



SiPearl Outlook

Teratec

Jean-Marc Denis  
Chief Strategy Officer  
[jean-marc.denis@sipearl.com](mailto:jean-marc.denis@sipearl.com)



# Company Update



# — SiPearl corporate overview

## The European Server Processor Solution

**HQ: Maisons-Laffitte (Paris), France**

**Incorporated in June 2019**

**CEO and Founder, Philippe Notton**


### **Design centers:**

- France: Maisons-Laffitte, Massy Palaiseau, Sophia Antipolis, Grenoble
- Germany: Duisburg (Düsseldorf)
- Spain: Barcelona

**Key Personnel from Intel, Atos, ST, Marvell, Nokia, Mstar-Mediatek**

**HPC Targeted Architecture based on Arm Neoverse V1 cores**

**+100 employees today, targeting >1,000 in 2025**




# SiPearl offices

We are close to our partners and customers



# Sipearl extensions 2022 and 2023







HPC is entering in its  
Cambrian era

## CPU



## ACC



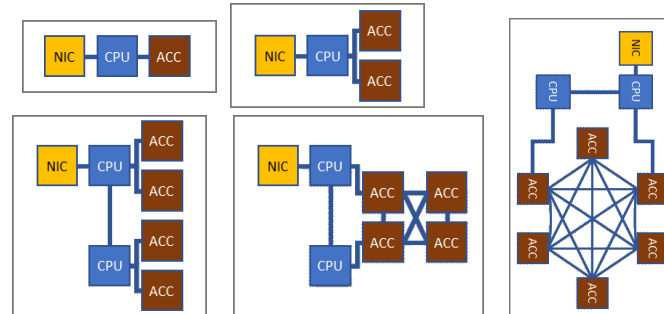
## Intranode network



## HPC network



## Intranode topology






## Cloud



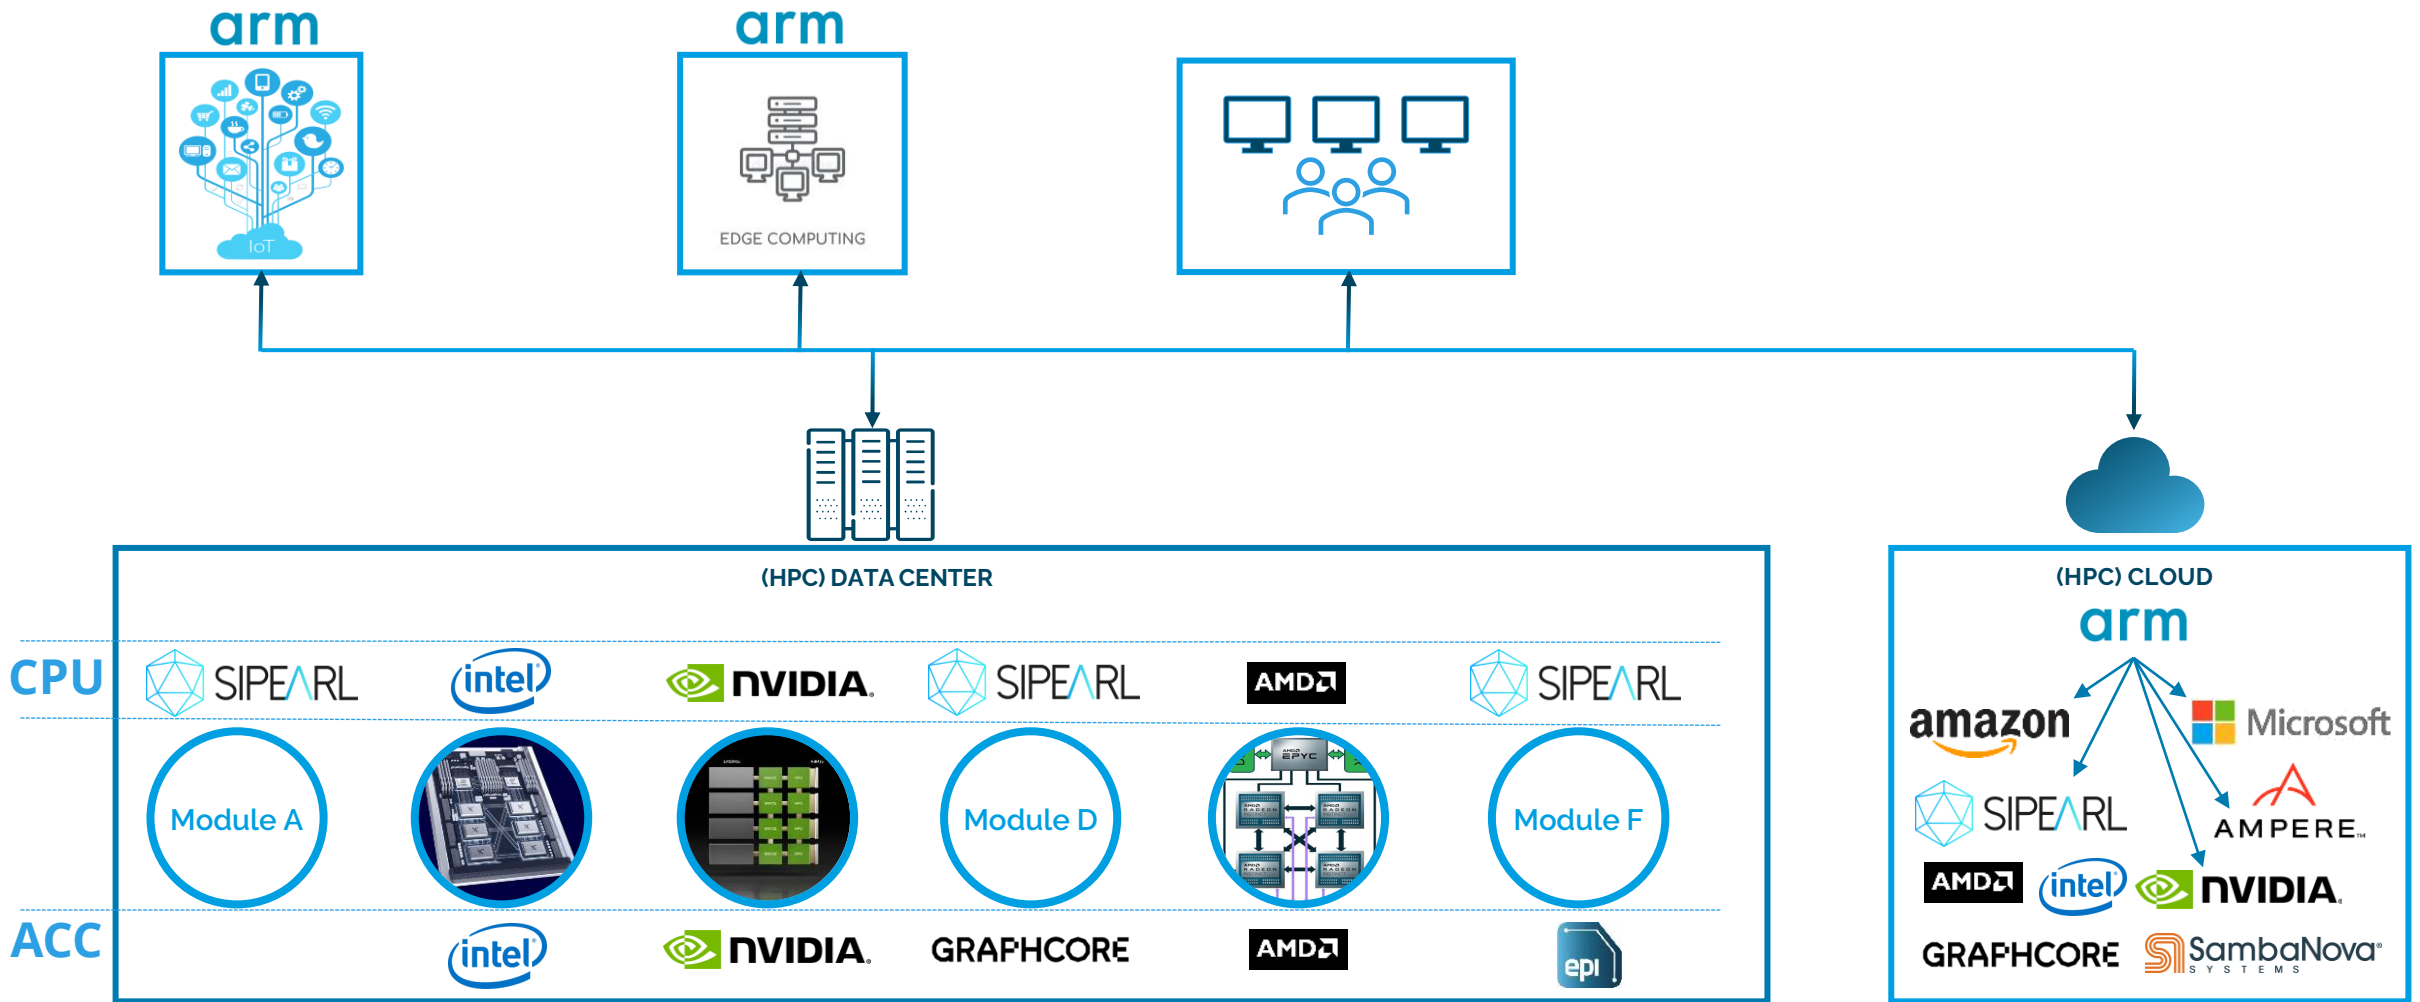


# Many announcements - all "Super-SOC" based

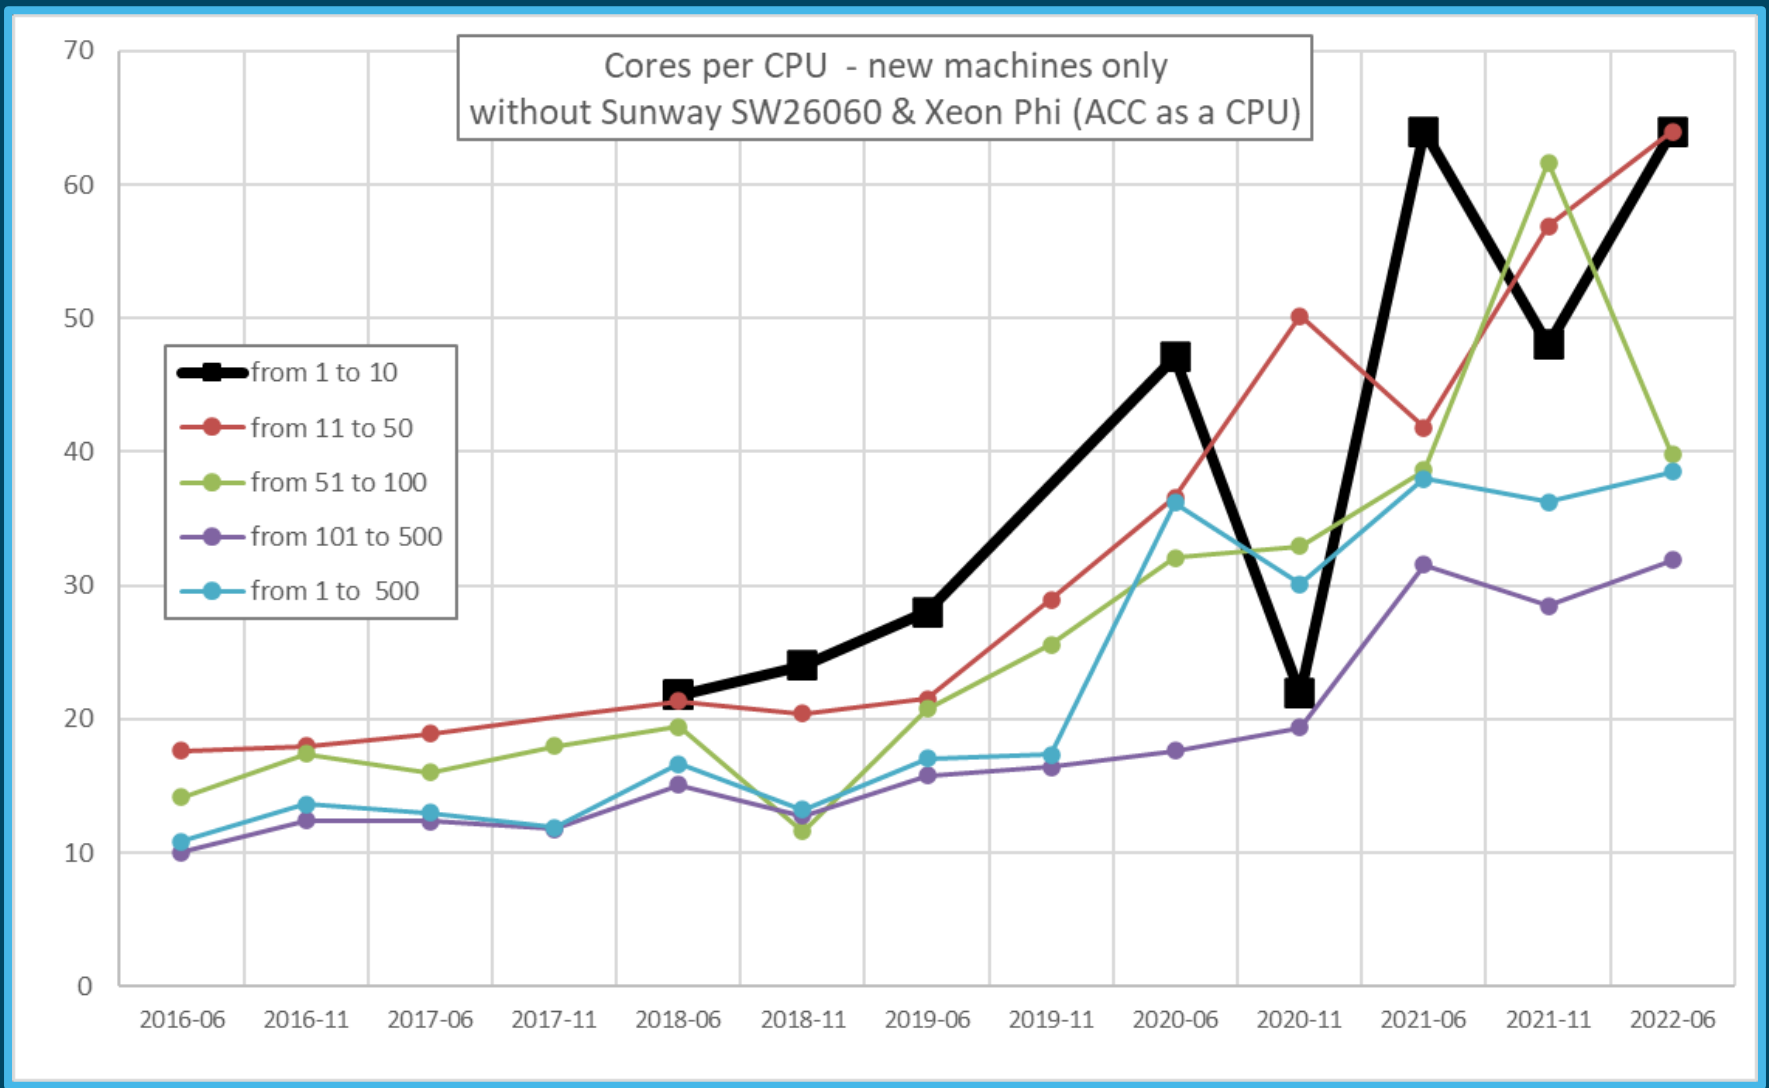




# HPC modular and hybrid architecture





# Core count per CPU socket over time (from Top500)





# Supercomputers (top500) – how many GPU per CPU? The historical perspective



**In top10:** bias because of Fugaku and Tianhe-2 machines. GPU part is artificially low. Next release (Nov.2022) should show >0.5GPU per CPU. Very variable ratio over time because Top10 is where new, innovative and sometimes non conventional technologies are introduced in first place.

