Meteor Lake Architecture Overview

intel

Advancing Moore's Law
5 NODES IN 4 YEARS

Meteor Lake Pillars

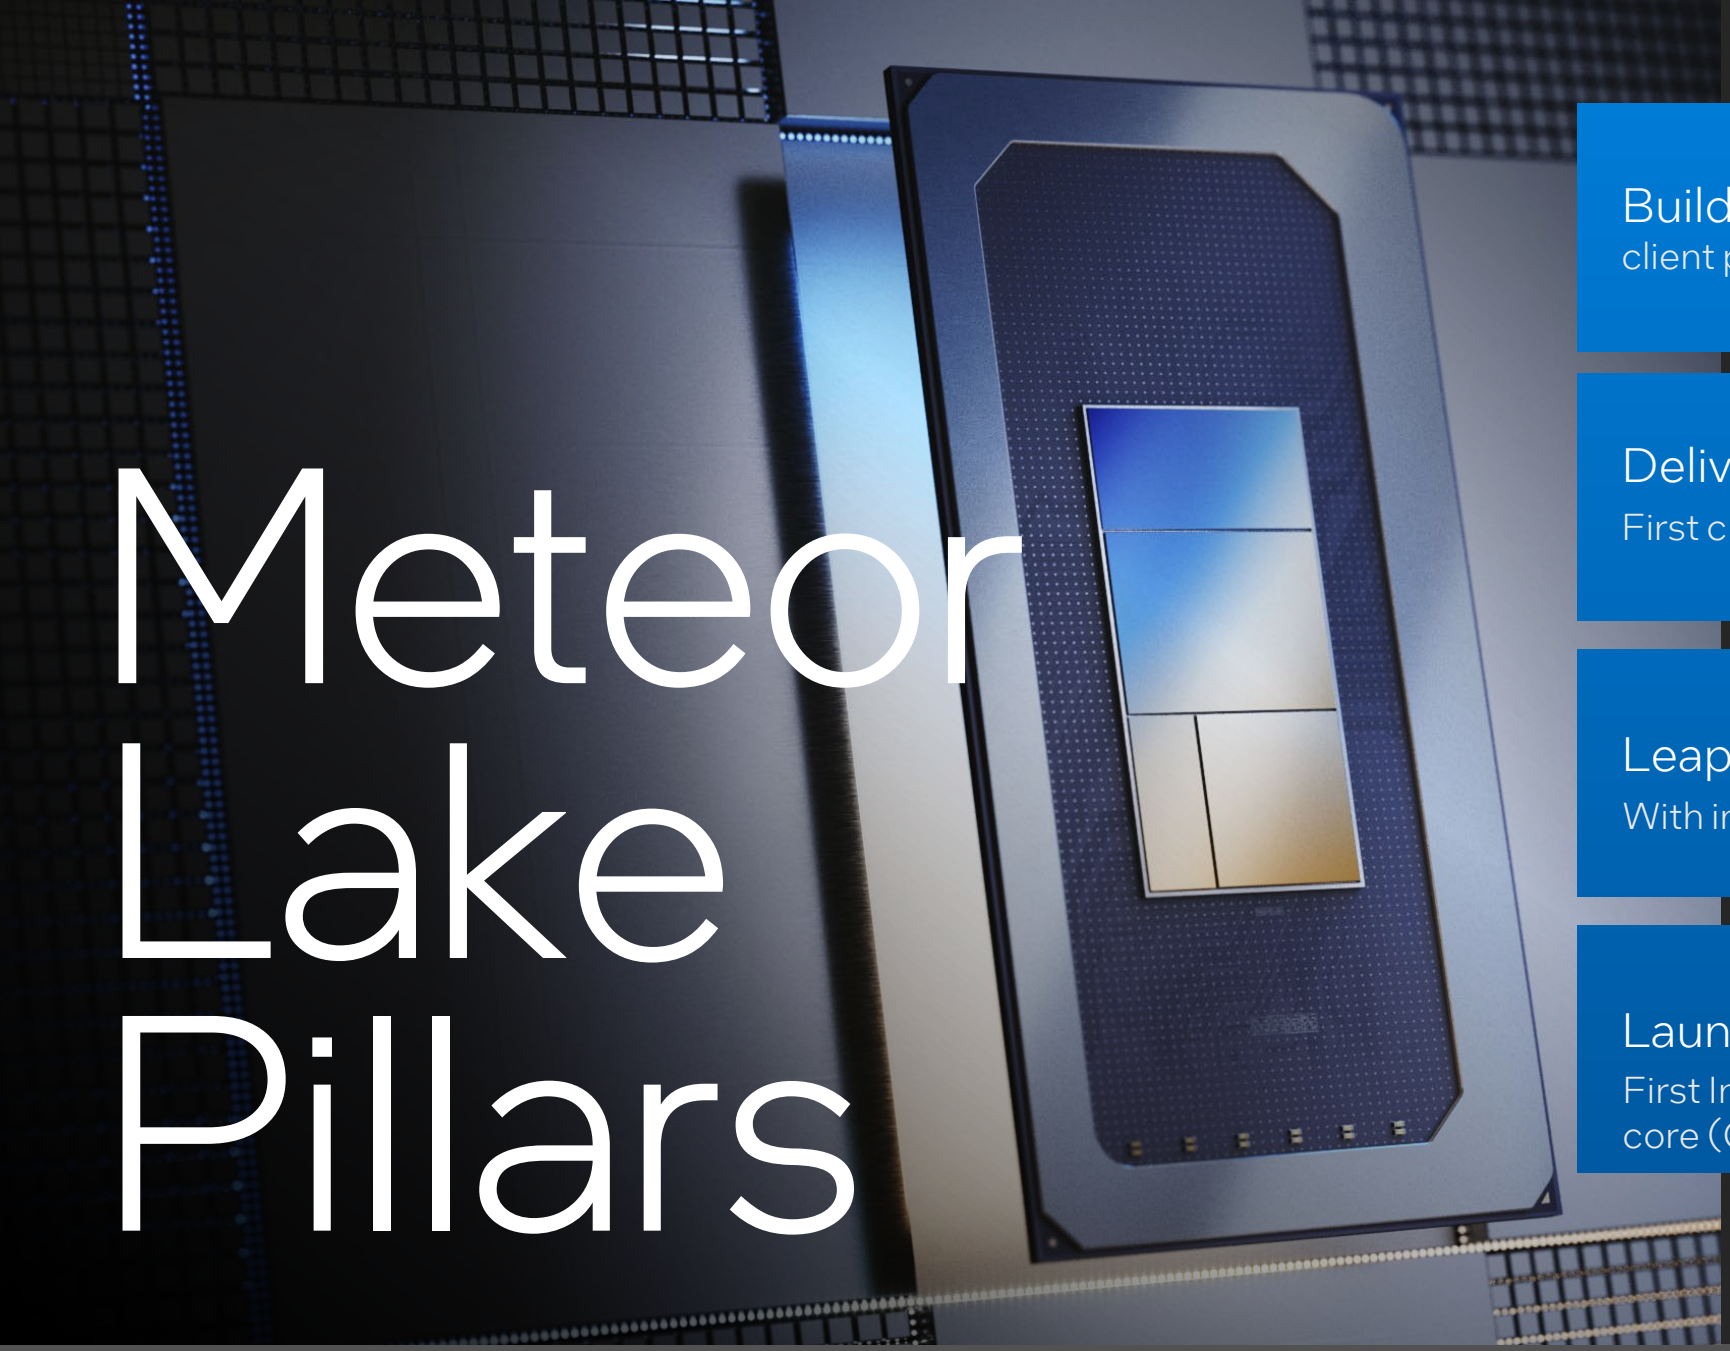

Build our most **power-efficient**
client processor in history

Deliver **AI at Scale**
First client integration of AI engine (NPU)

Leap in **graphics** performance
With increased power efficiency

Launch IA on **Intel 4**
First Intel 4 P-core (Redwood Cove) & E-core (Crestmont)

INTRODUCING

Meteor Lake

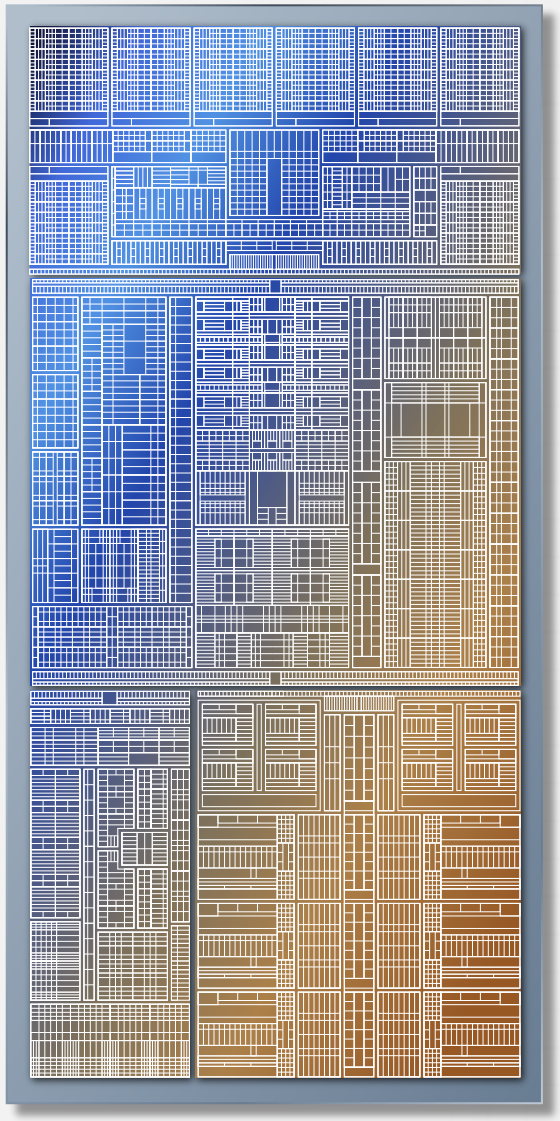

INTRODUCING

Meteor Lake

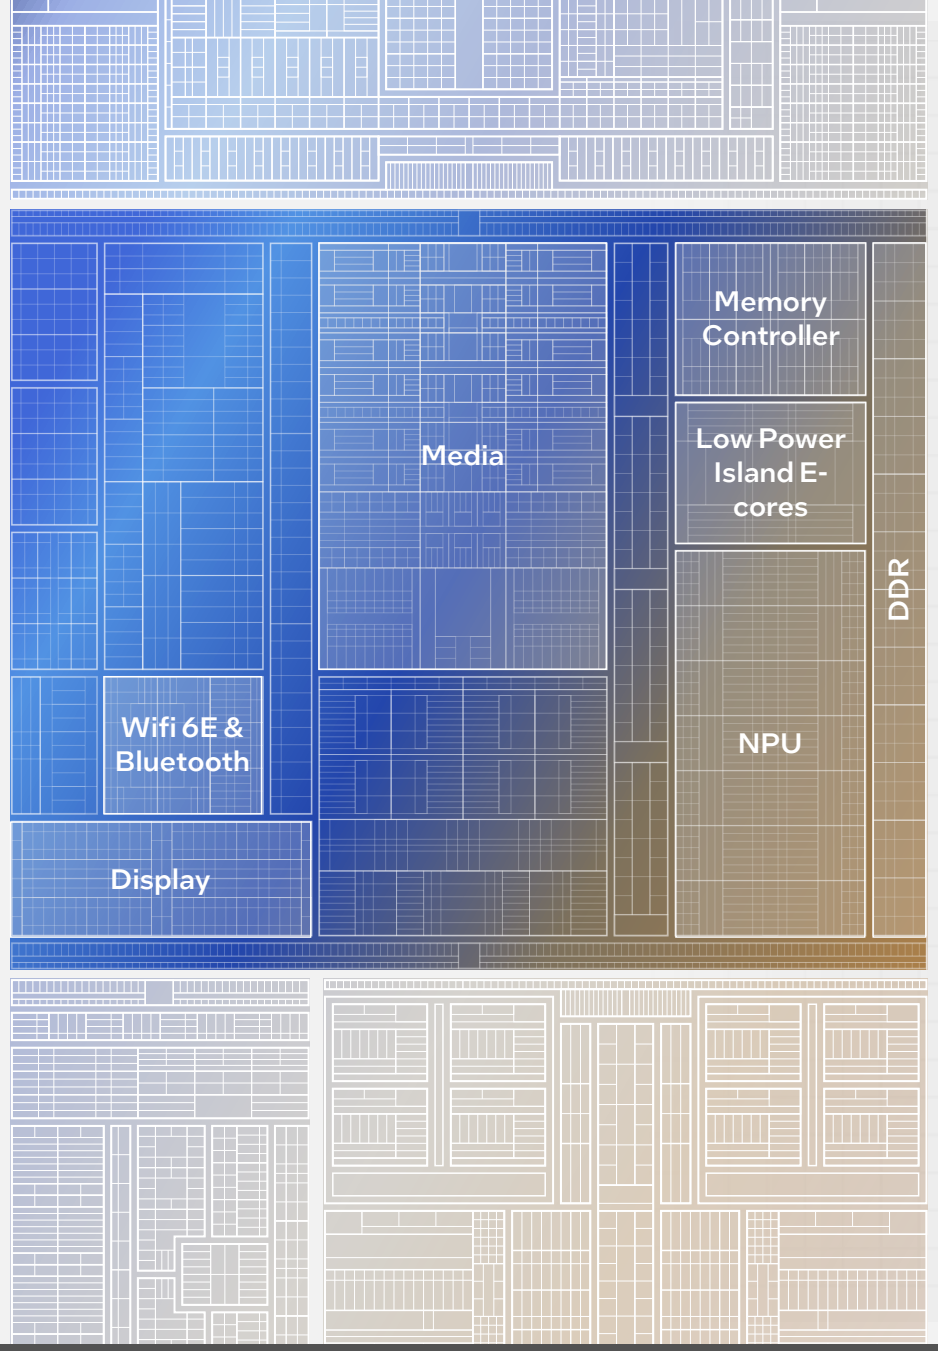

SOC Tile

- New **low power island** E-cores
- First built-In **NPU AI Engine**
- Leading **Wi-Fi 6E & Wi-Fi 7** support
- **8K HDR & leading AV1** support
- Native **HDMI 2.1 and DP 2.1** standards
- Integrated **memory controller & DDR**

