intel

TECH tour.TW

Lunar Lake Architecture Session Highlights

Robert Hallock
Vice President Client AI Marketing

Lunar Lake

Flagship SoC for the next gen of AI PCs

Breakthrough x86 power efficiency

Up to 40% lower SoC power*

Exceptional core performance

Similar ST perf at half the power*

Massive leap in graphics

Up to 1.5X better graphics*

Unmatched AI compute

Up to 120 platform TOPS

*Versus Intel's previous gen. For more information, go to Intel.com/PerformanceIndex.

Lunar Lake Construction

Built with advanced packaging

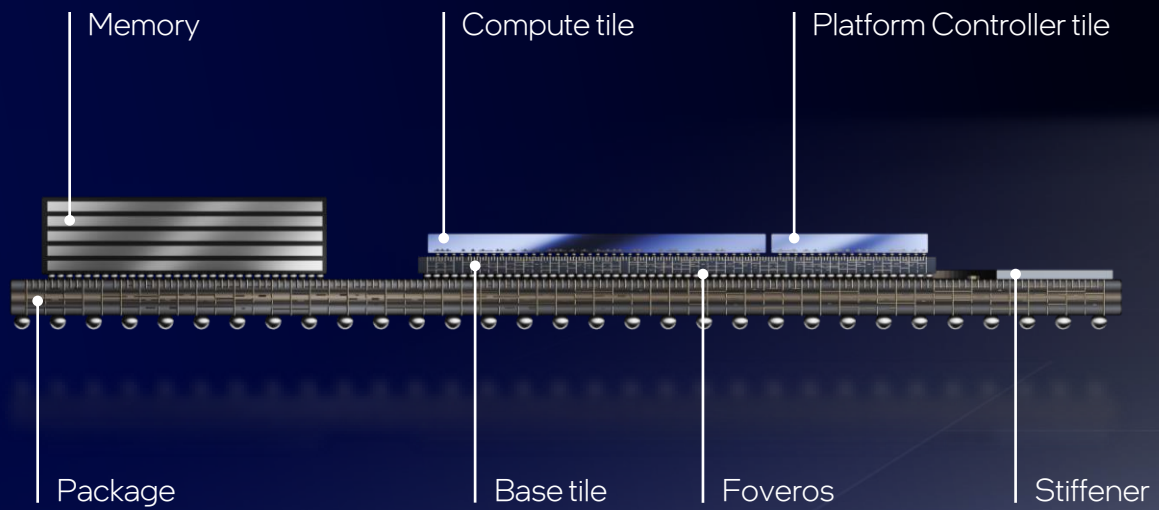

Memory on Package

First ever Intel integration onto package

Up to **32GB** with 2 ranks	Support for **LPDDR5x** DRAM	Up to **8.5GT/s** per chip
Support for **16b x4** channels	Up to **40%** lower PHY power	Up to **250mm²** area savings