**In top50:** growing very fast (x2 during the last 3Y)

**In top100:** growing fast

**Over the complete list:** 1) CPU still massively dominant. But GPU part is growing. Inertia is dominant. GPU cannot do everything everywhere. 2) even with Fugaku bias, total #socket is decreasing Reason: more powerful CPUs & GPUs

# — Cambrian era – Consequences

## Must have

Open/Standard  
HW interfaces  
(UCIe, CXL, PCIe)

## Must Have

Open/Standard  
SW interfaces  
(UCIe, CXL, PCIe)

## Must have

On-premise & in  
the Cloud

CPU's still matter. GPU's  
and ACCs don't fit all  
needs

CPU Core count is  
growing (fast).  
What's the trade-off?




An aerial photograph of a river network, overlaid with a teal color gradient. The rivers are shown as a dense, branching pattern of lines, with the largest river in the center. The background is a dark teal color.

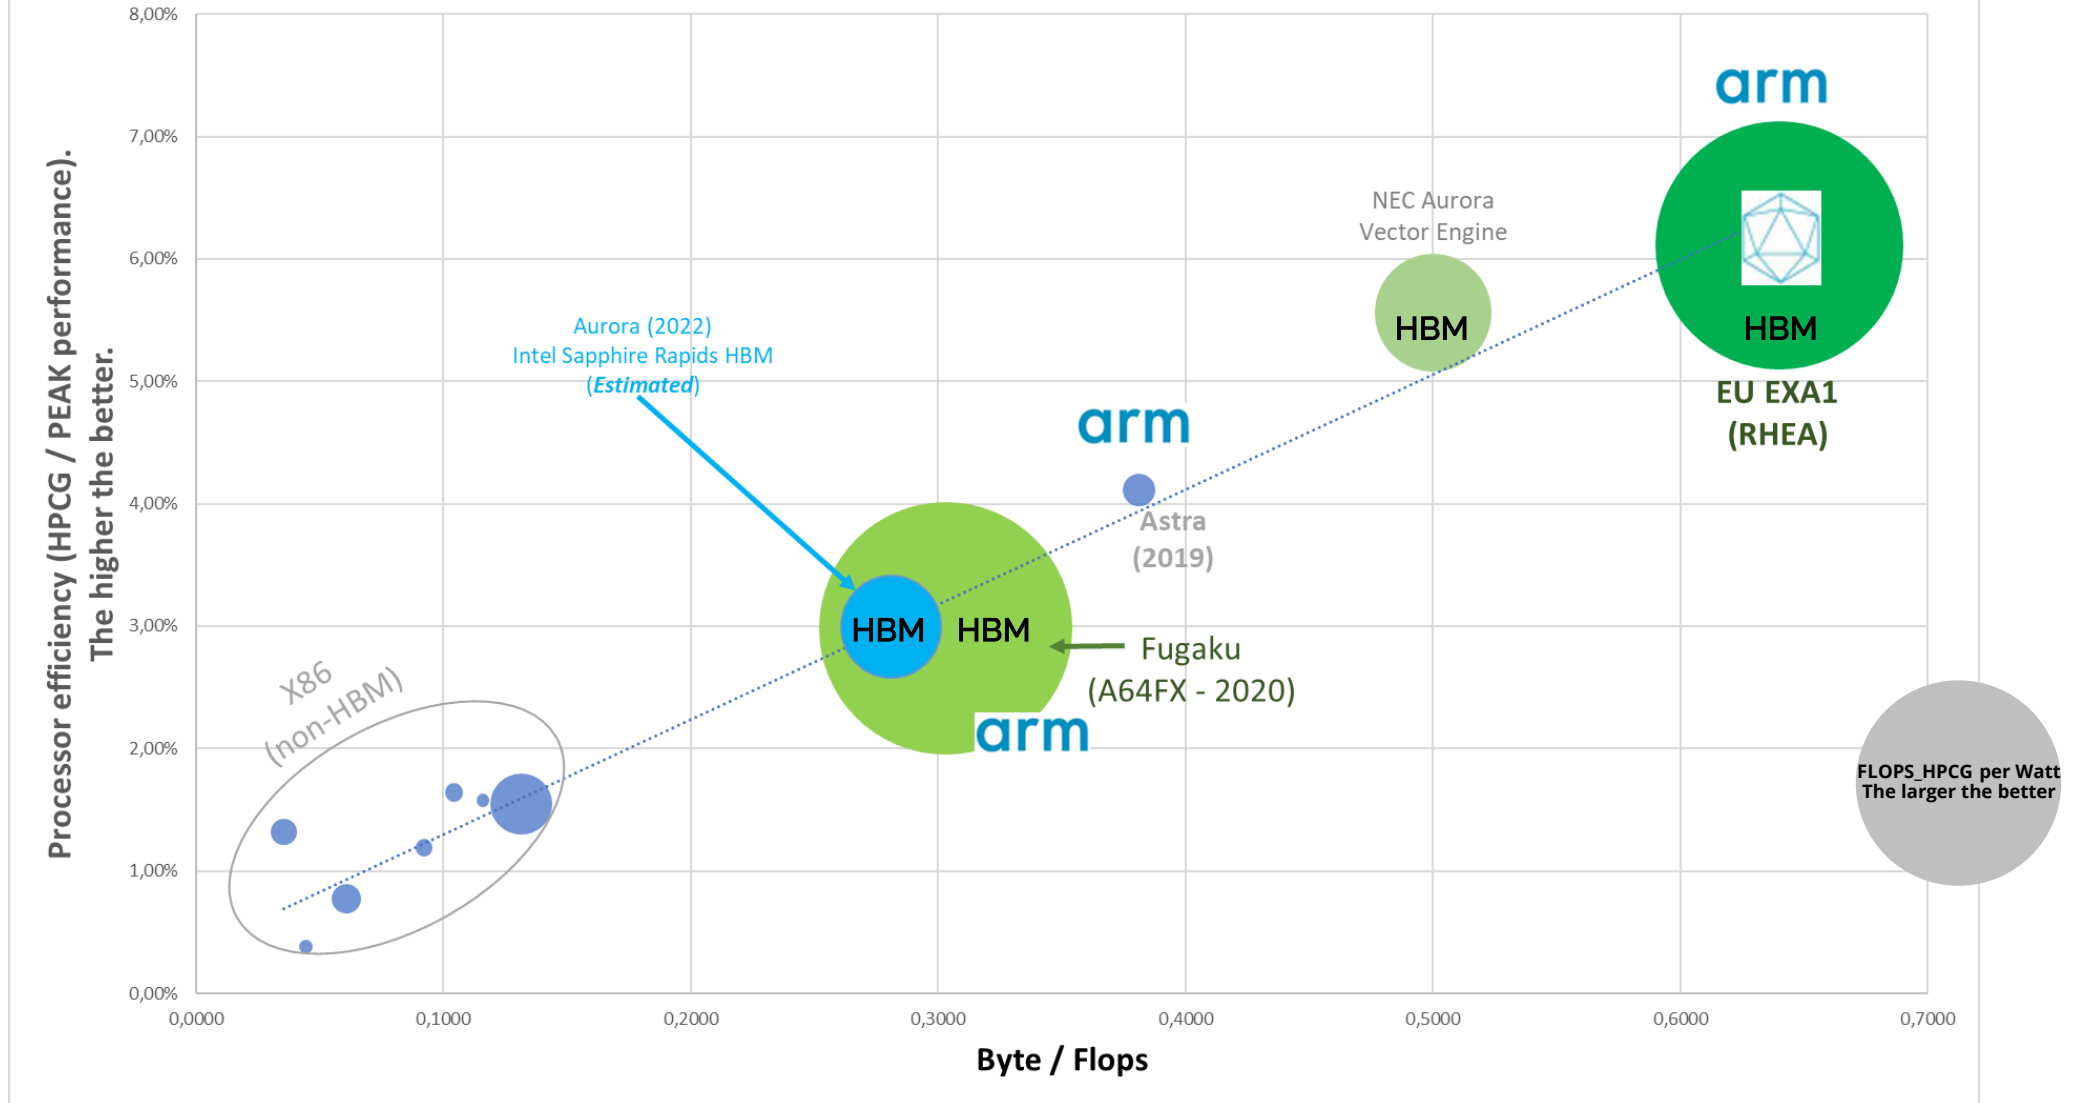
Performance?