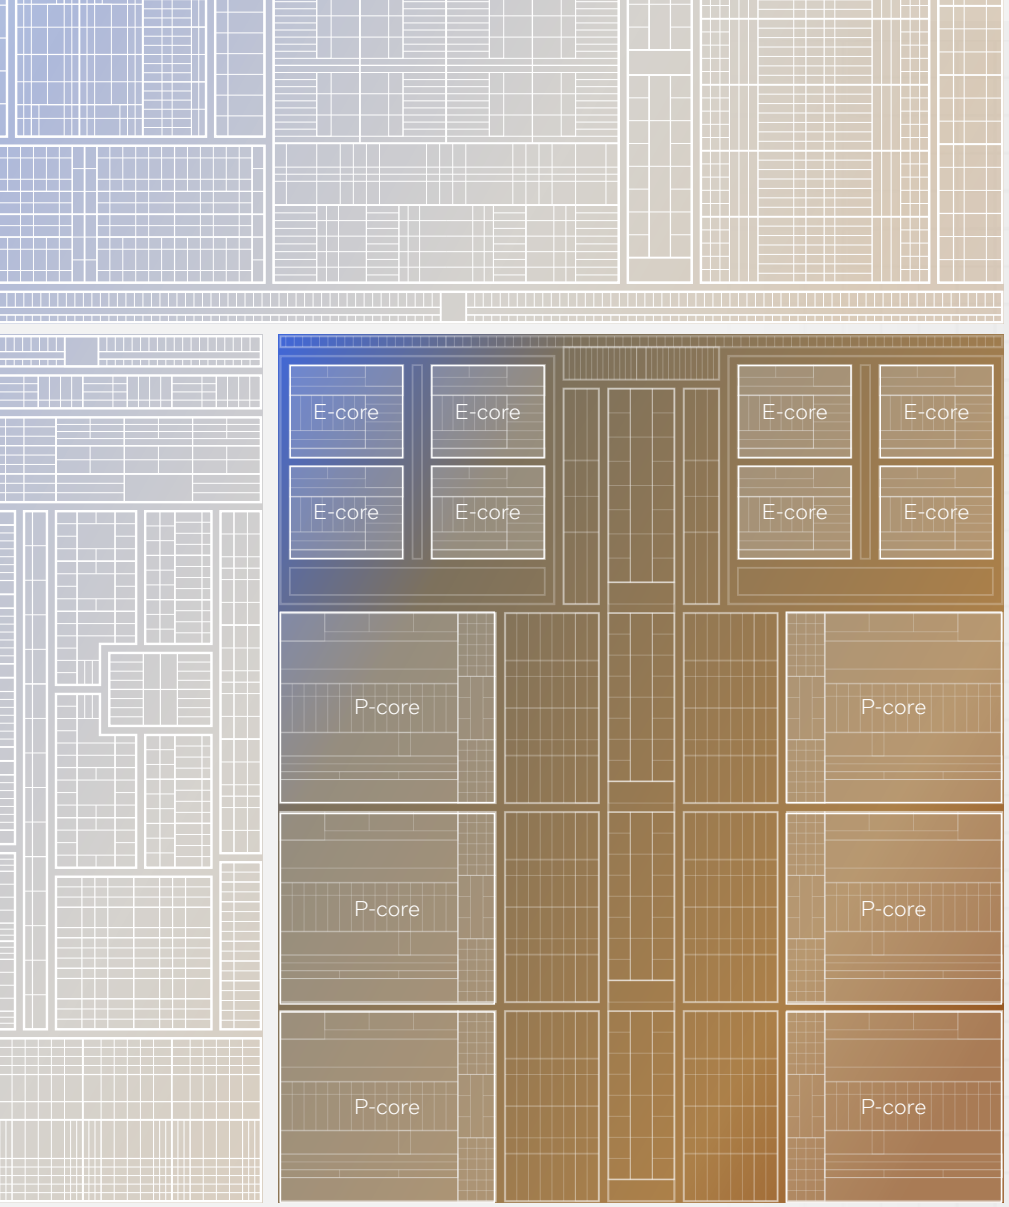

Compute Tile

New **E-core** microarchitecture

New **P-core** microarchitecture

First on **Intel 4** process technology

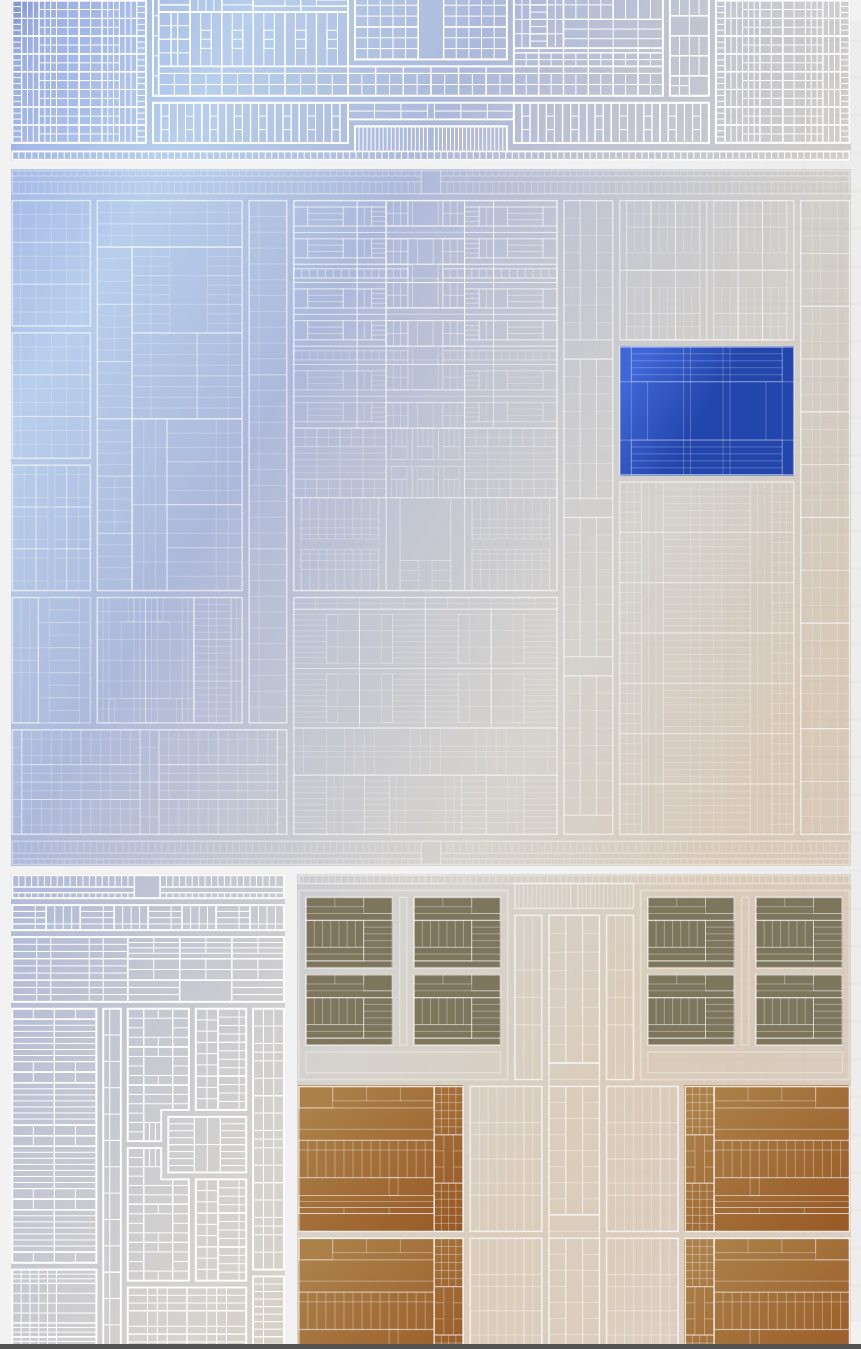

3D performance hybrid architecture

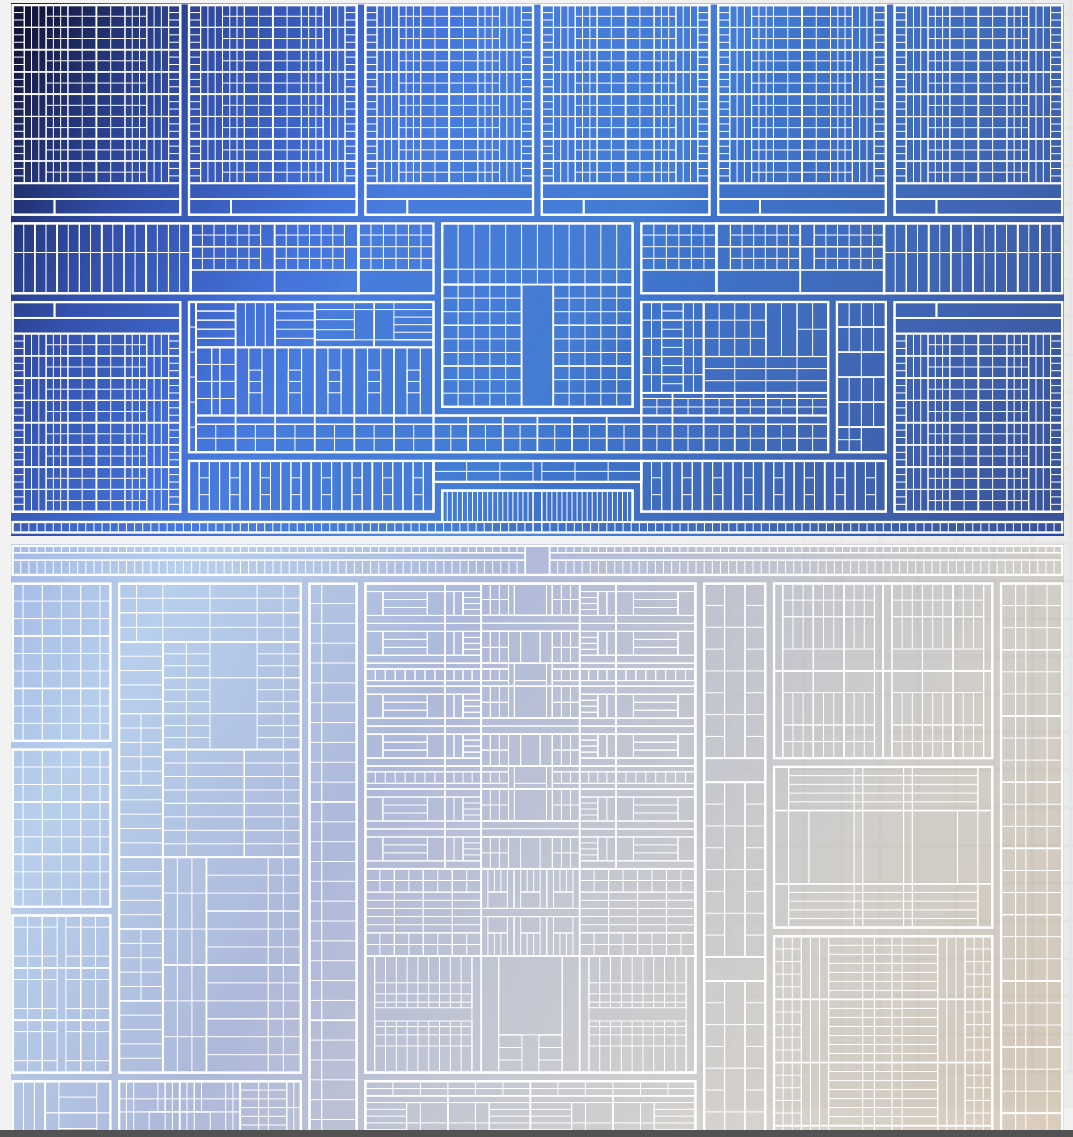

intel®
ARC™
Graphics

Intel® Arc™ graphics only available on select MTL processor-powered systems with dual-channel memory.

IO Tile

Industry leading connectivity with integrated **Thunderbolt 4** & **PCIe Gen5**

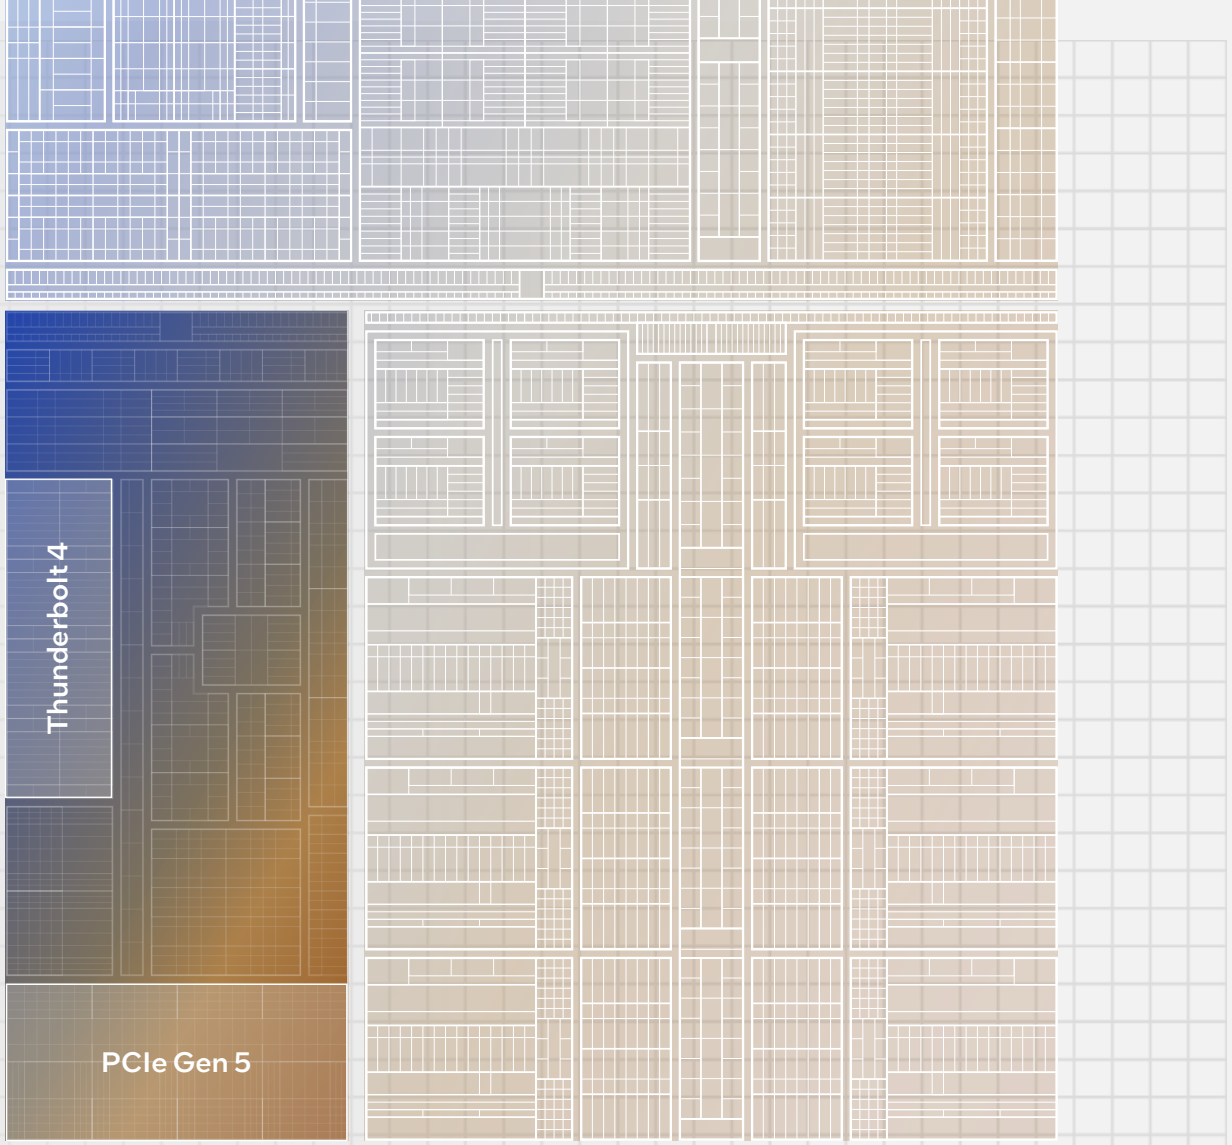

Meteor Lake

- Industry leading **Wi-Fi 6E & Wi-Fi 7** support
- Industry leading **FOVEROS 3D packaging**
- Industry leading **Thunderbolt 4**
- Latest connectivity **PCIe Gen5**