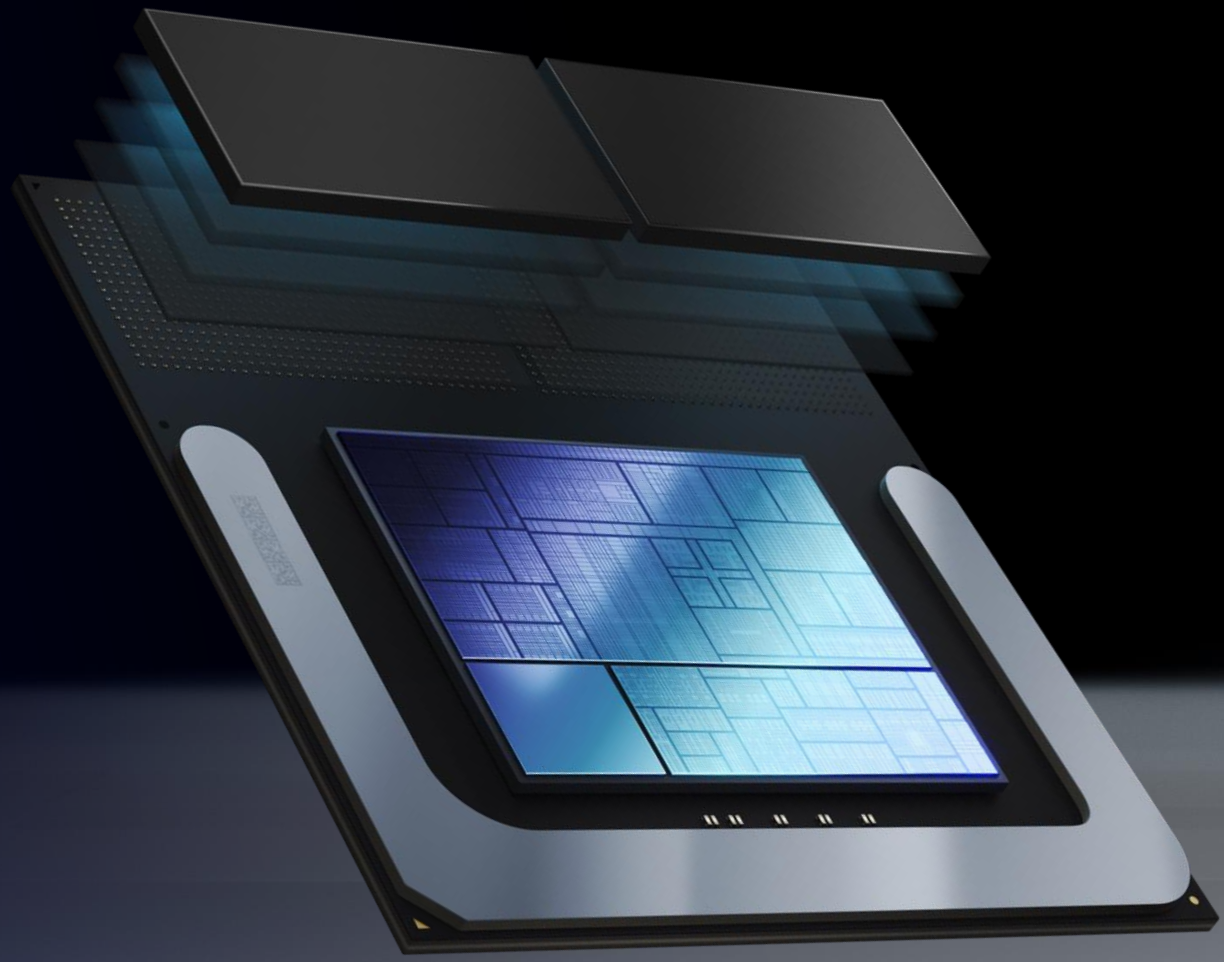

Lunar Lake

Architecture overview

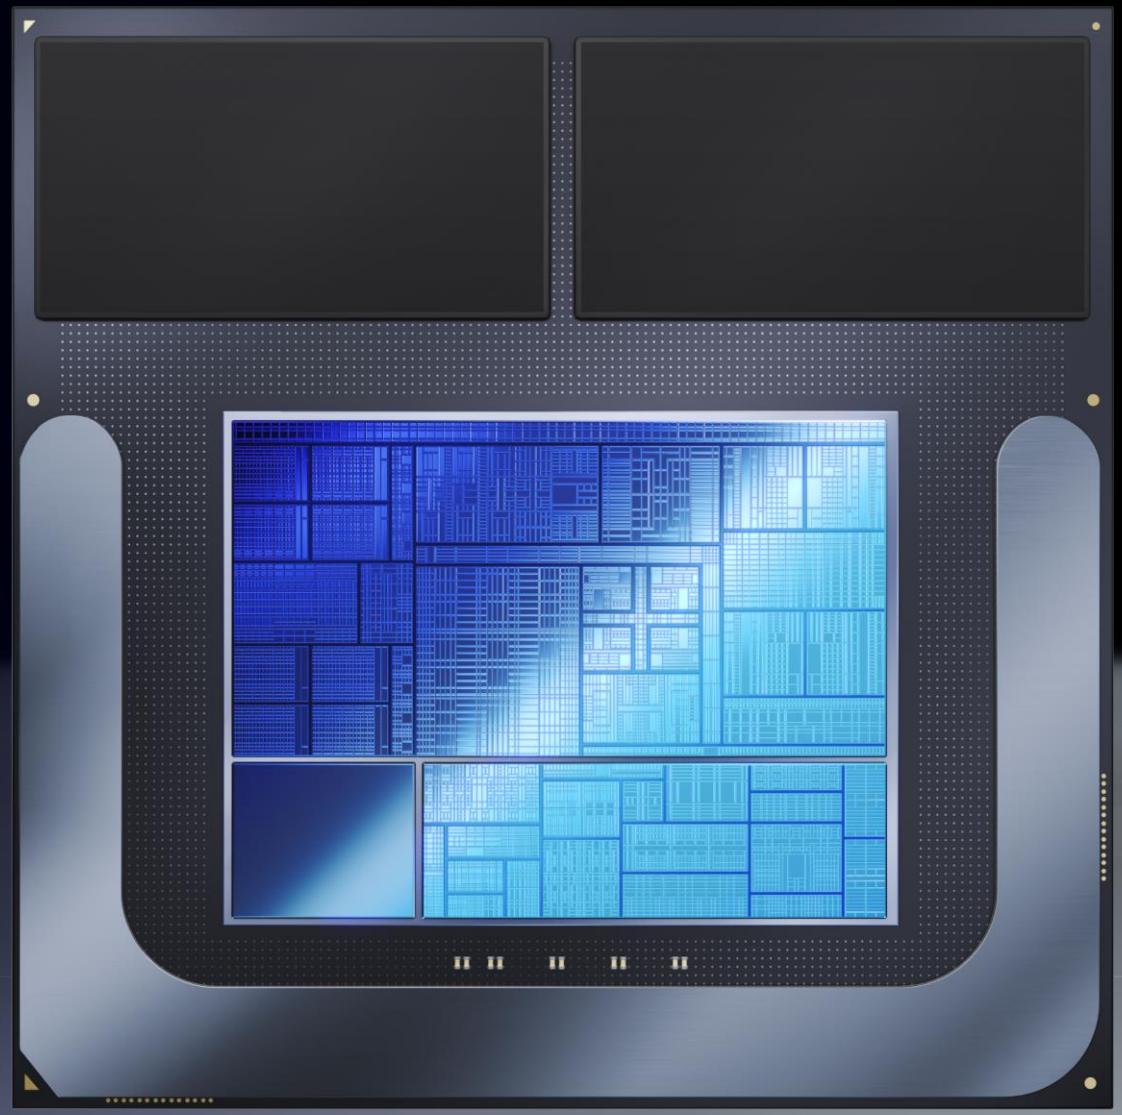

Lunar Lake

Performance cores

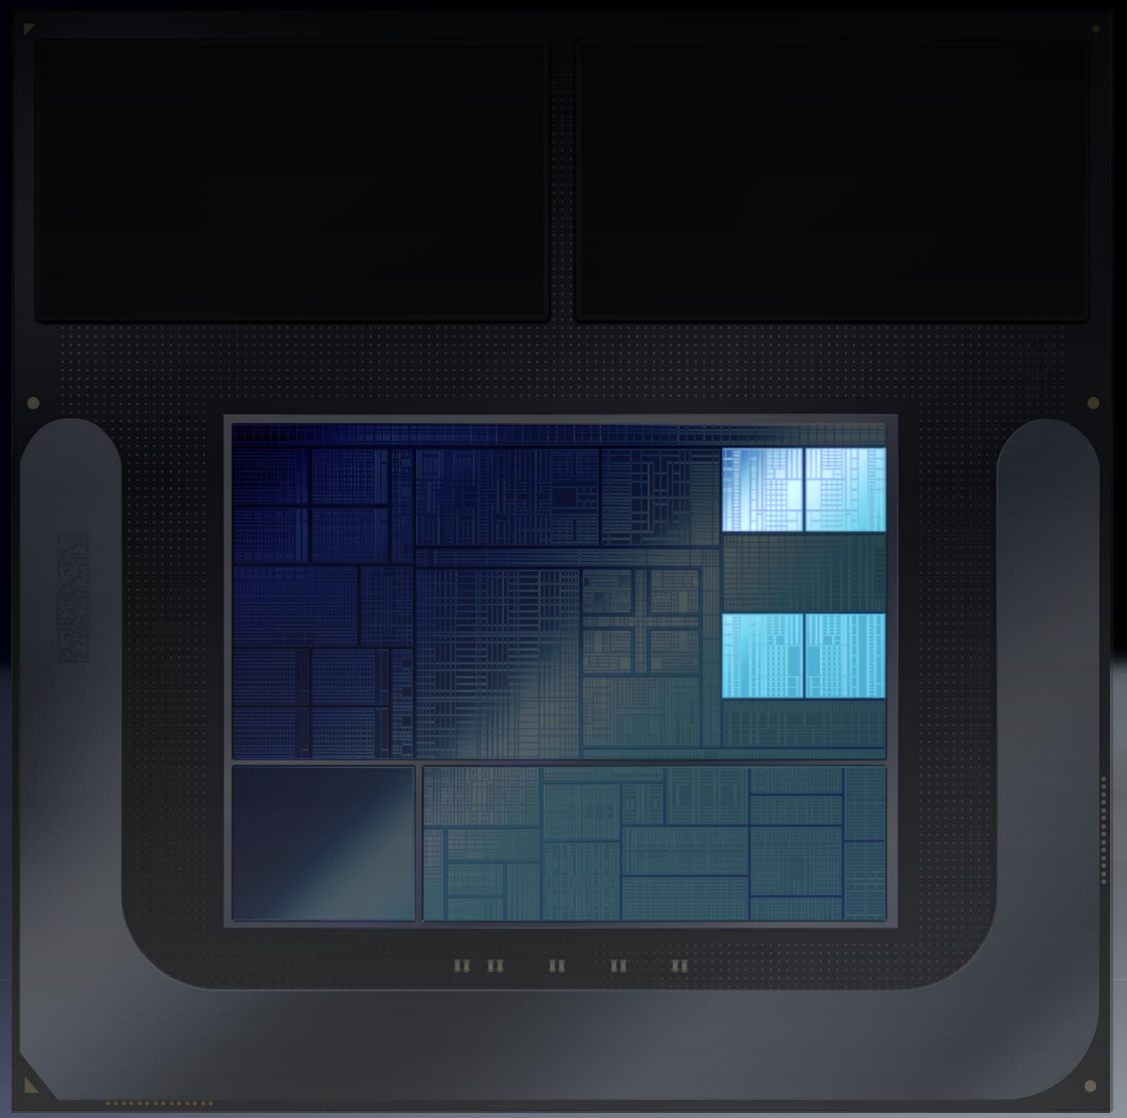

Lion Cove P-core

Architecture goals

Performance & area efficiency
Optimize ST perf/watt and perf/area for client SoCs