# Why HPCG matters





# Core count per CPU socket for new machines in top10 over time



More cores in HBM-less than in HBM processors → (much) **reduced memory BW per core in HBM-less processors**

Overtime, (2025 and beyond), HBM processors should close the gap in core count with HBM-less processors while preserving memory BW per core

# — Performances – Consequences

arm  
HBM

CPU bound → GPUs  
MEM bound → CPUs

Bytes/Flops  
Max mem BW → **HBM**

Slow memory only when  
required  
→ **DDR**

Power efficiency  
(perf on real apps per watt)  
really matters  
→ **HPCG**

Need new specific  
memory and energy  
profiling tools


The background of the slide is a composite image. The upper portion shows a dark night sky filled with stars, with the bright, hazy band of the Milky Way galaxy stretching diagonally from the top left towards the center. Below the sky, a rugged mountain range is visible, with peaks and ridges covered in snow. The overall color palette is dominated by deep blues, teals, and whites, creating a serene and expansive atmosphere.

# The developer/end-user perspective




# Processing elements and Development time


## Processing elements in a system



Back in the day

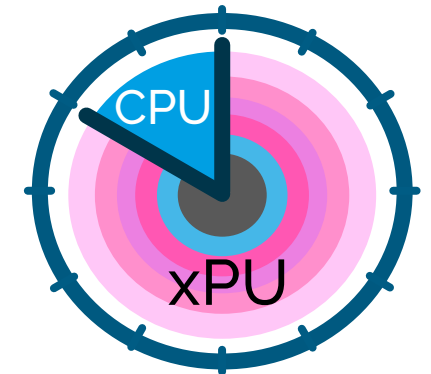
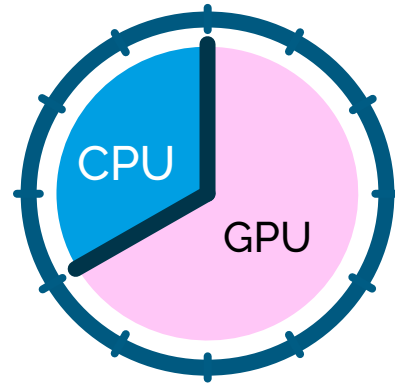



Now



Next gen (2024+)

## Development time



Developers will spend the bulk of their time programming accelerators.

**CPUs that are efficient « out of the box » offer an advantage**

# Software Ecosystem



**Only SiPearl covers the complete software ecosystem**

# – Developer & end-user perspective – Consequences

arm

xPU: (toward) a unified programming model

CPU: No time spent on optimization


Must Have  
Global and unified  
memory

Must Have  
Open/standard SW  
stacks


Develop once, run on many




Conclusion  
Rhea: a Processor for the  
Exa era




<b>Must have</b> Open/Standard HW interfaces (UCIe, CXL, PCIe)	<b>Must Have</b> Open/Standard SW interfaces (UCIe, CXL, PCIe)	<b>Must have</b> On-permise & in the Cloud
CPU's still matter, GPU's and ACC's don't fit all needs	CPU Core count is growing (fast). What's the trade-off?	



<b>arm</b>	xPU: (toward) a unified programming model	CPU: No time spent on optimization
	Must Have Global and unified memory	Must Have Open/standard SW stacks



<b>arm</b> <b>HBM</b>	CPU bound → GPU's MEM bound → CPU's	Bytes/Flops Max mem BW → <b>HBM</b>
Slow memory only when required → DDR	Power efficiency (perf on <u>real</u> apps per watt) really matters → HPCG	Need new specific memory and energy profiling tools