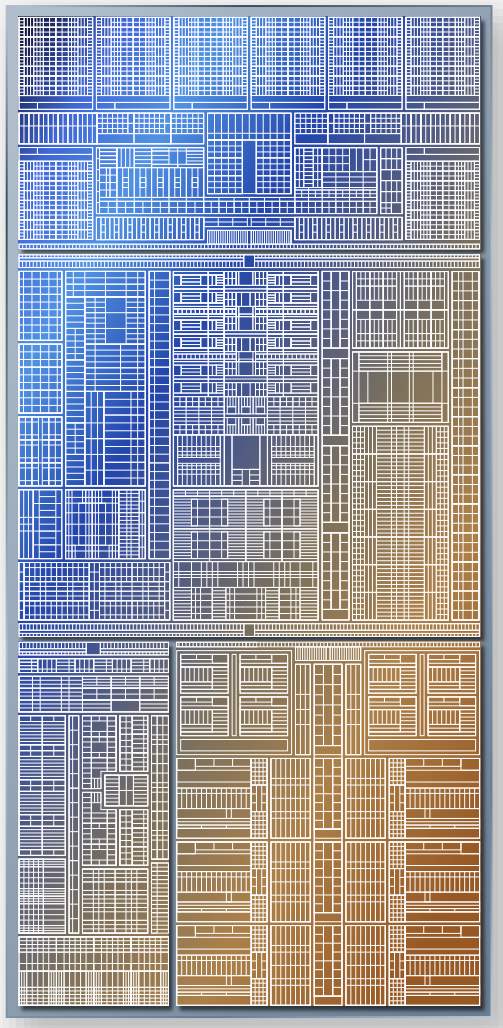

- **New** **Intel Arc** graphics
- **New** **Media & Display** standards
- **New** **Low power island** E-cores
- **First** Built-In **NPU AI engine**
- **First** **3D performance hybrid** architecture
- **New** **P-core & E-core** microarchitectures
- **First** on **Intel 4** process technology

Intel's largest client SoC architectural shift in 40 years

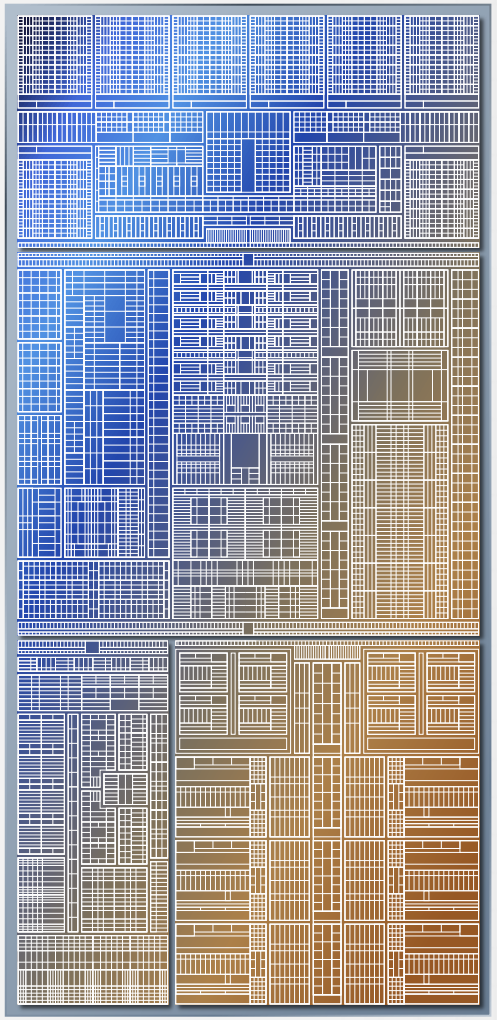


Industry leading
FOVEROS 3D
packaging

Next-gen Uncore Guiding Principles

1. Repartition compute intensive IPs for **power optimization**
2. Enable IO bandwidth **scalability**
3. Extend hybrid architecture with the addition of **low power** IA cores
4. Re-construct **power management**

Prior Gen

Repartition Compute Intensive IP

Media IP is embedded in graphics IP

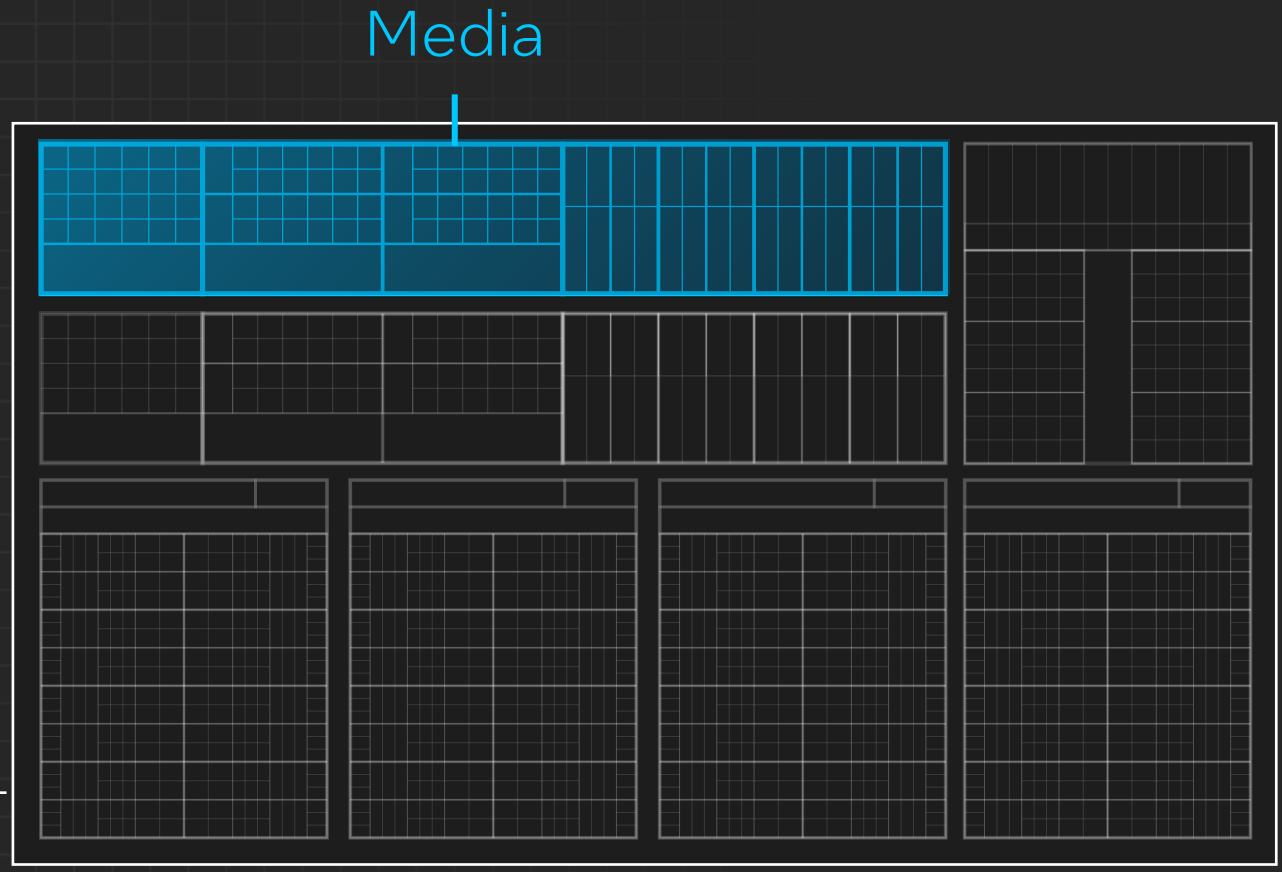

Prior Gen

Repartition Compute Intensive IP

- Media IP is embedded in graphics IP
- Graphics attached to core complex

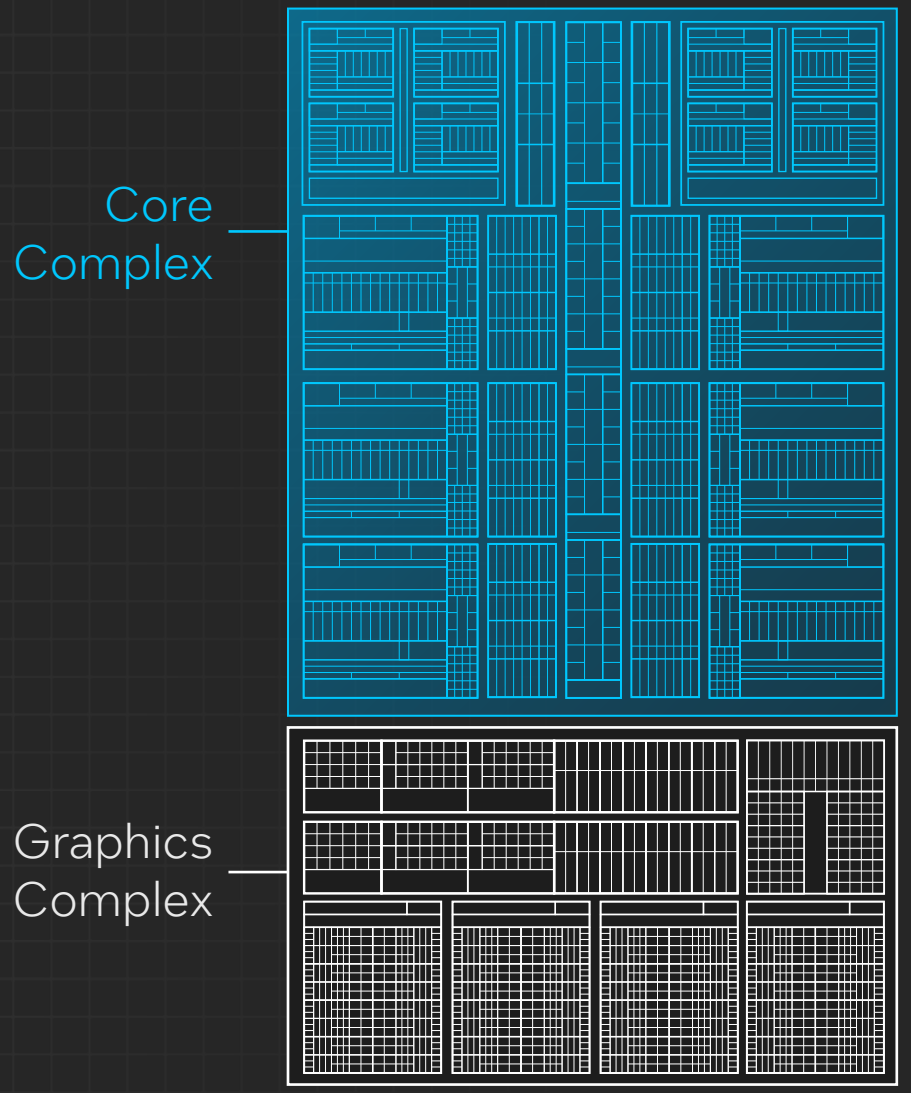

Prior Gen

Repartition Compute Intensive IP

- Media IP is embedded in graphics IP
- Graphics attached to core complex
- All use same ring fabric
- Ring fabric only way to access mem

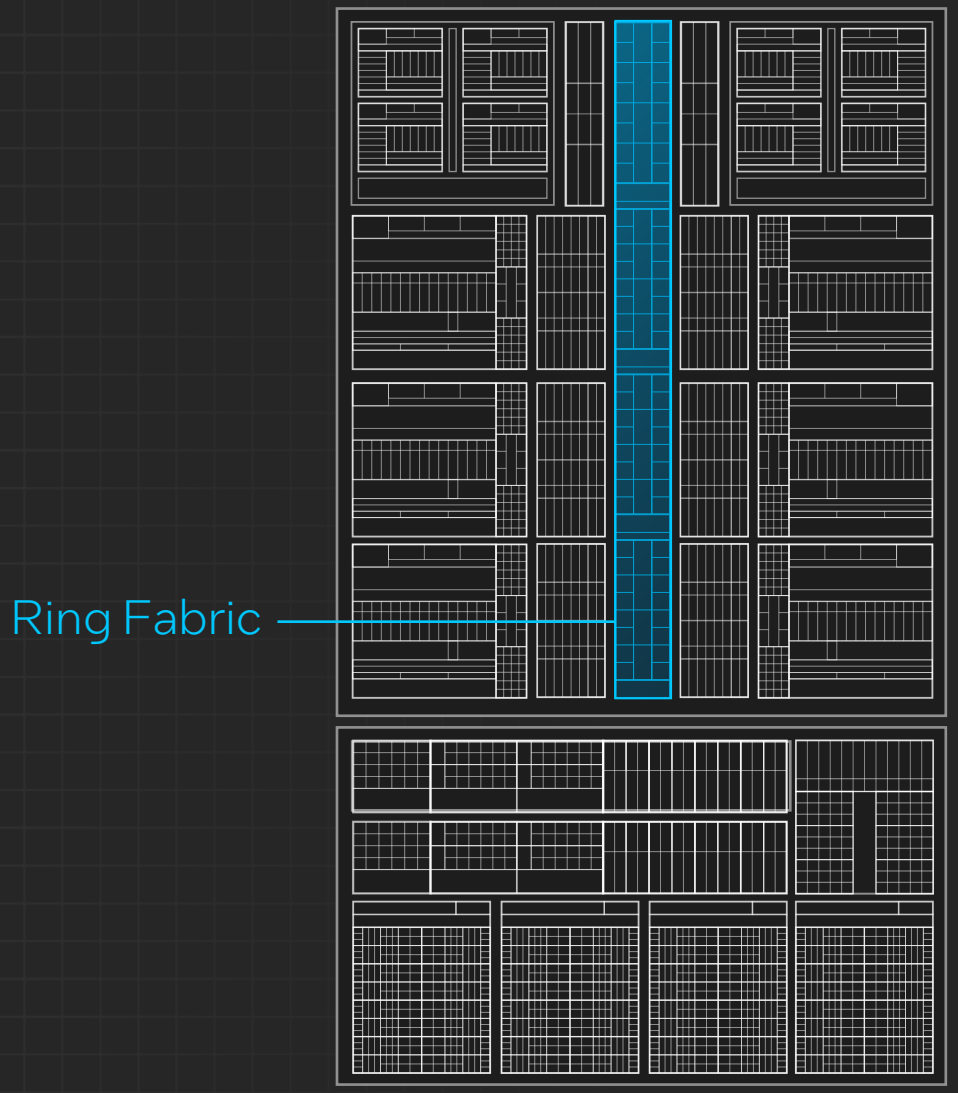

Repartition Compute Intensive IP

- Media IP is embedded in graphics IP
- Graphics attached to core complex
- All use same ring fabric
- Only way to access mem & cache