Overhaul microarchitecture
Generational IPC improvement and future scalability

Modernize design database
Accelerate innovation going forward

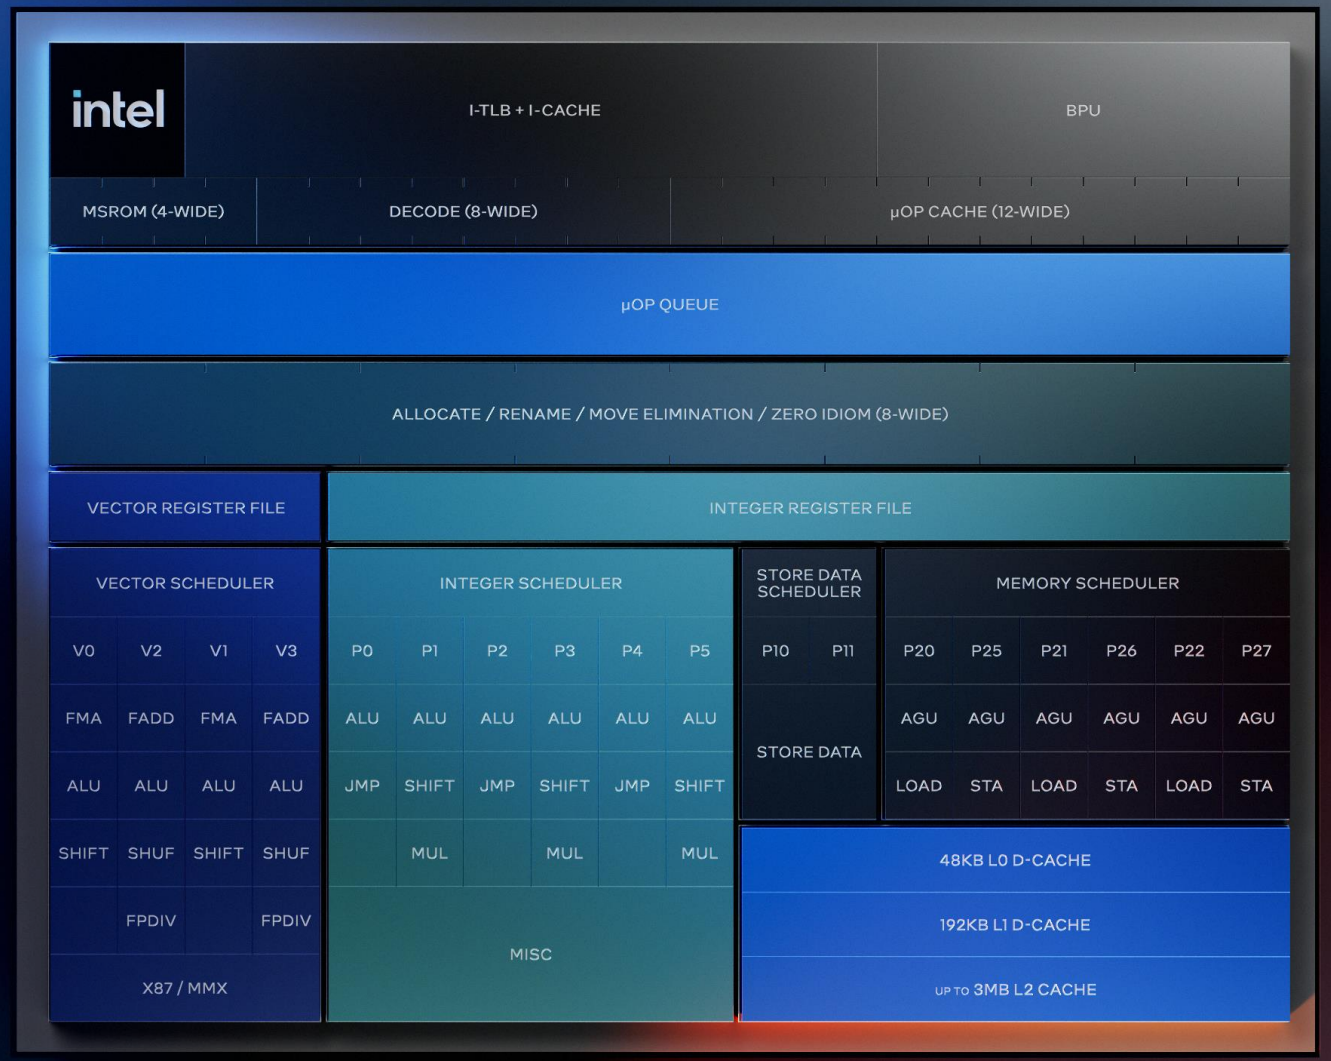

Lion Cove P-core

Rearchitected for efficient performance

18 Execution ports

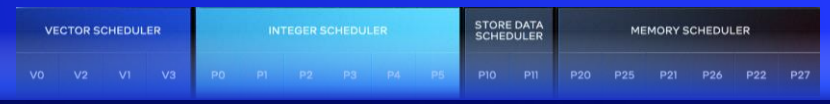

Up to 12MB
shared L3 cache
on Lunar Lake

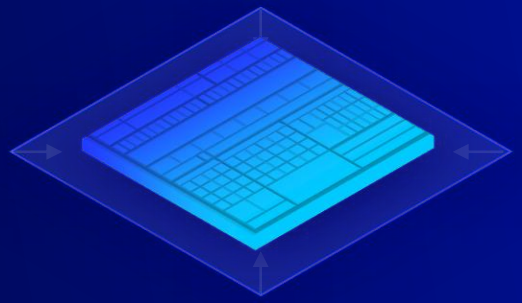

Optimized for **PPA**

16.67MHz
Finer clock intervals

Wider scheduling

across both allocation/rename & retire

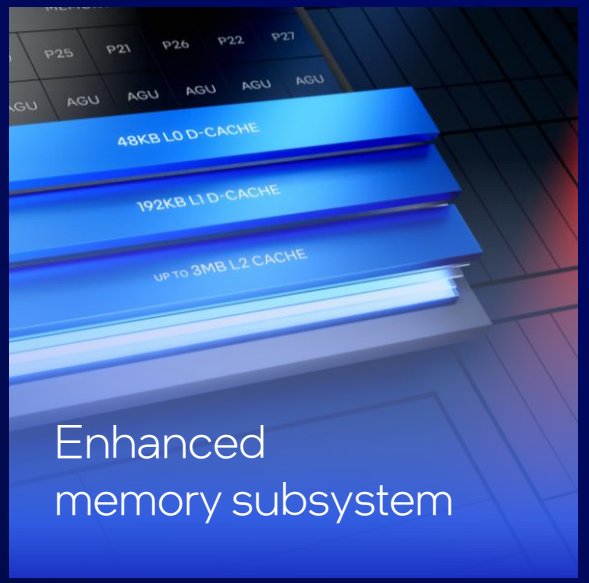

VEC ◀ **Split** out of order engine ▶ INT

intel

Lion Cove P-core

Double-digit performance gains over prior generation

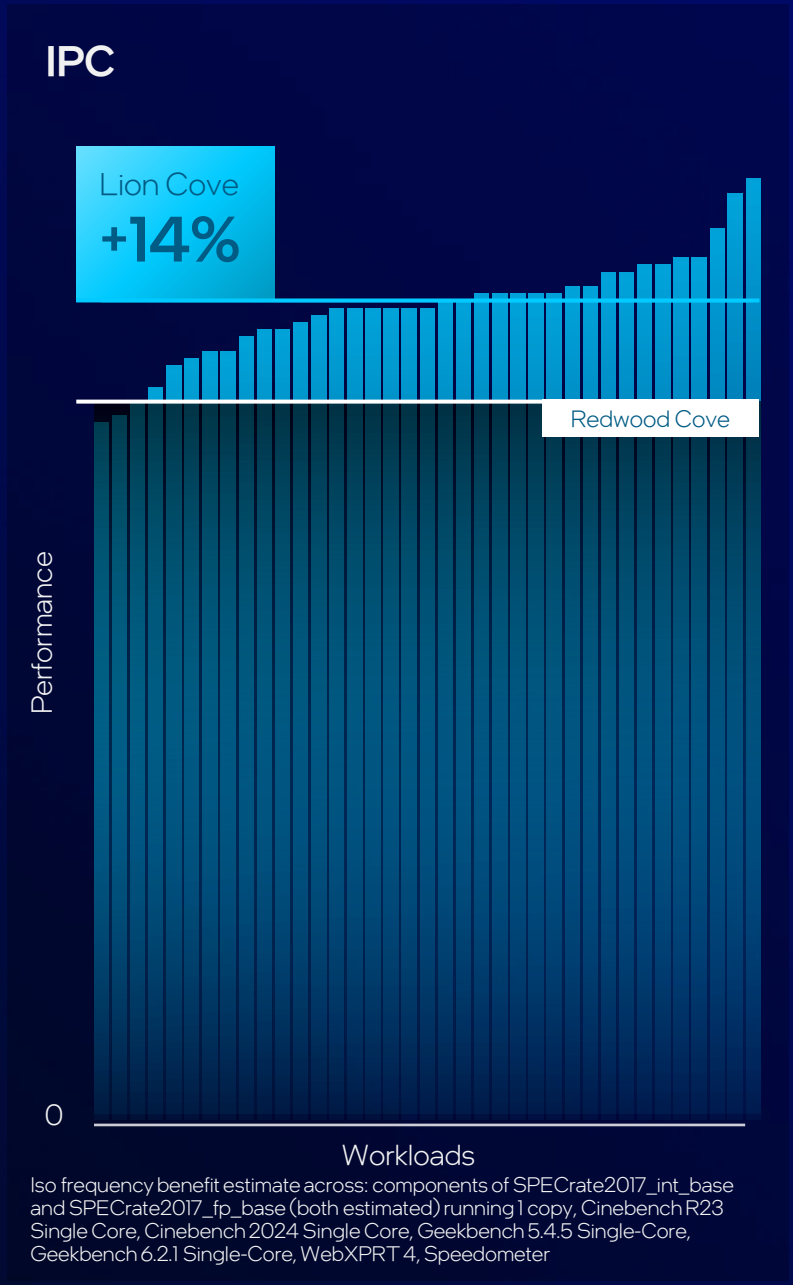

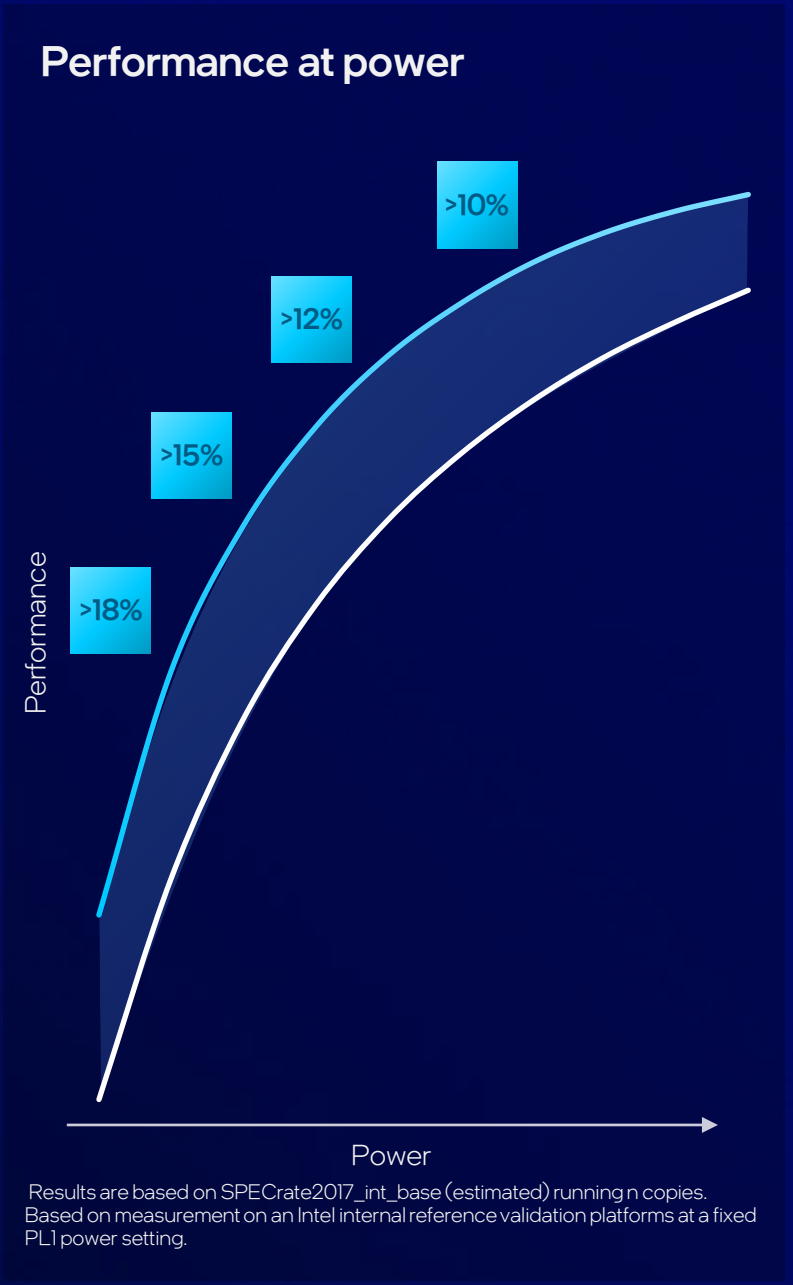

Lion Cove in Lunar Lake

Redwood Cove in Meteor Lake

Iso frequency benefit estimate across: components of SPECrate2017_int_base and SPECrate2017_fp_base (both estimated) running 1 copy, Cinebench R23 Single Core, Cinebench 2024 Single Core, Geekbench 5.4.5 Single-Core, Geekbench 6.2.1 Single-Core, WebXPRT 4, Speedometer

Results are based on SPECrate2017_int_base (estimated) running n copies. Based on measurement on an Intel internal reference validation platforms at a fixed PL1 power setting.

