European Processor Initiative

MPSoC'24

Will HPC be a next decade disruptor, or will it be disrupted?

Eric Monchalin
Chair of the European Processor Initiative
Vice President at Eviden, Head of Machine Intelligence
09/07/2024

EVIDEN

the uncertain certainties

1945

Thomas J.Watson (CEO of IBM)

« World market for may be five computers »

1977

Ken Olson (CEO of DEC)

« No reason for anyone to have a computer at home »

1980

IBM study

« Only about 50 Cray-1 class computers will be sold per year »

So , let be cautious

The views expressed on the following slides are those of the presenter.

They do not necessarily represent any view of Eviden, EPI or EuroHPC JU organizations, affiliates or employees.

European Processor Initiative

MPSoC'24

1 What's numerical simulation?

2 The ground truth

3 What's at stake?

4 The threats

5 The weaknesses

6 Cloud convergence

7 Making HPC future a reality

8 The European initiatives

9 Conclusion

EVIDEN

1 What's numerical simulation

Continuum to discretization

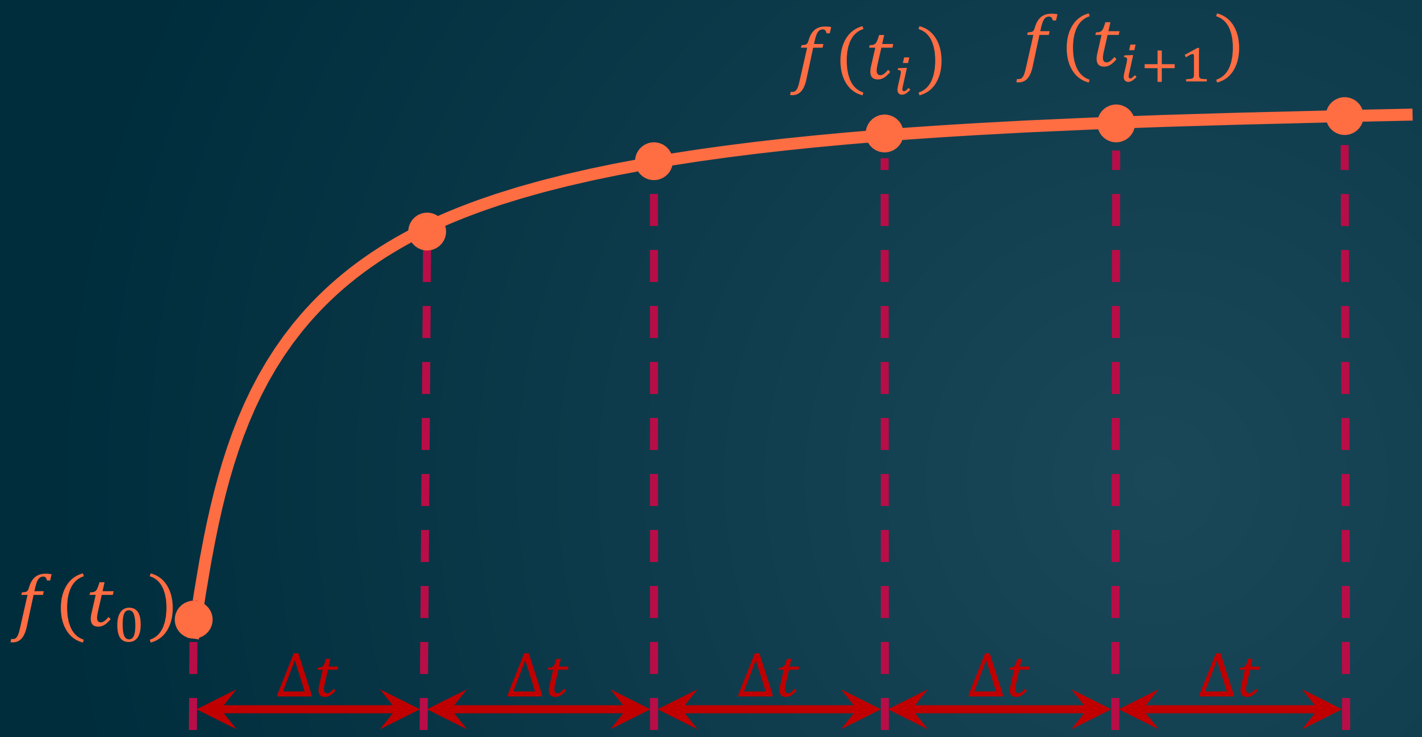

Derivative at t_i

$$\frac{df}{dt}(t_i) \cong \frac{f(t_{i+1}) - f(t_i)}{t_{i+1} - t_i} = \frac{f(t_{i+1}) - f(t_i)}{\Delta t}$$

Smaller the time stamps, more accurate the simulation and heavier the computation

Mathematical Model to Numerical Solution

$$F_u \cong -cv$$

$$F_d = mg$$

$$F_d + Fu = m\,\gamma \Rightarrow mg - cv = m\,\frac{dv}{dt}$$

$$\frac{dv}{dt} = g - \frac{c}{m}\,v$$

Analytical: $v(t) = \frac{mg}{c}\left(1 - e^{-\frac{c}{m}t}\right)$

$$\frac{dv}{dt} \cong \frac{\Delta v}{\Delta t} = \frac{v(t_{i+1}) - v(t_i)}{t_{i+1} - t_i} = g - \frac{c}{m}\,v(t_i)$$

Numerical: $v(t_{i+1}) = v(t_i) + \left[g - \frac{c}{m}\,v(t_i)\right](t_{i+1} - t_i)$

European Processor Initiative

MPSoC'24
EVIDEN

Decompose the compute domain to parallelize the computation

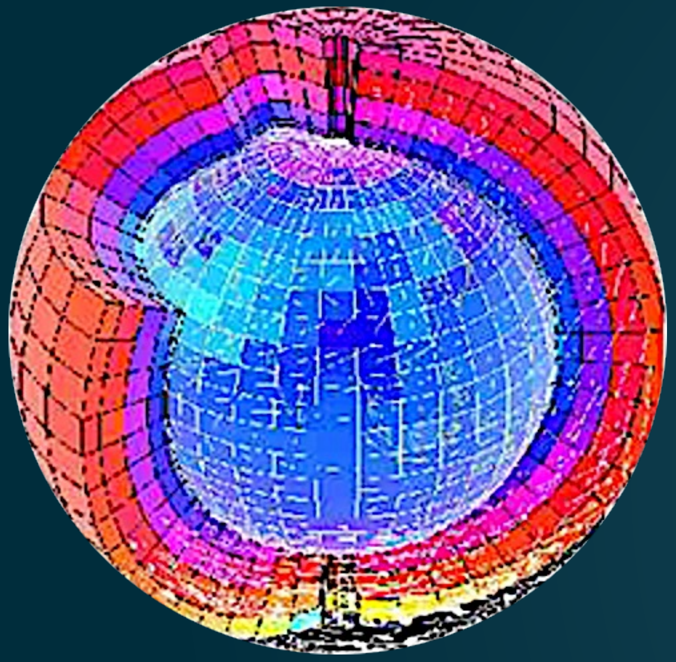

Parallel computation on mesh elements

And

Surface data exchanges with neighbors

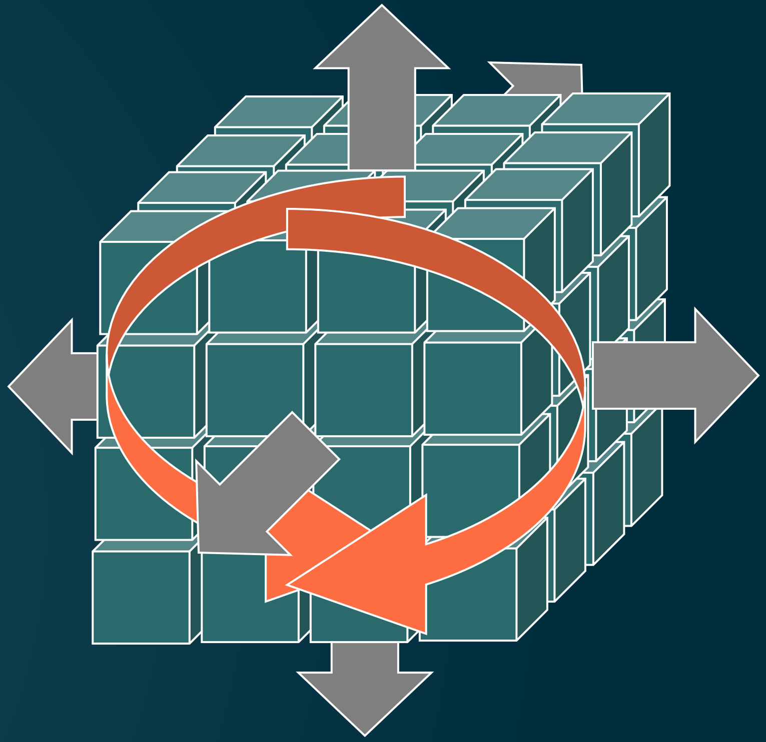

Refine a 3D mesh by a factor of 10 on each axis

- X 1,000 mesh elements
- 15 years of technology enhancements for the same solving budget

Almost a regular execution scheme

1st Mesh construction

Few CPUs
Large Memory

2sd Time stamps computation

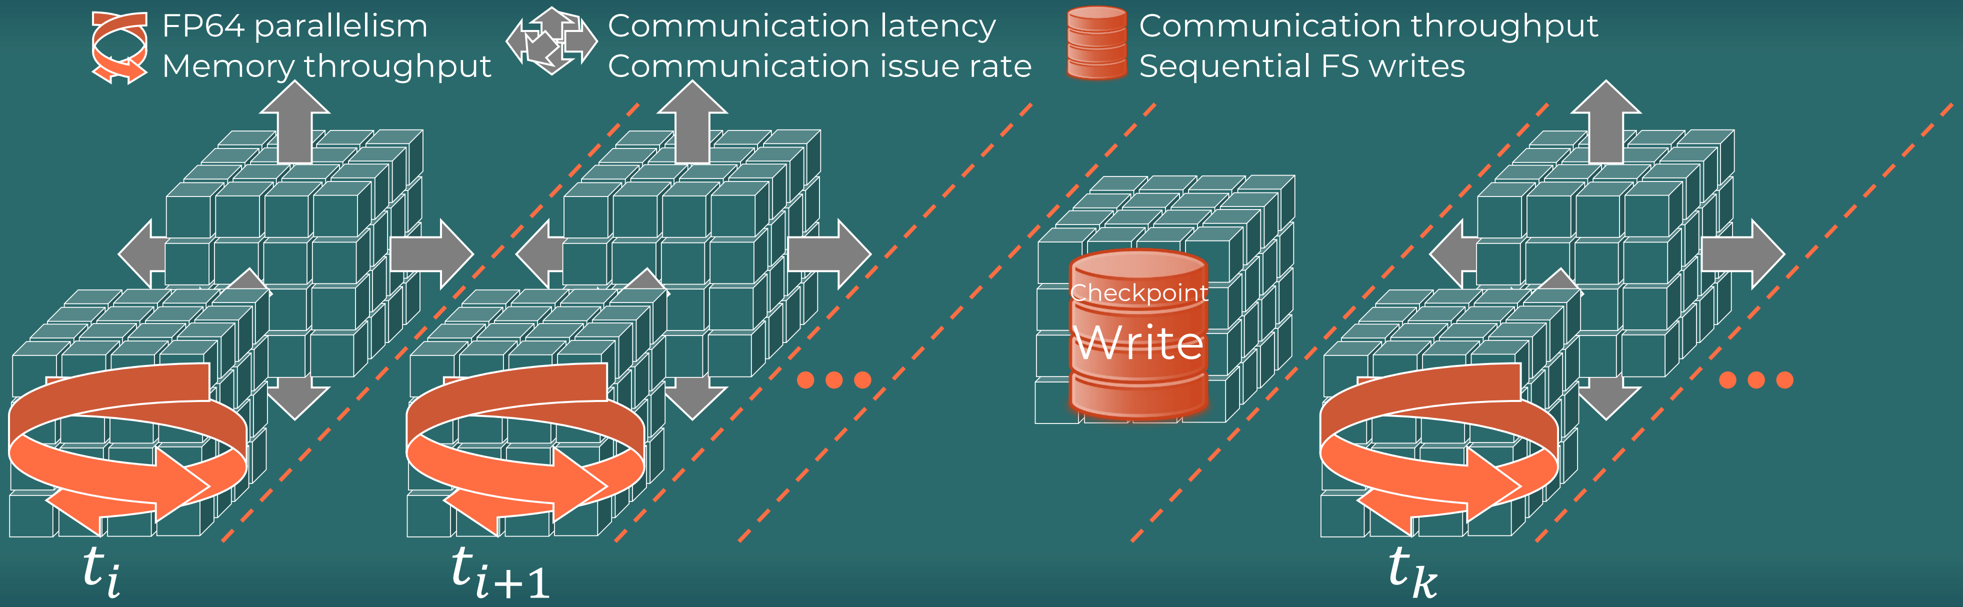

3rd Result saving

Communication throughput
Sequential FS writes

The harsh reality of the Amdahl's maw
Any sequential code or computing / communication delay matters

$\propto$: % of none parallelizable execution time

t: unparalyzed execution time

$\propto t + \frac{(1-\propto)t}{N}$: execution time for N tasks

$\frac{1}{\propto N+(1-\propto)}$: efficiency of the parallelism

Nvidia A100: 3456 FP64 cores

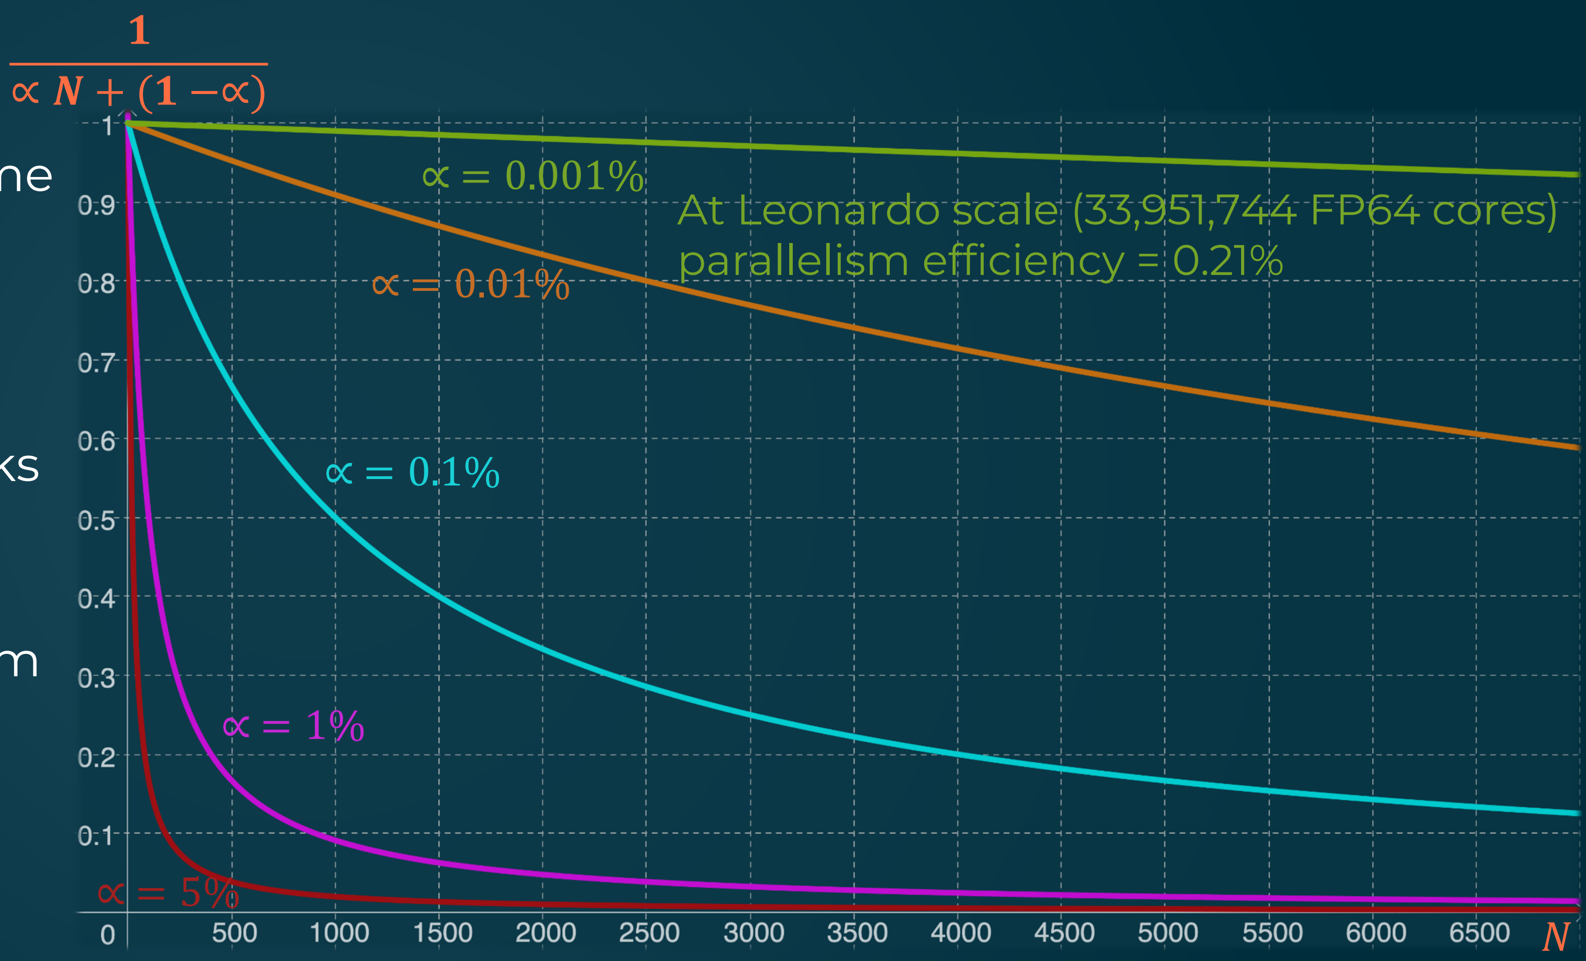

European Processor Initiative

MPSoC'24
EVIDEN

European Processor Initiative

2 The ground truth

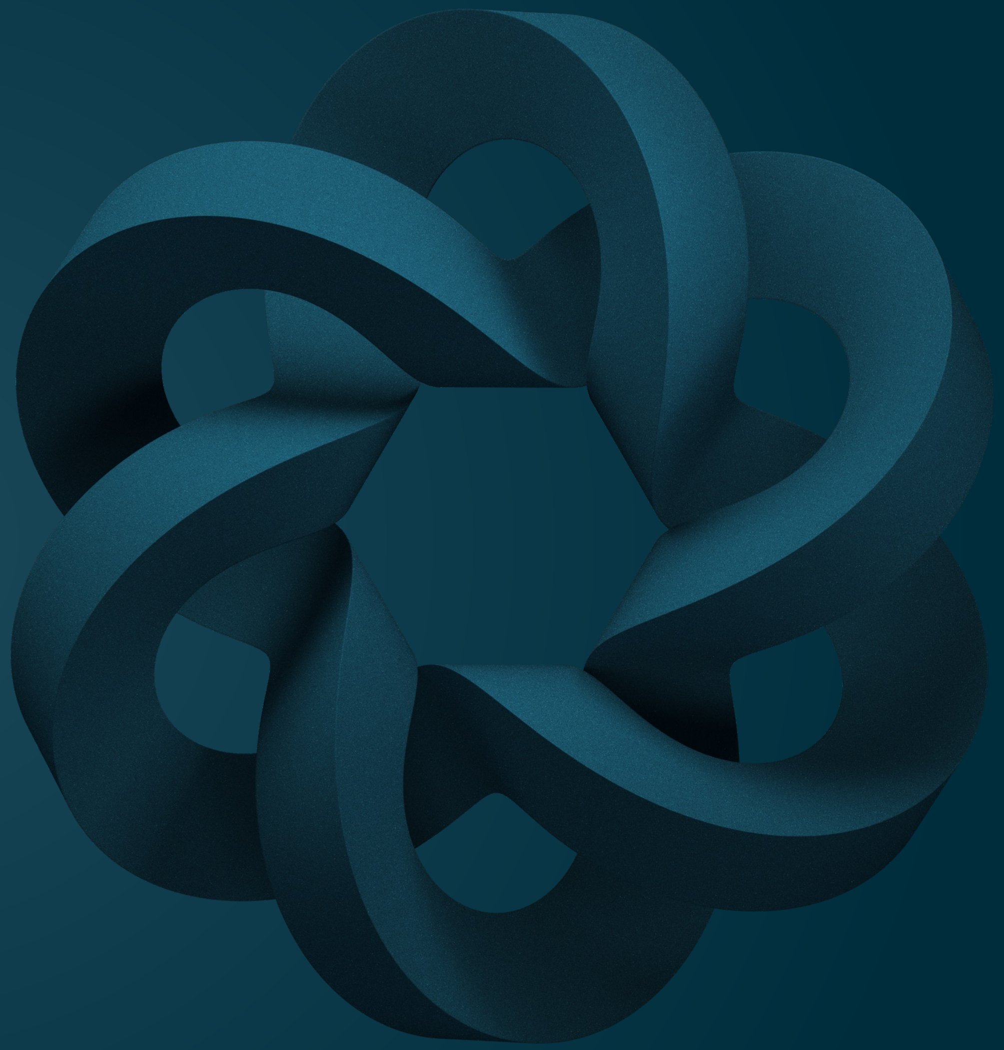

Software is eating the world (Marc Andreessen)
Yet hardware is still shaping it

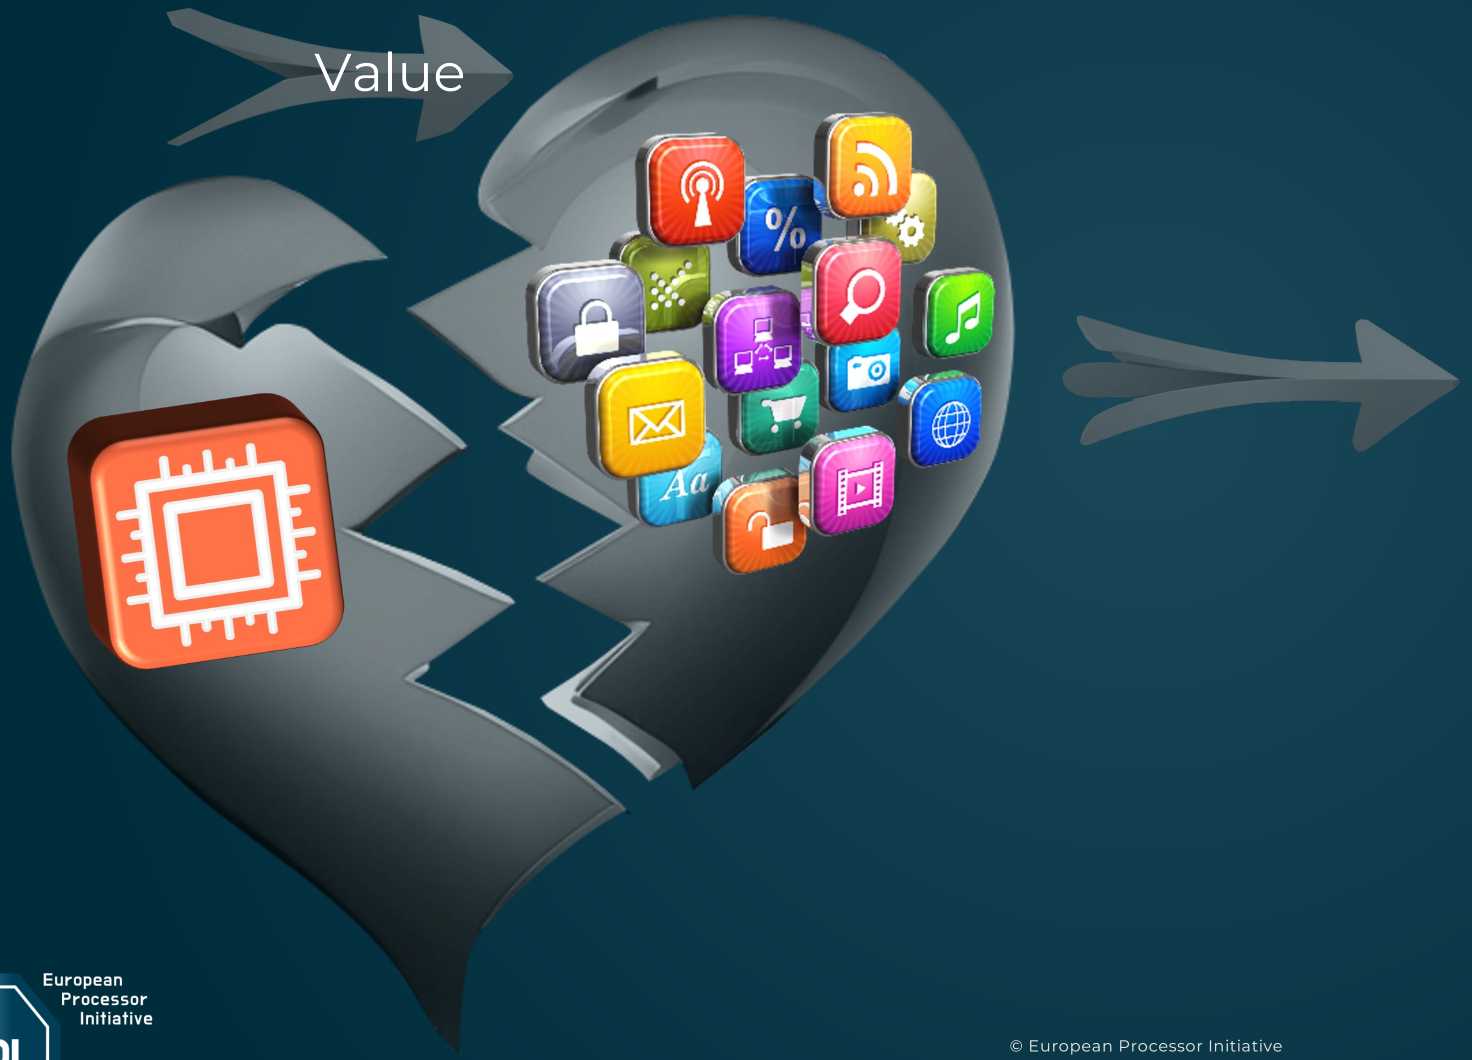

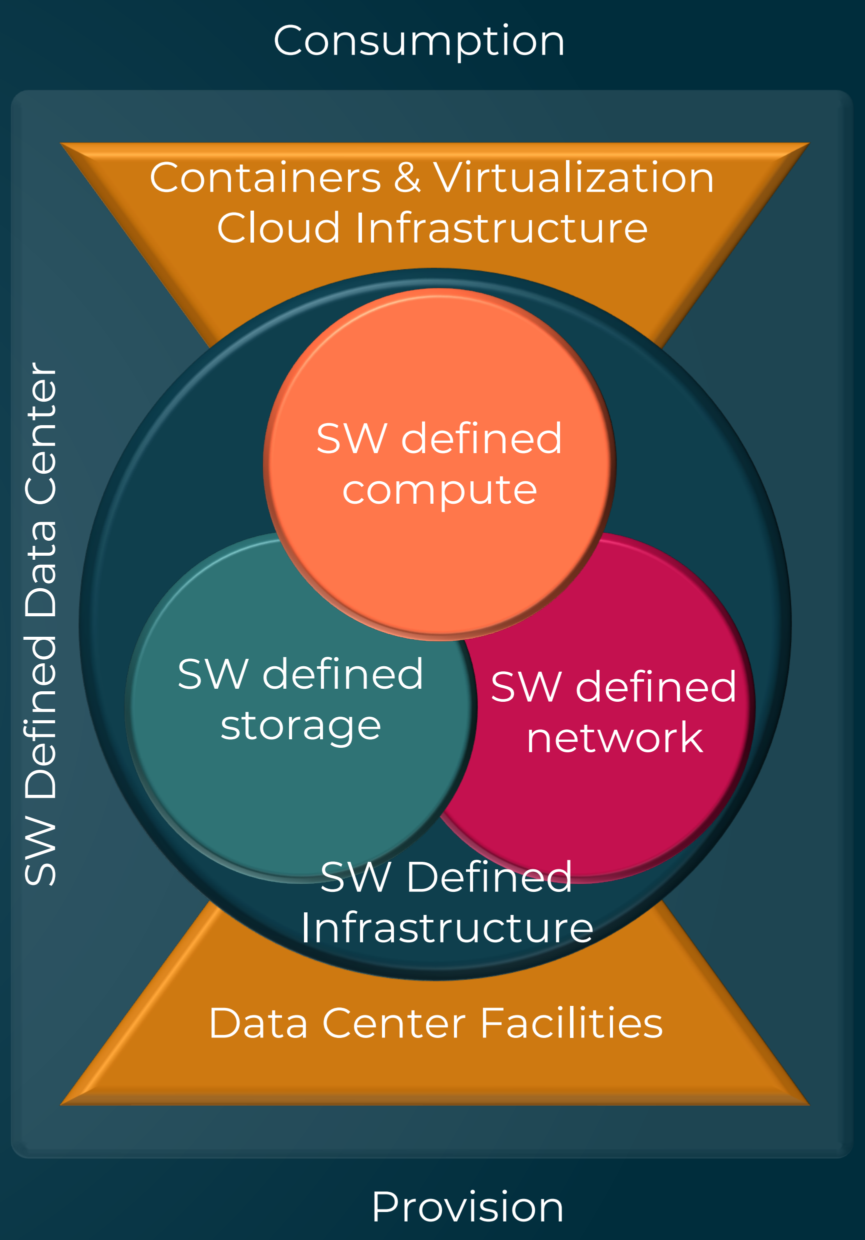

May be still few isolated Jurassic worlds
A pure construct of my mind or a reality for serious reasons?