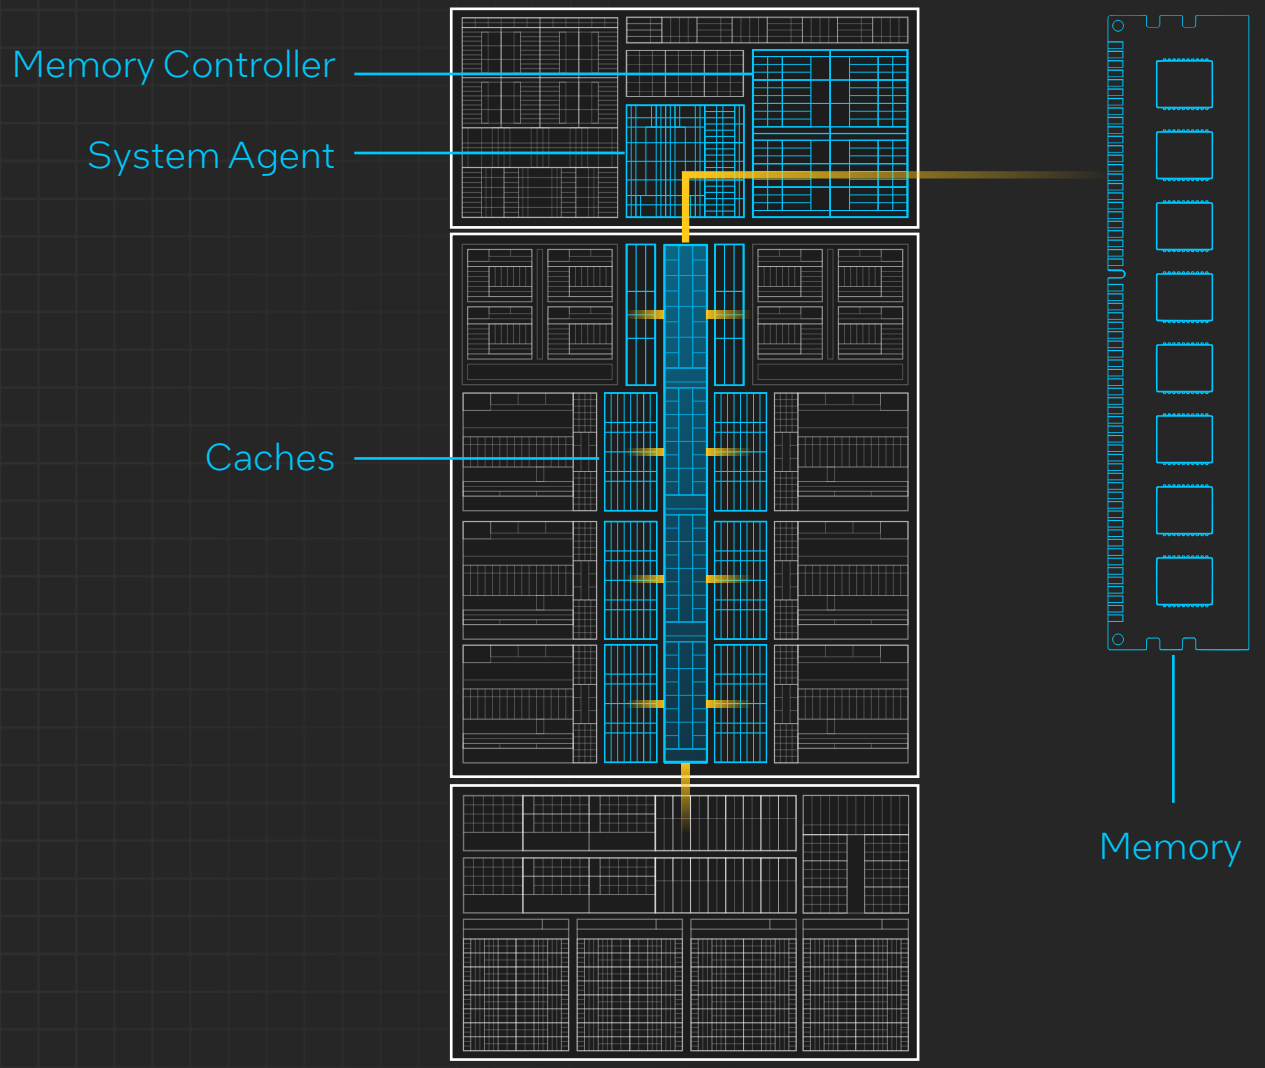

Repartition Compute Intensive IP

Core complex stays on if either graphics or media need to access memory

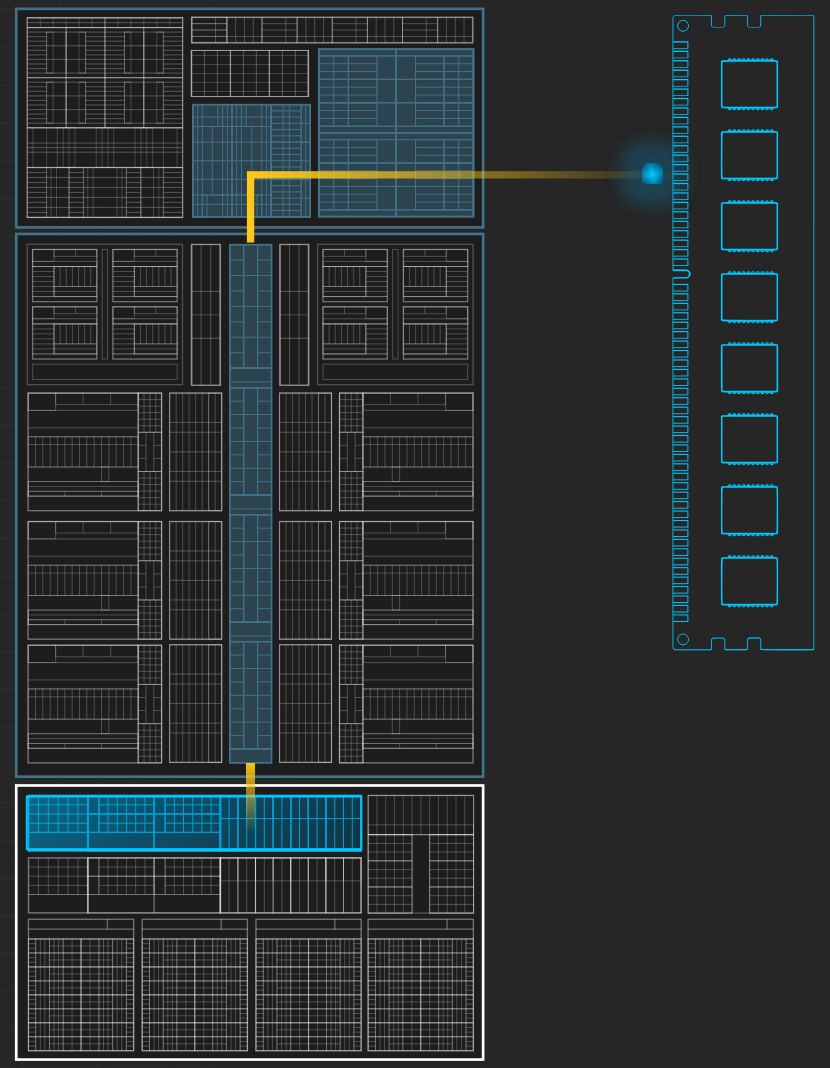

Compute Intensive IP in Meteor Lake

Media separated from graphics

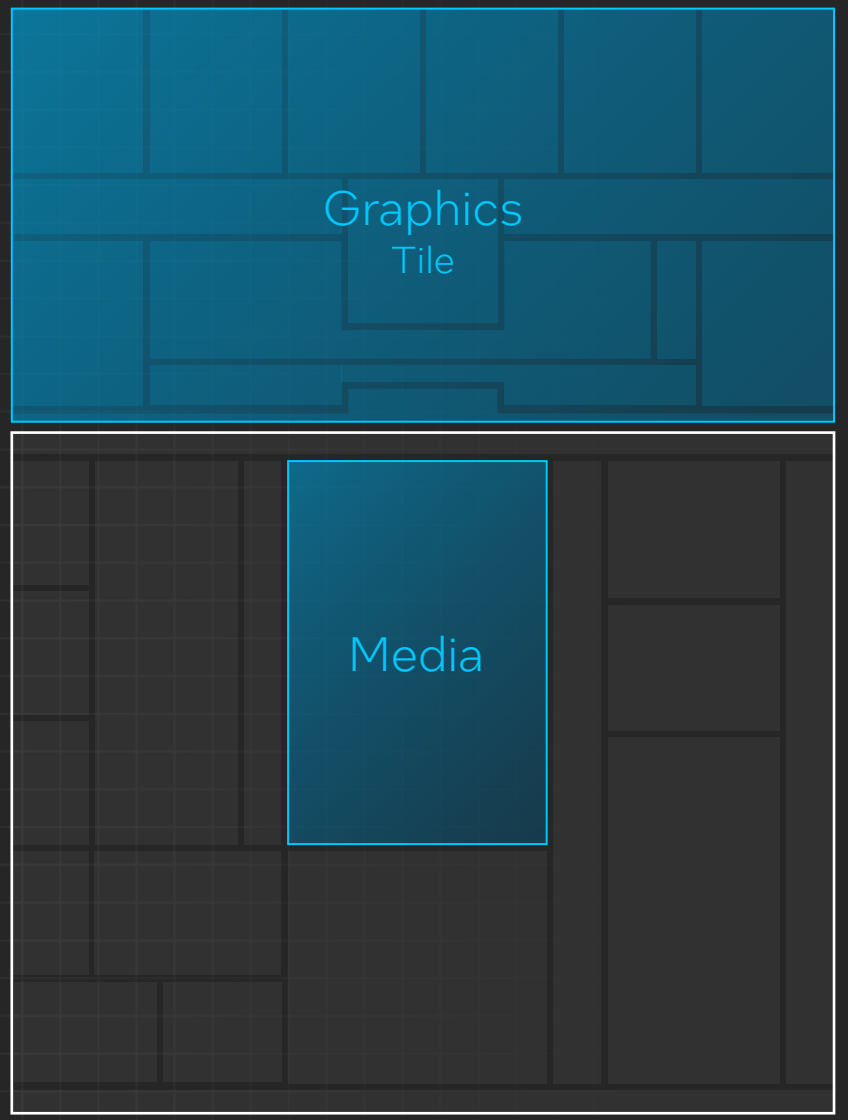

Compute Intensive IP in Meteor Lake

- Media separated from graphics
- Both independently attached to SOC

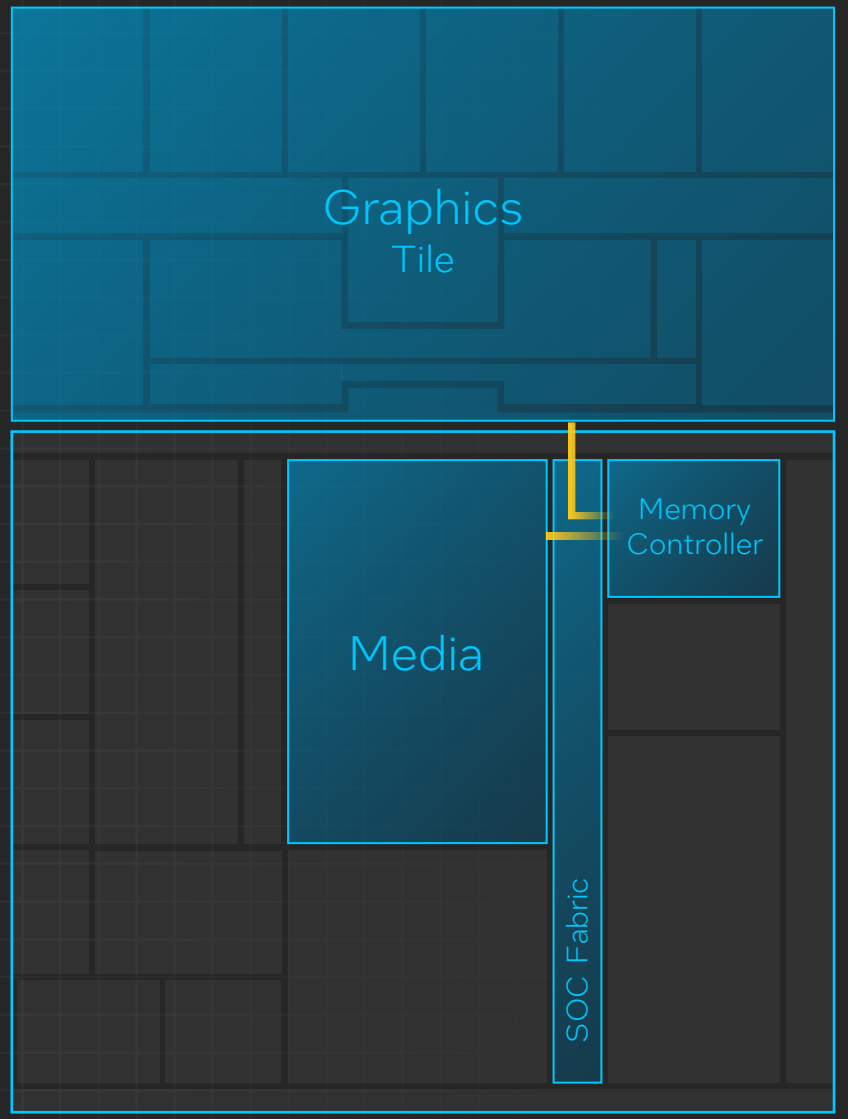

Compute Intensive IP in Meteor Lake

- Media separated from graphics
- Both independently attached to SOC
- Independent core complex