Lunar Lake

Efficient cores

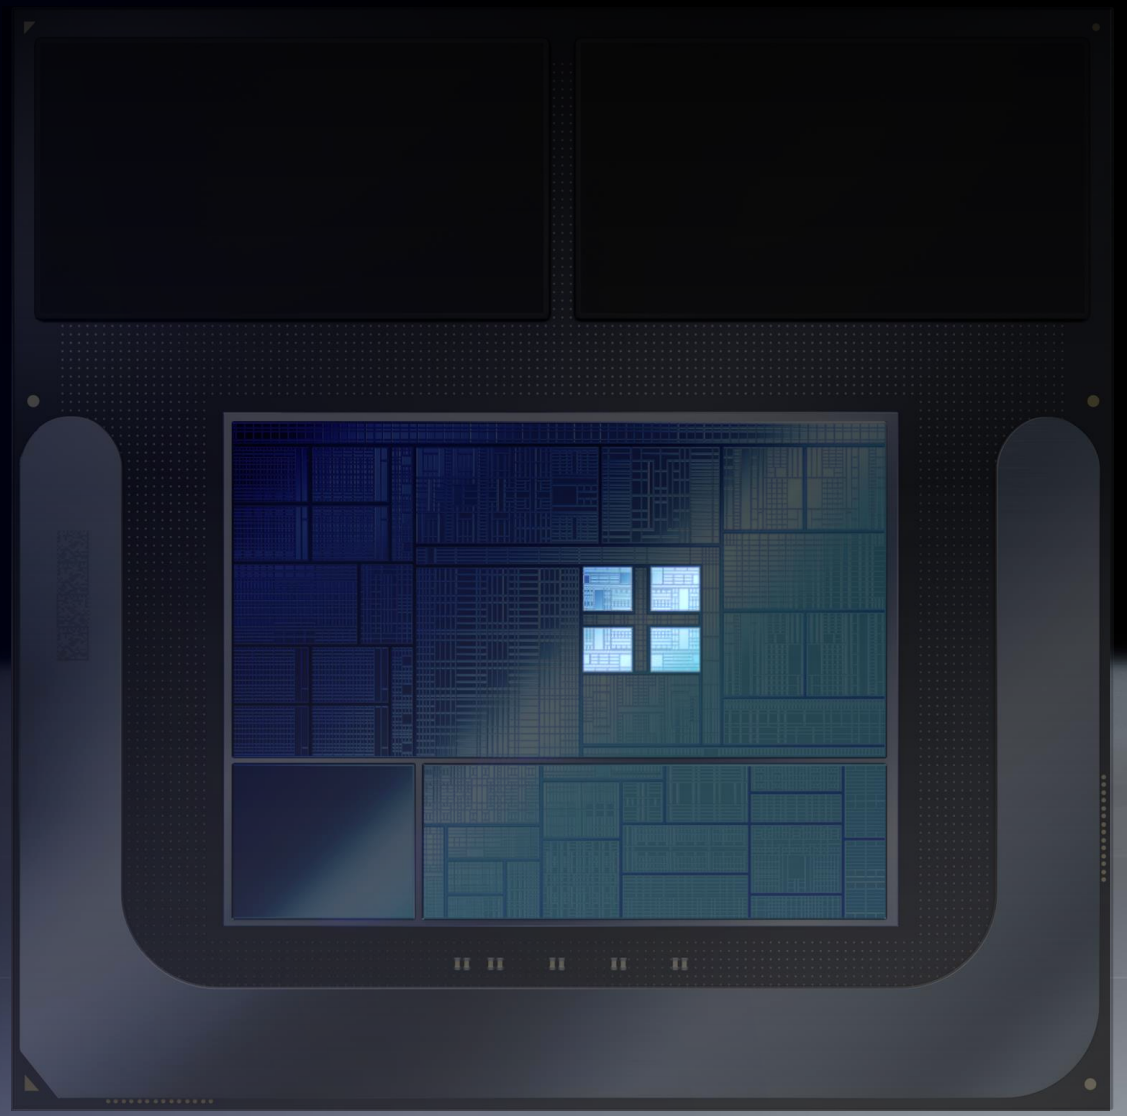

Skymont E-core

Architecture goals

Increase workload coverage
Increase range of low power island & MT perf

Double vector & AI throughput
For increased VNNI capability support

Increase scalability
For overall performance uplift

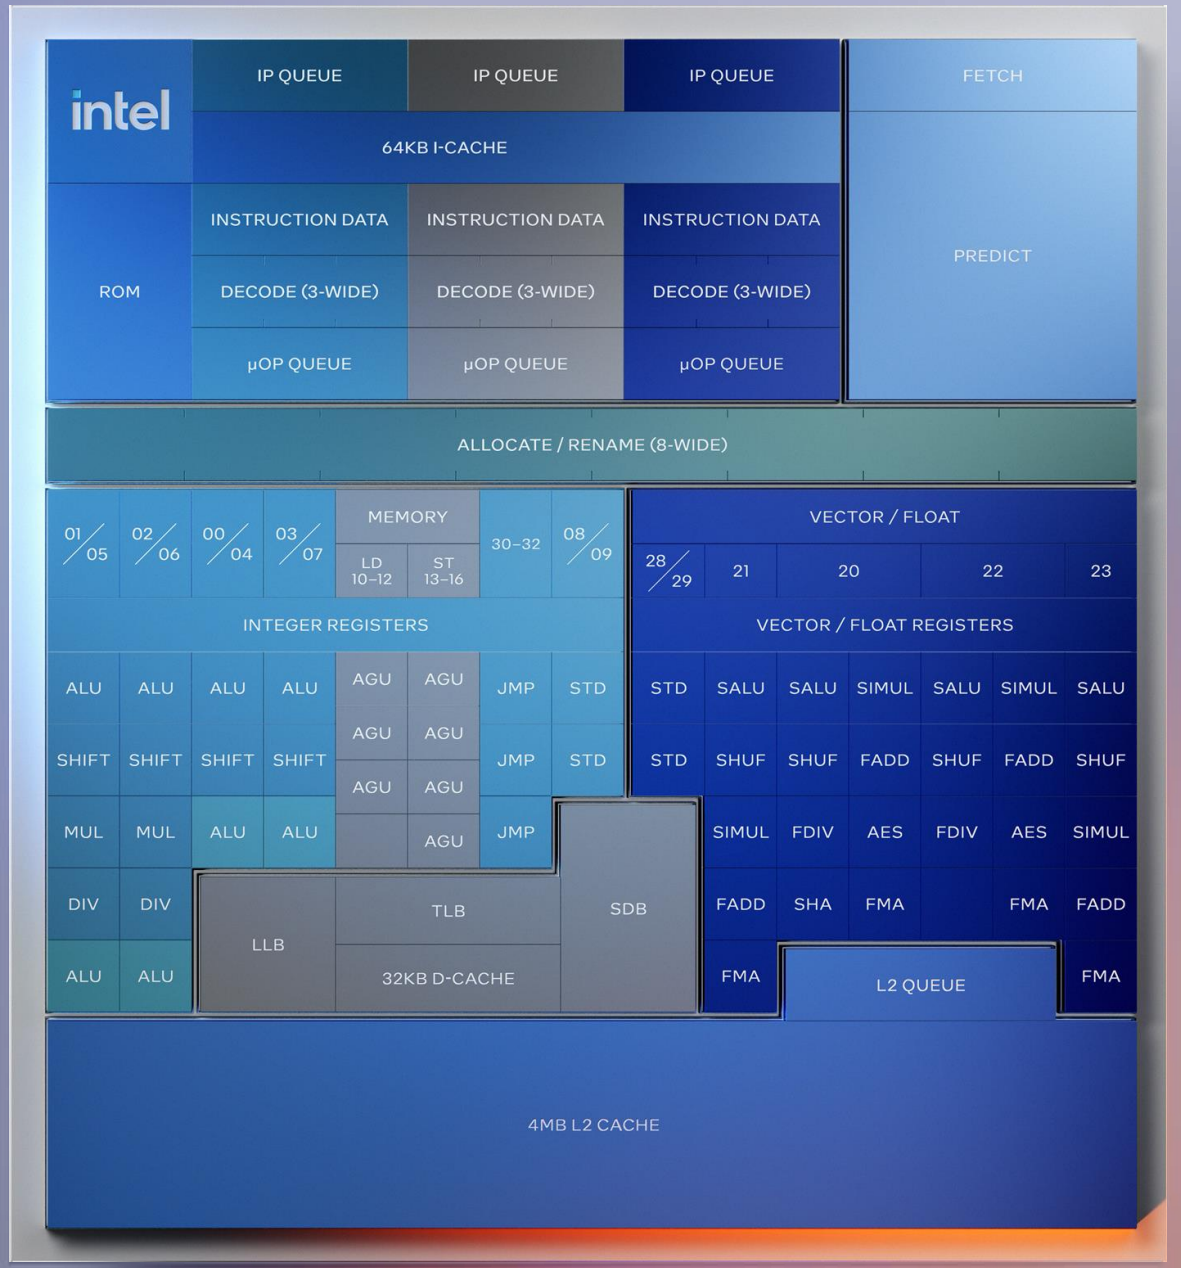

Skymont E-core

Our most efficient performant architecture

26 dispatch ports

Scalable & flexible

across multiple implementations

2x

L2 cache bandwidth

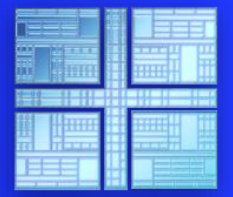

4MB

shared L2 cache

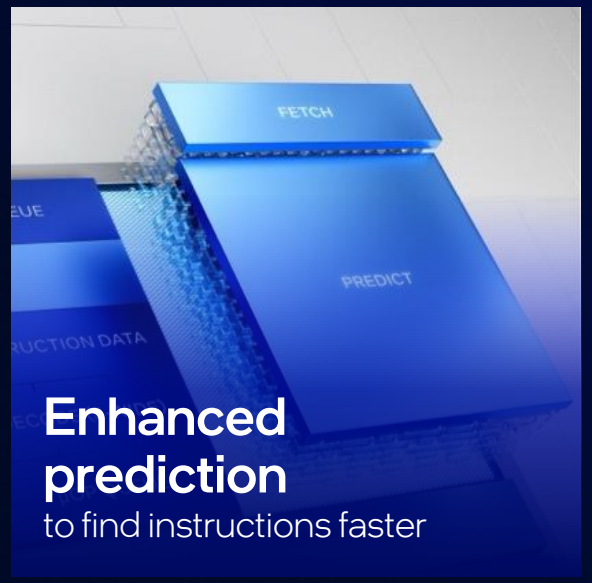

Enhanced prediction

to find instructions faster

96B

parallel fetching

Deeper queueing

for better parallelism

Wider allocation & retire

ALLOCATE

RETIRE

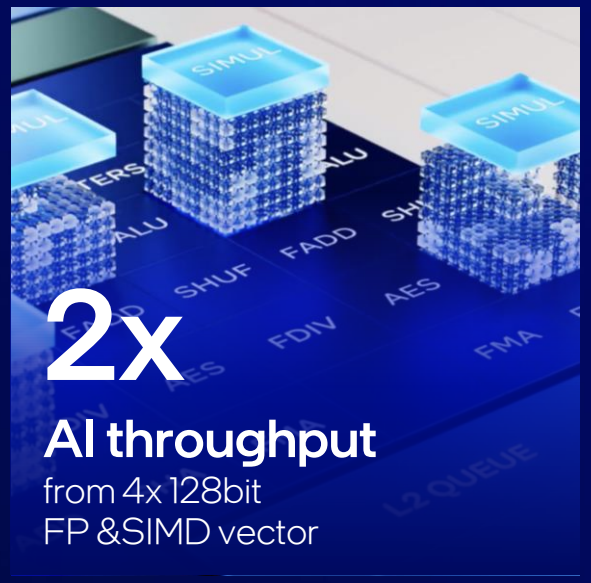

2x

AI throughput

from 4x 128bit FP &SIMD vector

intel

Skymont E-core

Leaps in performance & capability vs. previous implementation

Skymont in Lunar Lake

Crestmont in Meteor Lake

Single threaded FP improvement

SPECrate2017_fp_base est/ GCC

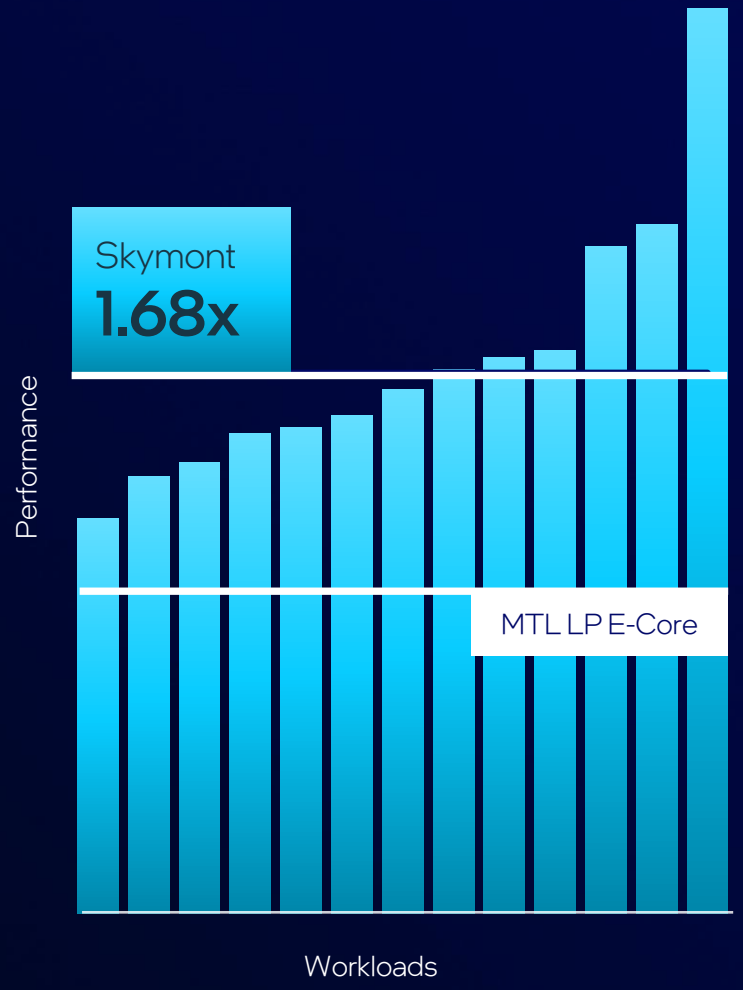

Multi threaded INT

Meteor Lake LP E-Core 2C vs. Skymont 4C

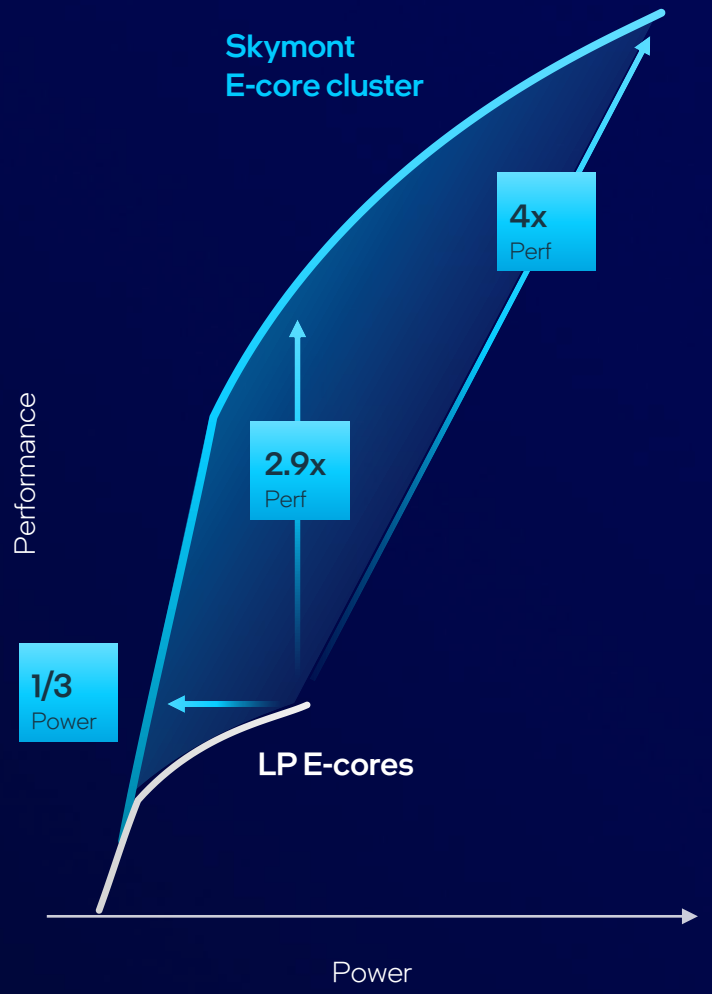

Lunar Lake

E-core cluster & P-core cluster

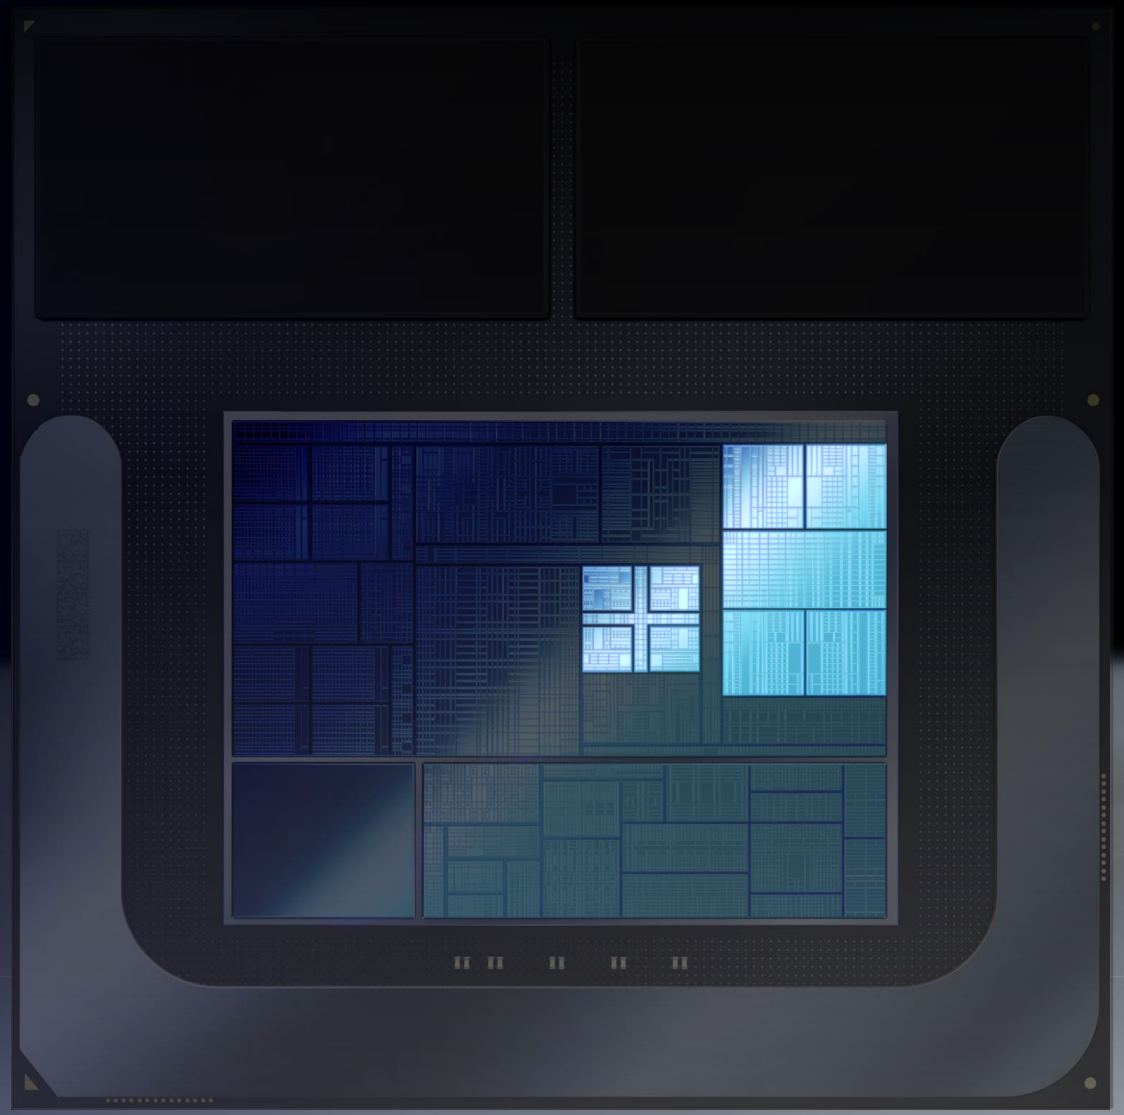

Breakthrough x86 Efficiency