The divergences of the microelectronics performances

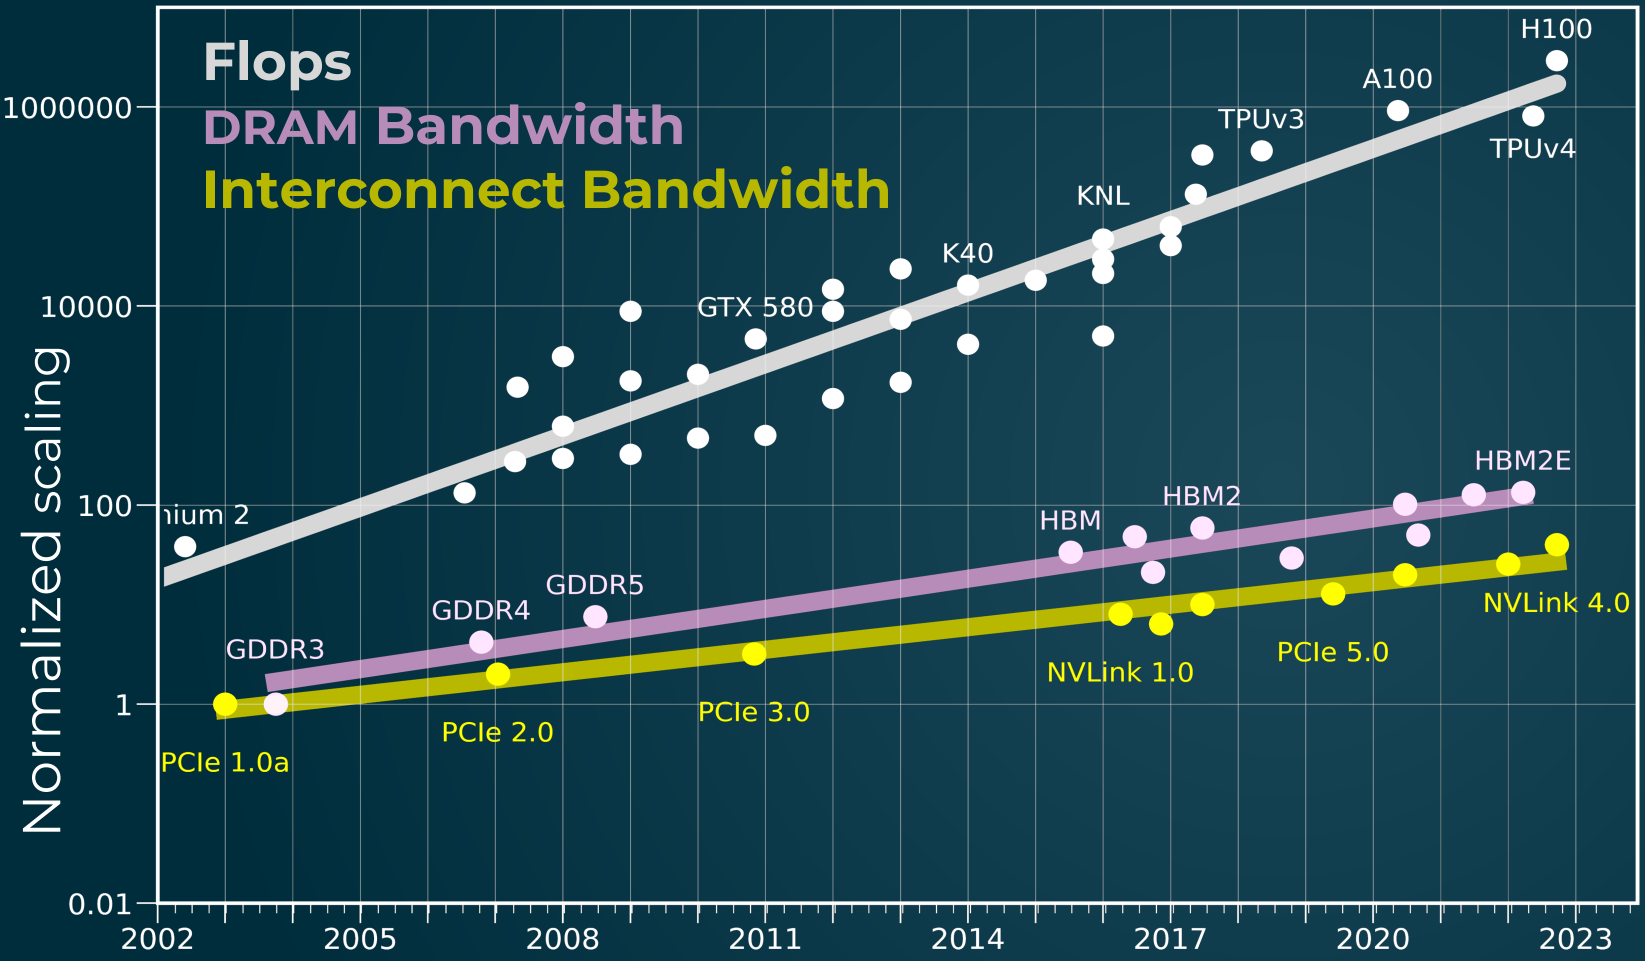

Source: medium.com (AI and Memory Wall)

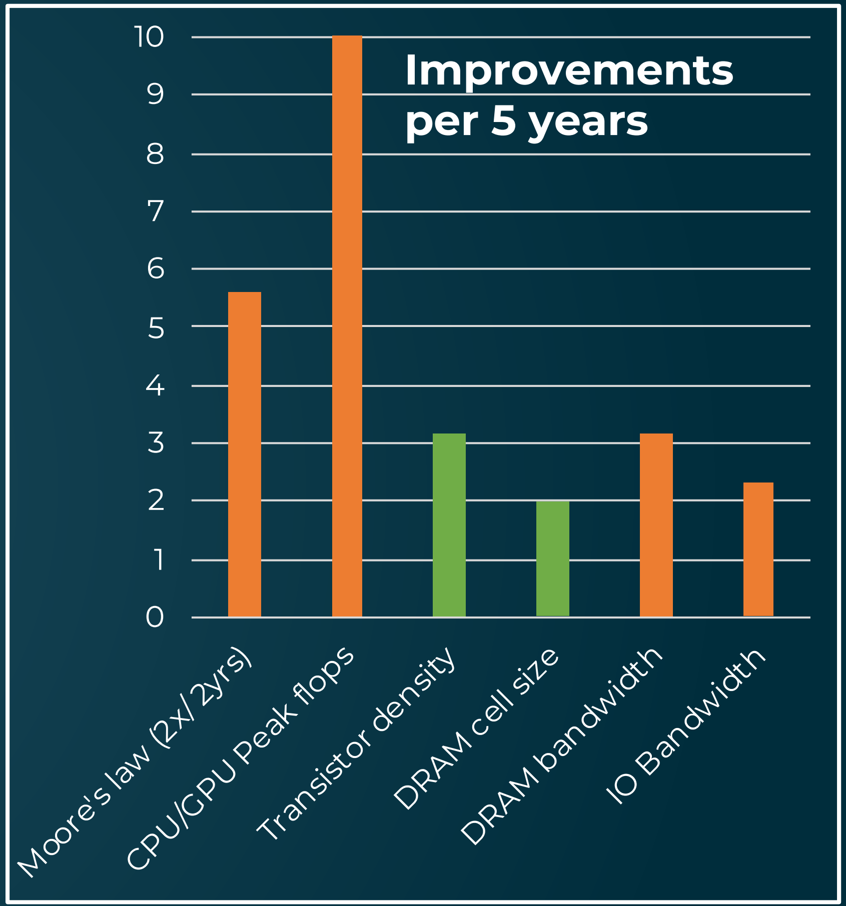

European Processor Initiative

MPSoC'24 EVIDEN

Supremacy of data generation over computation

Large Hadron Collider
Tens of Petabytes per year

Square Kilometer Array
400 exabytes per year

Climate modeling
400 Petabytes per year

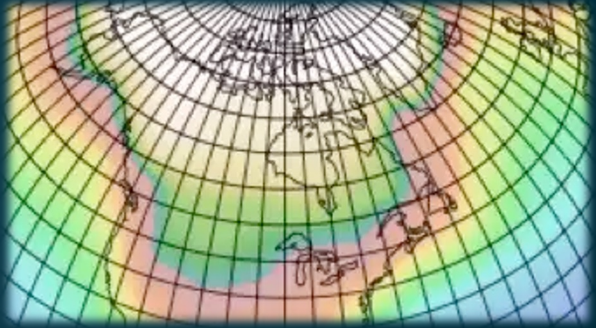

Advanced Light Source
7 Terabytes per hour

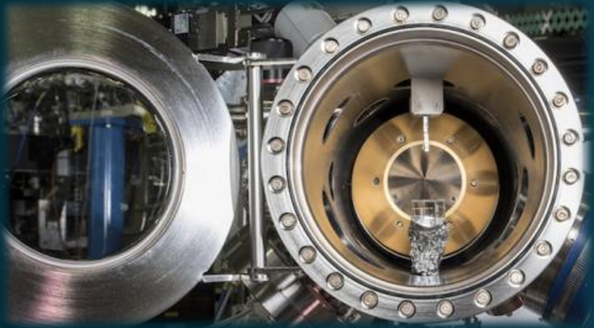

Genomics
Exabytes per

Data volumes are
growing far beyond

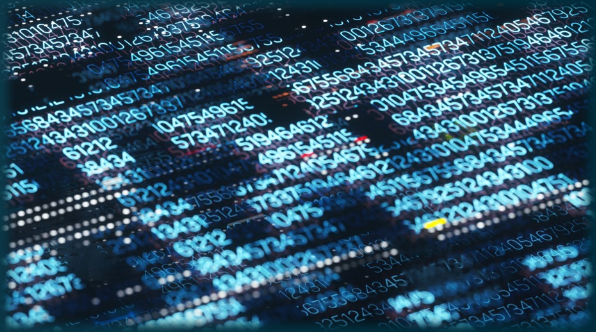

The need of leading edge HPC applications is far broader then ever

European Processor Initiative

MPSoC'24
EVIDEN

New era of the cooperation in compute and data continuum

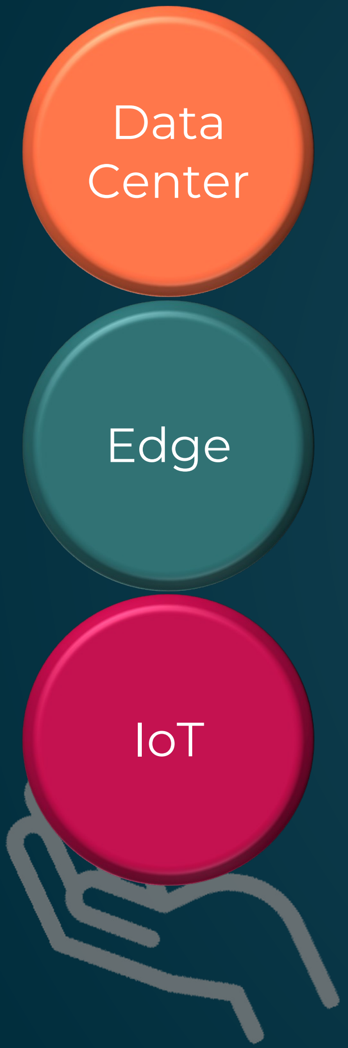

From
Hierarchy
To
Swarm

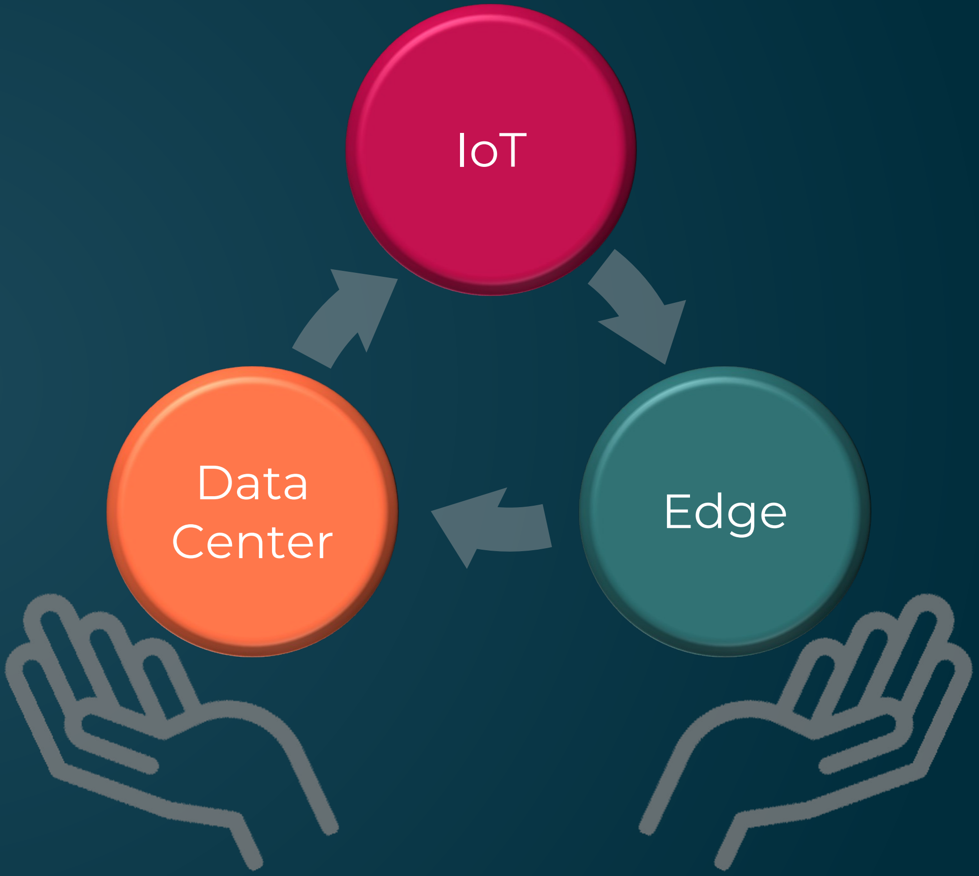

3 What’s at stake?

Tackle research, economic, industry, and societal challenges

Anticipate climate changes

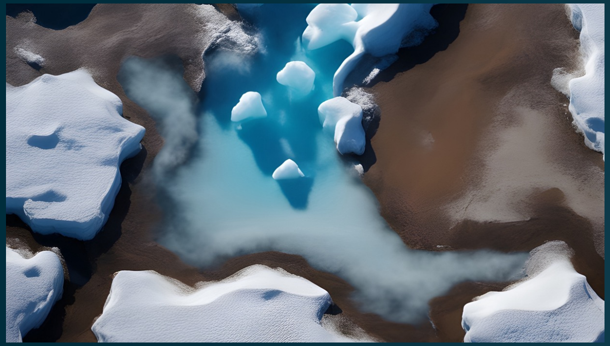

Study earthquakes

Control epidemics

Care for aging population

Innovate without limit

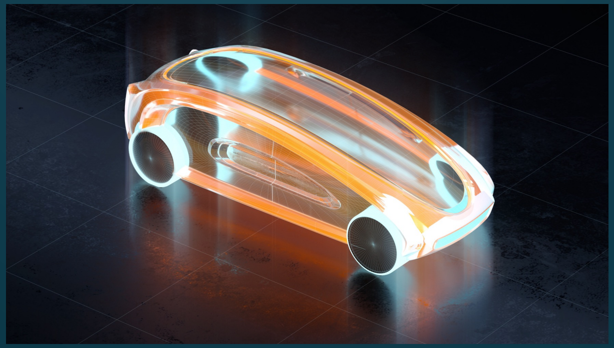

Secure energy resources

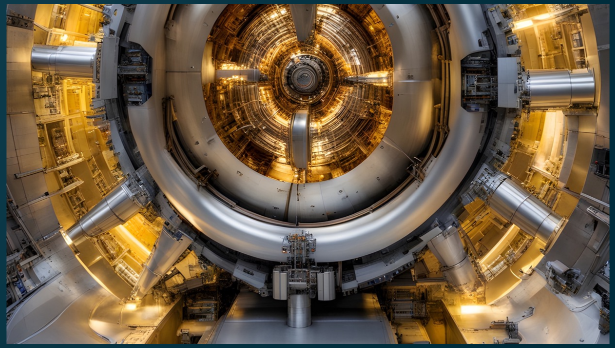

European Processor Initiative

MPSoC'24 EVIDEN

Who knows what's HPC and numerical simulation?

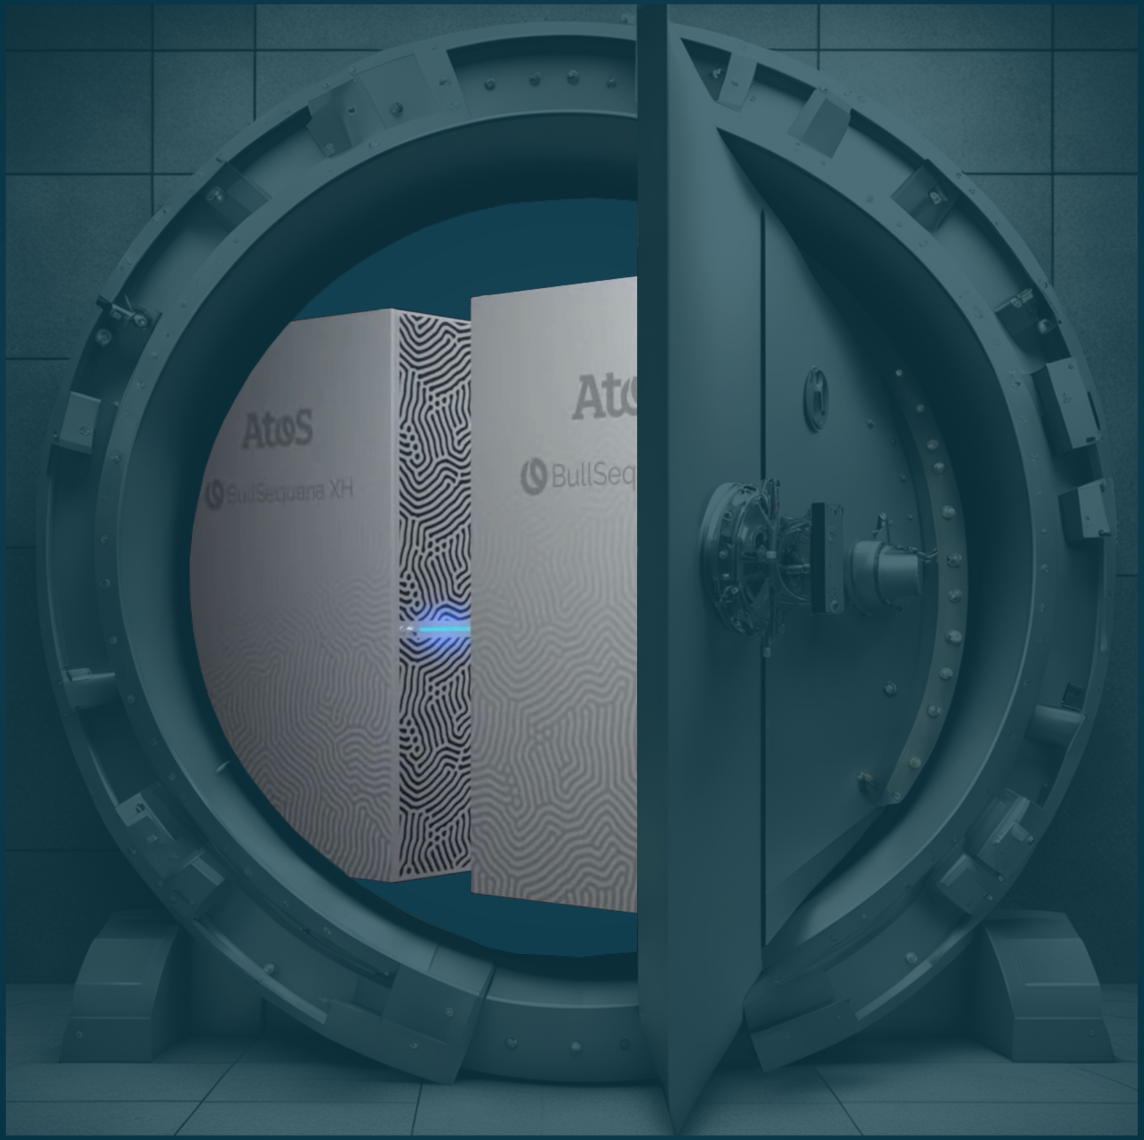

Keeping HPC so confidential is at risk

Reinvent HPC for a brilliant phygital future

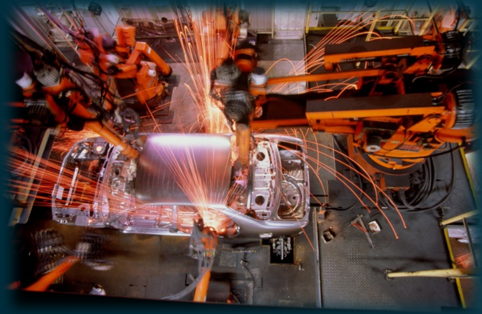

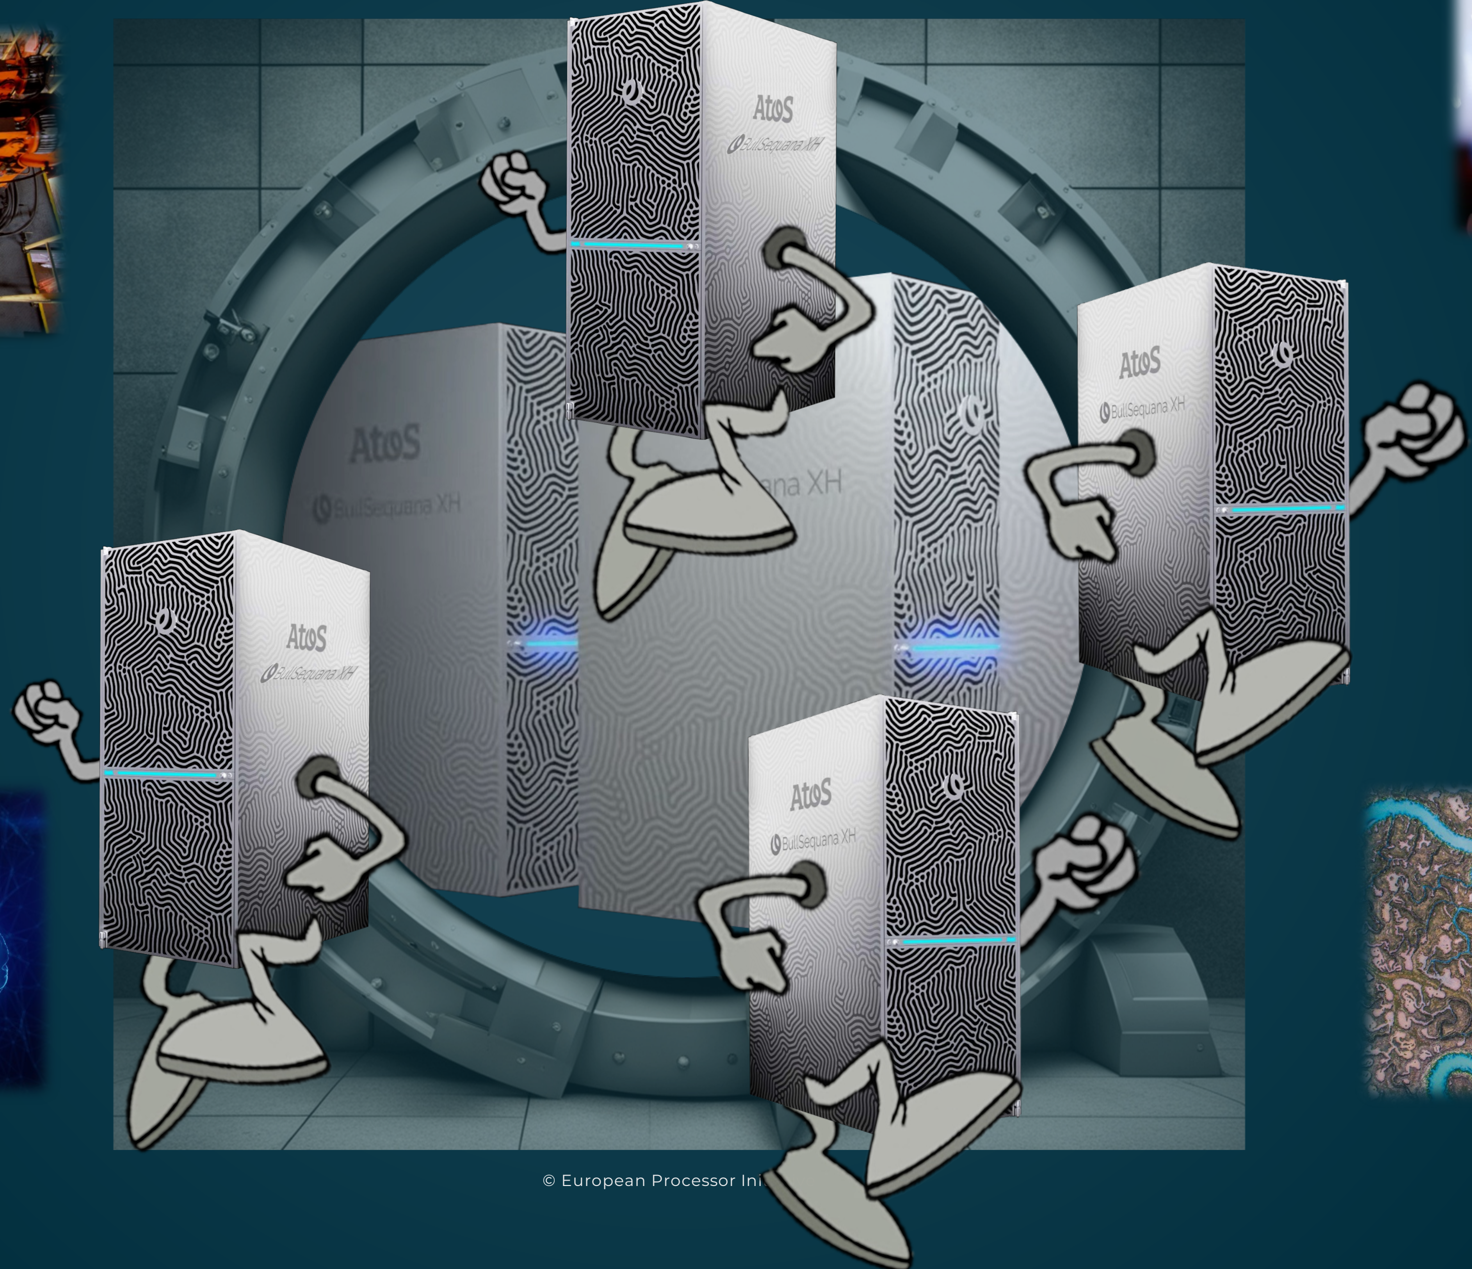

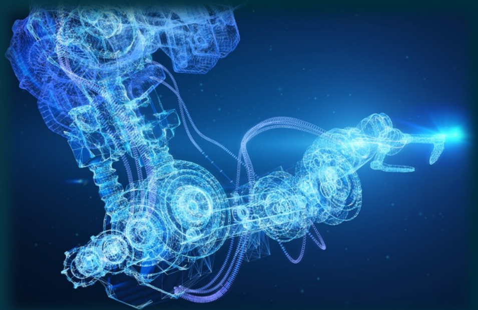

4 The threats

Gen Z moto: You Only Live Once

Leave an impact on the world

Align job and interests

Entrepreneur

Tech-savvy

Multi-taskers

Few focus

work-life balance

Digital interactive

Freedom

Make scientific education and jobs attractive again

Foreseen mismatch between electricity supply and demand

2020

91% Impact of energy on global CO_2 emissions

14% Energy that is green Electricity

7% Electricity weight of Information & Communication Technology

2030

+50% Increase in demand for electricity

20% Electricity weight of Information & Communication Technology

European Processor Initiative

epi

Energy, a sensitive subject

The paper that forced Timnit Gebru out of Google in 2020

Common carbon footprint benchmarks

in lbs of CO2 equivalent

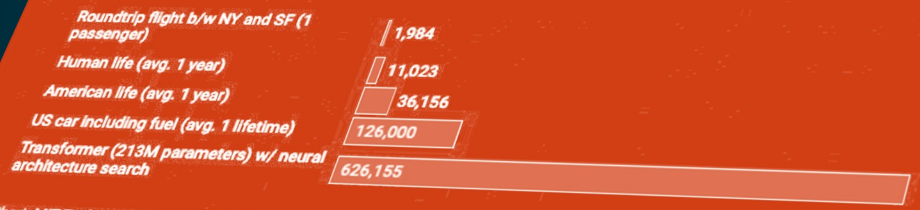

Chart: MIT Technology Review • Source: Strubell et al. • Created with Datawrapper

Strubell's study found that one language model with a particular type of "neural architecture search" (NAS) method would have produced the equivalent of 626,155 pounds (284 metric tons) of carbon dioxide—about the lifetime output of five average American cars. A version of Google's language model, BERT, which underpins the company's search engine, produced 1,438 pounds of CO2 equivalent in Strubell's estimate—nearly the same as a roundtrip flight between New York City and San Francisco.