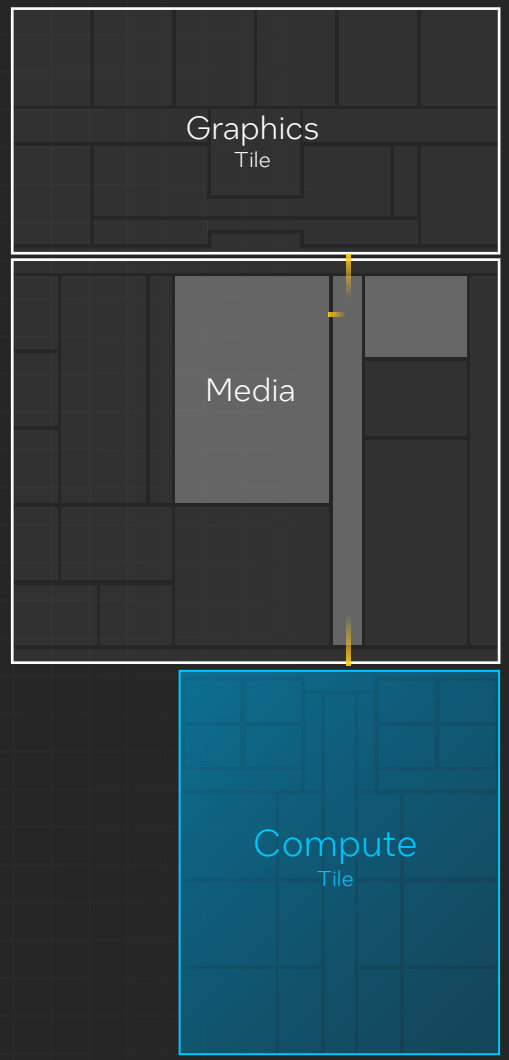

Compute Intensive IP in Meteor Lake

- Media separated from graphics
- Both independently attached to SOC
- Independent core complex
- All IPs have independent paths to memory

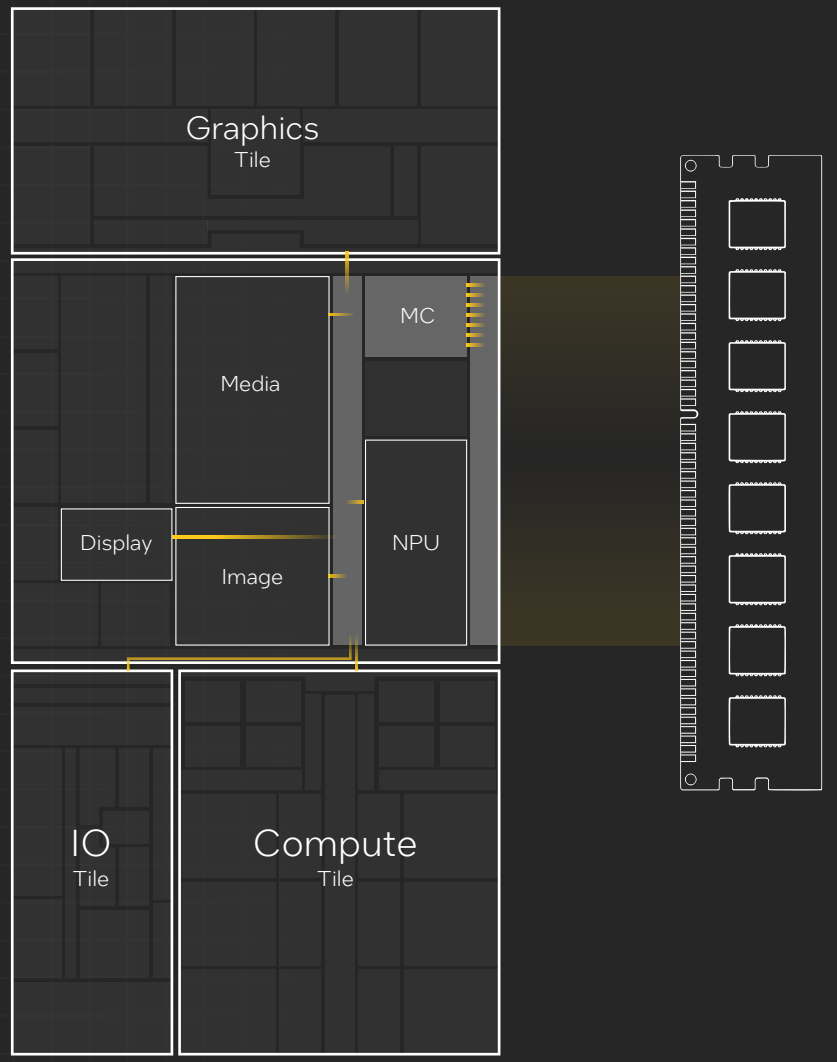

Compute Intensive IP in Meteor Lake

- Media separated from graphics
- Both independently attached to SOC
- Independent core complex
- All IPs have independent paths to memory
- All IPs can be independently powered on/off

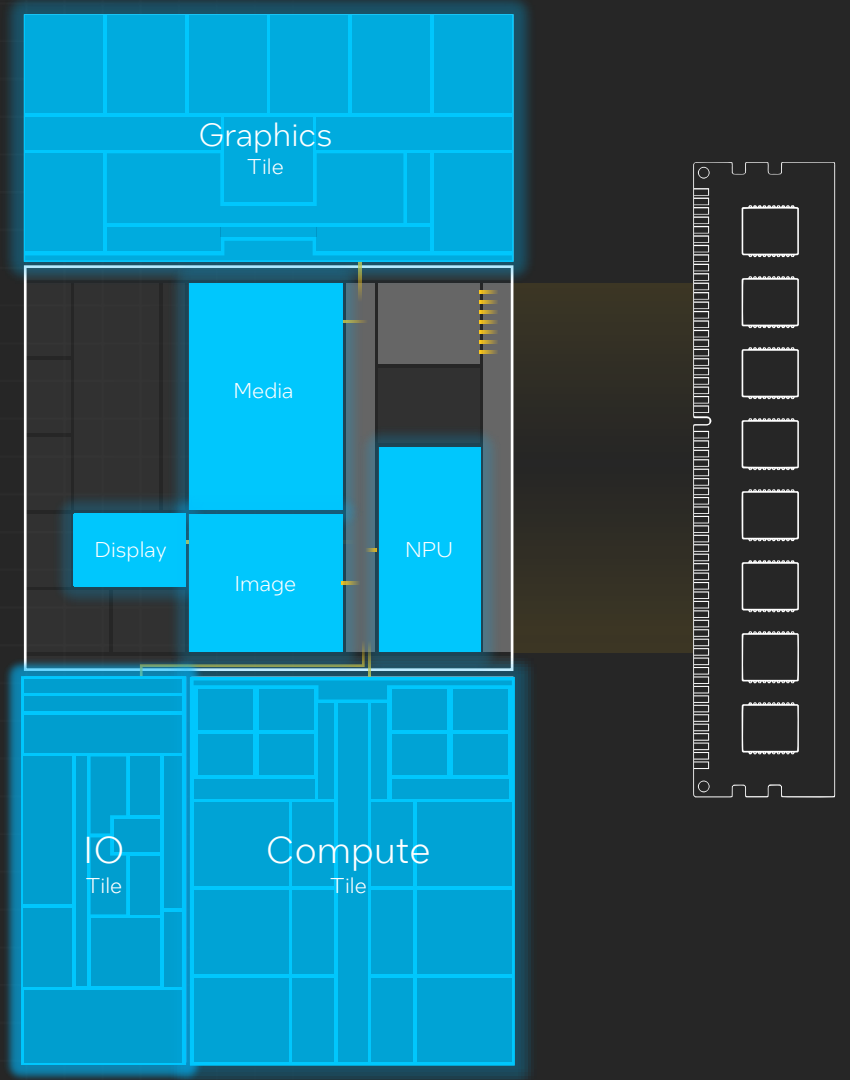

Compute Intensive IP in Meteor Lake

- Media separated from graphics
- Both independently attached to SOC
- Independent core complex
- All IPs have independent paths to memory
- All IPs can be independently powered on/off → **media playback**

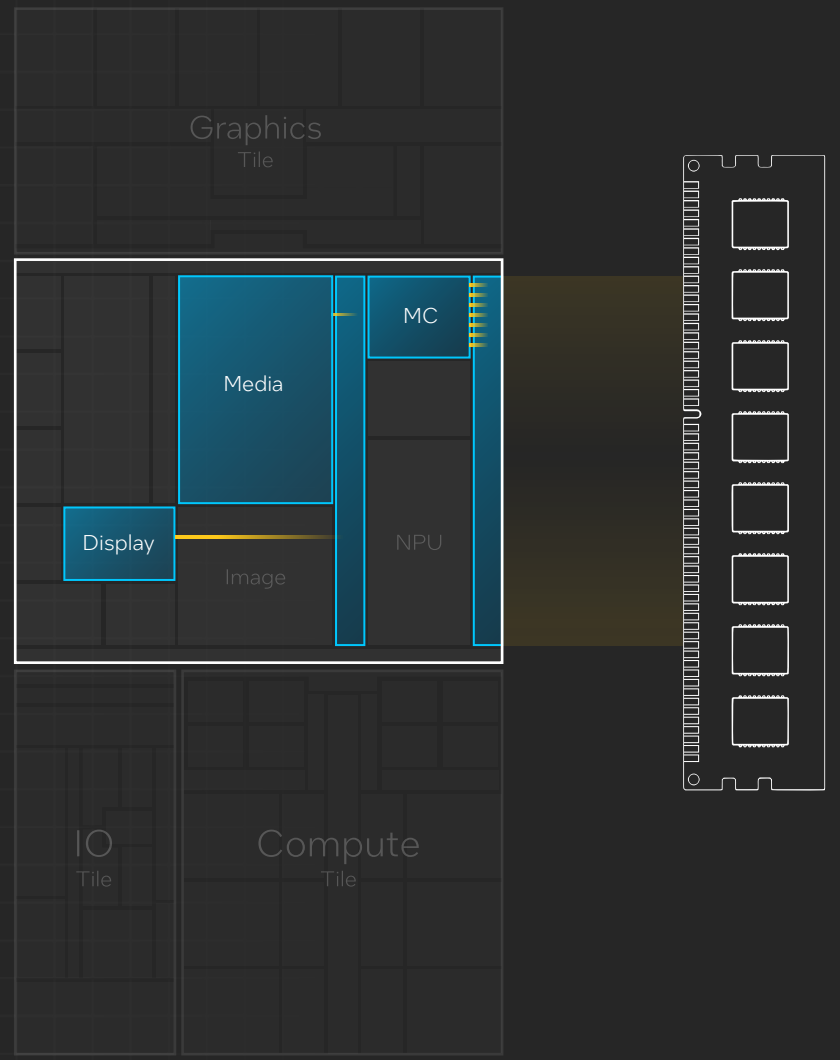

I/O Bandwidth Scalability

- Media/Graphics/NPU/IO tile add significant traffic
- IO Port now critical bottleneck