Whilst covering the full CPU performance range

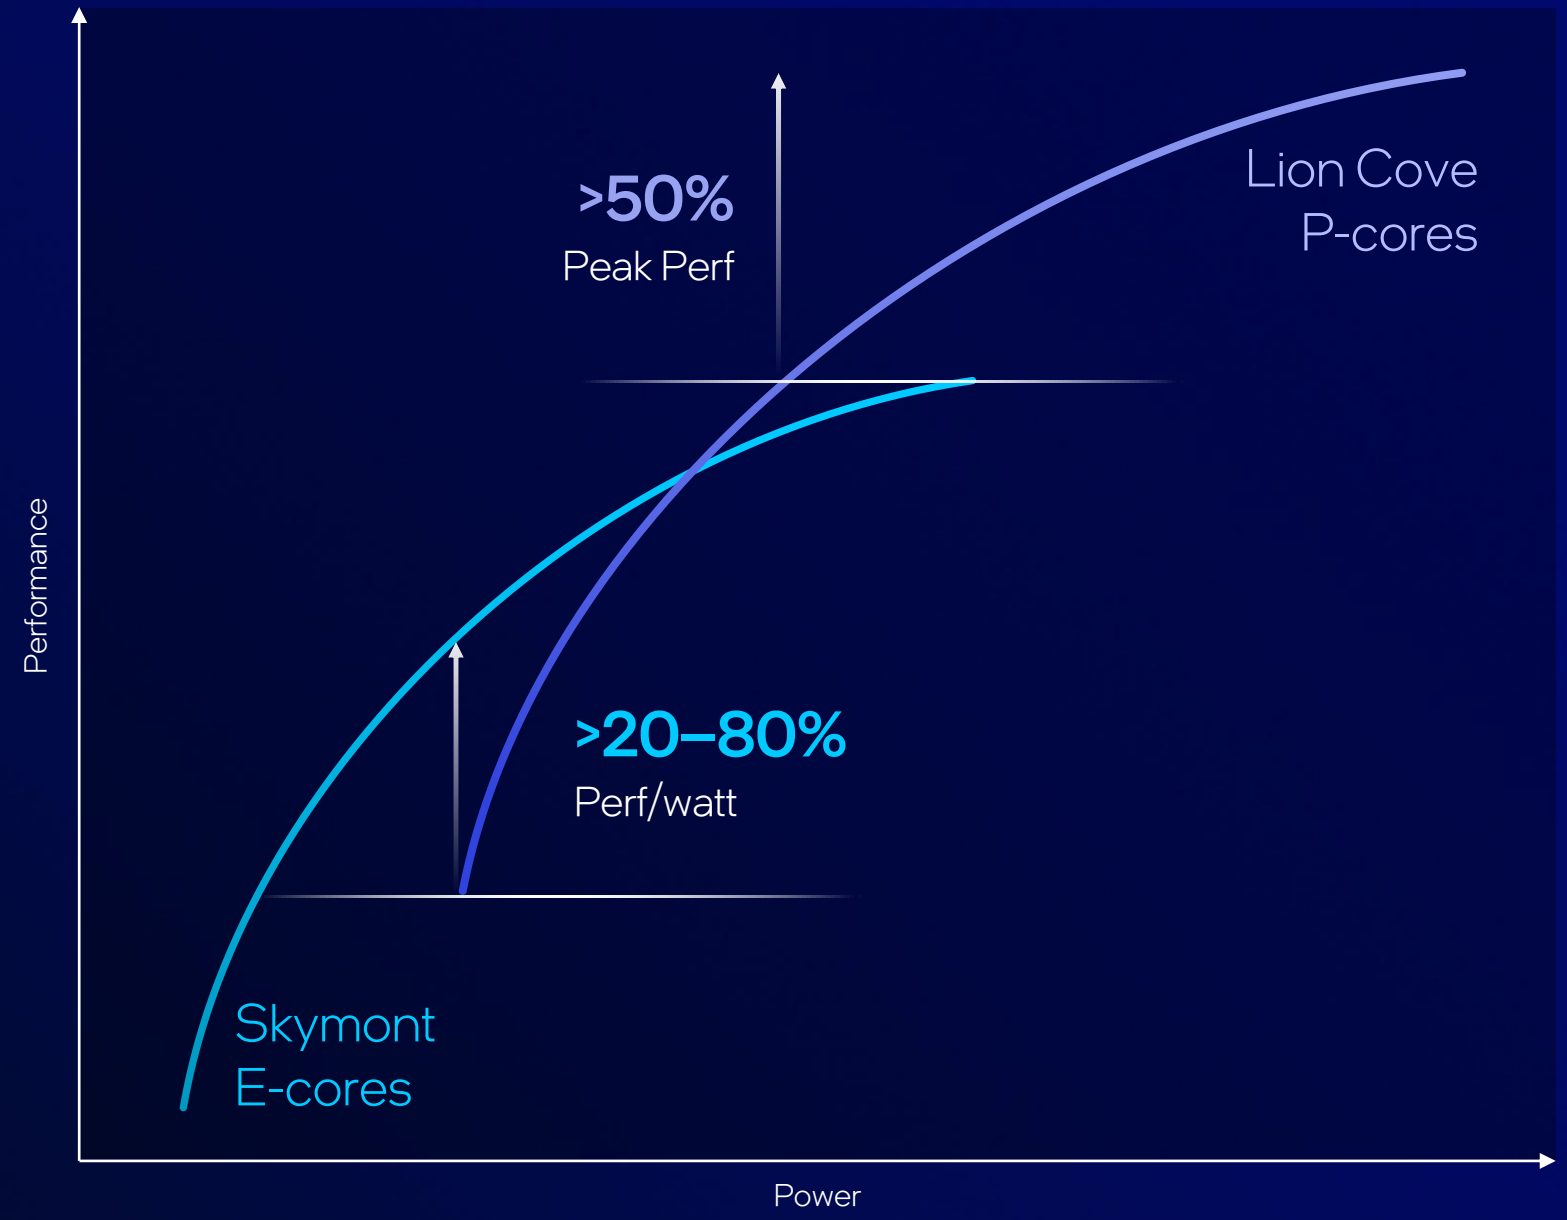

Illustration of the relative Lunar Lake P-core and E-core performance across the SoC power range.
See Intel.com/PerformanceIndex for more details.

Intel Thread Director

Architecture goals

Increase intelligence
To optimize workload to core matching

Improve OS and OEM integration
For more informed and controlled scheduling

Expand efficiency capabilities
Driving overall better battery life

Intel Thread Director

Our next generation intelligent thread direction for advanced hybrid architectures

OS containment zones

Efficiency

Hybrid / compute

Zoneless

Enhanced power management integration

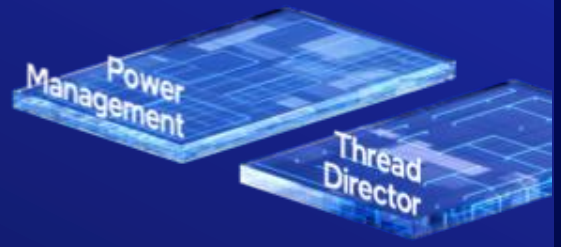

Upgraded foundations

Enhanced algorithms

Finer granularity

Experience continuity

OEM mode selection

Efficiency — Performance

Lunar Lake Scheduling

With Intel Thread Director

- Dynamic scheduling policy used
- First single E-core as long as work fits

Lunar Lake Scheduling

With Intel Thread Director

- Dynamic scheduling policy used
- First single E-core as long as work fits
- Expand to other E-cores for MT

Lunar Lake Scheduling

With Intel Thread Director

- Dynamic scheduling policy used
- First single E-core as long as work fits
- Expand to other E-cores for MT
- Move to P-cores based on demand ITD guidance