The estimated costs of training a model

	Date of original paper	Energy consumption (kWh)	Carbon footprint (lbs of CO2e)	Cloud compute cost (USD)
Transformer (65M parameters)	Jun, 2017	27	26	$41-$140
Transformer (213M parameters)	Jun, 2017	201	192	$289-$981
ELMo	Feb, 2018	275	262	$433-$1,472
BERT (110M parameters)	Oct, 2018	1,507	1,438	$3,751-$12,571
Transformer (213M parameters) w/ neural architecture search	Jan, 2019	656,347	626,155	$942,973-$3,201,722
GPT-2	Feb, 2019	-	-	$12,902-$43,008

Note: Because of a lack of power draw data on GPT-2's training hardware, the researchers weren't able to calculate its carbon footprint.

Table: MIT Technology Review • Source: Strubell et al. • Created with Datawrapper

European Processor Initiative

MPSoC'24 EVIDEN

Is next step owning its electricity production plant ?

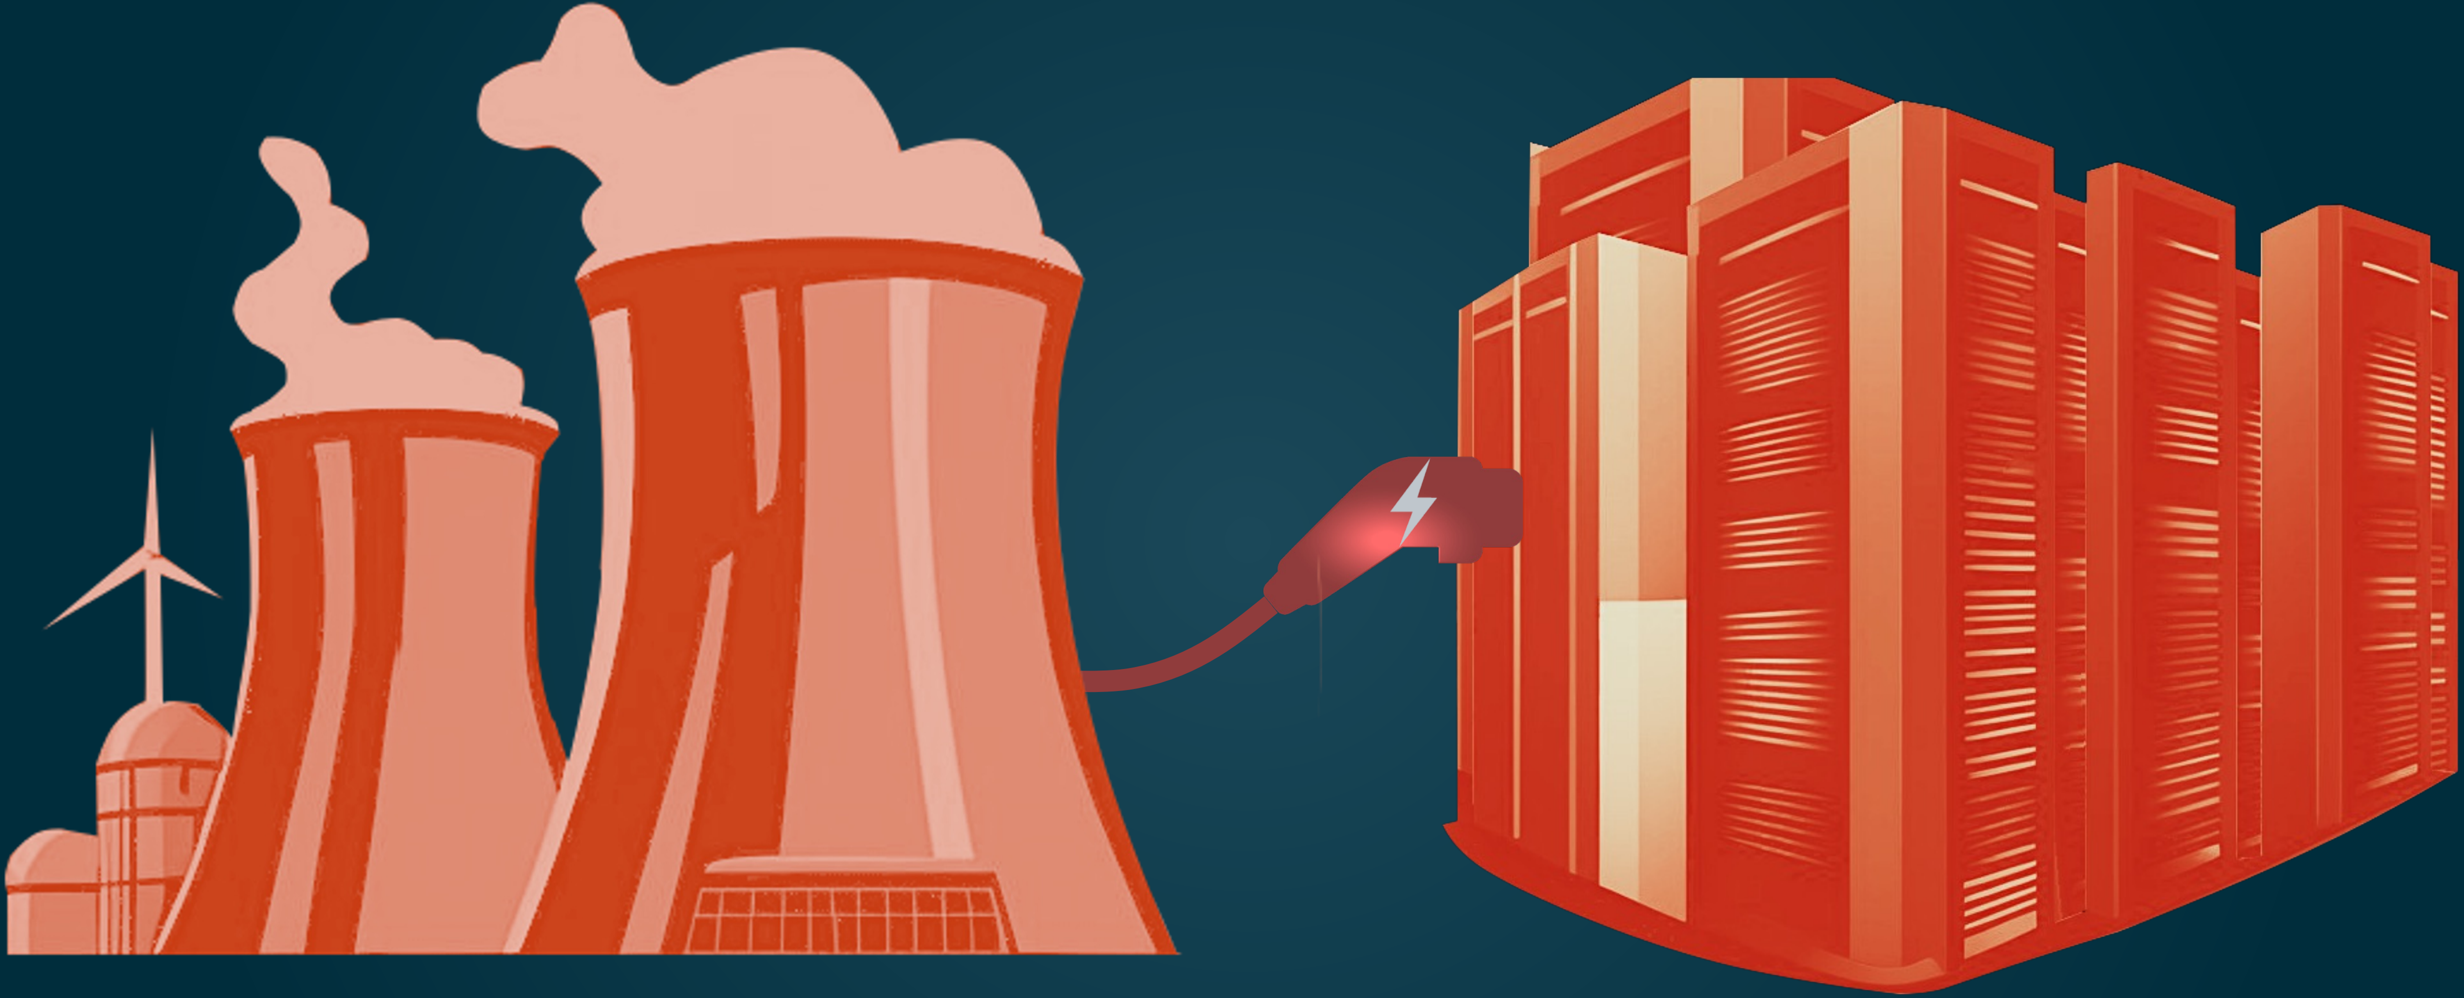

Example: Amazon data center (960 MW) directly connected to Susquehanna Steam Electric Station (2.5 GW)

5 The weaknesses

The most secret number of every HPC datacenter

Percentage of data generated, archived and never exploited

- Compute time if often free of charge for users
- More application runs than time and tooling to explore the results
- Applications not behaving well and not stopped earlier enough
- Mis of old data garbage collectors
- …

European Processor Initiative

epi

MPSoC'24

EVIDEN

From Flops to Bytes per Flop

Direct solver
Dense Matrix
$O(N^3)$

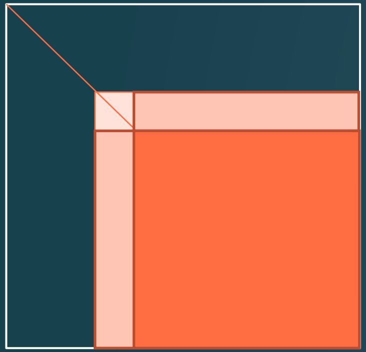

Linpack benchmark

Direct solver
Banded Matrix
$O(N^{2.7})$

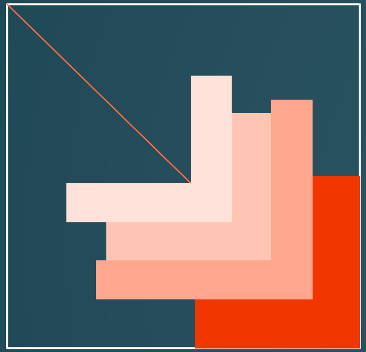

Iterative solver (CG)
Sparse Matrix
$O(N^{1.5})$

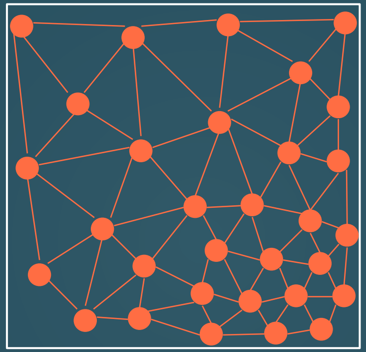

HPCG benchmark

Iterative solver (AMG)
Sparse Matrix
$O(N.log(N))$

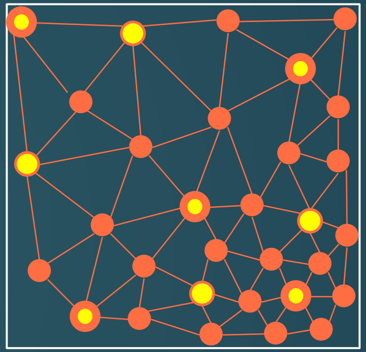

Spectral solver (FFT)
Regular grid
$O(N.log(N))$

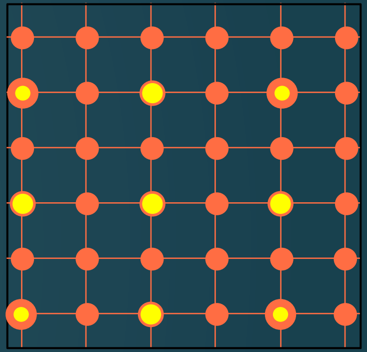

Dense blocks

n^2 memory access - n^3 computation

Core intensive

Sparse data

n memory access - n computation

Memory intensive

Lack of memory throughput

Top 5 2023 06	Tflop/s peak	Linpack efficiency	HPCG efficiency
Frontier	1 679 819	71%	0,8%
Fugaku	537 212	82%	3,0%
LUMI	428 704	72%	0,8%
Leonardo	304 466	78%	1,0%
Summit	200 795	74%	1,5%

	FP64 (Tflop/s)	Bandwidth (TW*/s)	BW / FP64
NVIDIA A100	9.70	0.485	5%
Millan 64C 2Ghz	2.05	0.051	2.5%

W* = 4 bytes

European Processor Initiative

MPSoC'24
EVIDEN

So energy expensive memory throughput

Operation		Picojoules per operation		
		45 nm	7 nm	Ratio
Add. (IEEE)	FP 16	0.4	0.16	2.5
	FP 32	0.9	0.38	2.4
Mult. (IEEE)	FP 16	1.1	0.34	3.2
	FP 32	3.7	1.31	2.8
SRAM 64b access	32 KB	20	**8.5**	2.4
	1 MB	100	**14**	7.1
DRAM 64b access		Circa 45 nm	Circa 7nm	
	DDR 3/4	1 300	**1 300**	1.0
	HBM2	**250-450**		

Source: Ten Lessons From Three Generations Shaped Google's TPUv4i

7nm - Picojoules per operation (log_{10} scale)

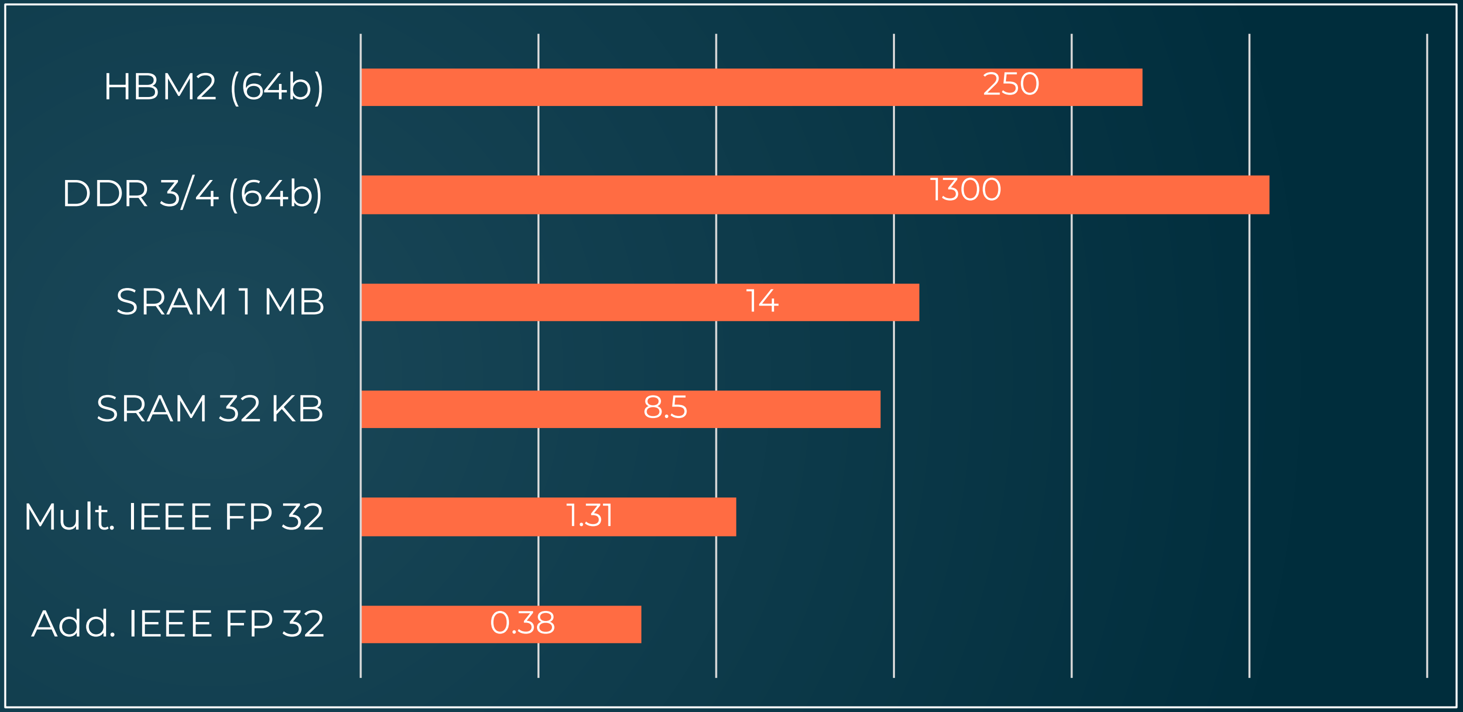

The uncertain future of monolithic chips

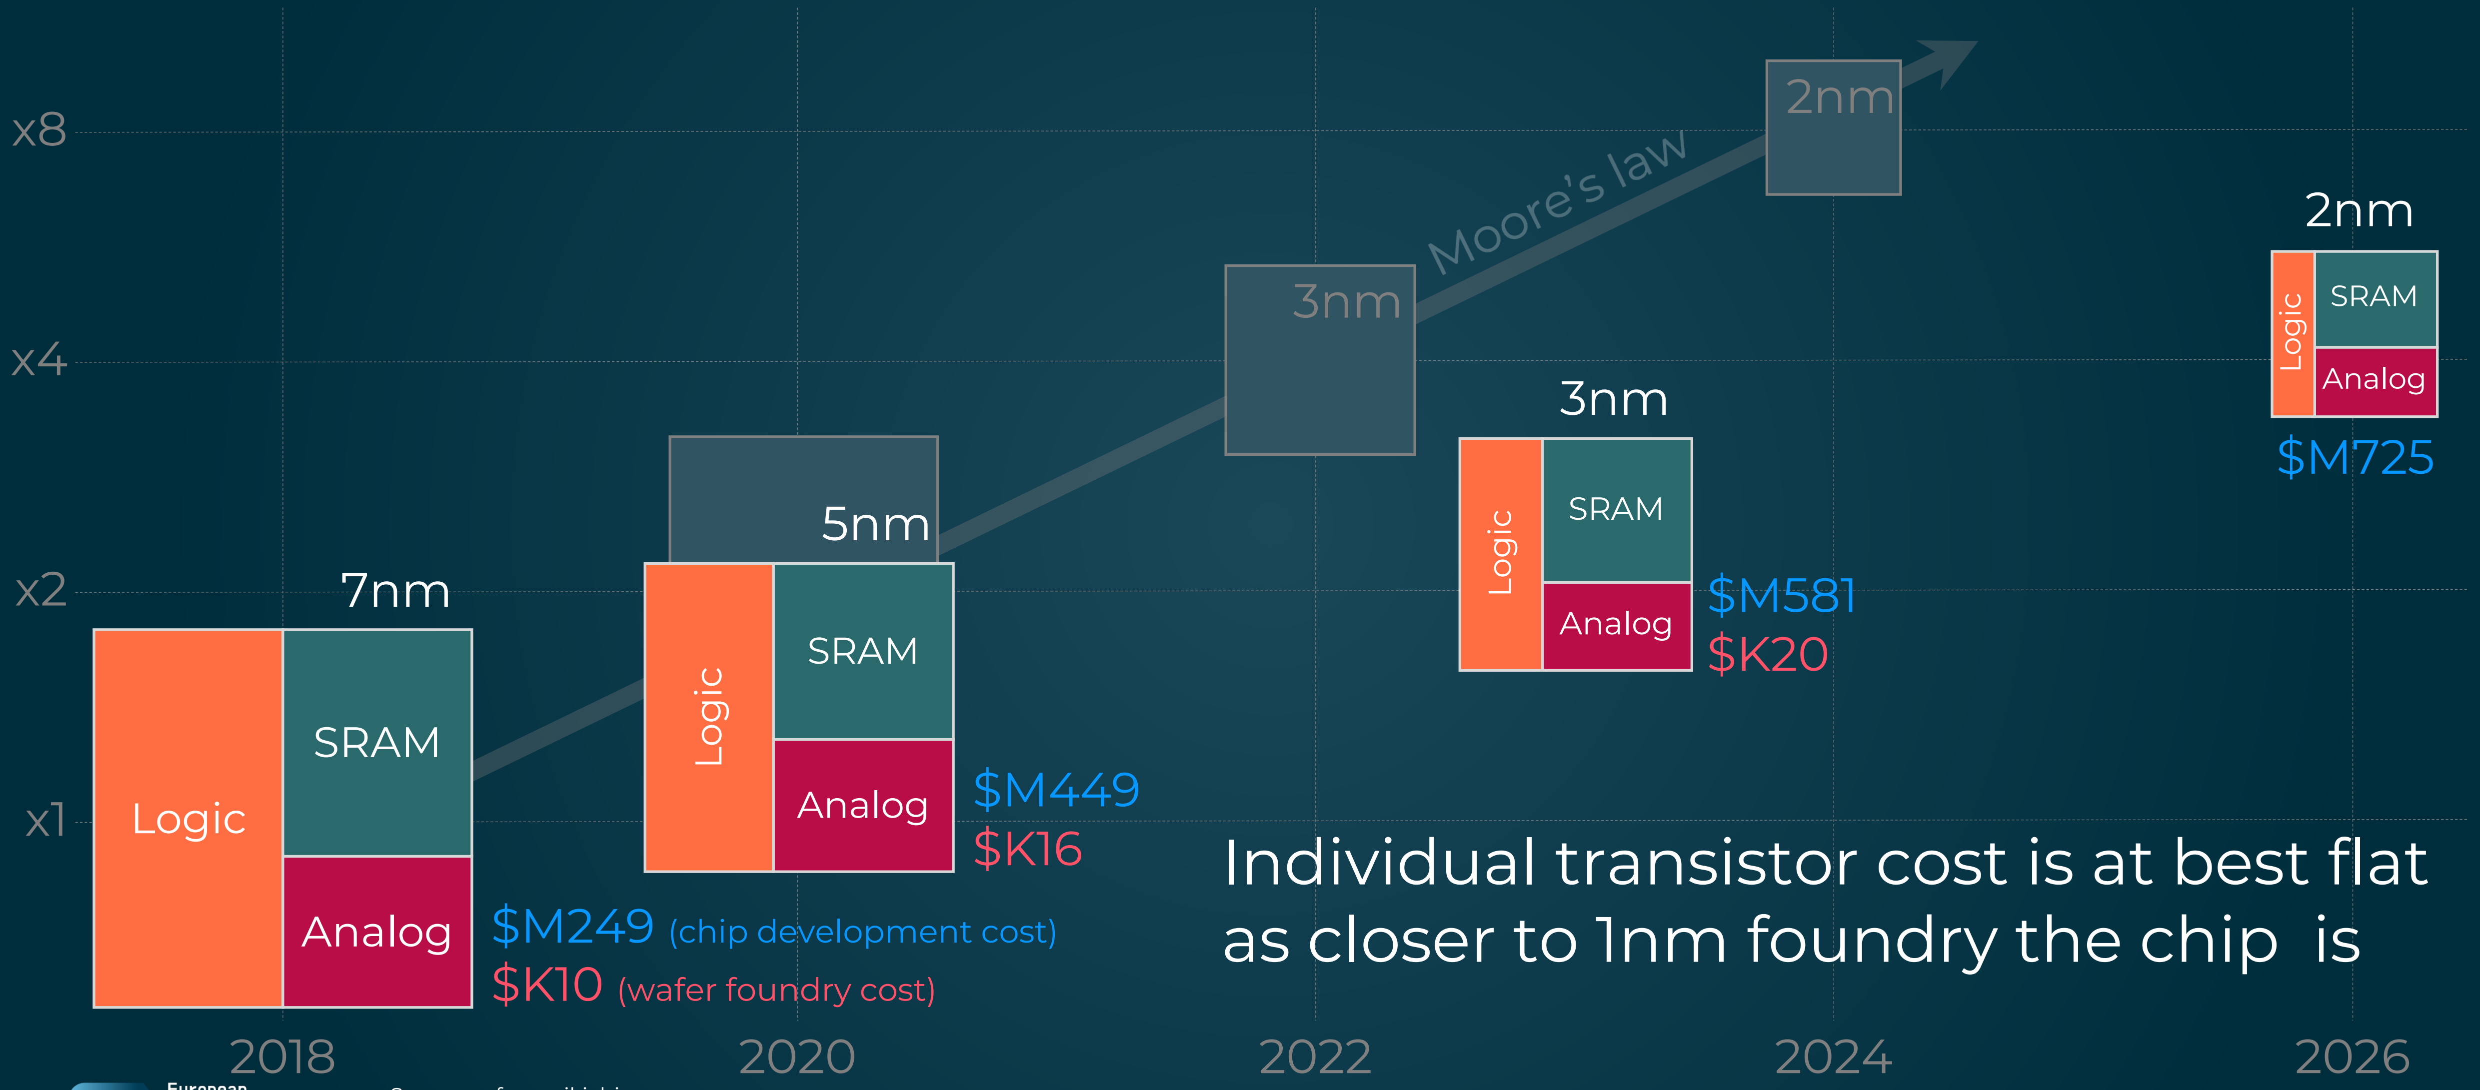

Individual transistor cost is at best flat as closer to 1nm foundry the chip is

Sources: fuse.wikichip.org
granitefirm.com (Andy Lin's blog)
TSMC
IBS, July 2022
Denis Dutoit (CEA LIST)