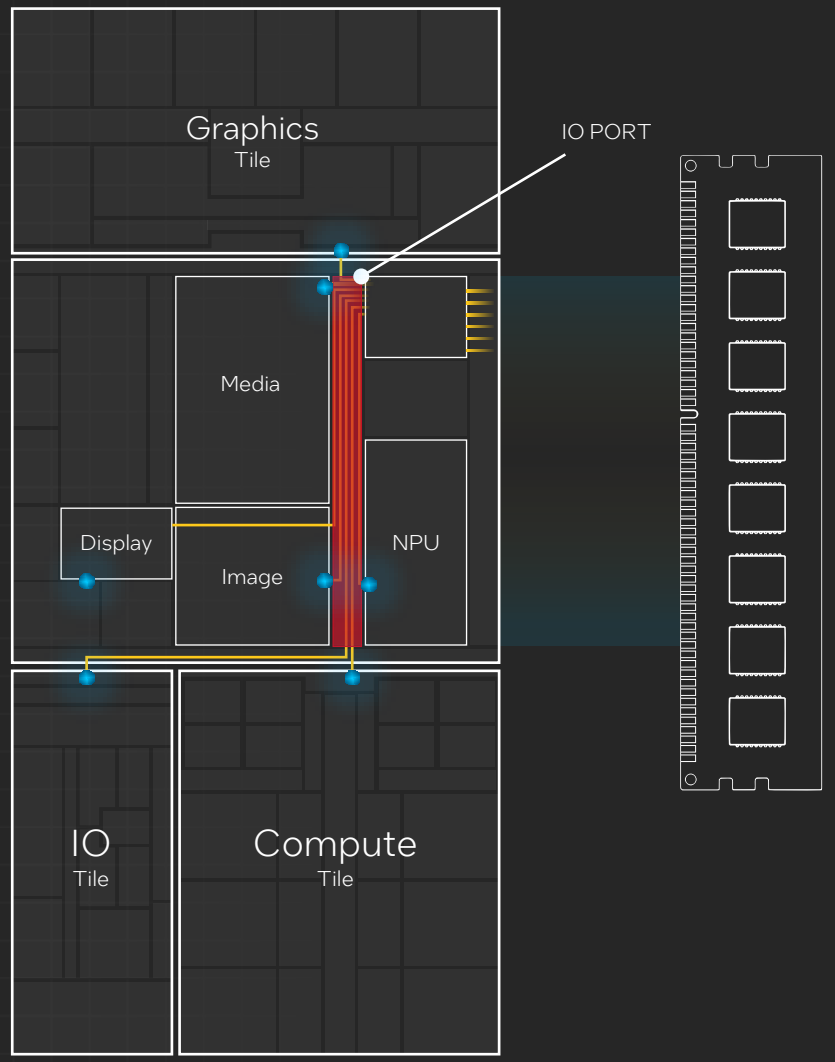

I/O Bandwidth Scalability

Media/Graphics/NPU/IO tile add significant traffic

IO Port now critical bottleneck

SOLUTION #1

Custom side channels for each IP → **Not scalable**

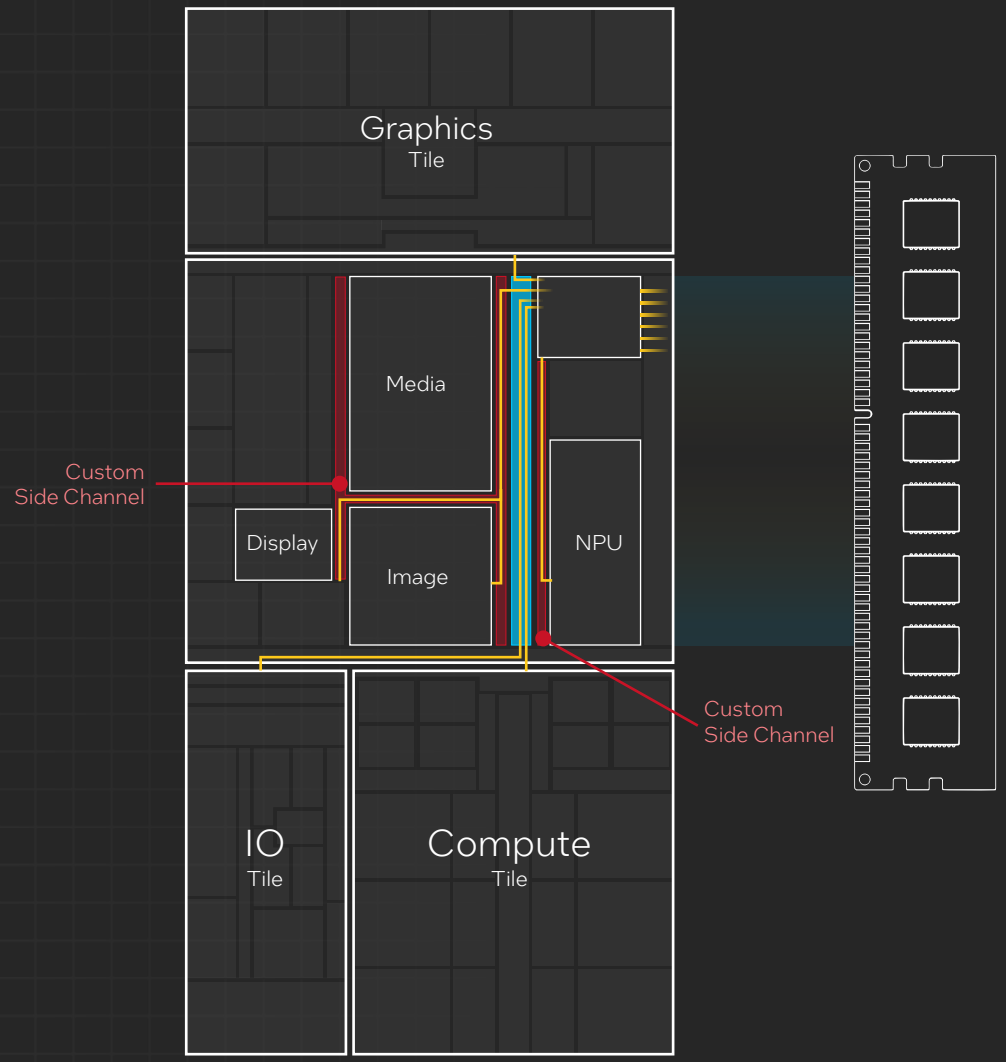

I/O Bandwidth Scalability

Media/Graphics/NPU/IO tile add significant traffic

IO Port now critical bottleneck

CHOSEN SOLUTION

New Scalable Fabric for high BW (128 GBs) connectivity

All IO ordering and address translation goes through IOC

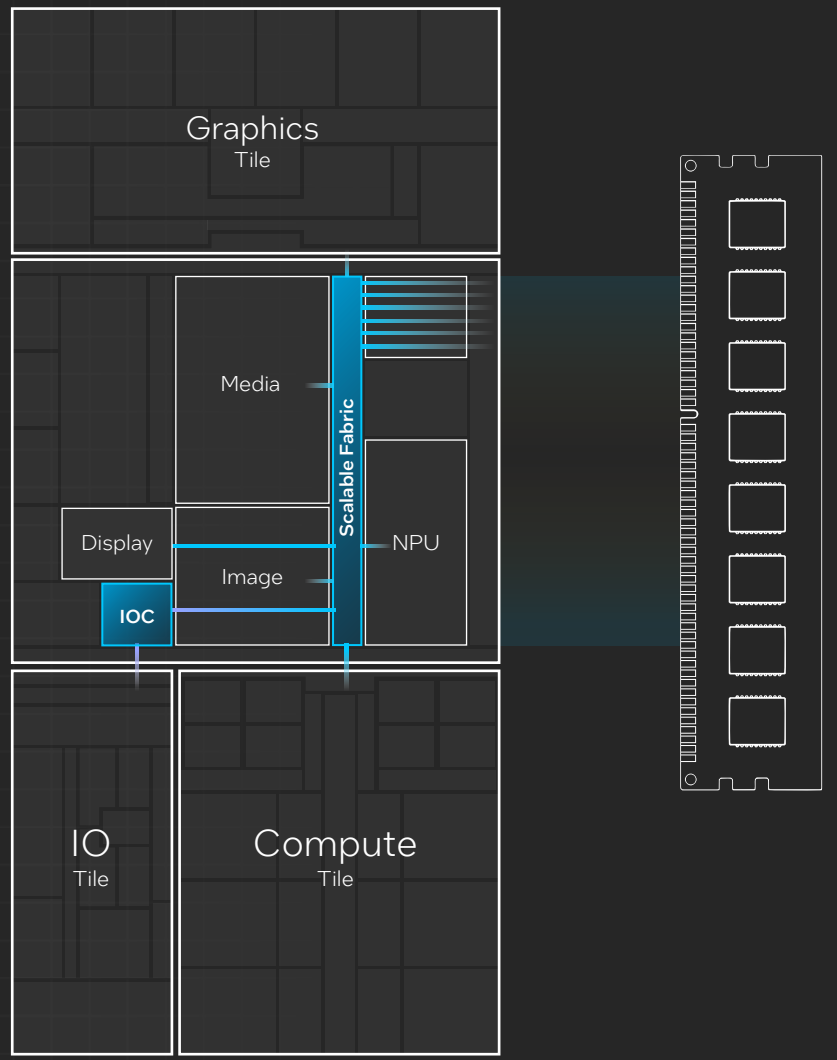

Next Evolution of Our Hybrid Architecture

OPPORTUNITY

IA complex is woken up even for low compute intensity workloads

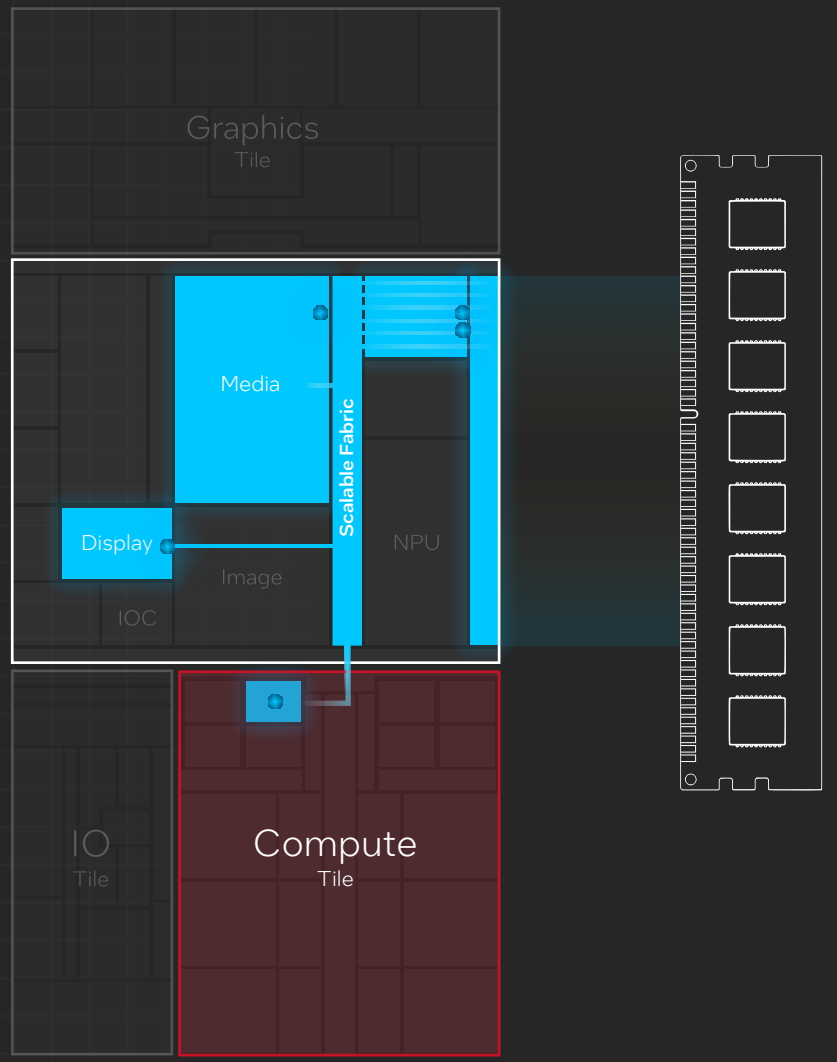

Next Evolution of Our Hybrid Architecture

OPPORTUNITY

IA complex are woken up even for low compute intensity workloads

SOLUTION

Lower power E-cores on SOC

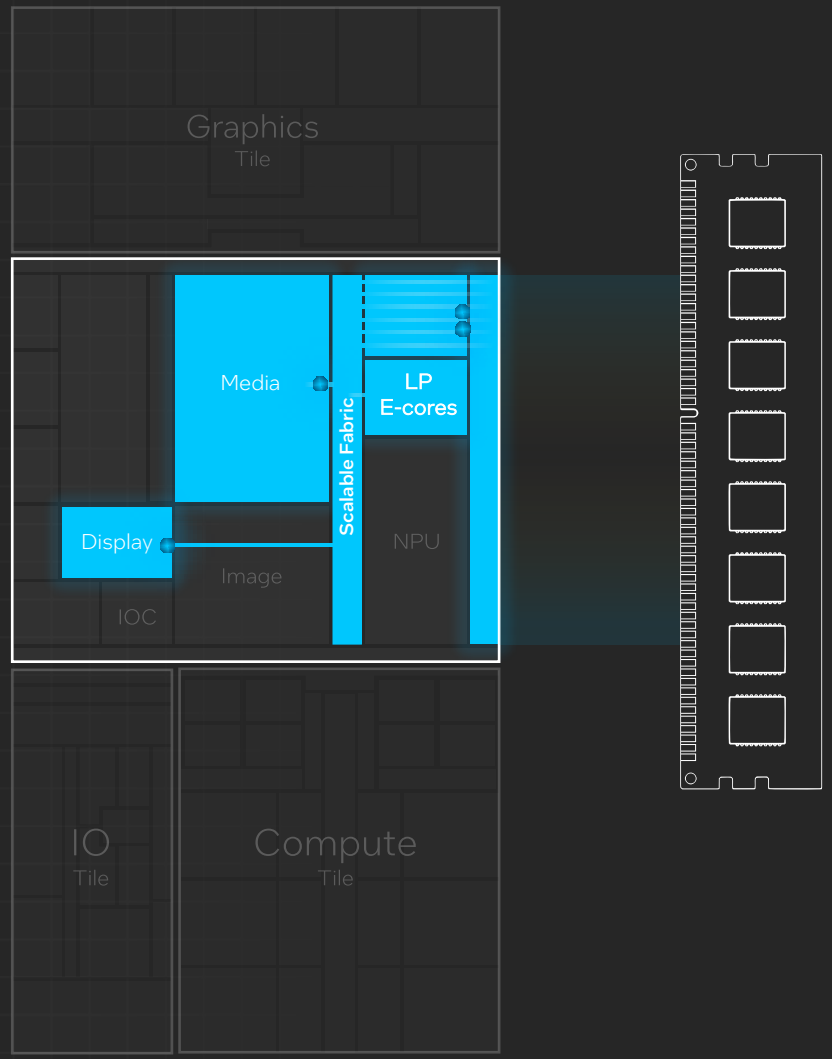

Re-constructing Power Management

- Grounds up **modular and scalable PM** architecture for disaggregation
- New **scalable fabric** for improved bandwidth and energy efficiency
- Coordination between **multiple PM controllers** on different tiles
- Coordination between SOC PM controllers and system software

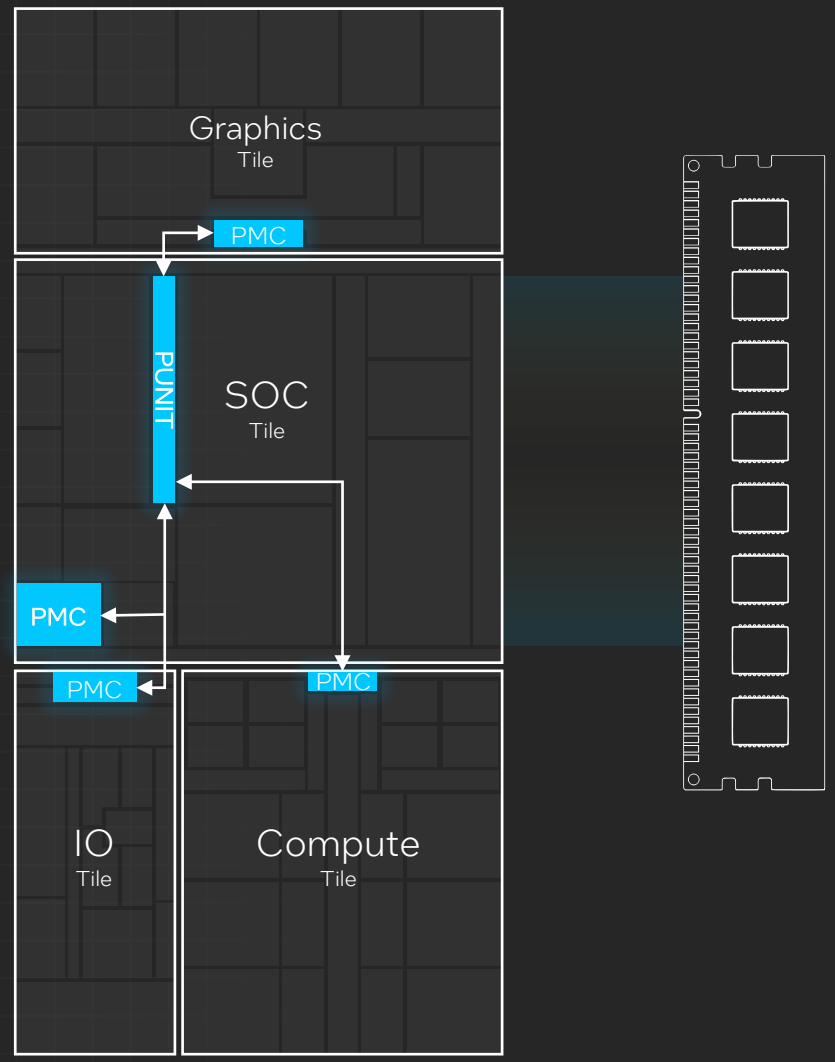

Next-gen Uncore Guiding Principles

1. Repartition compute intensive IPs for **power optimization**
2. Enable IO bandwidth **scalability**
3. Re-design of hybrid architecture with the addition of **low power** IA cores
4. Re-construct **Power Management**