Improved Experience

With OS containment & power management optimization on Lunar Lake

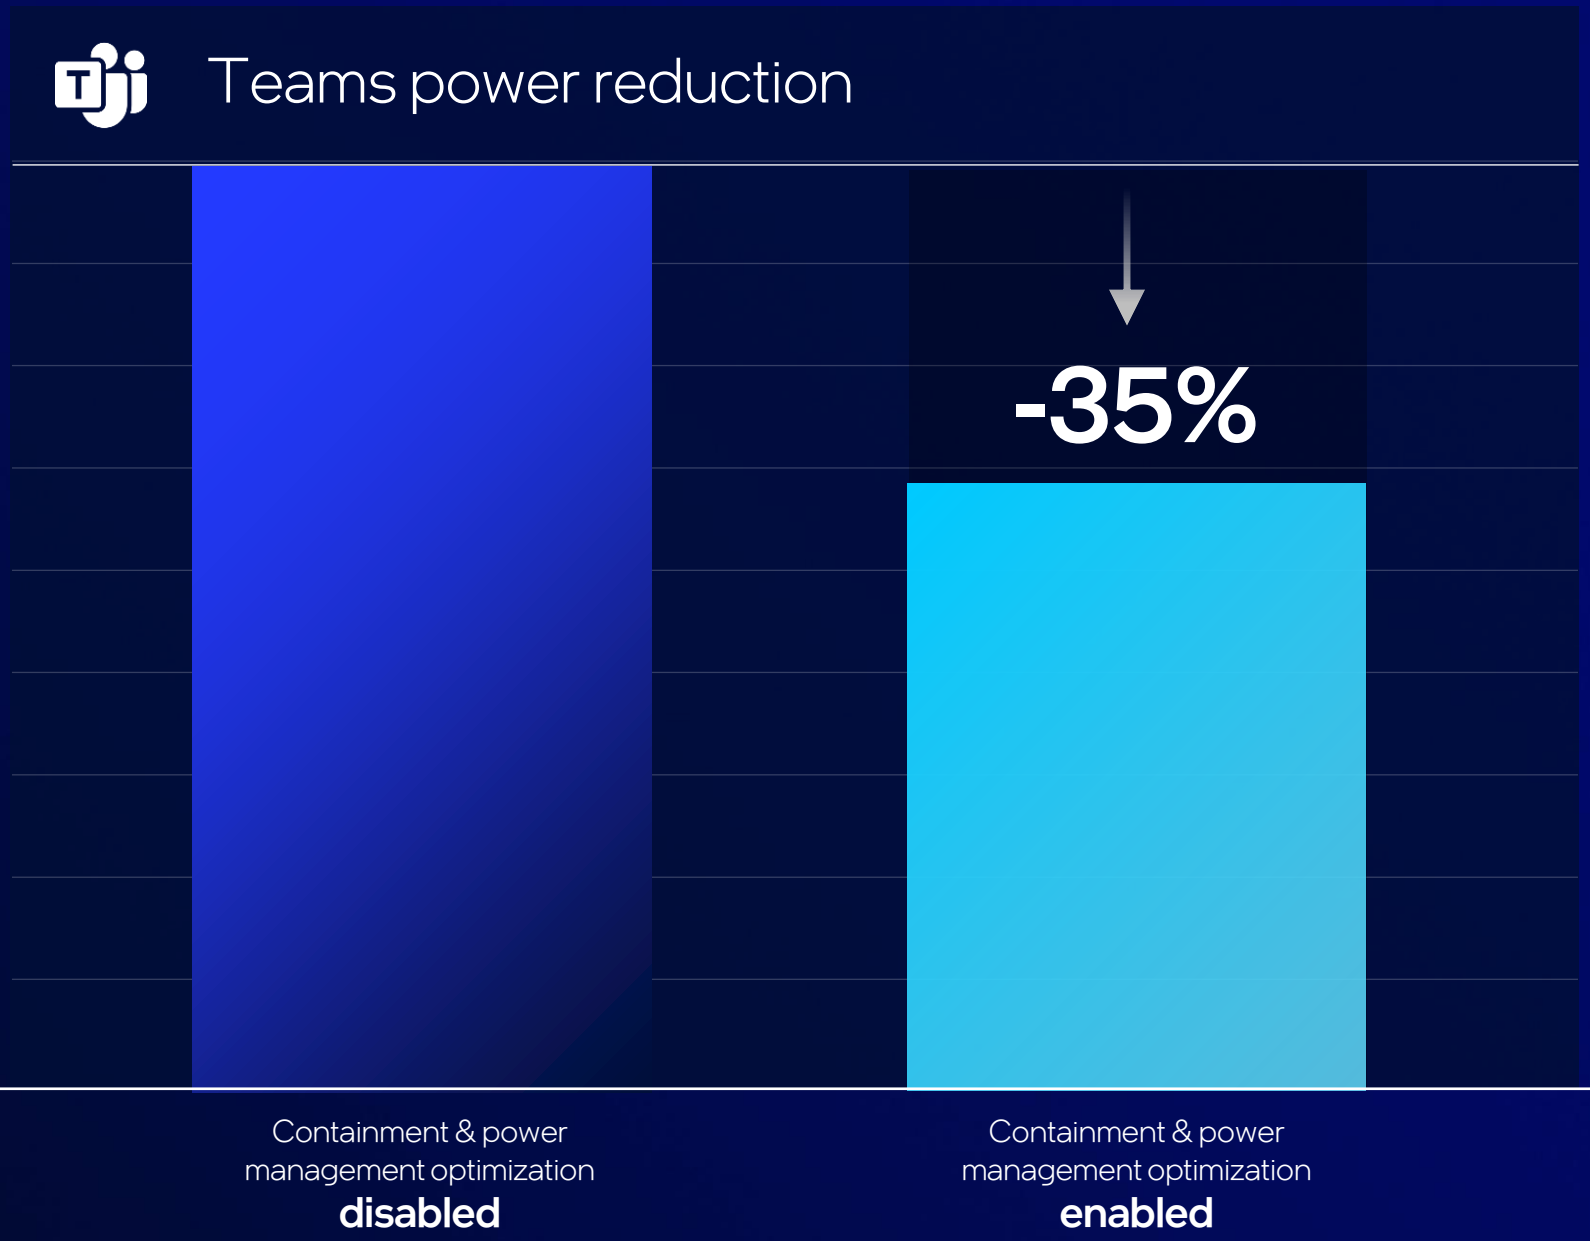

See backup for details. Results may vary.