the irresistible rise in energy consumption

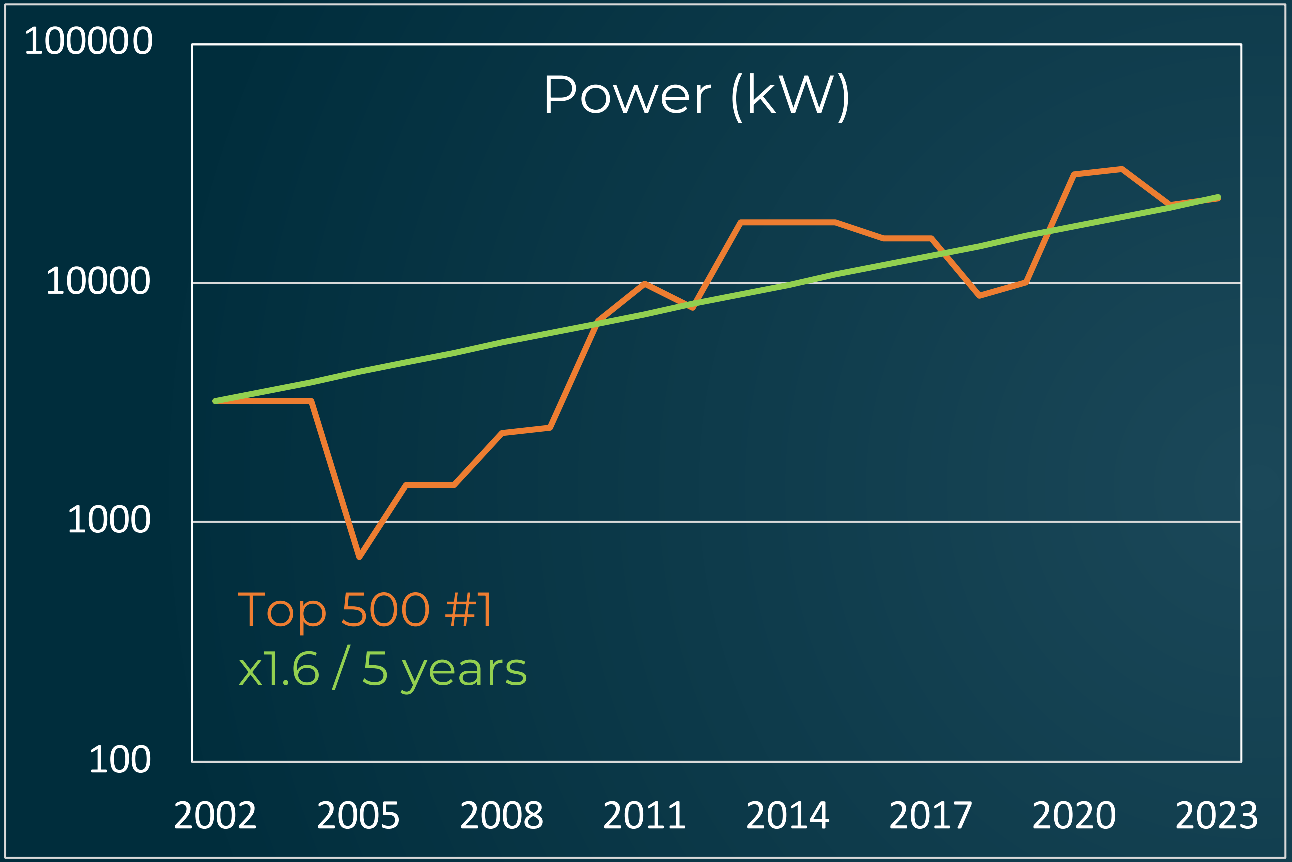

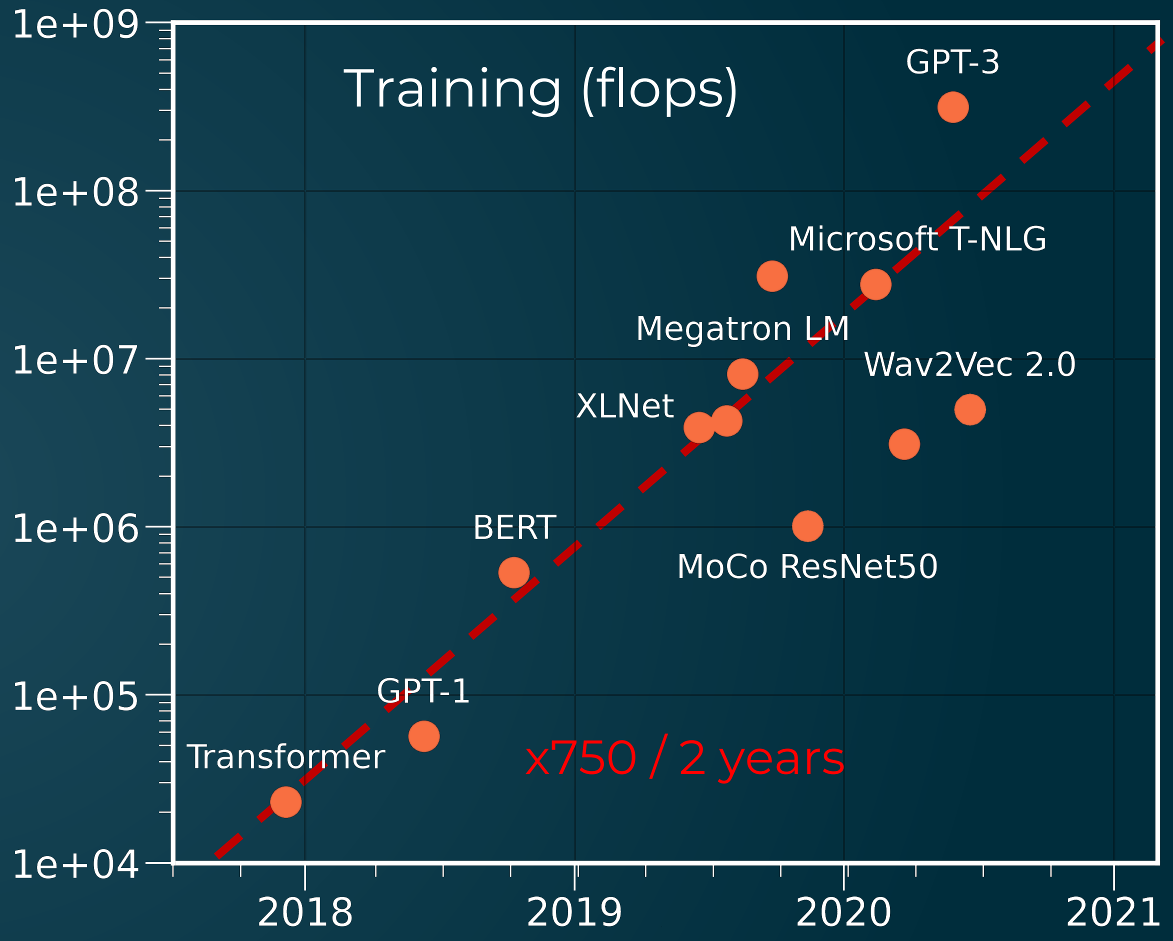

Source: medium.com (AI and Memory Wall)

Gate expensive FP64 computation not so useful for AI

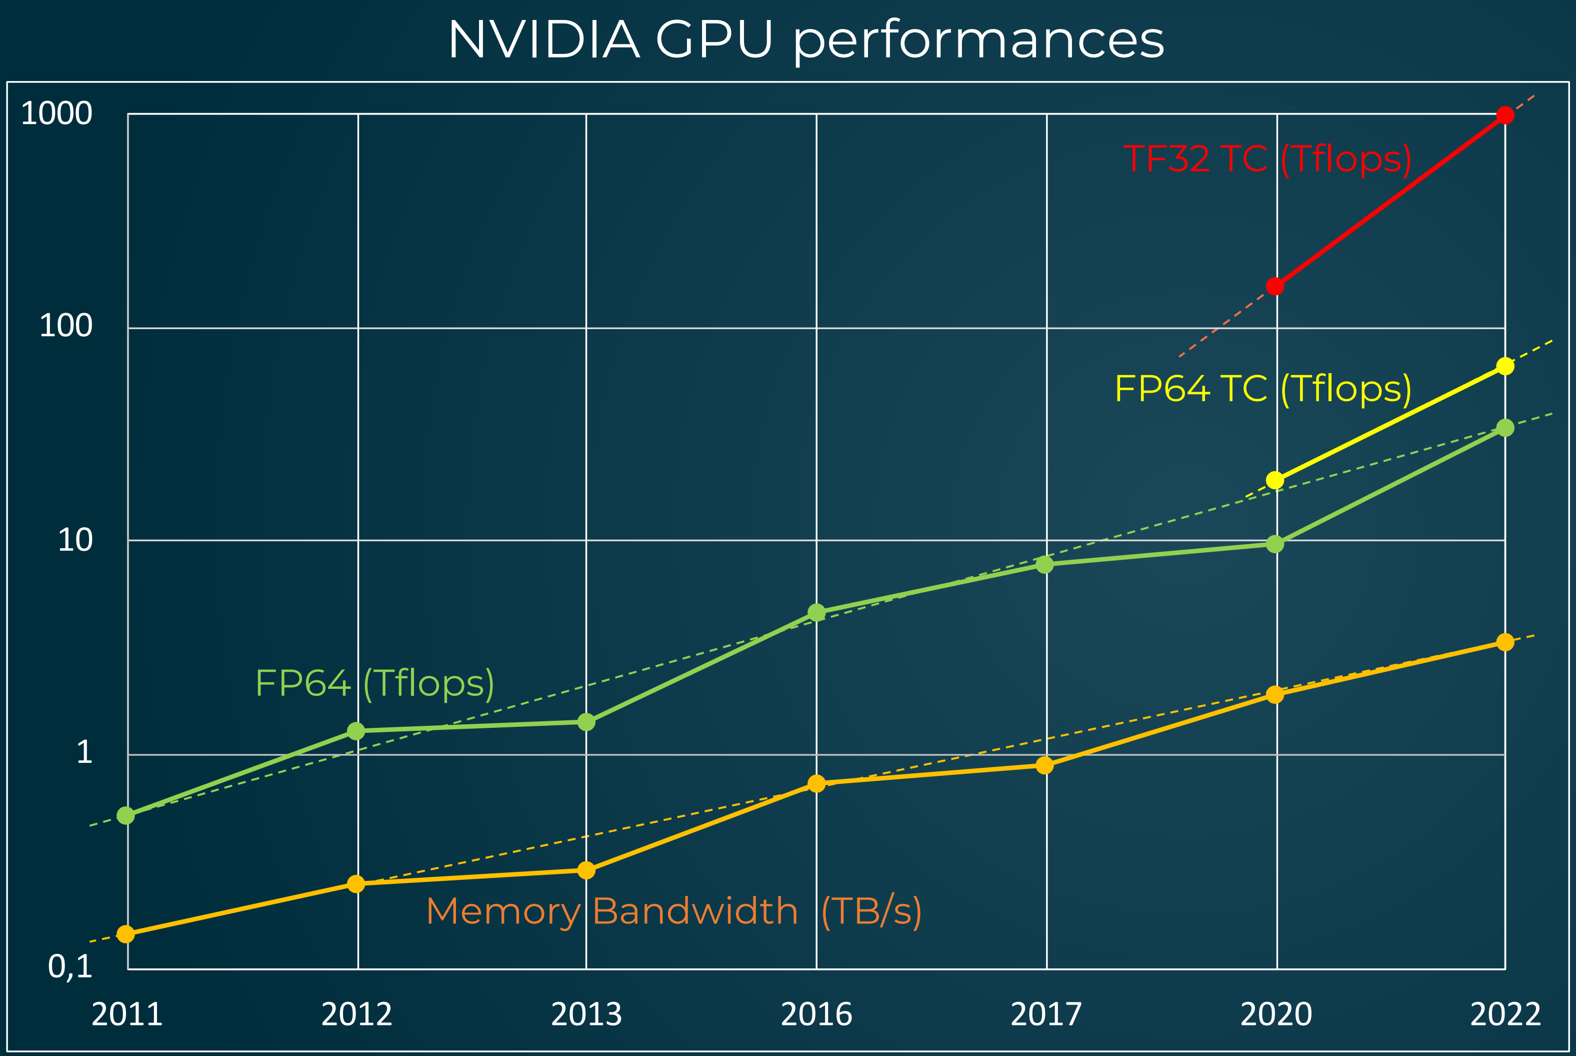

More expensive than ever Eflops Linpack in the coming years?

European Processor Initiative

MPSoC'24
EVIDEN

Flops are isolated, not IOs that are polluted by workflow concurrencies

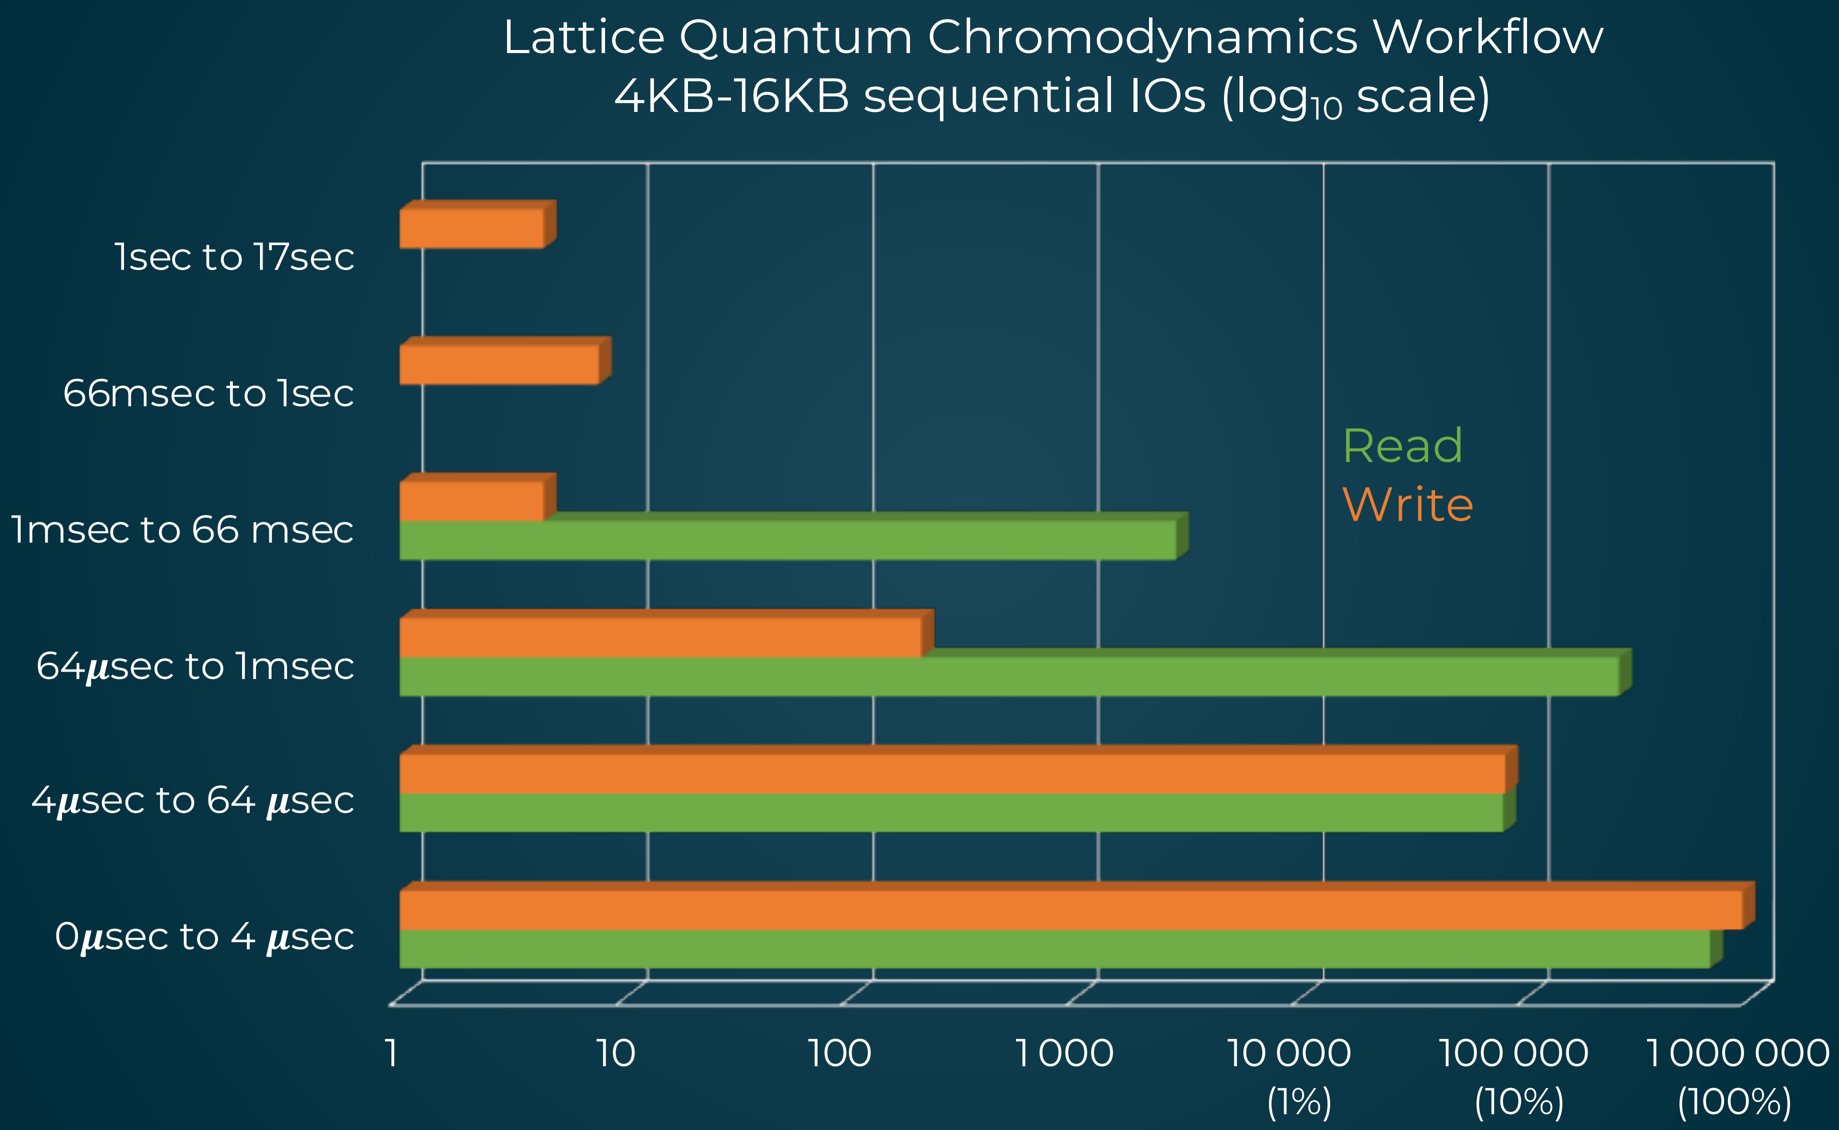

European Processor Initiative

MPSoC'24

EVIDEN

European Processor Initiative

6 Cloud convergence

Four flavors of Cloud Computing

Process	Business Process Outsourcing (BPO)	process
Consume	Software as a Service (SaaS)	Application Data
Build	Platform as a Service (PaaS)	Middleware OS
Host	Infrastructure as a Service (IaaS)	Compute Storage Network

One Cloud doesn't fit all

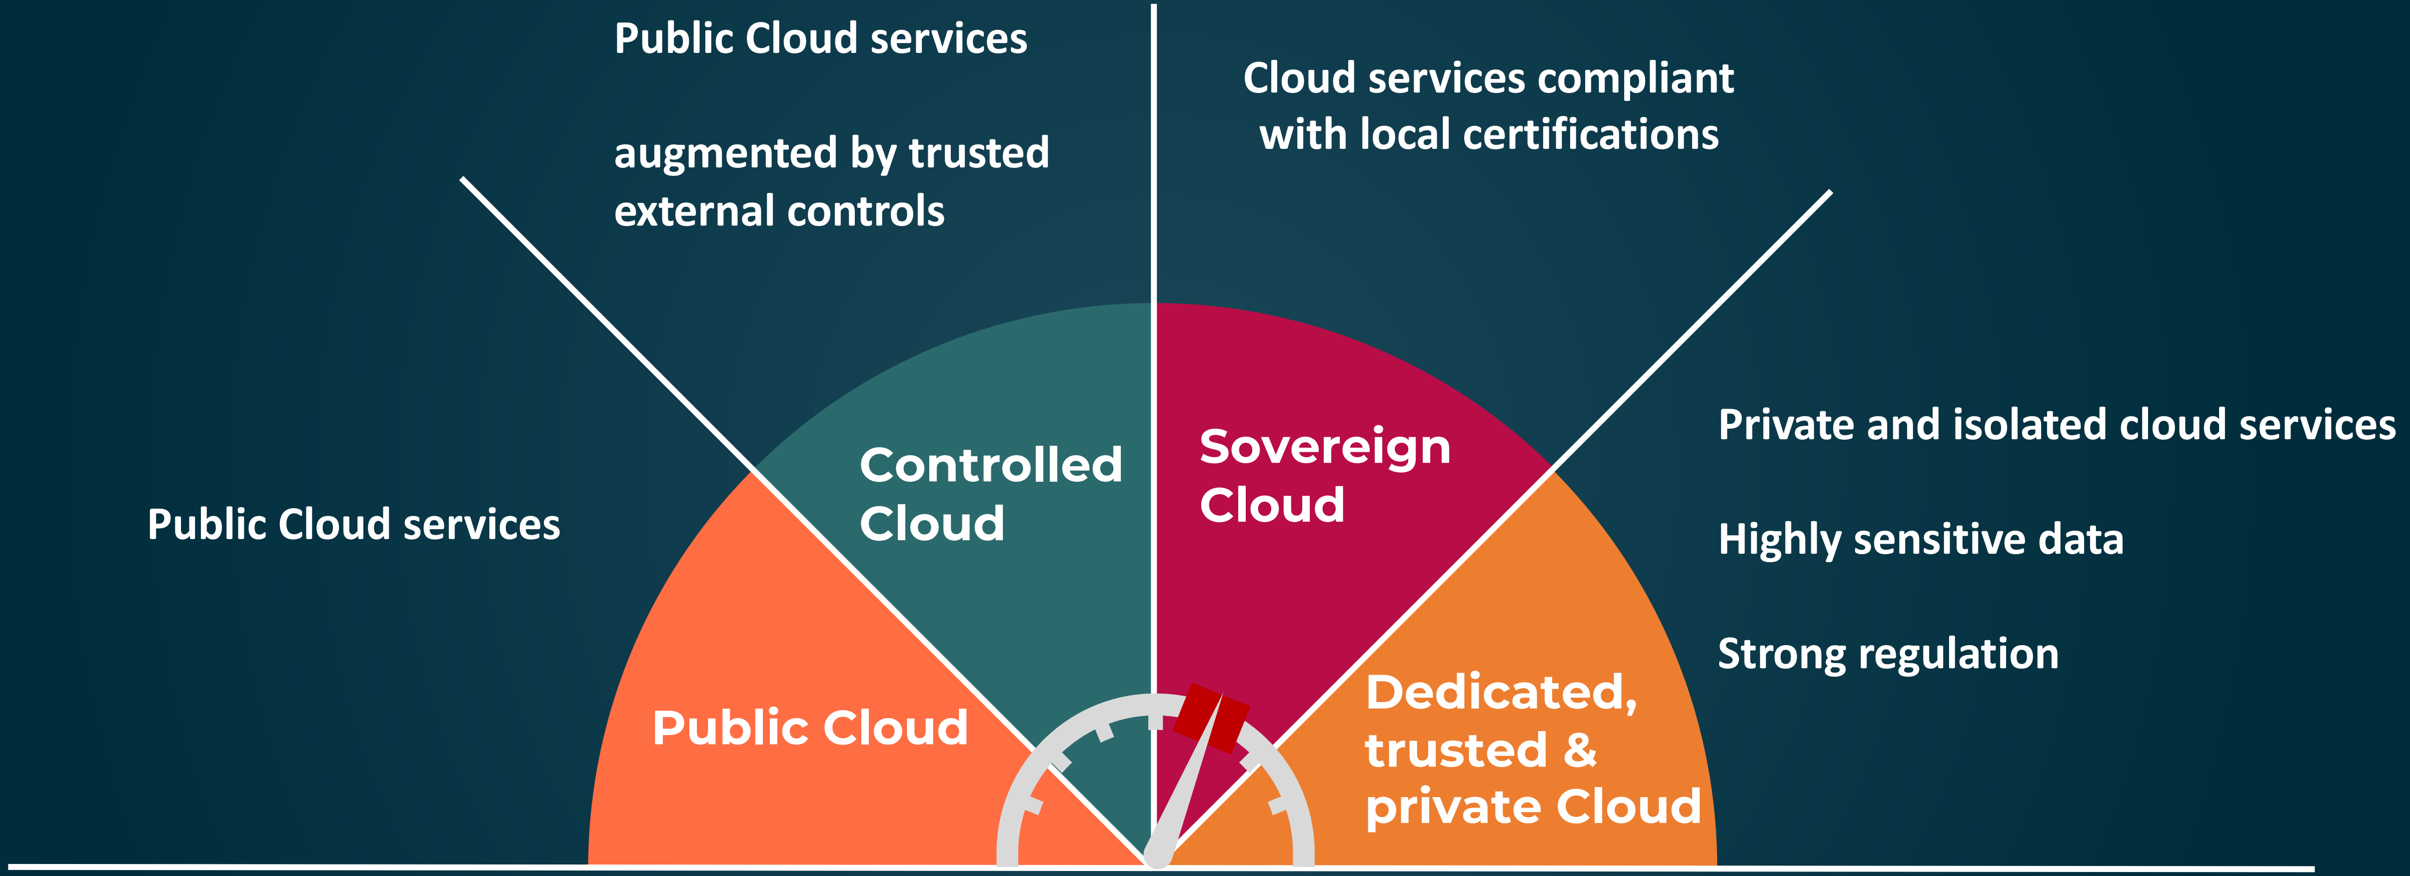

Sovereignty
Data sensitivity

Cloud computing tradeoffs

Benefits

- Pays as you use (CAPEX to OPEX)
- Easily scale resources up or down
- Faster time to market
- Access resources and applications from anywhere
- Advanced security
- Data loss prevention
- Seamless collaboration & data sharing

Limitations

- Risk of vendor lock-in
- Unforeseen costs, might be more expensive
- Relies on an internet connection
- Less control over underlying cloud infrastructure
- Concerns about security risks like data privacy and online threats
- Integration complexity with existing systems