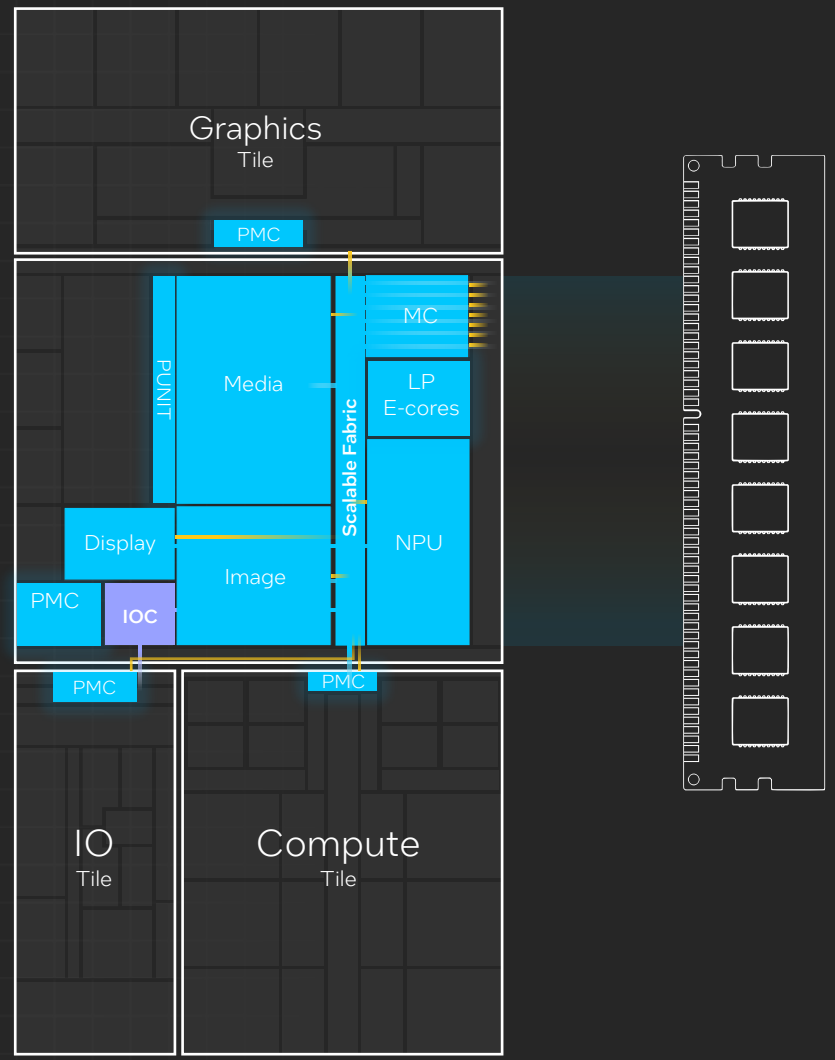

New Architectural Capabilities

AI

is Everywhere

INTRODUCING

Intel's First Integrated NPU

Dedicated AI Engine for Low Power Inference

Purpose built for efficient client AI

Ideal for sustained AI and AI offload

Standardized program interfaces

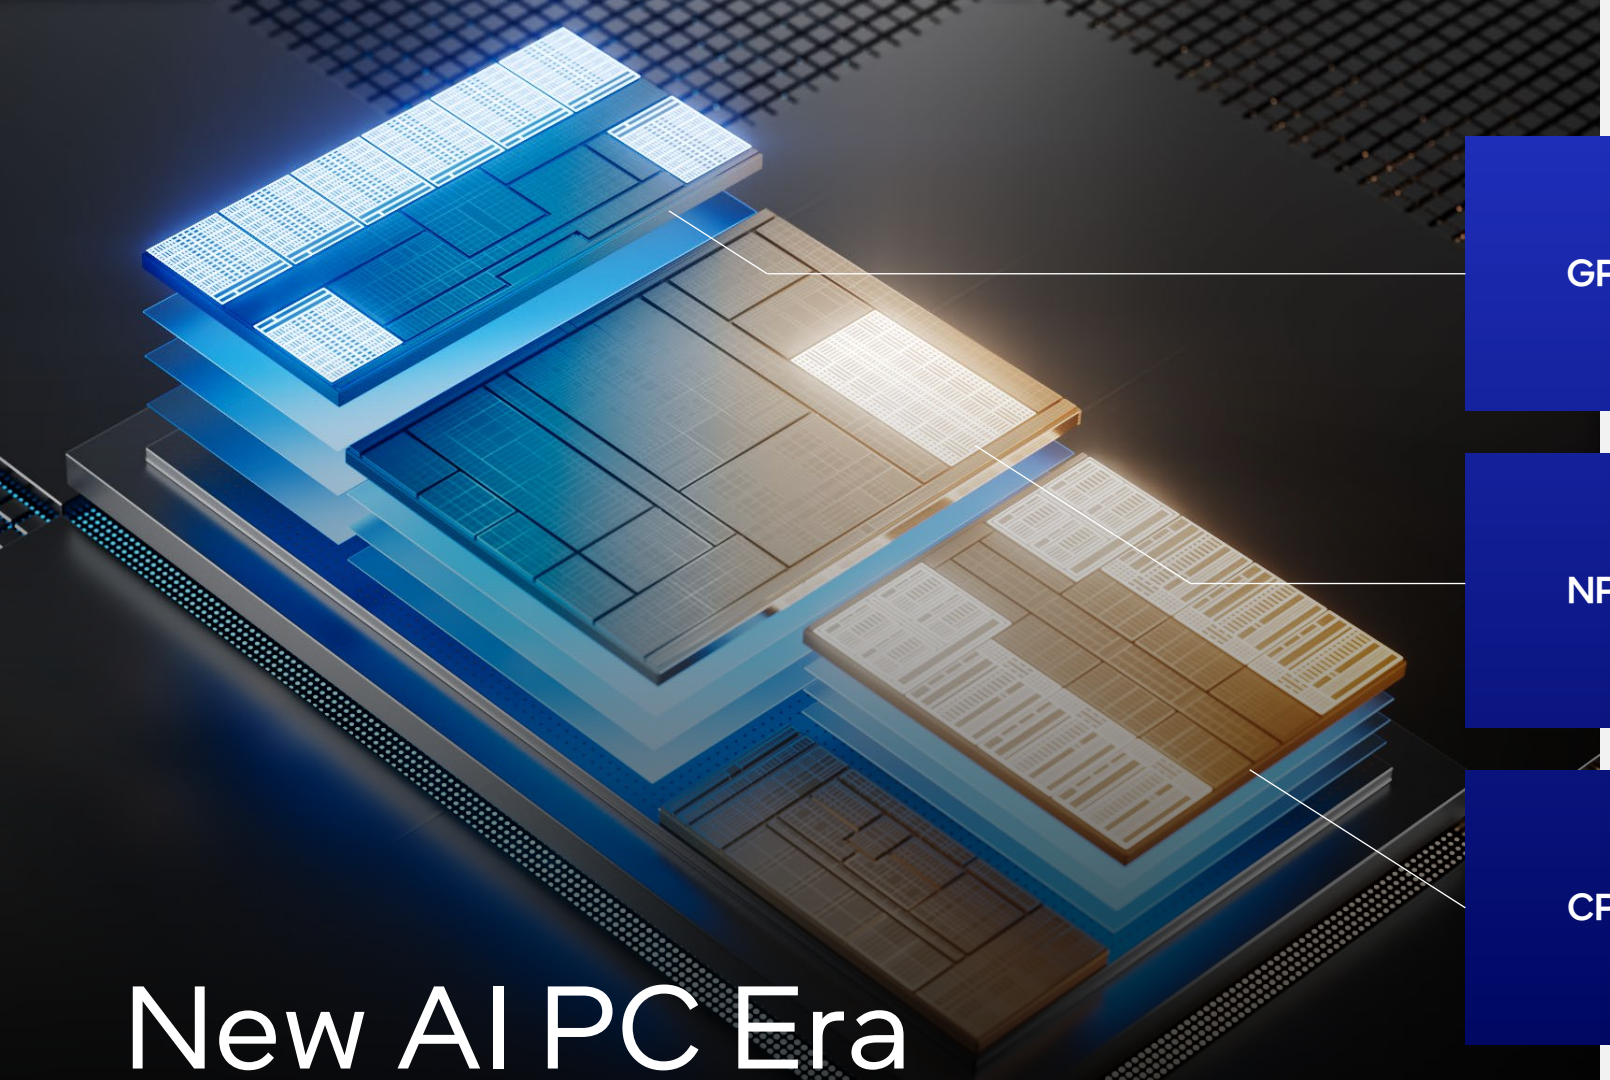

New AI PC Era
Powered By Meteor Lake

GPU

Performance Parallelism & Throughput
Ideal for AI infused in Media/3D/render pipeline

NPU

Dedicated Low Power AI Engine
Ideal for sustained AI and AI offload

CPU

Fast Response
Ideal for light-weight, single inference low-latency AI tasks

Xe LPG

Graphics IP

Scaling the graphics engine

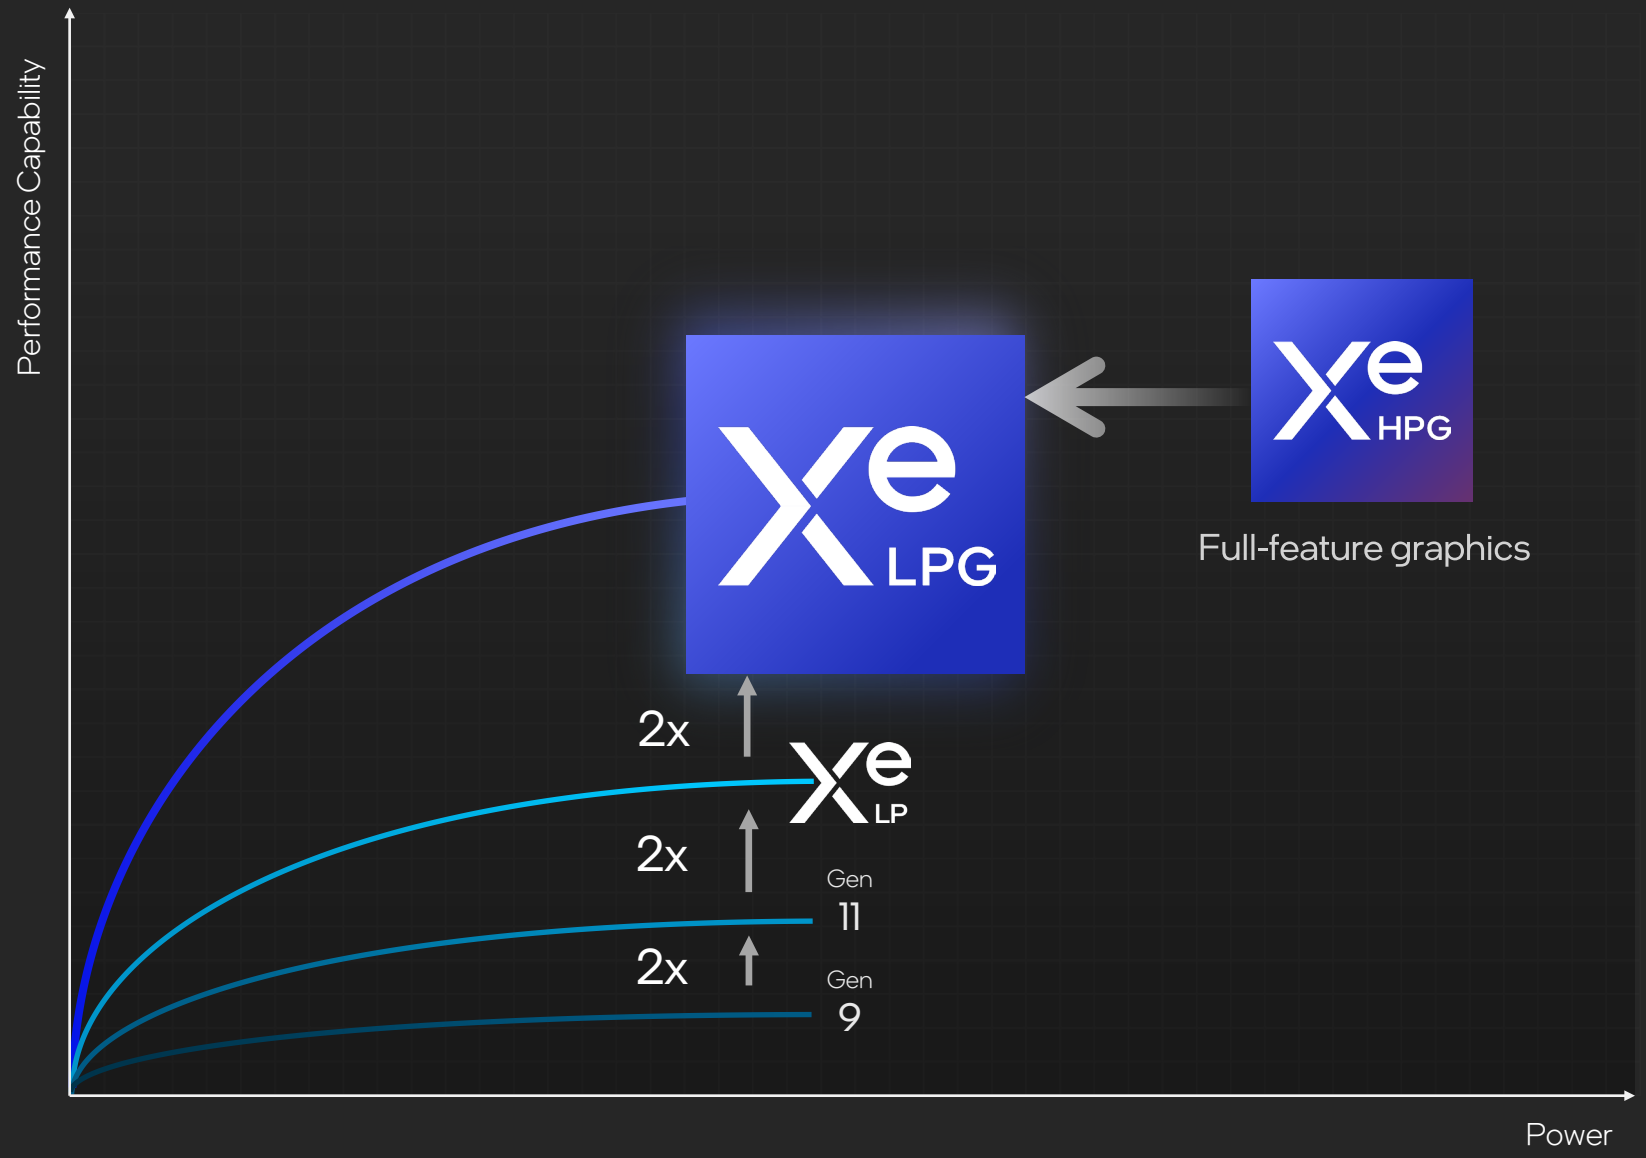

*See appendix for workloads and configurations. Results may vary.

Meteor Lake **GPU**

*Compared to prior generation. See appendix for more information. Results may vary.
**Intel® Arc™ graphics only available on select MTL processor-powered systems with dual-channel memory.

