grade CPU & Accelerators*
- *Software development tools & robust OS support*

USER-CENTRIC APPROACH

- *Customer-first: industrial user viewpoint*
- *Lower barriers for newcomers*
- *Proactive engagement with ISVs, ODMs, OEMs, product makers*

TOOLING & DOCUMENTATION

- *Search tools for extensions, options, CSRs*

CERTIFICATION

- *Solid RISC-V Certification Program*

TALENT

- *Academic curricula, mentorships & internships*
- *Build a new generation of RISC-V experts*

APPLICATIONS

- *Porting, optimization & certification of open-source and ISV packages*
- *Plug & play approach*

E4 COMPUTER ENGINEERING

THANK YOU

CONTACTS

info@e4company.com
support@e4company.com
sales@e4company.com

E4 Computer Engineering SpA
Via Martiri della Libertà, 66 . 42019 Scandiano (RE) - Italy
Tel. +39 0522 991811