HPC has started to jump in cloud arena

Some signals to move one step further

New
mainstream
for
technological
influence

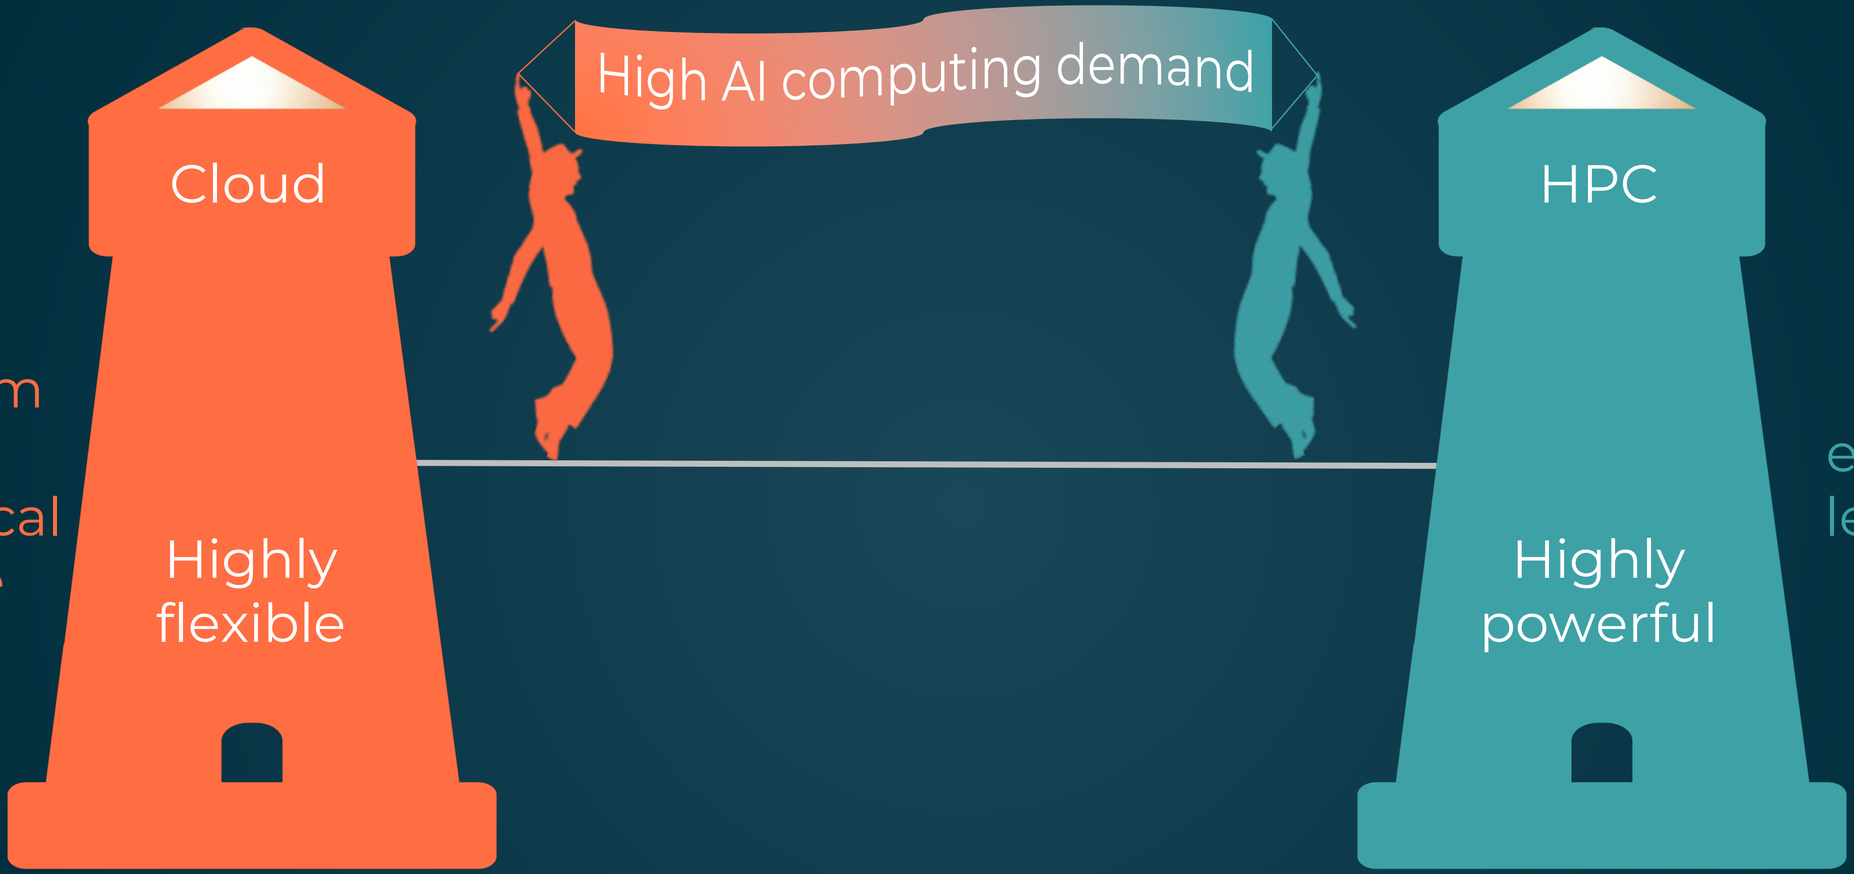

Small
and
endothermic
leading edge
business

Rich environment of container services
on top of a microservice architecture

Number crunching software optimized
for underlined hardware configuration

Fiscal year 2022, R&D investments ($Billions)

Sources: statista.com
HPE Fiscal Year 2022 Form 10-K

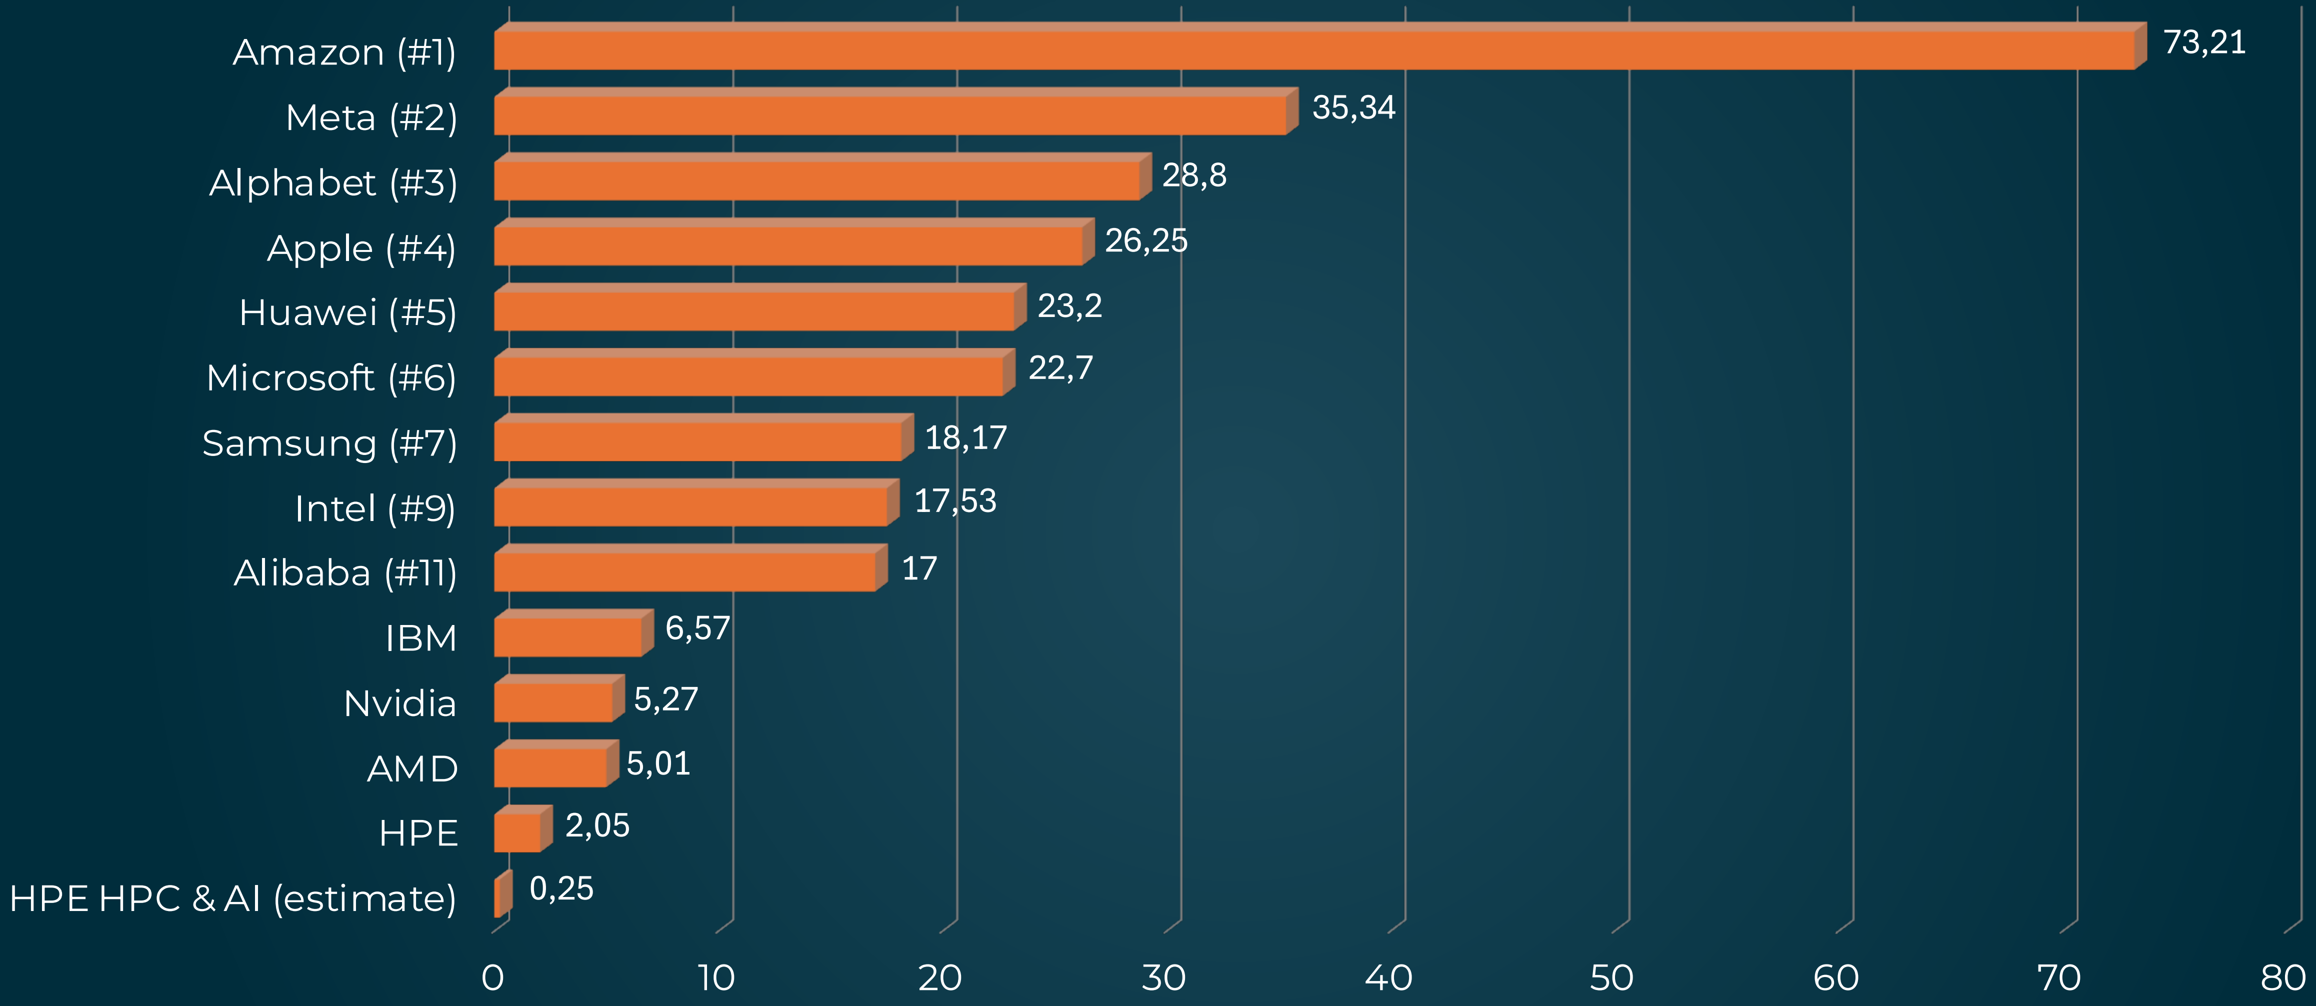

Ultra domination of top Cloud providers

European Processor Initiative

MPSoC'24
EVIDEN

One Cloud still doesn't fit all HPC

Generic Cloud / Commodity HPC

Expertise focuses on key business proposition

HPC friendly container management

Swarm data repositories

Specific Cloud / Tailor-made HPC

- Capacity computing
- Bursting workload
- Episodic large-scale computing
- Use case burn-in