3D Performance Hybrid Architecture

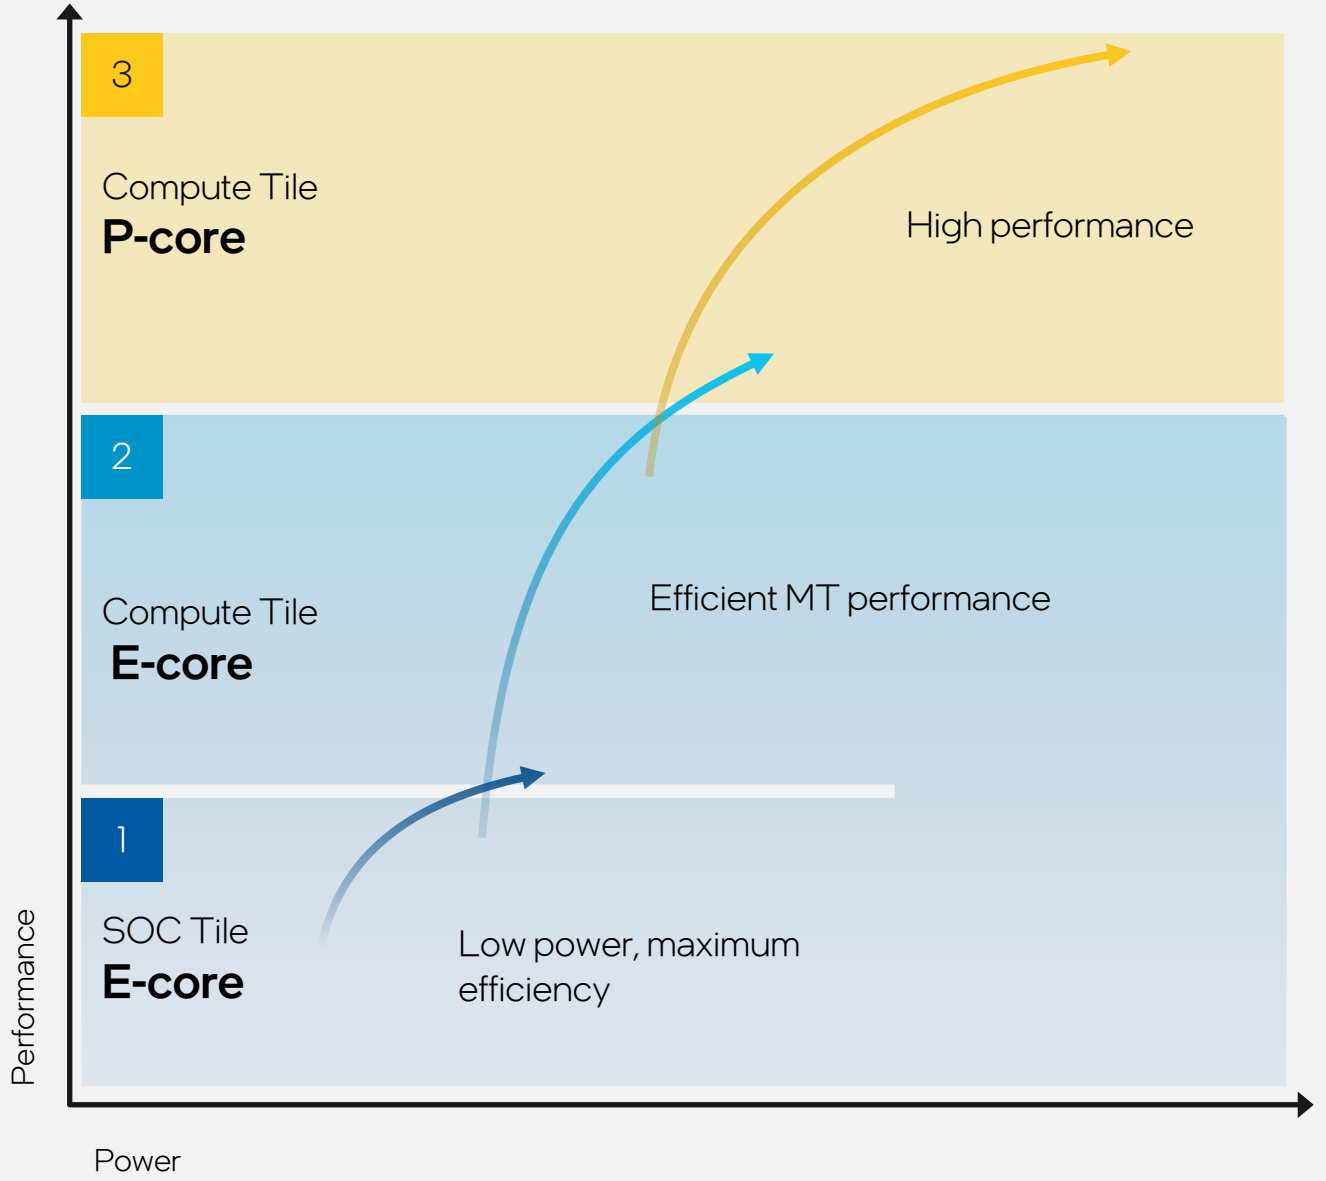

*Conceptual representation of 3D Perf Hybrid Arch

INTEL THREAD DIRECTOR

Architecture

OS Scheduler

Intel Thread Director

E-cores & P-cores


Leveraging
Disaggregation

Leveraging Disaggregation

“Experience First” Client
drives New Era of System level integration

Leveraging Disaggregation

“Experience First” Client
drives New Era of System level integration

Process, packaging & architecture
together make this possible

Leveraging Disaggregation

“Experience first” client
drives new era of system-level integration

Process, packaging & architecture
together make this possible

Extremely flexible architecture
that scales across design points and time

FOVEROS TECHNOLOGY ADVANTAGE

High Density
Wires / Area

Energy Efficient
pJ / Bit

Low Latency
Nanosecond / trip

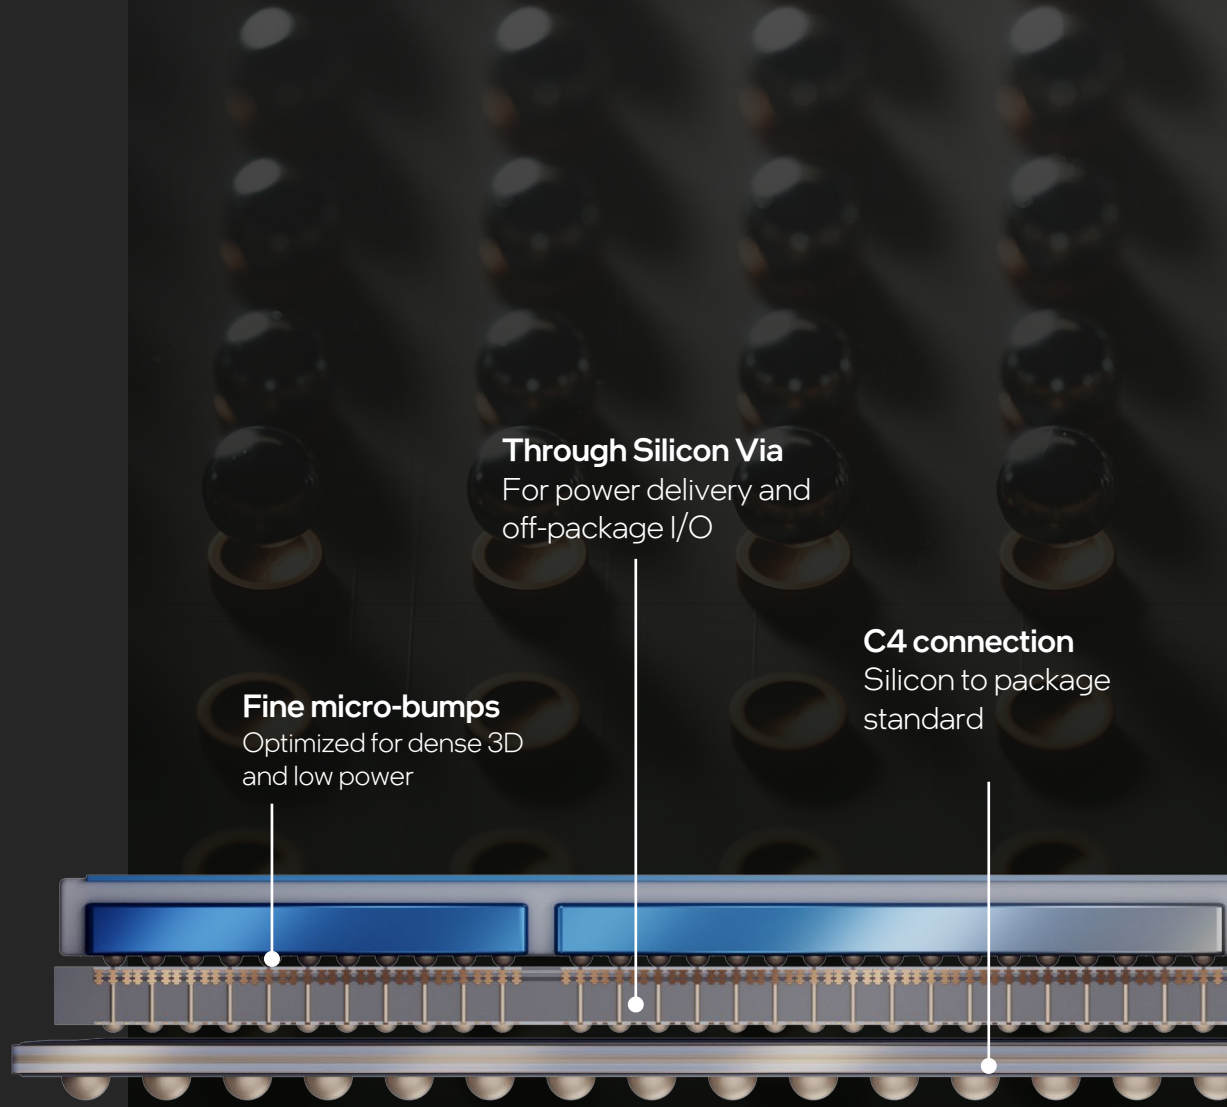

OVEROS FOVEROS FOVEROS FOVERO

Introducing

Intel 4

Logic process technology

2x
area scaling
for High Perf Logic library vs Intel 7*

EUV
lithography for process simplification

>20%
power efficiency vs Intel 7*

Compatible with
3D Foveros
Advanced Packaging

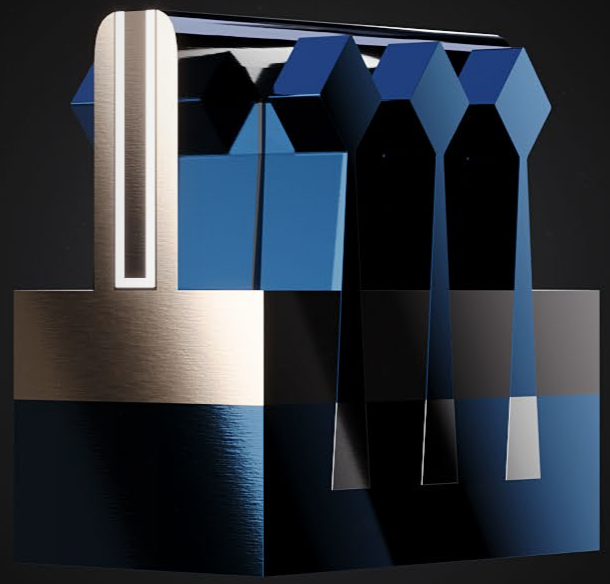

*Based on internal estimates.
Learn more at www.intel.com/PerformanceIndex.
Results may vary.

Meteor Lake

Most power-efficient processor we've ever built

Build our most **power-efficient** client processor in history

Launch IA on **Intel 4**
First Intel 4 P-core (Redwood Cove) & E-core (Crestmont)

Leap ahead on **graphics**
Up to 2x GFX performance/watt*

Deliver **AI at Scale**
First client integration of AI engine (NPU)

*Compared to prior generation. See appendix for more information. Results may vary.

Meteor Lake

New **P-core & E-core** microarchitectures

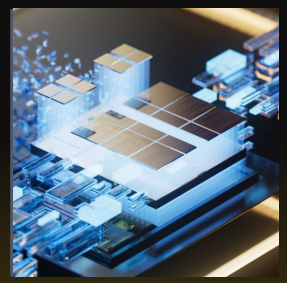

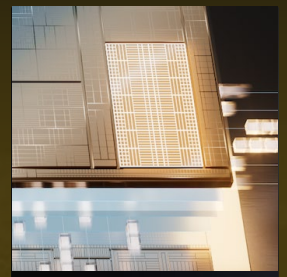

3D Performance Hybrid Architecture

Built-In **NPU AI Engine**

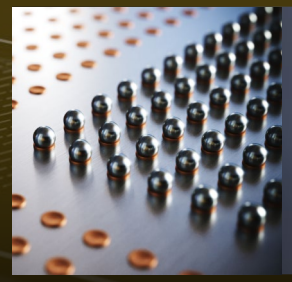

FOVEROS 3D packaging

Latest **Media & Display** Standards

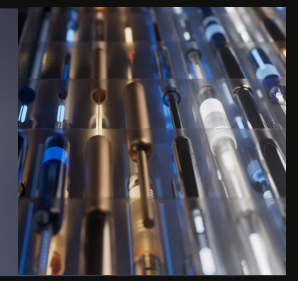

Thunderbolt 4

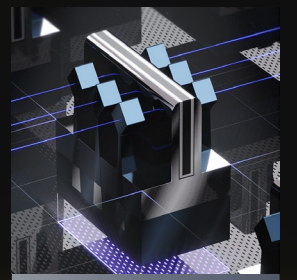

WiFi 7

Low Power Island E-cores

First on **Intel 4**

intel ARC™

Power efficiency & AI at scale

Notices and Disclaimers

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details.

Results that are based on pre-production systems and components as well as results that have been estimated or simulated using an Intel Reference Platform (an internal example new system), internal Intel analysis or architecture simulation or modeling are provided to you for informational purposes only. Results may vary based on future changes to any systems, components, specifications or configurations.

All product and service plans, and roadmaps are subject to change without notice. Any forecasts of products, services or technologies needed for Intel's operations are provided for discussion purposes only. Intel will have no liability to make any purchase in connection with forecasts published in this document. Code names are often used by Intel to identify products, services or technologies that are in development and usage may change over time. Product, service and technology performance varies by use, configuration and other factors. No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. Learn more at www.Intel.com/PerformanceIndex and www.Intel.com/ProcessInnovation.

Reference to research results, including comparisons to products, services or technology performance are estimates and do not imply availability. The products and services described may contain defects or errors which may cause deviation from published specifications. Current characterized errata are available on request. Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade. Statements in this document that refer to future plans or expectations are forward-looking statements. These statements are based on current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in such statements. For more information on the factors that could cause actual results to differ materially, see our most recent earnings release and annual report on form 10-K and other SEC filings at www.intc.com.

Intel contributes to the development of benchmarks by participating in, sponsoring, and/or contributing technical support to various benchmarking groups, including the BenchmarkXPRT Development Community administered by Principled Technologies.

Altering clock frequency or voltage may void any product warranties and reduce stability, security, performance, and life of the processor and other components. Check with system and component manufacturers for details.

Intel is committed to the continued development of more sustainable products, software, and processes and working with suppliers and customers as we strive to lower the greenhouse gas footprint of our entire value chain. Where applicable, environmental attributes of a product family or specific SKU will be stated with specificity. Refer to the 2022 Corporate Responsibility Report for further information.
Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Your costs and results may vary. No product or component can be absolutely secure. Intel technologies may require enabled hardware, software or service activation.


intel