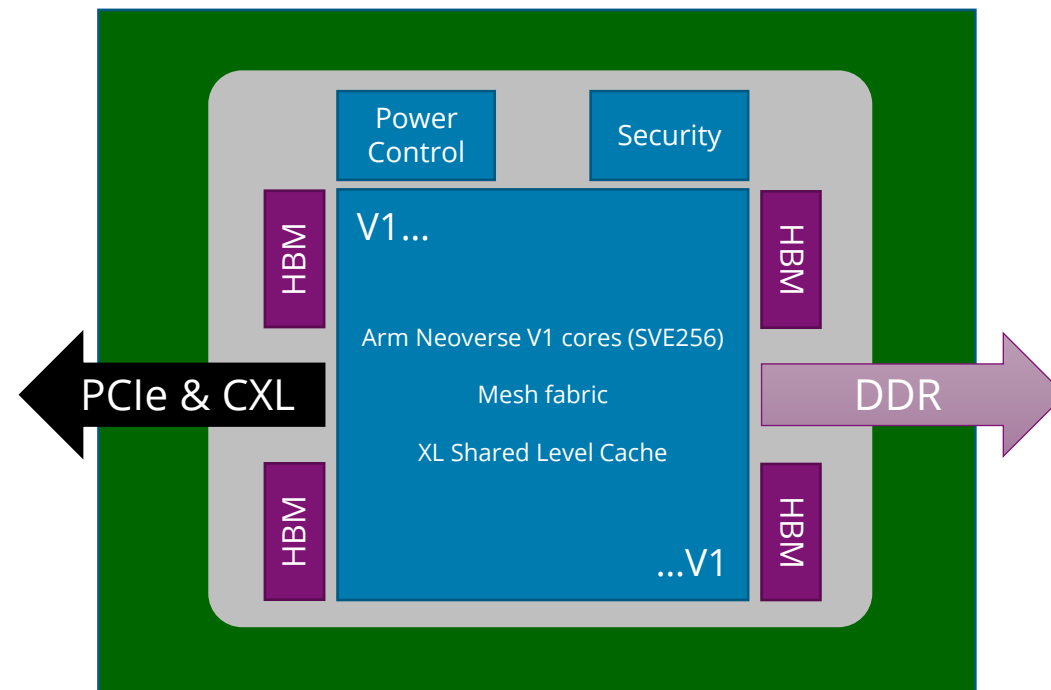


# At the heart of Rhea

With its high-performance, low-power Arm Neoverse V1 architecture, Rhea will meet the needs of all supercomputing workloads.

## Key features

Core	<ul style="list-style-type: none"> <li>- Arm architecture</li> <li>- Neoverse V1 cores</li> <li>- SVE 256 per core supporting 64/32/BF16 and Int8</li> <li>- ArmVirtualization extensions</li> </ul>
SoC	<ul style="list-style-type: none"> <li>- Arm mesh fabric</li> <li>- Advanced RAS support including Arm RAS extensions</li> <li>- Link protection for NoC &amp; high-speed IO</li> <li>- ECC support for selected memory</li> </ul>
Cache	<ul style="list-style-type: none"> <li>- Large L3 (Shared Level Cache)</li> <li>- RAS supported for all cache levels</li> </ul>
Memory	<ul style="list-style-type: none"> <li>- HBM2e</li> <li>- And DDR5</li> <li>- ECC for memory and link protection for controllers</li> </ul>
High Speed I/O	<ul style="list-style-type: none"> <li>- PCIe, CCIX &amp; CXL</li> <li>- Root and endpoint support</li> </ul>
Other I/O	<ul style="list-style-type: none"> <li>- USB, GPIO, SPI, I<sup>2</sup>C</li> </ul>
Power Management	<ul style="list-style-type: none"> <li>- Power management block to optimize perf/watt accross use cases and workloads.</li> </ul>
Security Block Support	<ul style="list-style-type: none"> <li>- Secure boot and secure upgrade</li> <li>- Crypto</li> <li>- True random number generation</li> <li>- Made in Europe</li> </ul>



Rhea will deliver extraordinary real compute performance and efficiency with an unmatched Bytes/Flops ratio.



## About SiPearl

Created by Philippe Notton, SiPearl is designing the high-performance, low-power microprocessor for European exascale supercomputers. This new generation of microprocessors will enable Europe to set out its technological sovereignty in strategic high performance computing markets such as artificial intelligence, medical research or climate modelling.

SiPearl is working in close collaboration with its 27 partners from the European Processor Initiative (EPI) consortium - leading names from the scientific community, supercomputing centres and industry - which are its stakeholders, future clients and end-users.

SiPearl employs 109\* people in France, Germany and Spain. Its first range of microprocessors, Rhea, will be launched at the end of the year.

The company is supported by the European Union (funding from the European Union's Horizon 2020 research and innovation program under specific grant agreement no.826647).

\* as of June 15th 2022

---

Contact  
Jean-Marc Denis  
Chief Strategy Officer  
[jean-marc.denis@sipearl.com](mailto:jean-marc.denis@sipearl.com)