- Large-scale computing
- Defense missions

ISV licenses terms & operations

Government investments & private sector partnerships

Gen Z view of computing

E2E innovation & co-design

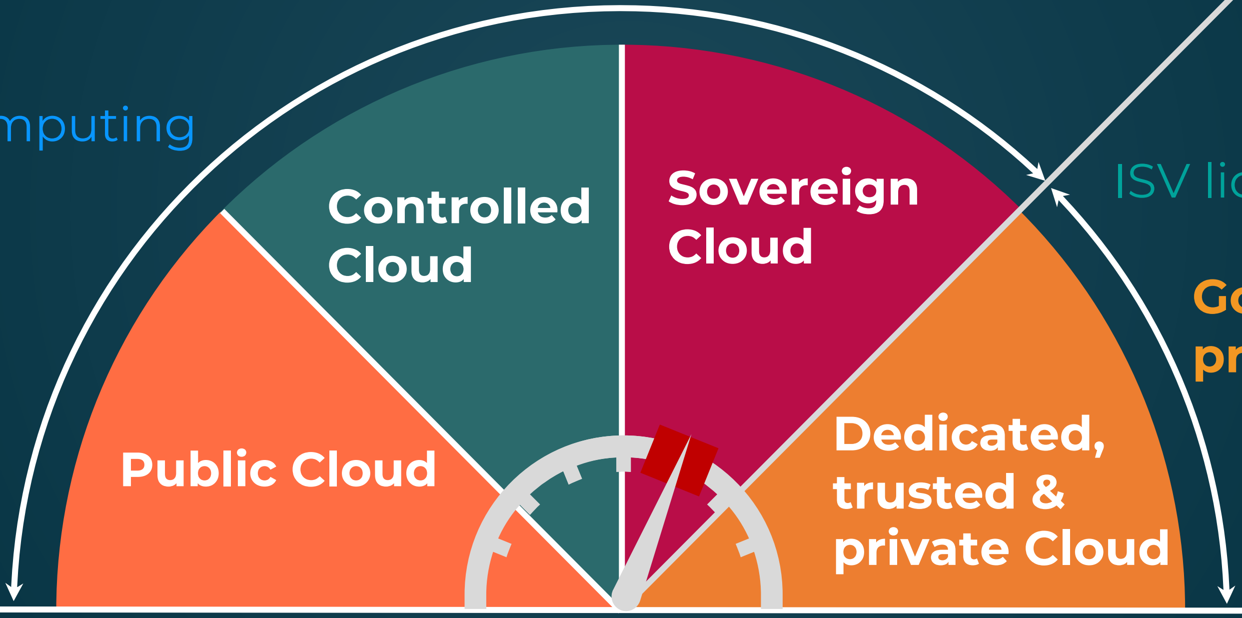

Parallel file system as a service, a cloud opportunity

Data moving through the continuum

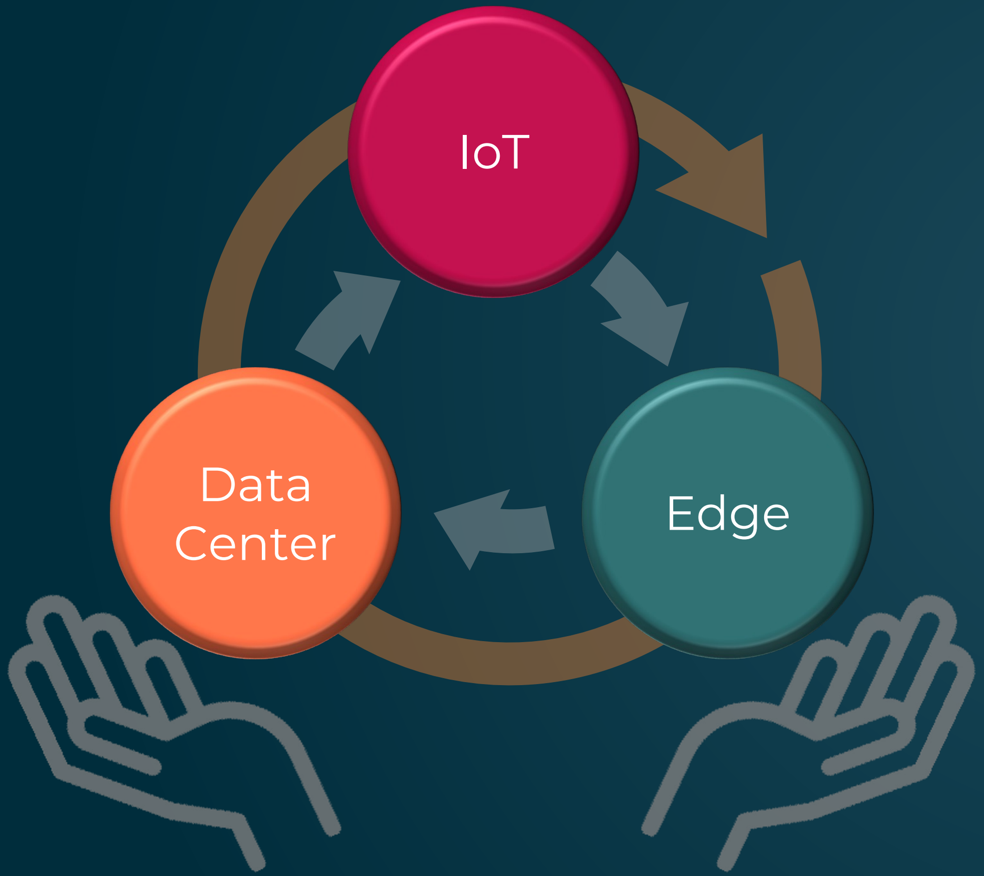

Application workflow pressure on file system

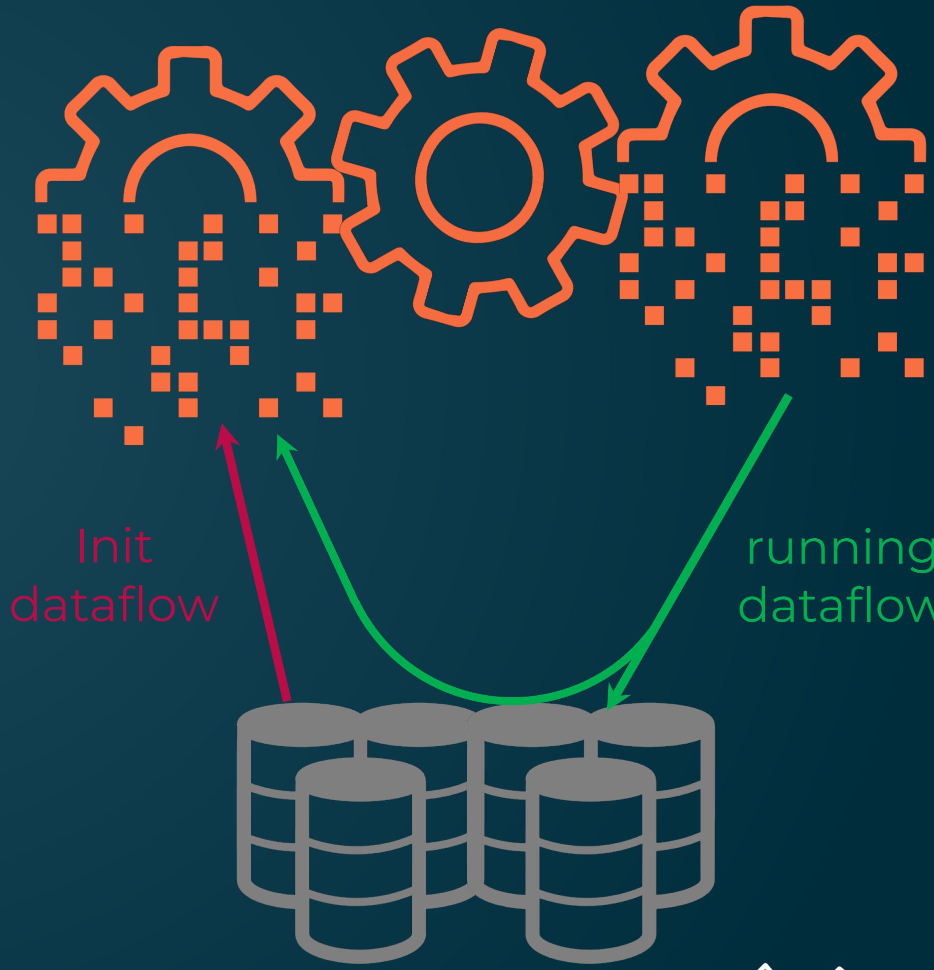

Polymorphic & Disaggregated IOs

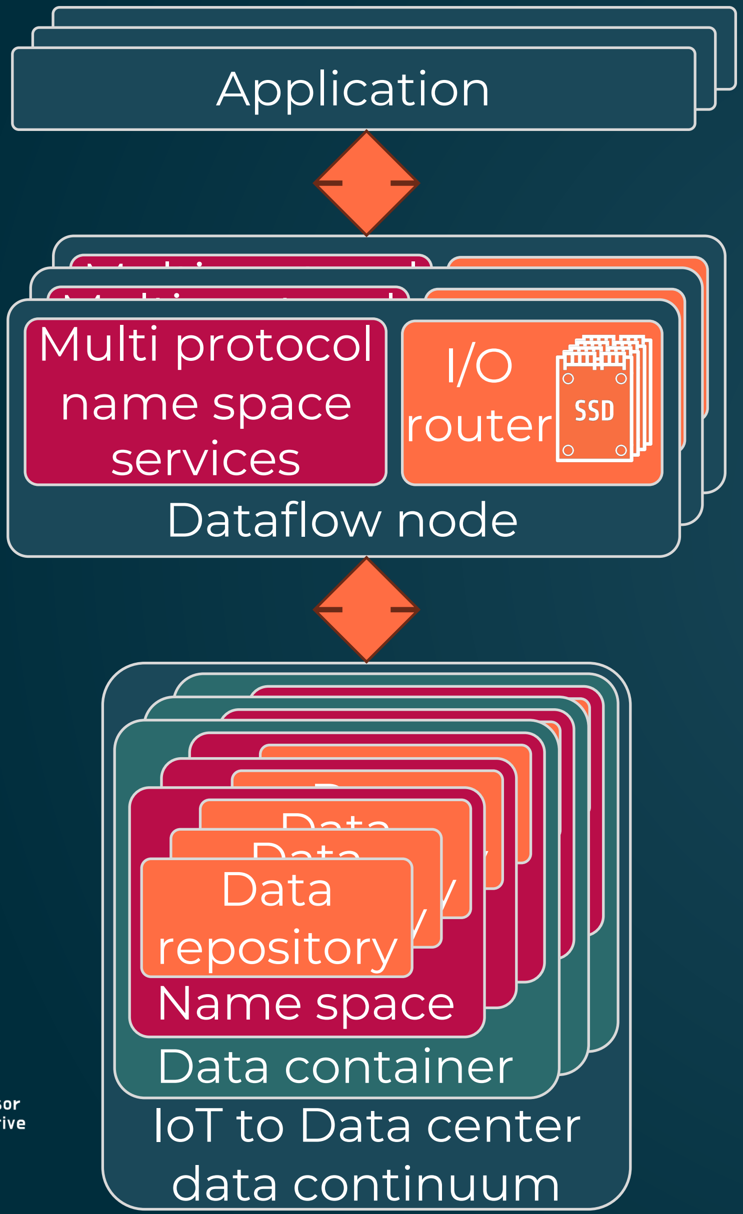

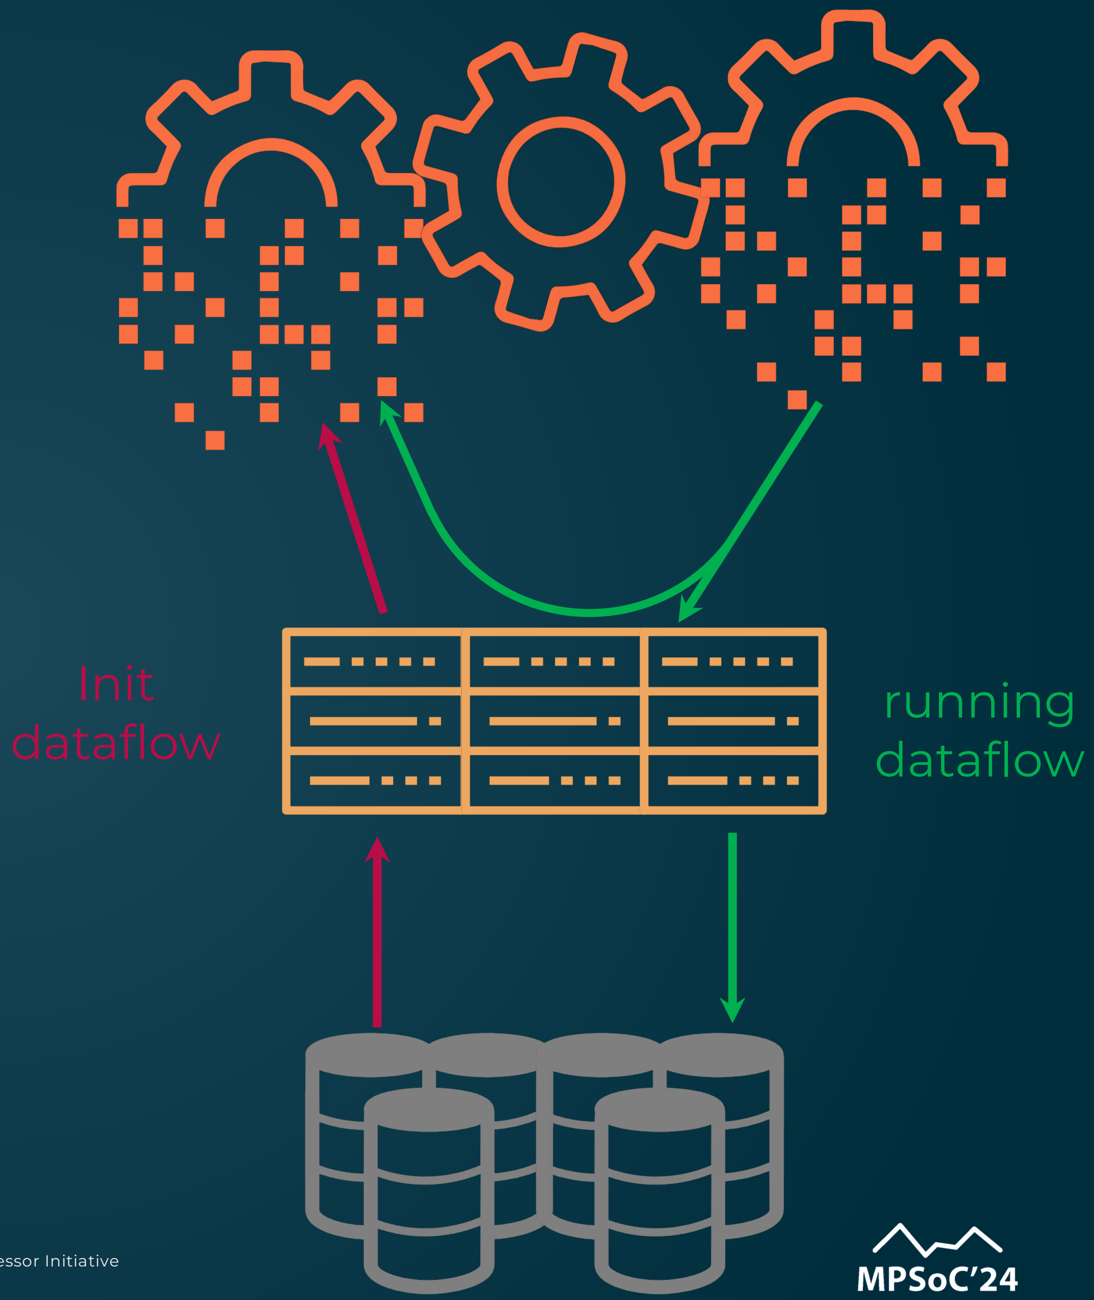

Dataflow nodes in action: write elapsed time divided by 2

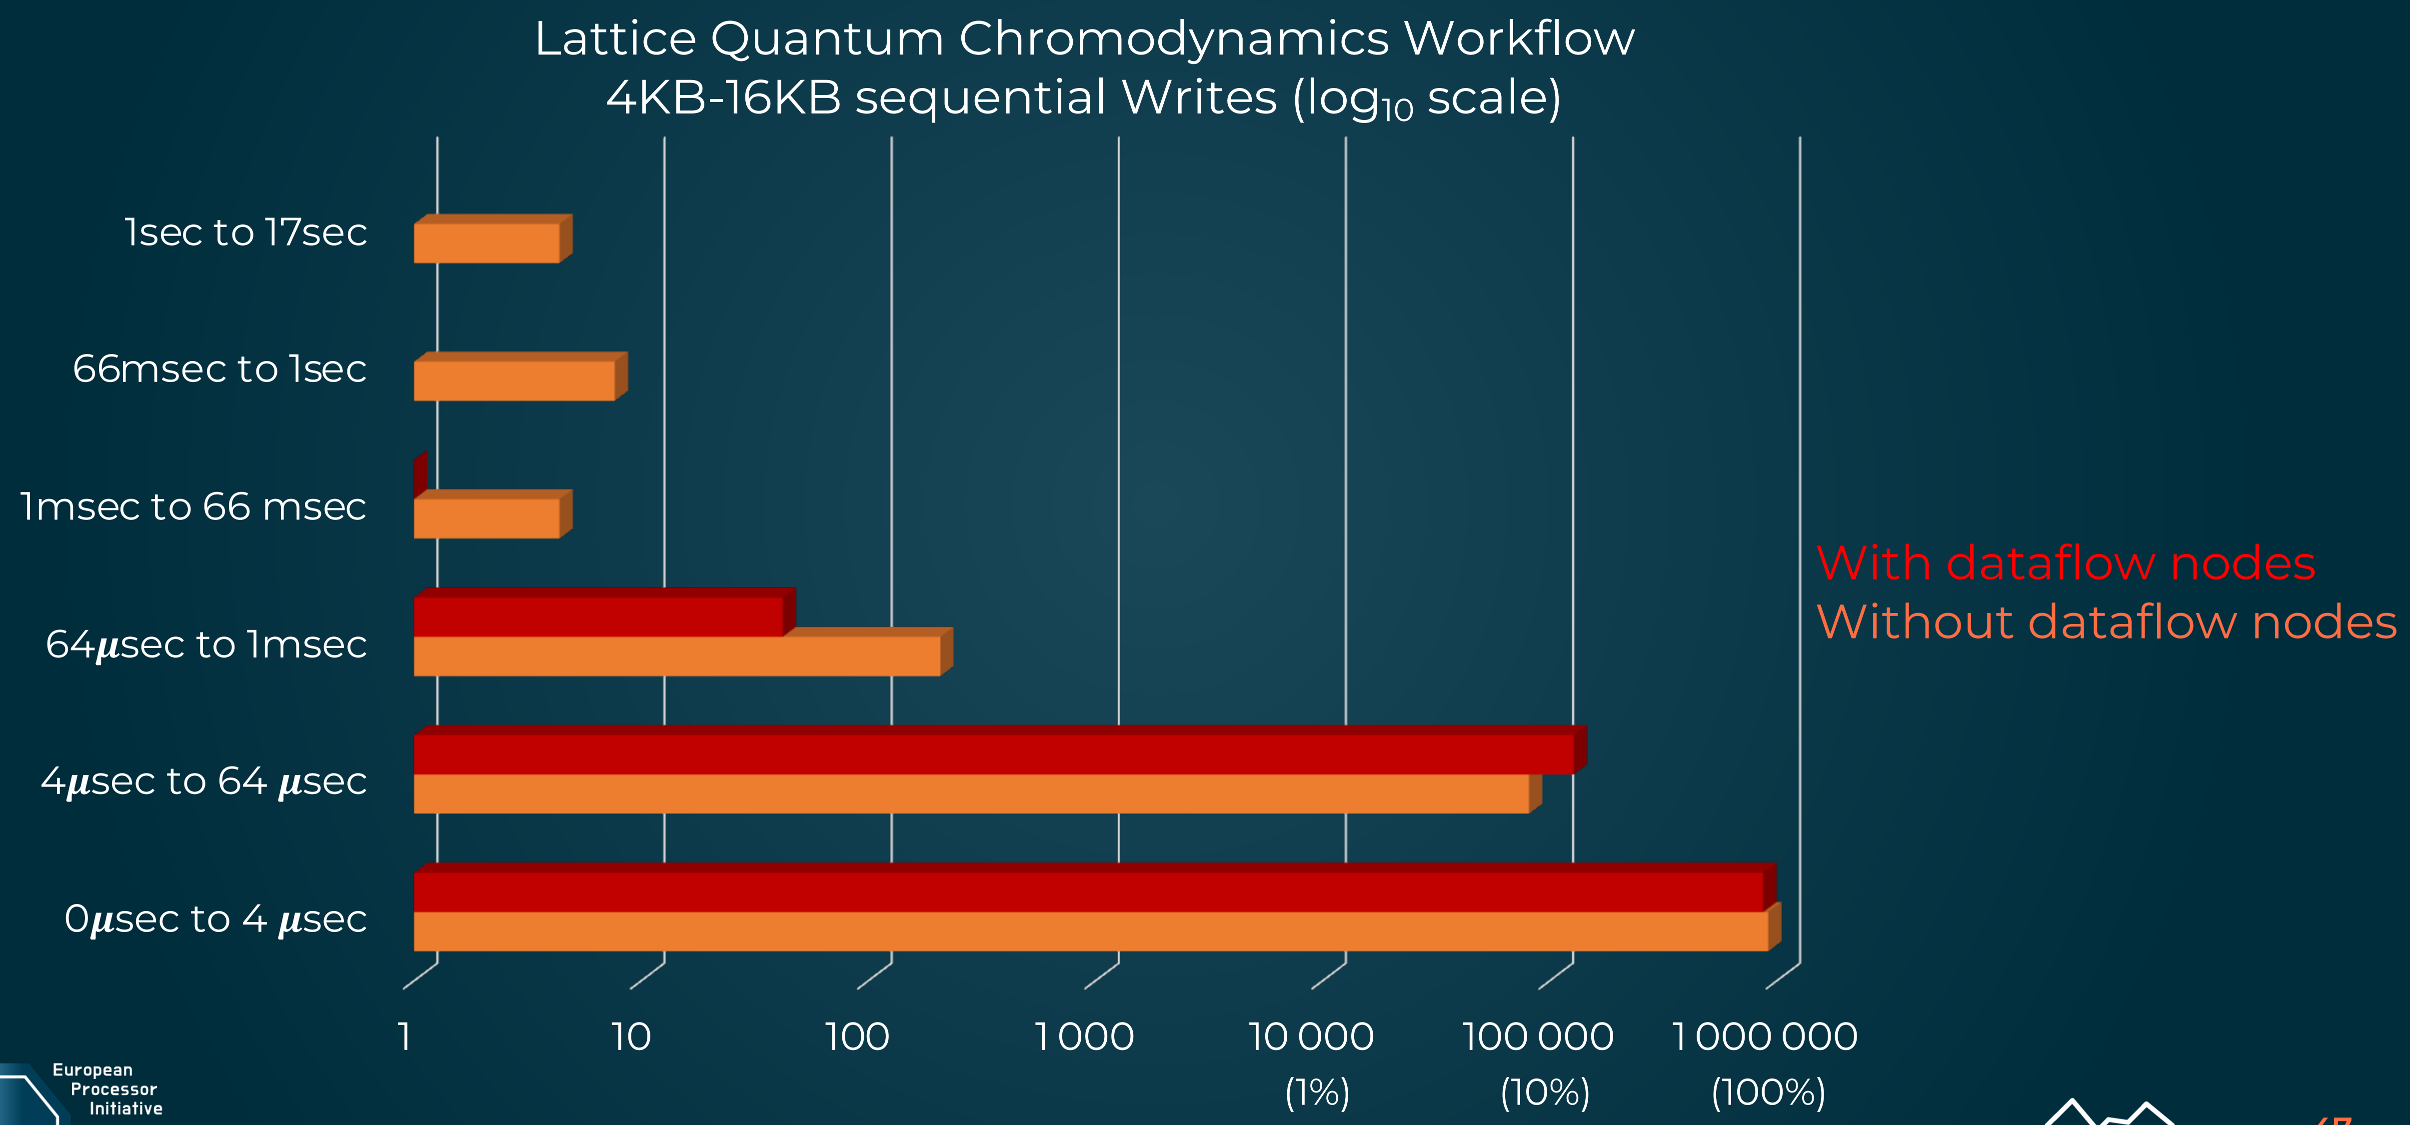

European Processor Initiative

MPSoC'24
EVIDEN

7 Making HPC future a reality

The basic ingredients

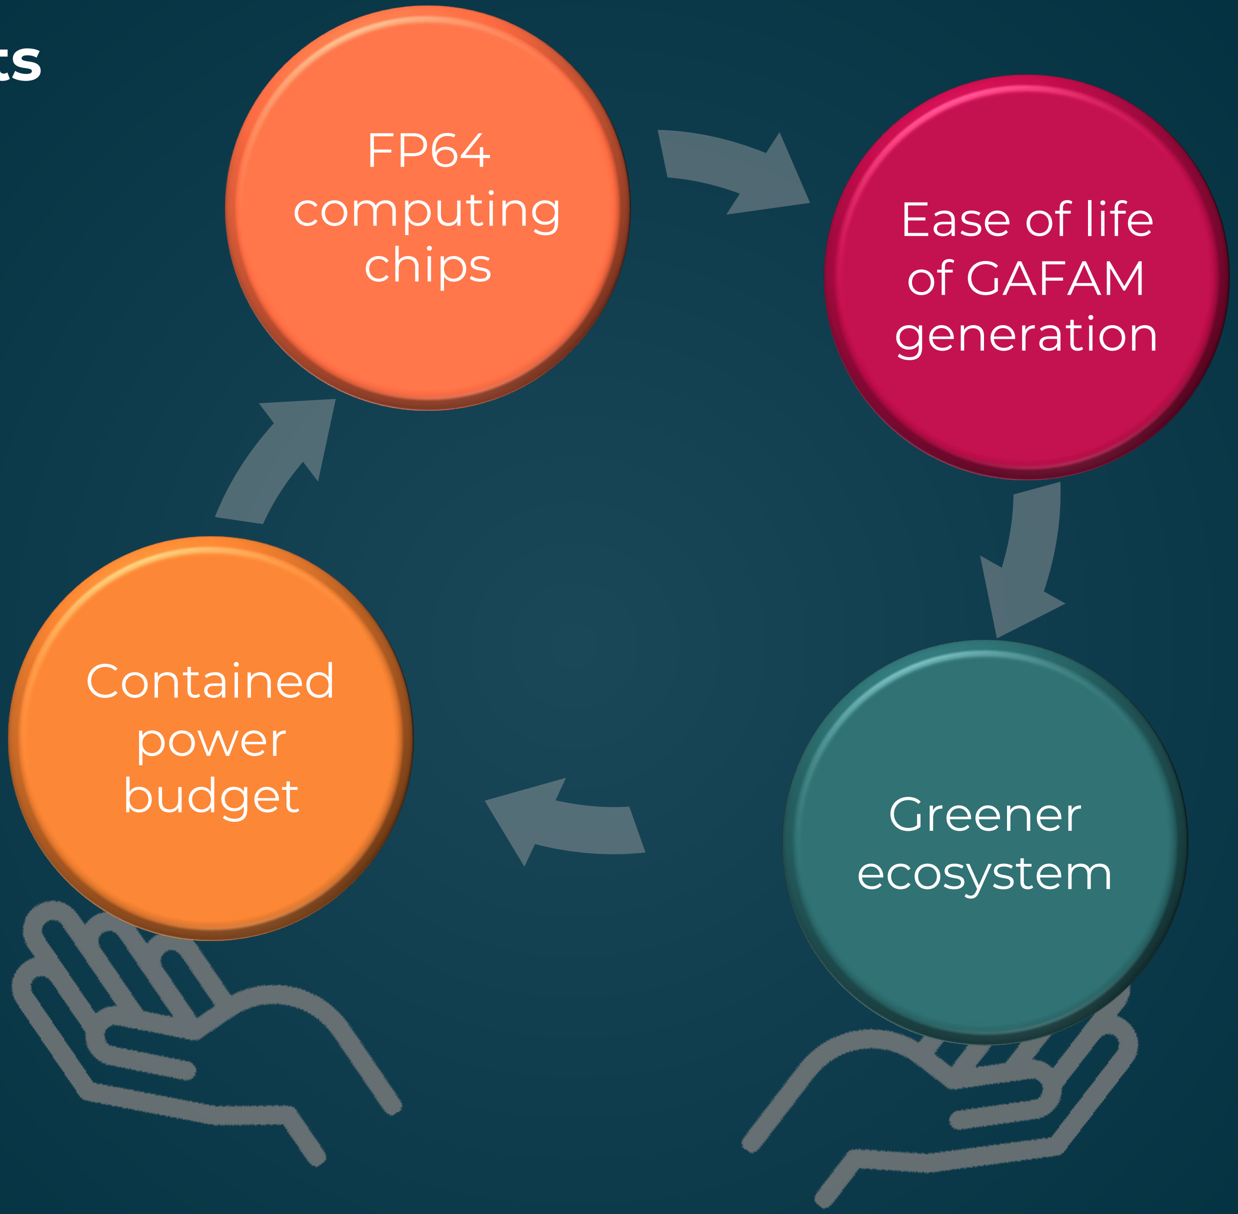

I had a dream

HPC/Cloud user environment convergence

Short term

- Generate/store key data only
- Programming language abstraction
- Monolithic simulations to workflows
- AI augmented algorithms
- Augmented orchestration
- Augmented Management
- Predictive maintenance
- E2E & unified security of a federation
- FP64 profiler
- Dynamic Frequency Management
- Polymorphic & Disaggregated IOs
- MPI offload

Mid term

- Standardized accelerator interface
- Less FP64 sensitive algorithms
- Prescriptive maintenance
- Data life cycle management
- End user tooling
- Standardized carbon footprint mngt
- Unified CPU/accelerator memory
- Ultra Ethernet for scale up/out
- Near memory Computing
- Standardized chiplet foundation
  - 3D memory
  - Inner / outer photonics
  - Disaggregated SoC
- Reconditionable infrastructure

Long term

- Photonic computing
- System on wafer
- Classical/Quantum convergence
- Superconducting CMOS
- Disaggregated architecture
- Full VHV to LV DC power

End to end co-design

Whole ecosystem alliances

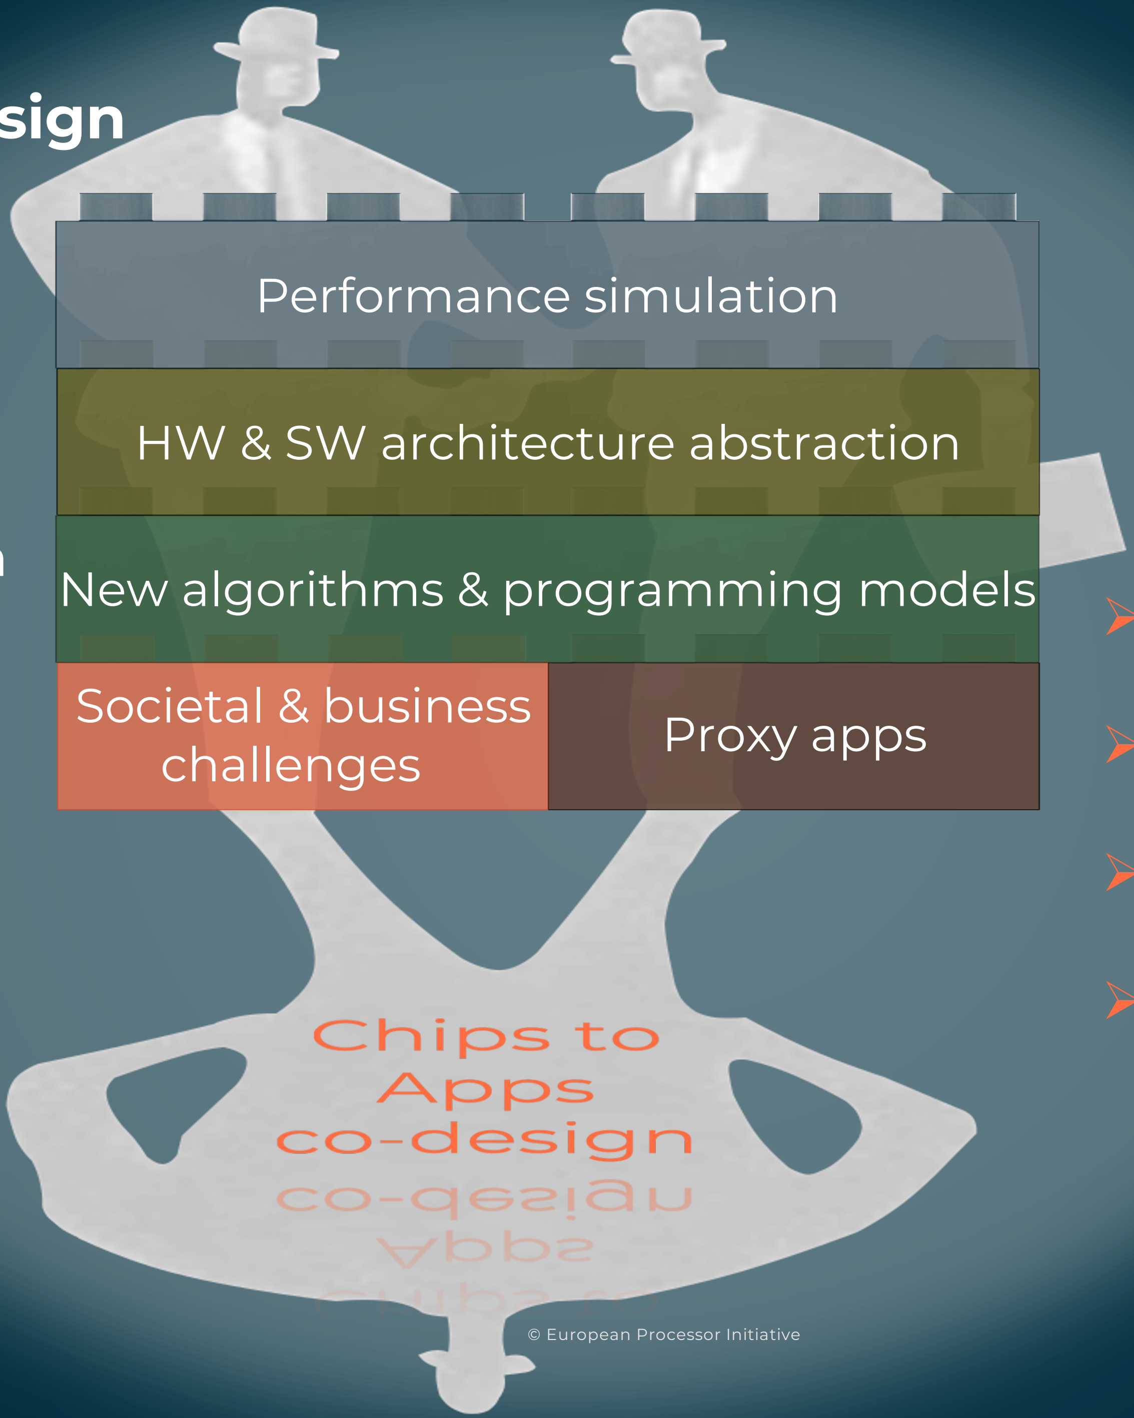

- Accelerate Time-to-Market
- Mitigate Technical Risk
- Optimize Design efficiency
- Leverage Rapidly Evolving Technologies

European Processor Initiative

MPSoC'24
EVIDEN

From free lunch to responsible computing and data generation

End user tool box

Data life cycle management

Relieve FP64 sensitivity

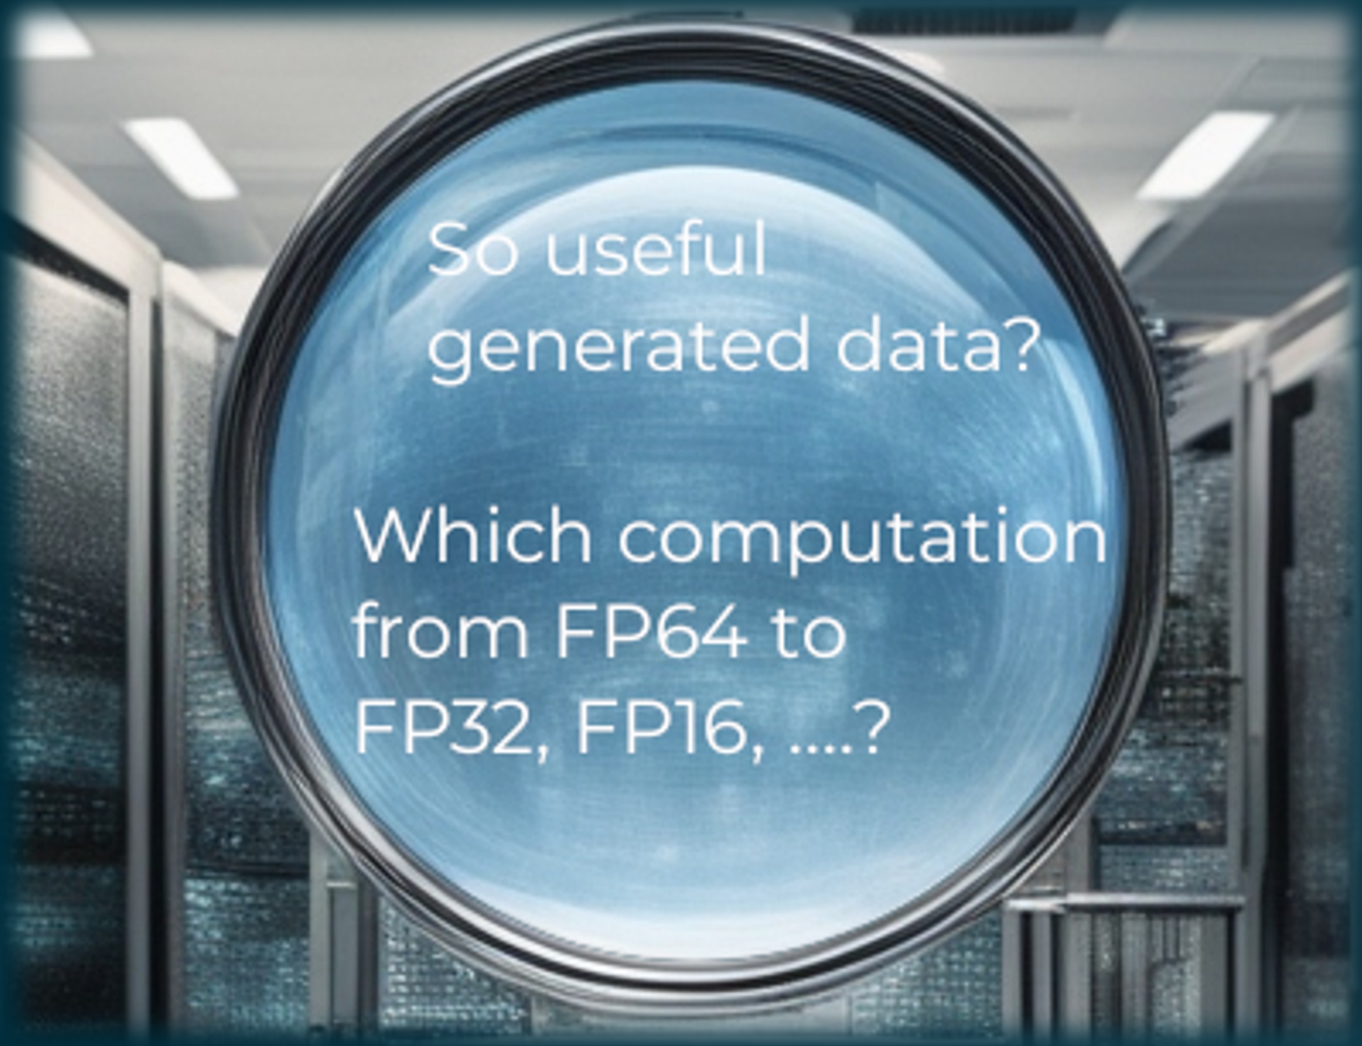

FP code profiler

Work around the Babel tower of languages

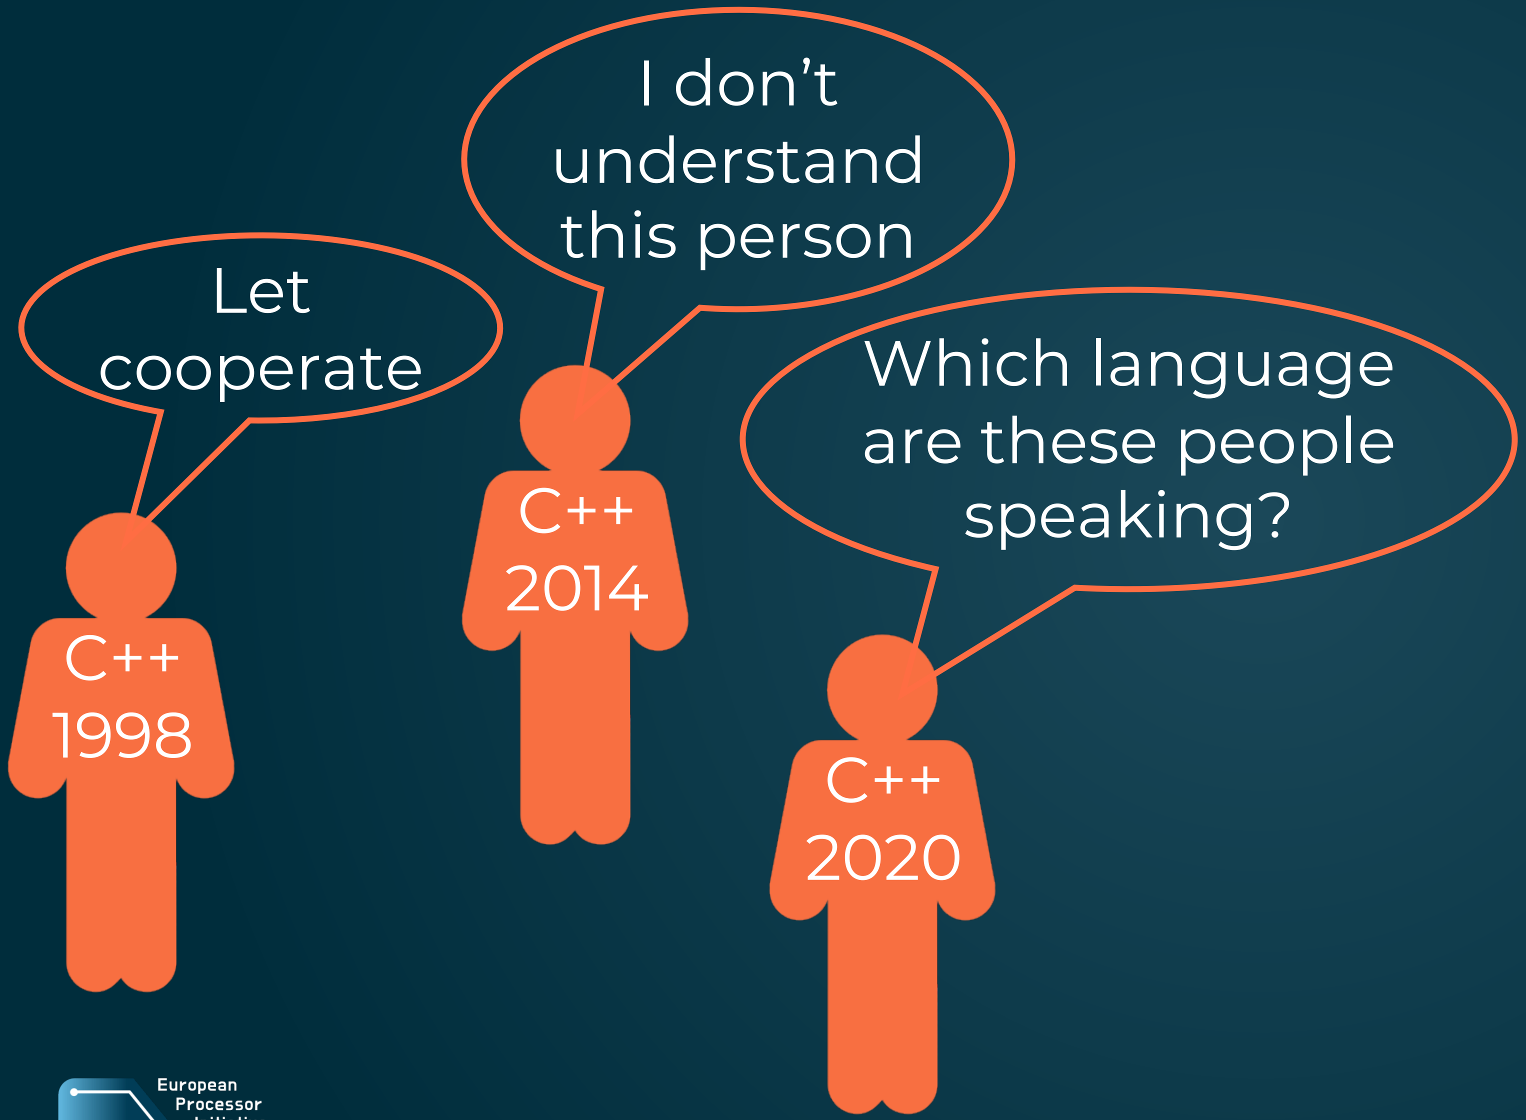

Programming Language abstraction

Gain the freedom to choose your accelerator

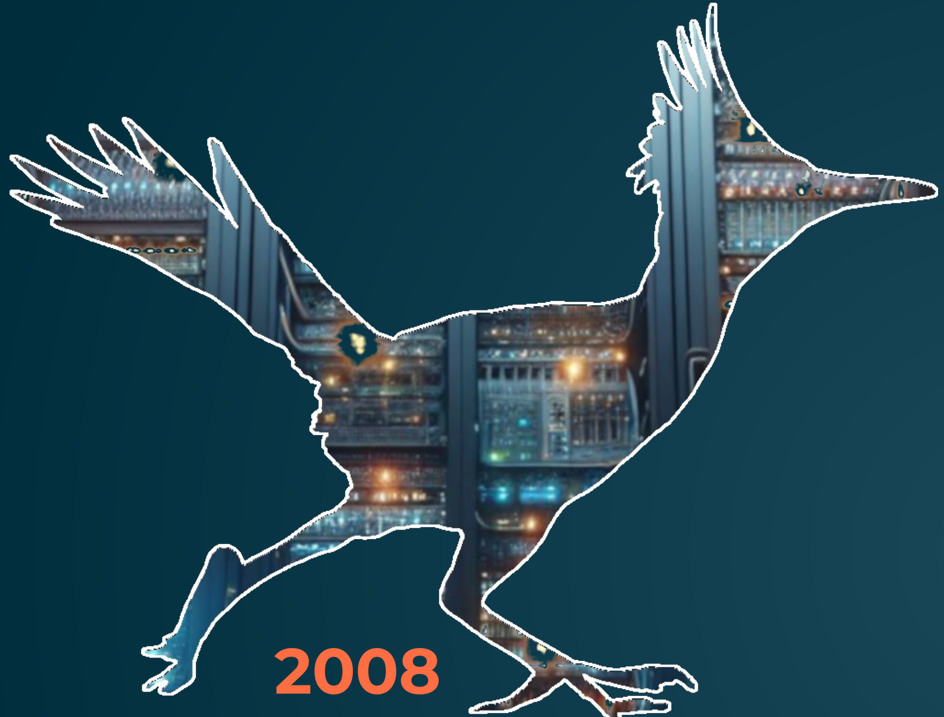

2008

Roadrunner, 1st supercomputer with GPUs

2024

Still challenging to program GPUs
Numerous codes not GPU compatible yet
Strong dependency from Nvidia CUDA API