Intel Thread Director

Office Productivity Example

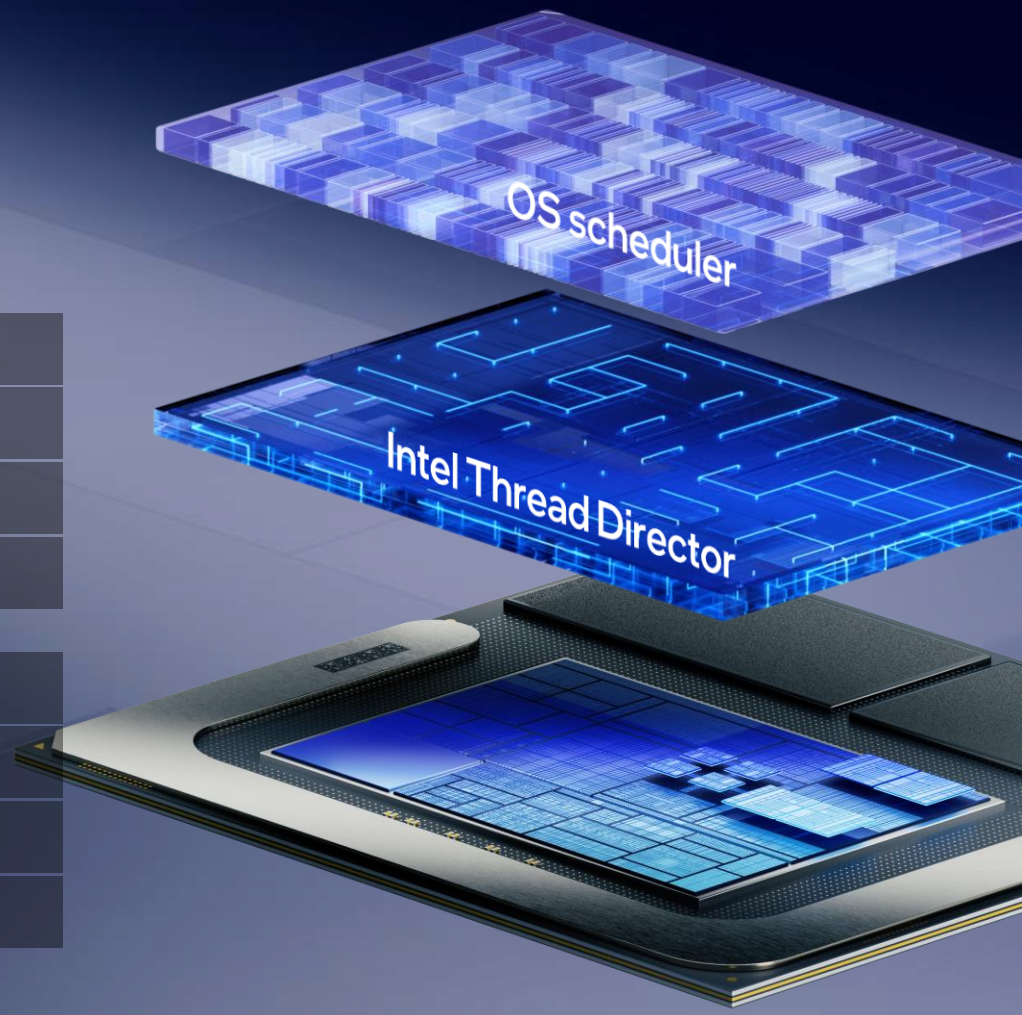

Intel Thread Director

Office Productivity Example

Moves to P-cores as soon as more performance is needed

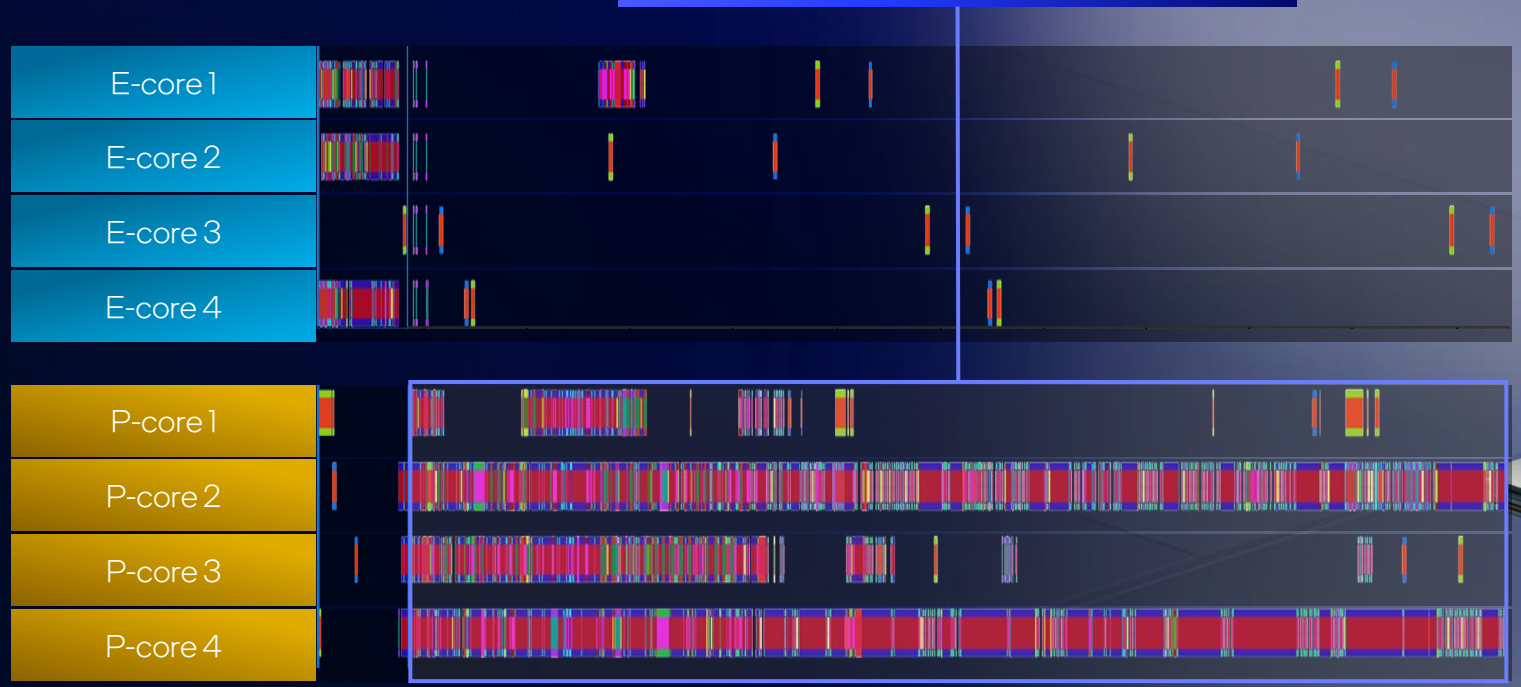

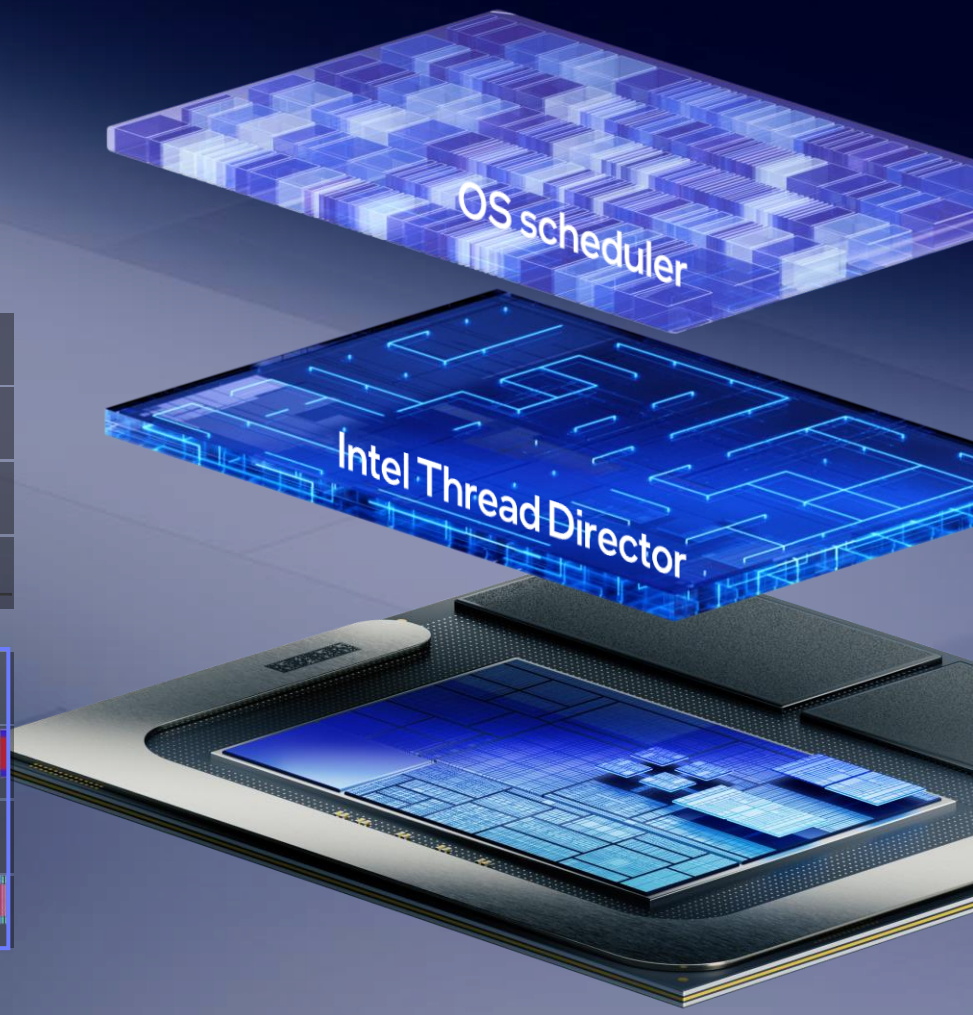

Lunar Lake

New Xe 2 GPU

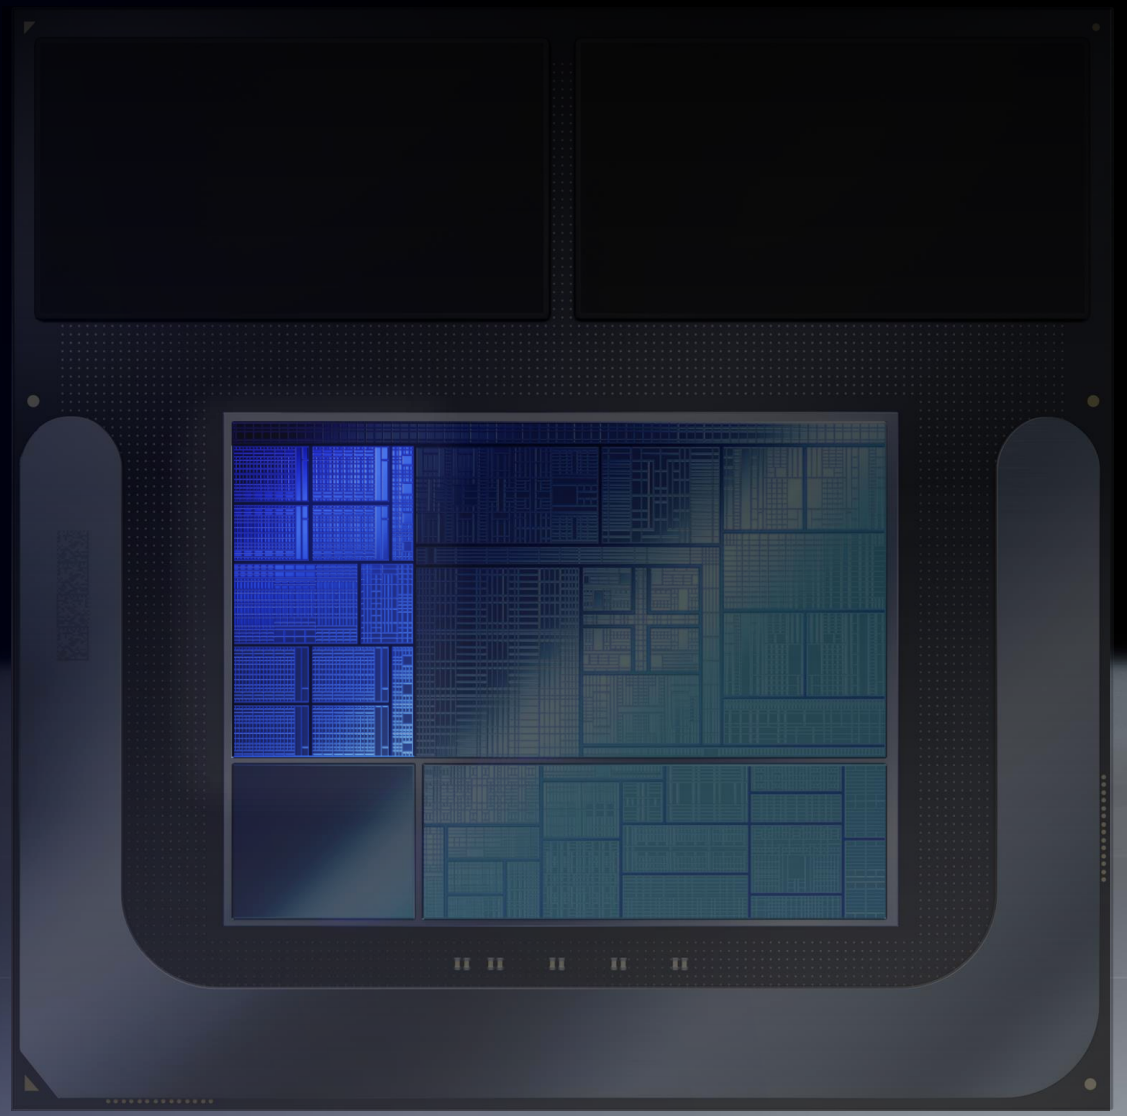

Next Gen Xe2 GPU

Architecture goals

Improved utilization
of hardware functions

Improved distribution
of workload across architecture

Improved integration
Of hardware & software

Next Gen Xe2 GPU

Major leap in graphics performance

up to **67 TOPS**

New XMX engines