"Standardized" programming API for accelerators

Unlash memory throughput for an efficient computing

	Package BW (GB/s)	Power efficiency (pJ/b)
HBM2	**410**	**5**
Darpa PIPES	100,000 **(244x)**	1 **(5x)**
AYAR LABS	32,000 **(78x)**	1 **(5x)**

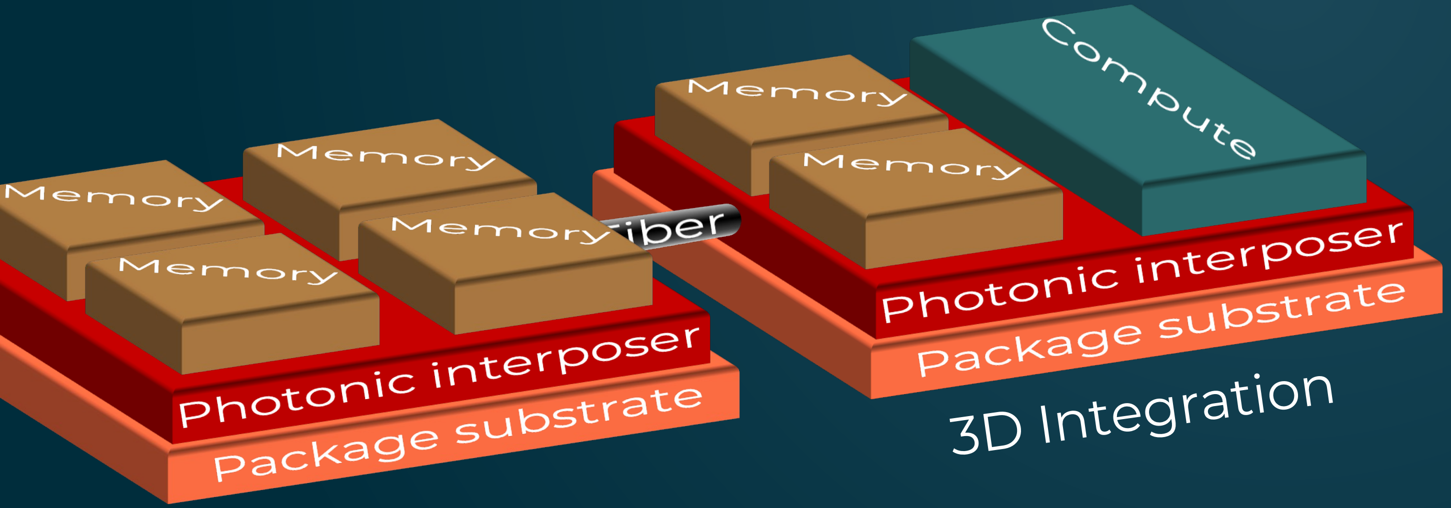

Compute to Memory photonic connections

Make tailor-made HPC HW affordable

Common chiplet platform architecture

- Mechanical, electrical, protocols
- Published Interface specifications
- Test suite for validation/certification
- Likely foundry process agnostic
- Market place

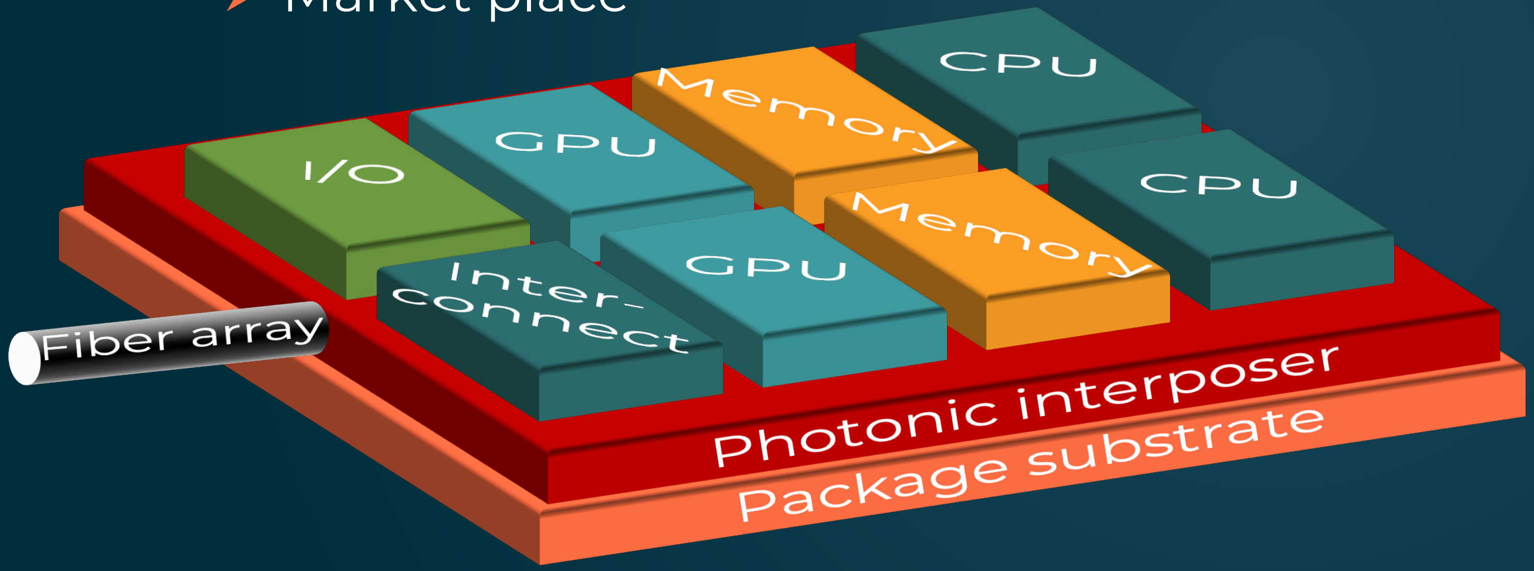

common
chiplet
platform
architecture

Superconducting drastically downsize the power consumption

Goal of Imec SCD (Super Conducting Digital) Project:

4 PFLOPs/cm² @ 120mW (cold) / 38W (wall plug)

	CMOS 7nm	Imec SCD 28nm
Speed	2 GHz	15x
Memory	500 MB/cm² (SRAM)	0.01x
Interconnects	1.6 GB/line @ 1 pJ/b	1000x
Power	1 TOPs/W	100x

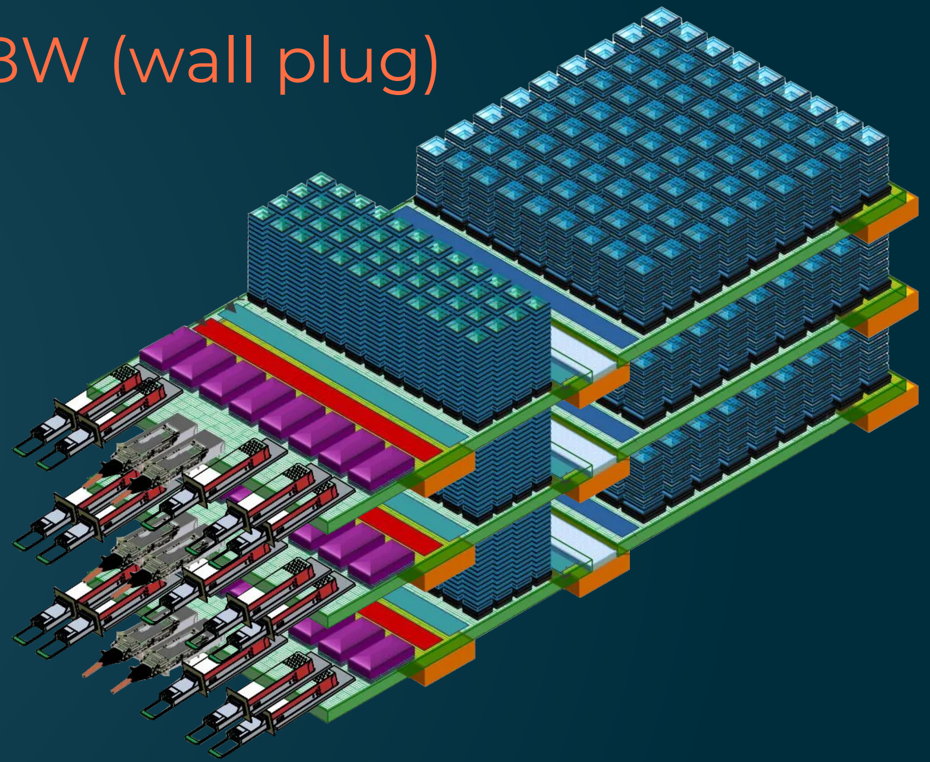

Data center in a shoe box:

- 100 blades
- 150x320x100mm

Sources: Imec lecture, SC23
superconducting-digital-technology-revolutionize-ai-and-machine-learning-roadmap (Imec)

Classical/Quantum computing convergence

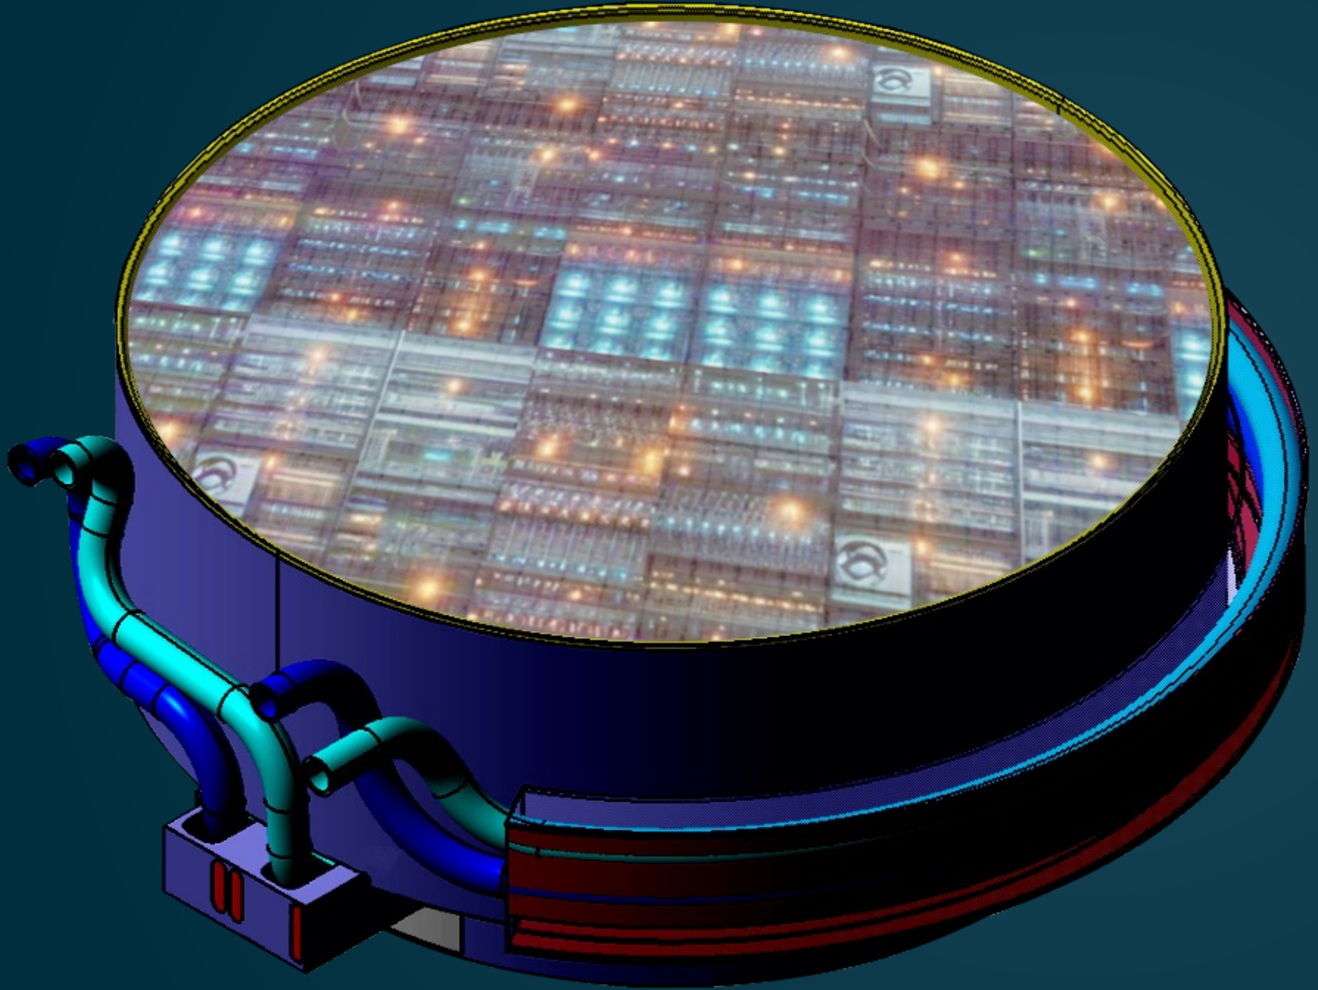

Datacenter on a wafer

Affordable cryocooler

European Processor Initiative

8 The European initiatives

Achieve autonomy in strategic processing technologies

20% Production (worldwide value) in Europe of cutting-edge and sustainable semiconductors including processors

5 to 2nm Manufacturing capacity

x10 Energy efficiency improvement from 2020

EU

Two pillars to make European ambition a reality

Empowering Europe's Semiconductor Future, uniting innovation and driving Progress

Leading the Way in European Supercomputing, developing a World Class Supercomputing Ecosystem

EU back in the race with EuroHPC JU

Deploy

Develop, deploy, extend & maintain a world-leading supercomputing, quantum computing, service & data infrastructure ecosystem in Europe

Innovate

Support the development of innovative supercomputing components, technologies, knowledge & applications to underpin a competitive European supply chain

Value

Widen the use of HPC & quantum infrastructures to a large number of public & private users wherever they are located in Europe and supporting the development of key HPC skills for European science and industry

Fuel European ambition (2021-2027)

Digital Europe Program	Horizon Europe Program	Connecting Europe Facility
1.98B Eur	**900M Eur**	**200M Eur**
Infrastructure	**Technology**	**Hyperconnectivity**
Federation of supercomputing services	**Application**	**Data connectivity**
Widening usage and skills	**International Cooperation**	

*Member states to match this with national contributions

Supercomputer deployment

NOV 2023	TOP500	Green500
LUMI	#5	#7
LEONARDO	#6	#18
MARENOSTRUM 5	#8	#6
MELUXINA	#71	#27
KAROLINA	#113	#25
DISCOVERER	#166	#216
VEGA	#198	#253

Underway	Year	Performance
JUPITER	2024	Exascale
DEADELUS	2024	Mide-range
JULES VERNES	2025	Exascale

European Processor Initiative

EuroHPC Joint Undertaking

Quantum deployment

Agreements with Six hosting entities

2 quantum simulators (100+ qubits) in
- ➢ Joliot Curie (GENCI / France)
- ➢ Juwels (JFZ / Germany)

Two procurements in progress
- ➢ EuroQCS-Poland (PSNC / Poland)
- ➢ Euro-Q-Exa (LRZ / Germany)

- Call in progress for 2 quantum Excellence Centers

Federate EuroHPC systems (2023+)

Authentication, Authorization and Identification services (AAI)

Computing services
- ➢ Interactive Computing
- ➢ Cloud access – Virtual Machines - Containers

Data services
- ➢ Archival Services and Data repositories
- ➢ Data mover / transport services

User and Resource management

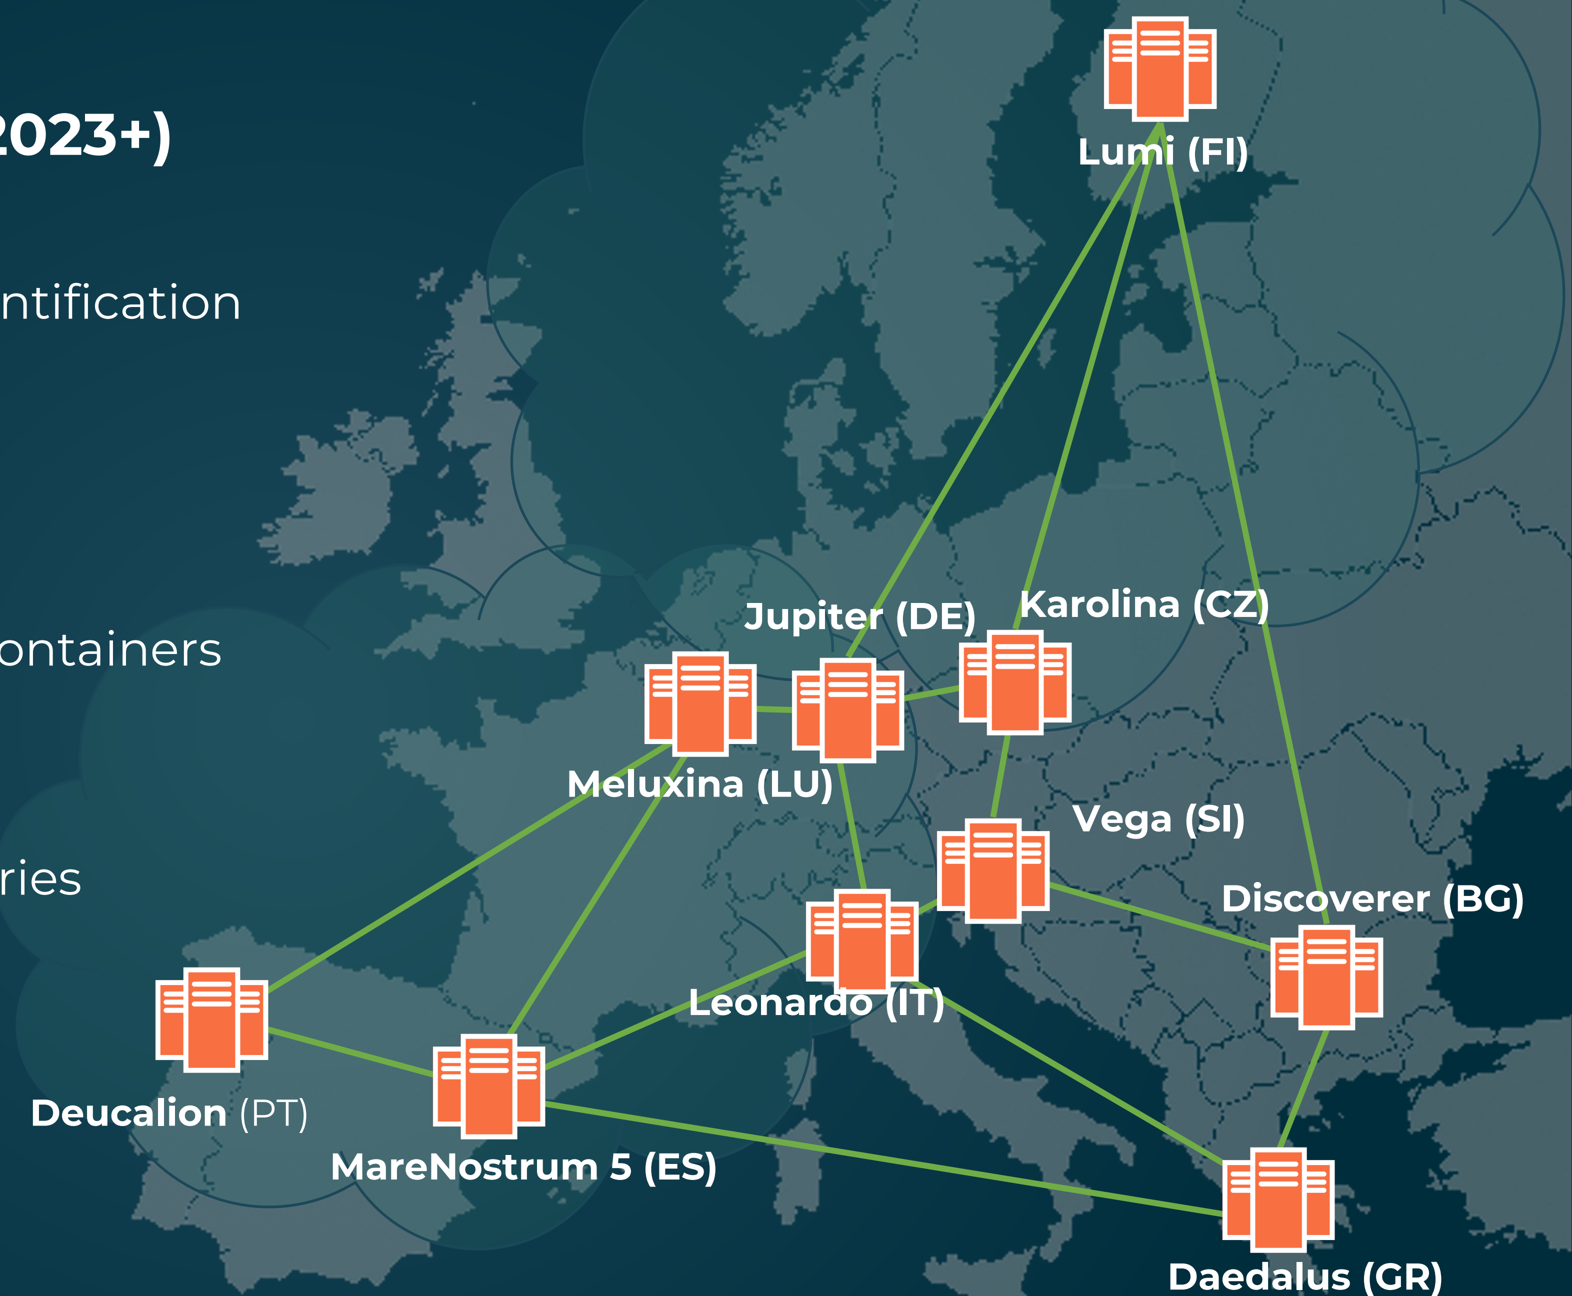

Infrastructure roadamp

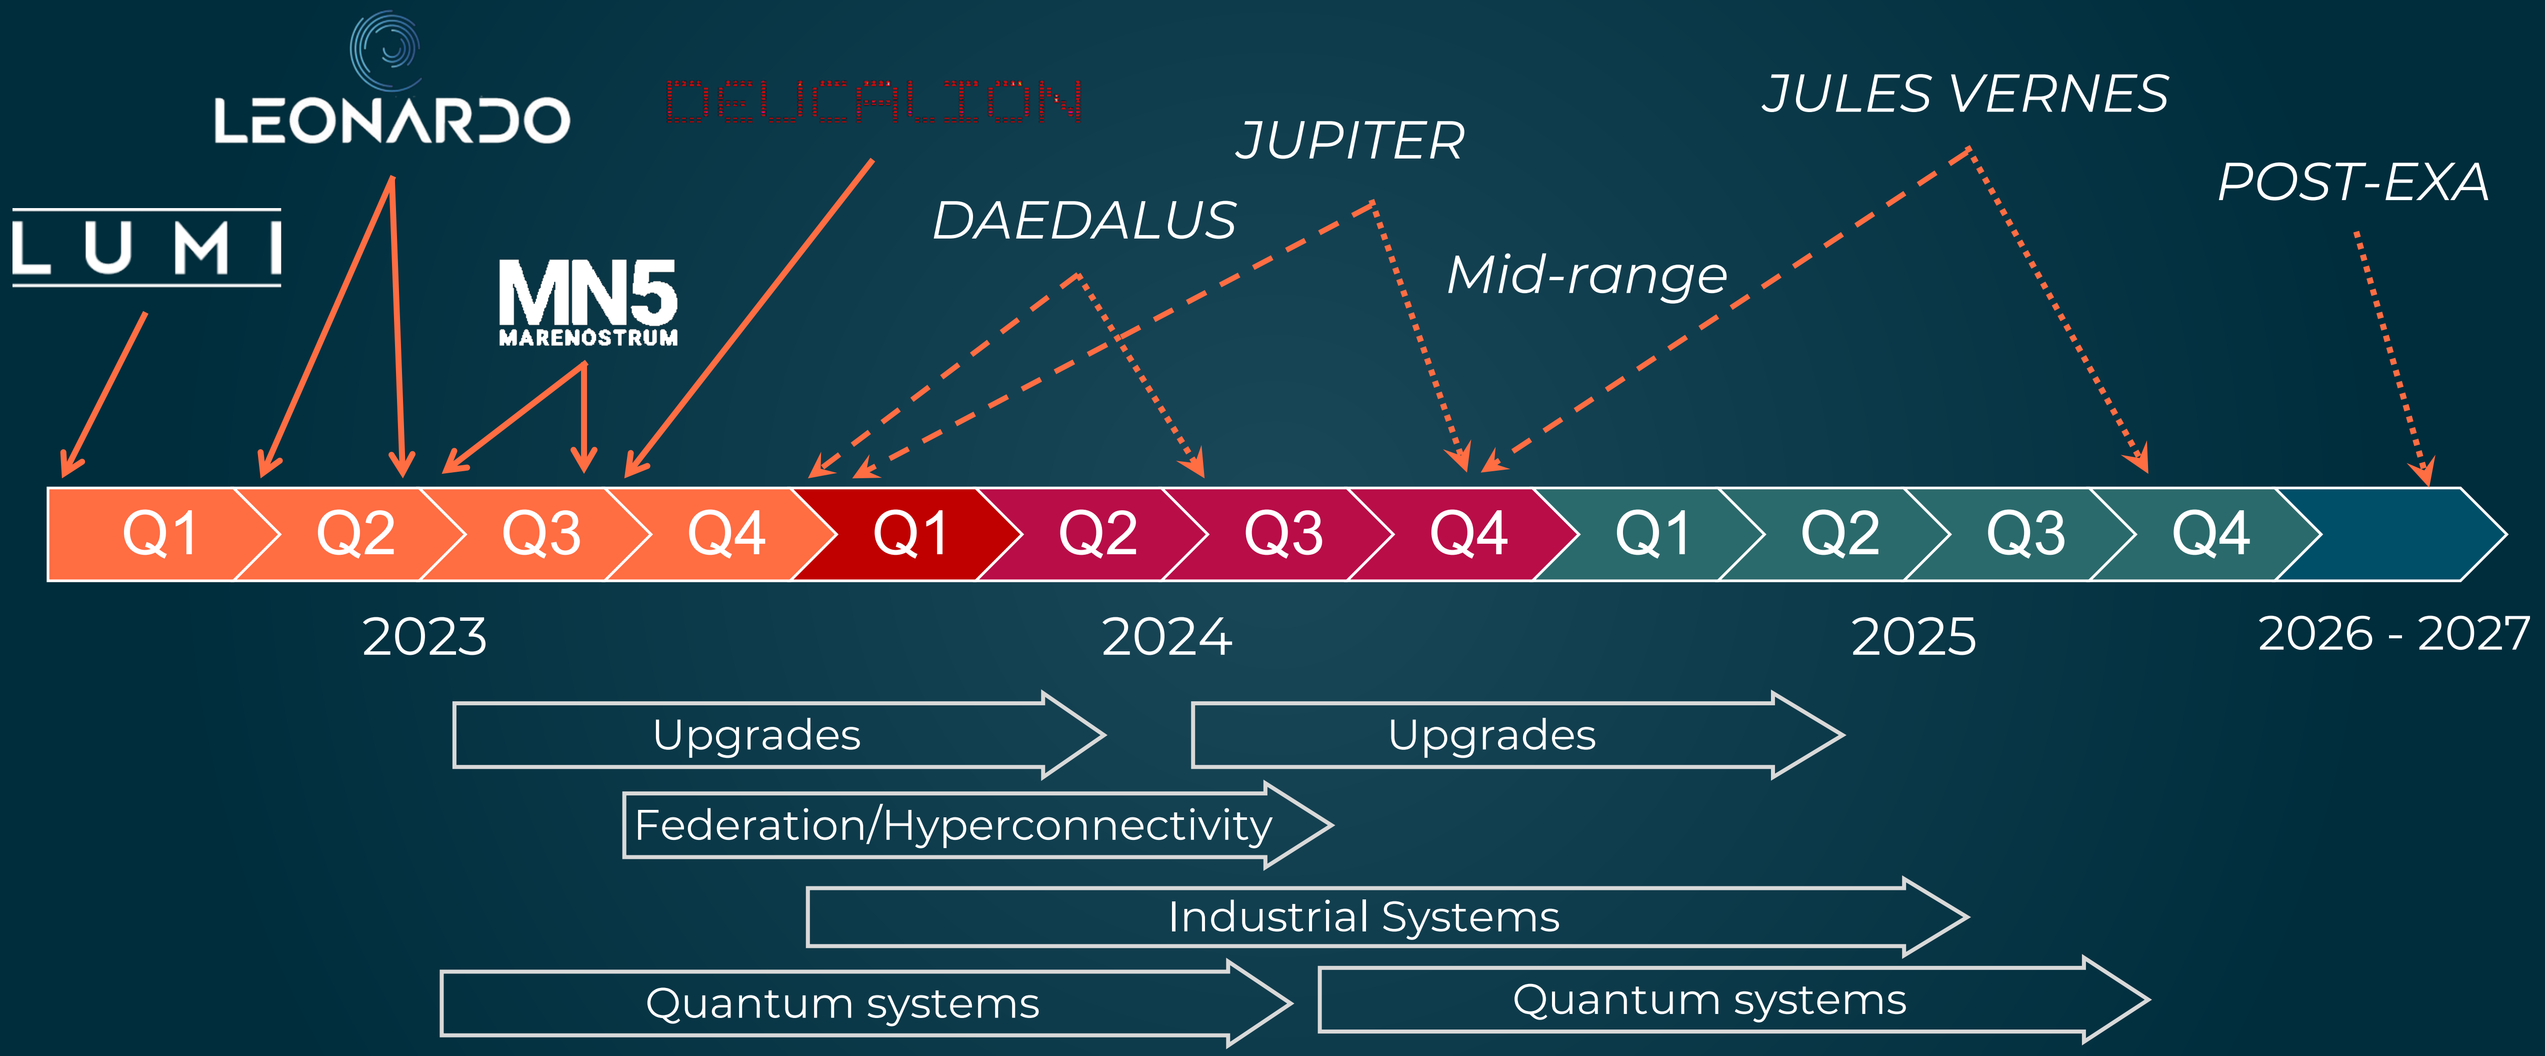

Areas of Strategic Research & Innovation

- **Leadership in Use & Skills**
  Competence Centres and training programs in HPC commensurate with the labor market.

- **Applications and Algorithms**
  Centres of Excellence for HPC Applications and new algorithms for European exascale technology.

- **European Software Stack**
  Software and algorithms, programming models and tools for exascale and post exascale systems.

- **European Open Hardware**
  Ecosystem for the low power high-end general purpose processor and accelerator.

Currently Around **40** running projects

(Partial) system technology roadmap

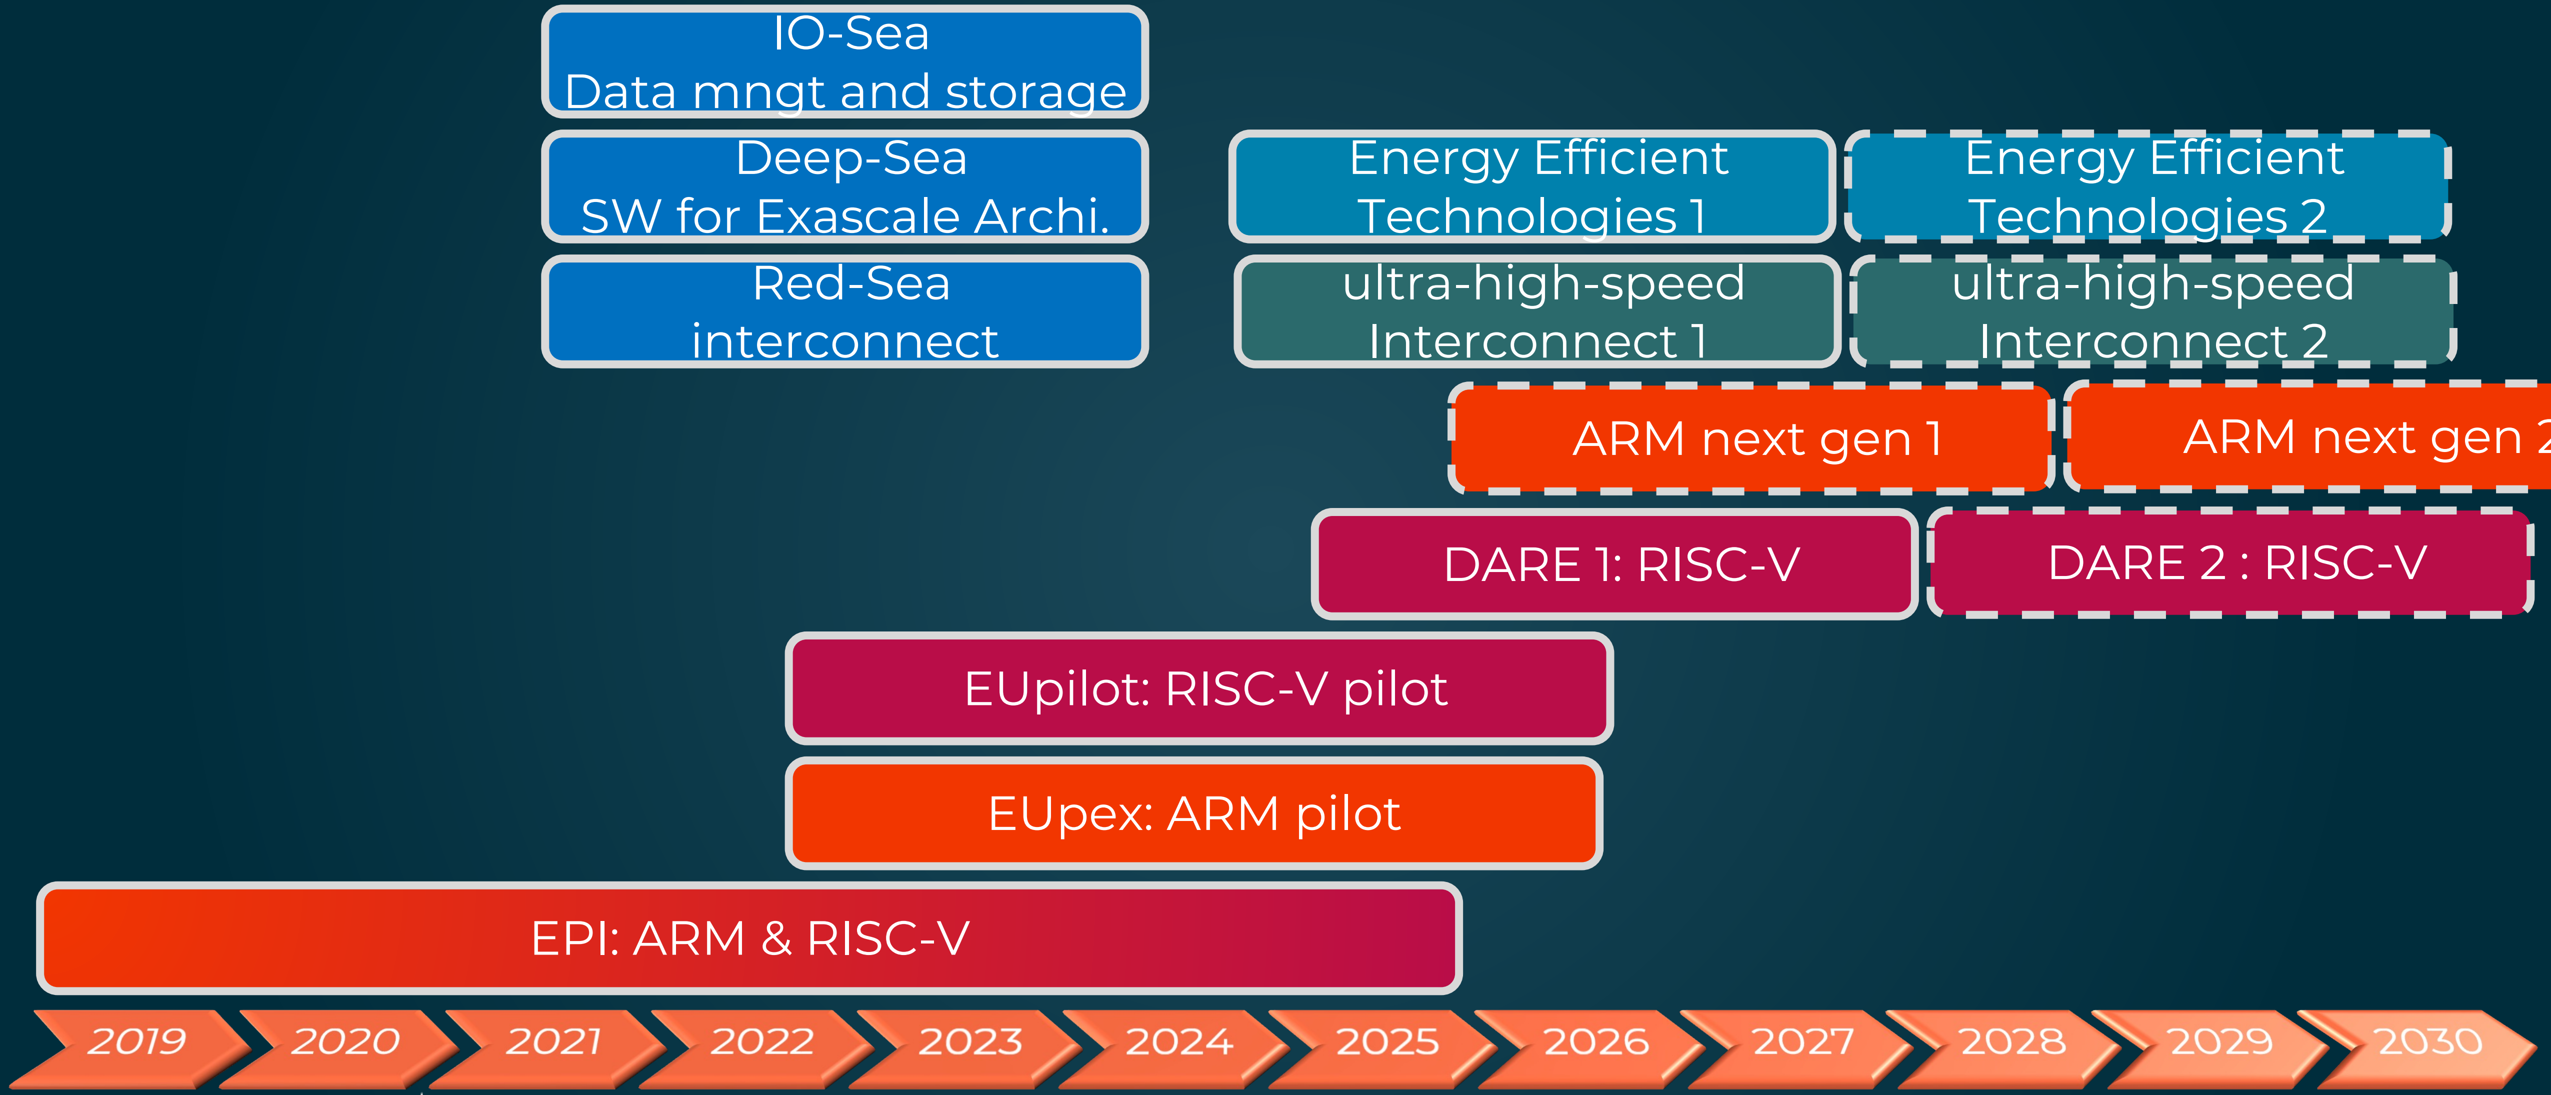

CPU and accelerator, a 400M€ Roadmap

European silicon proven IP

Flexible chiplets

Ready for state of the art AI

Deployment-ready through pilots

First IPs (2024-2026):

- Build on EPI efforts on ARM-based processor
- From test chips to TRL 9
- EuroHPC exascale systems as first customer

New architectures (2026-2028)

- Stand-alone competitive processors & accelerators
- Building on EU R&D in low power, security,…
- EuroHPC post-exascale system as first customer

Post-exascale(2028-)

- ARM/RISC-V systems based on EU R&D and IPs

A special attention to education and trainings

One-stop shop for HPC training offers of 33 National Competence Centers

Pan-European Master (2 years) for HPC

14 Training Centres with ~100 training events each year

Coming: virtual HPC academy

European Processor Initiative

9 Conclusion

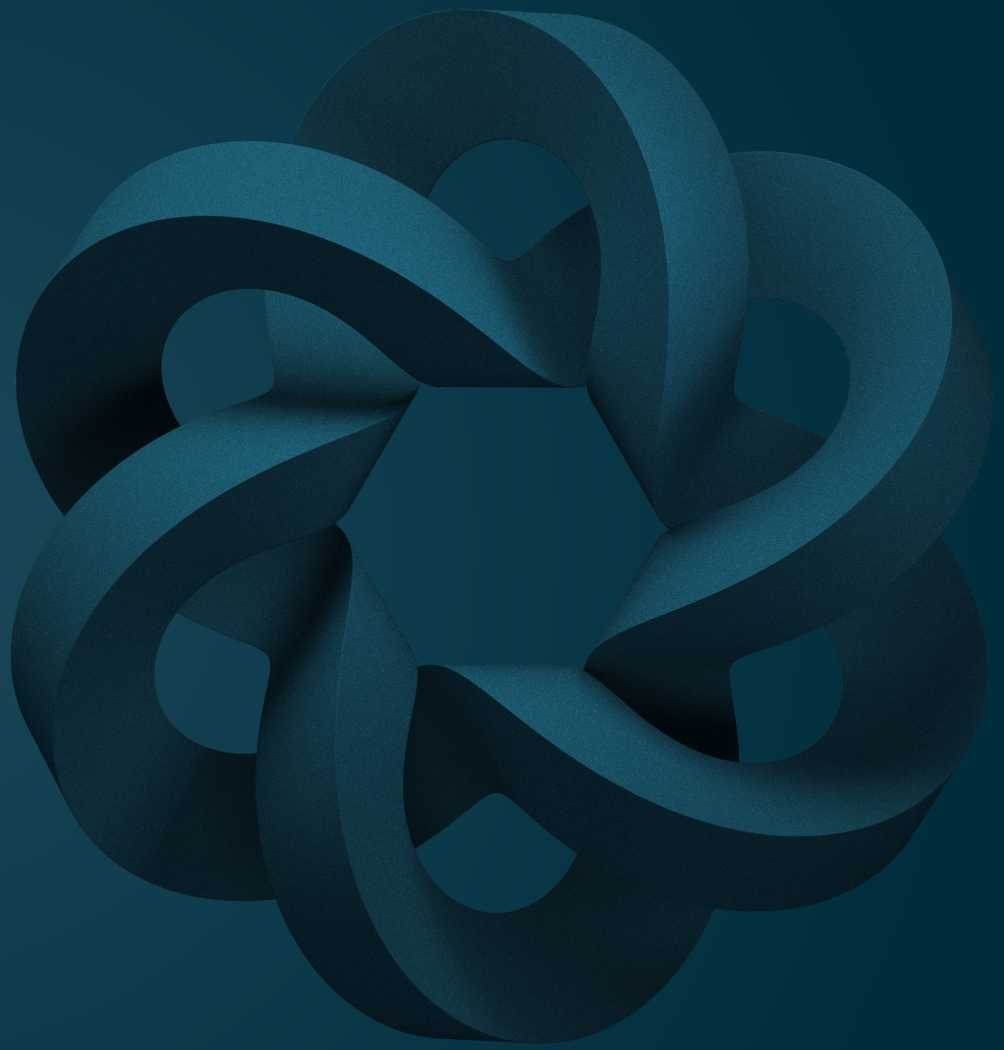

They did not know it was impossible…
So they did it” (Mark Twain)

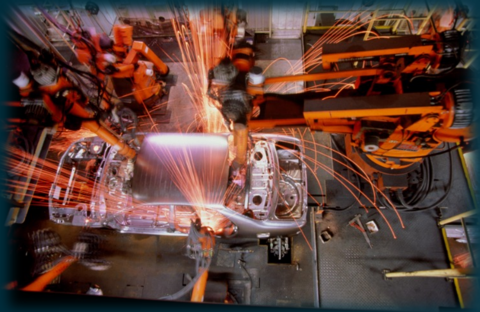

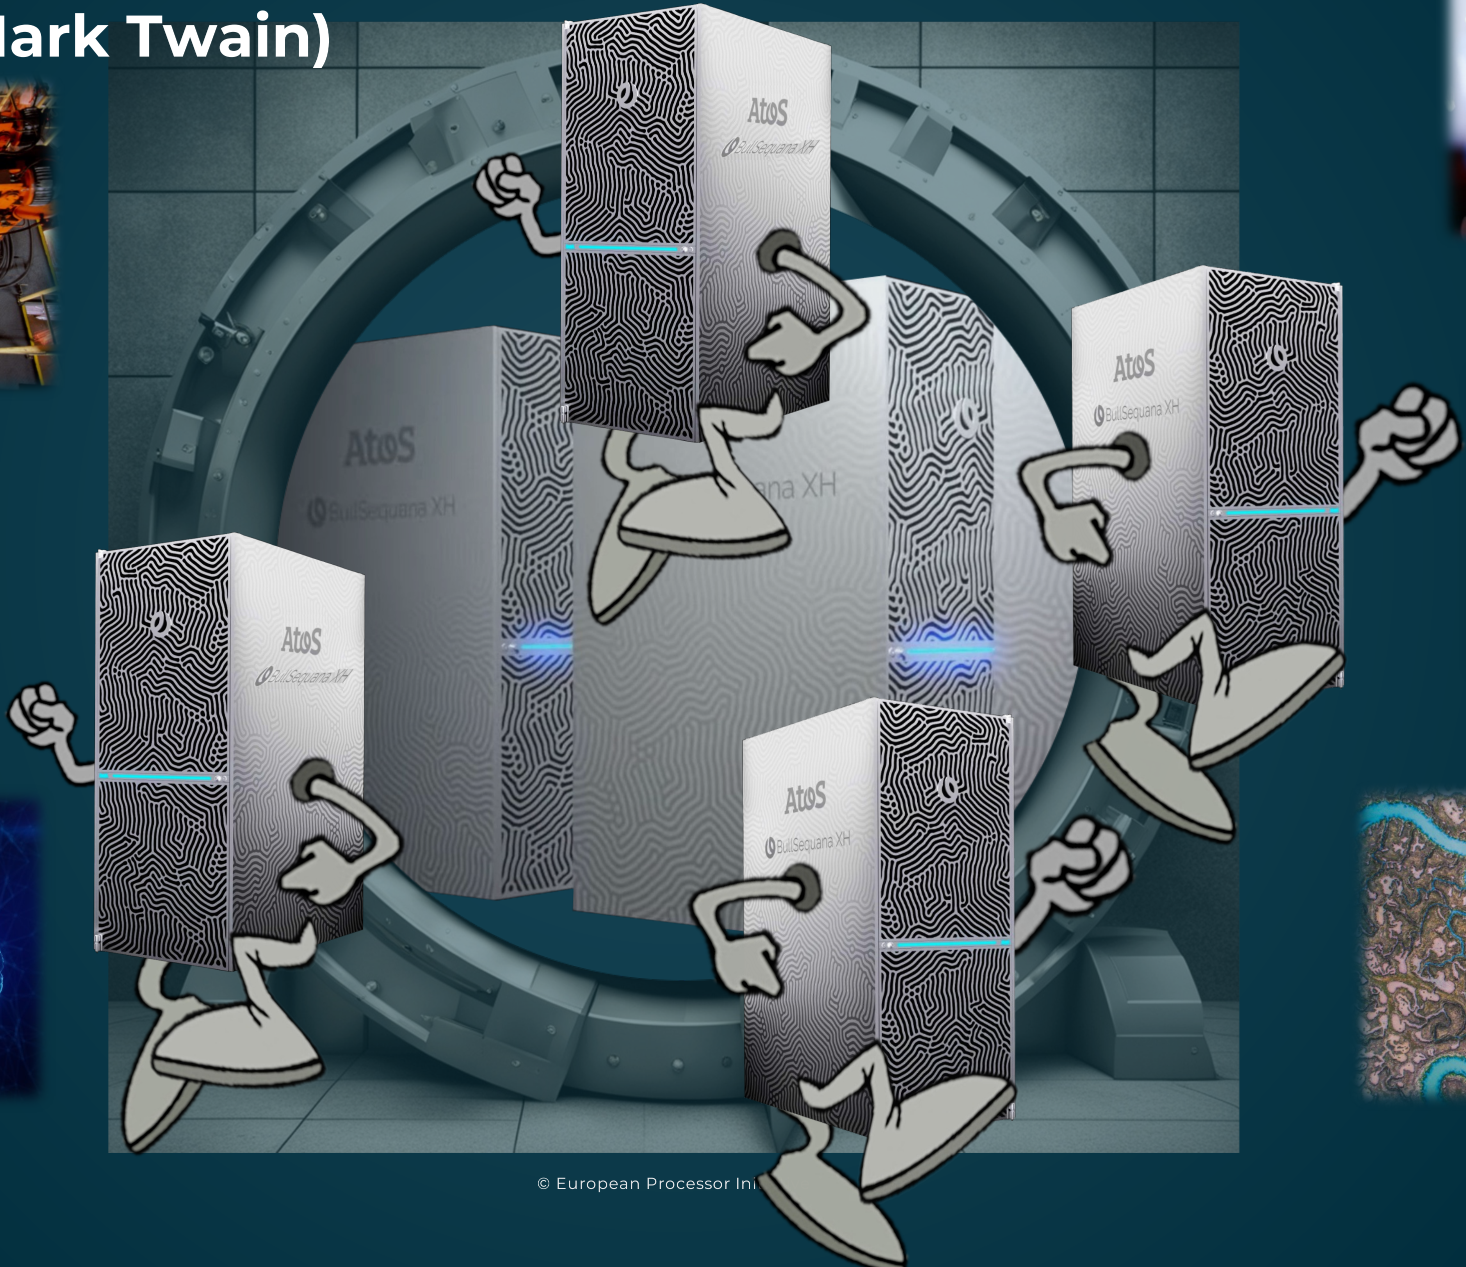

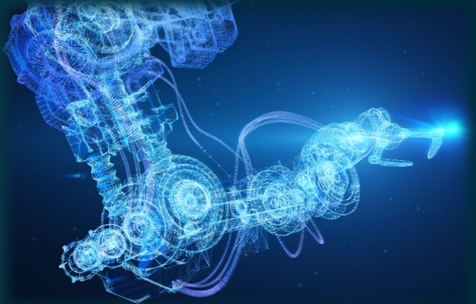

We are still moving data to compute!

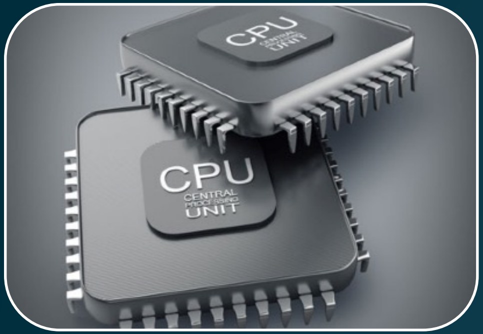

CPU

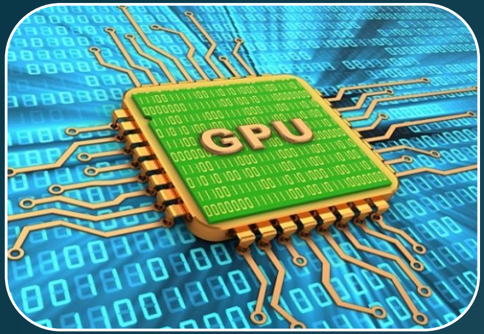

GPU

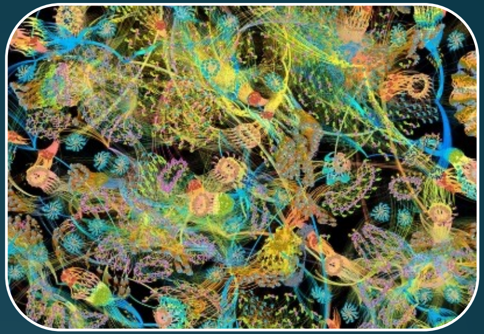

IPU

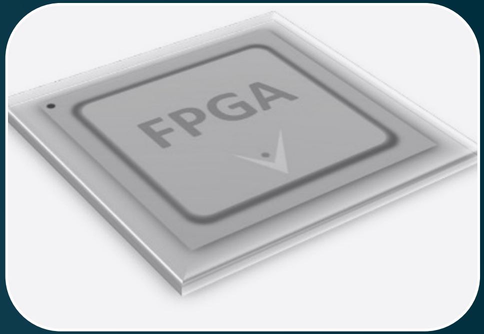

FPGA

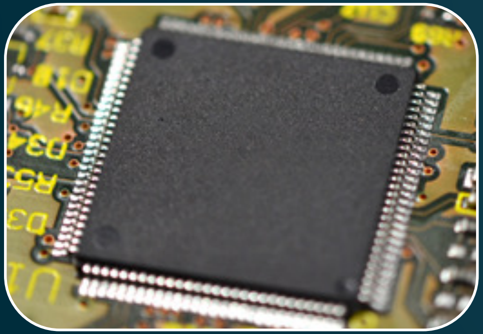

ASIC

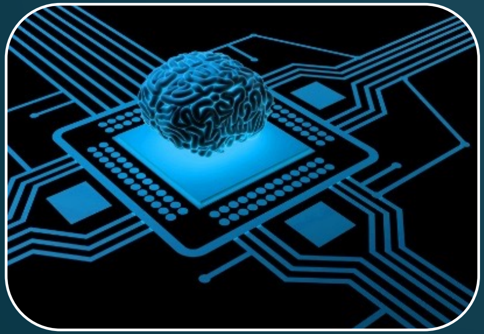

Neuromorphic

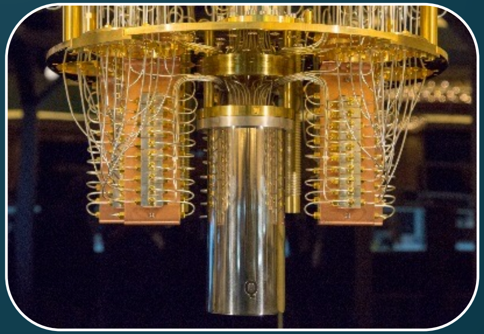

Quantum

DNA

European Processor Initiative

MPSoC'24
EVIDEN

we tend to shy away from simple and obvious solutions

Let be creative

Questions?

European Processor Initiative

MPSoC'24

eric.monchalin@eviden.com

EVIDEN