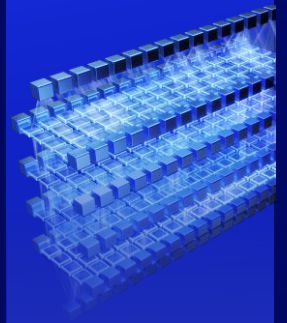

8 Larger ray tracing units

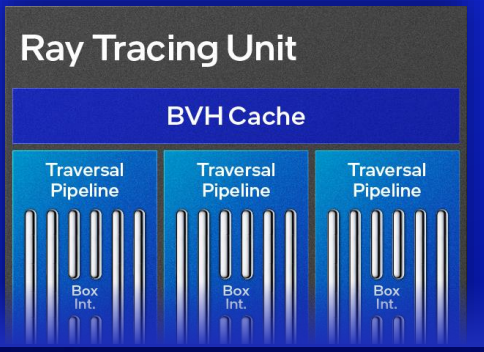

8

2nd gen Xe cores

Xe2 vector engines

1.5x

better vs. Meteor Lake GPU

intel ARC™ Software stack

eDP 1.5

Enhanced XeSS kernels

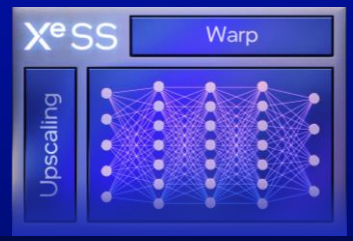

8 MB L2 cache

intel.

New Vector Engine

Optimized for efficiency & AI throughput

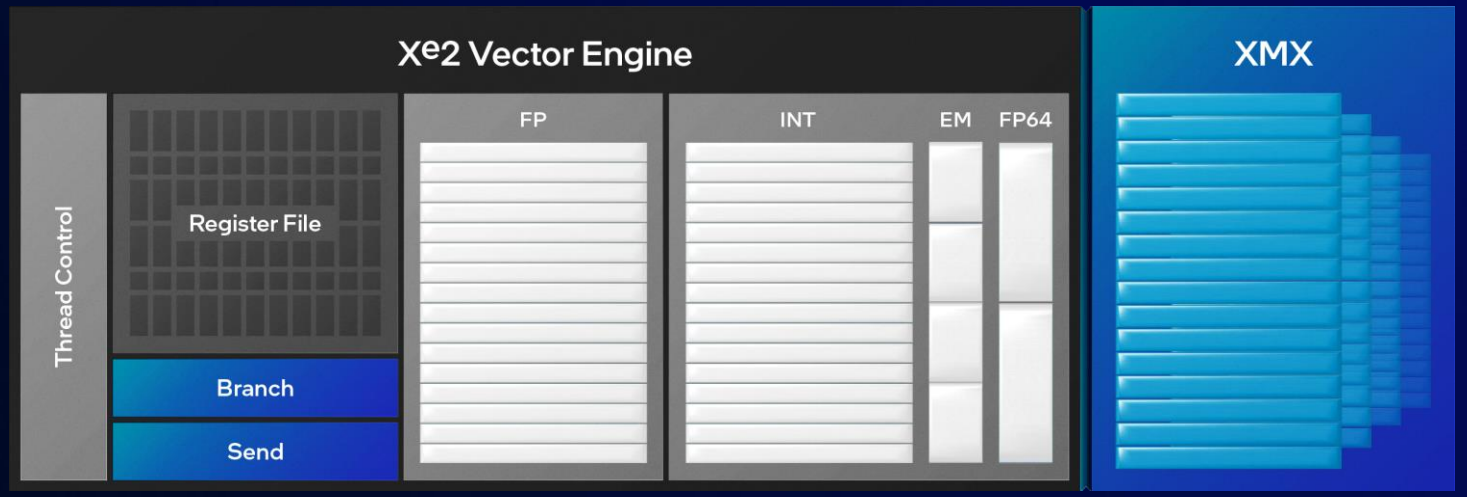

New Xe Matrix Extension Engines

INT8

4096 OPS/clock

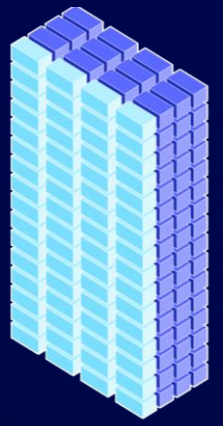

FP16

2048 OPS/clock

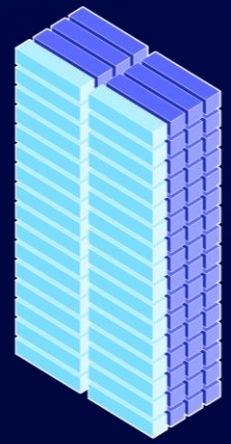

Next Gen Xe2 GPU

Major leap in graphics performance

~1.5x

vs. previous gen

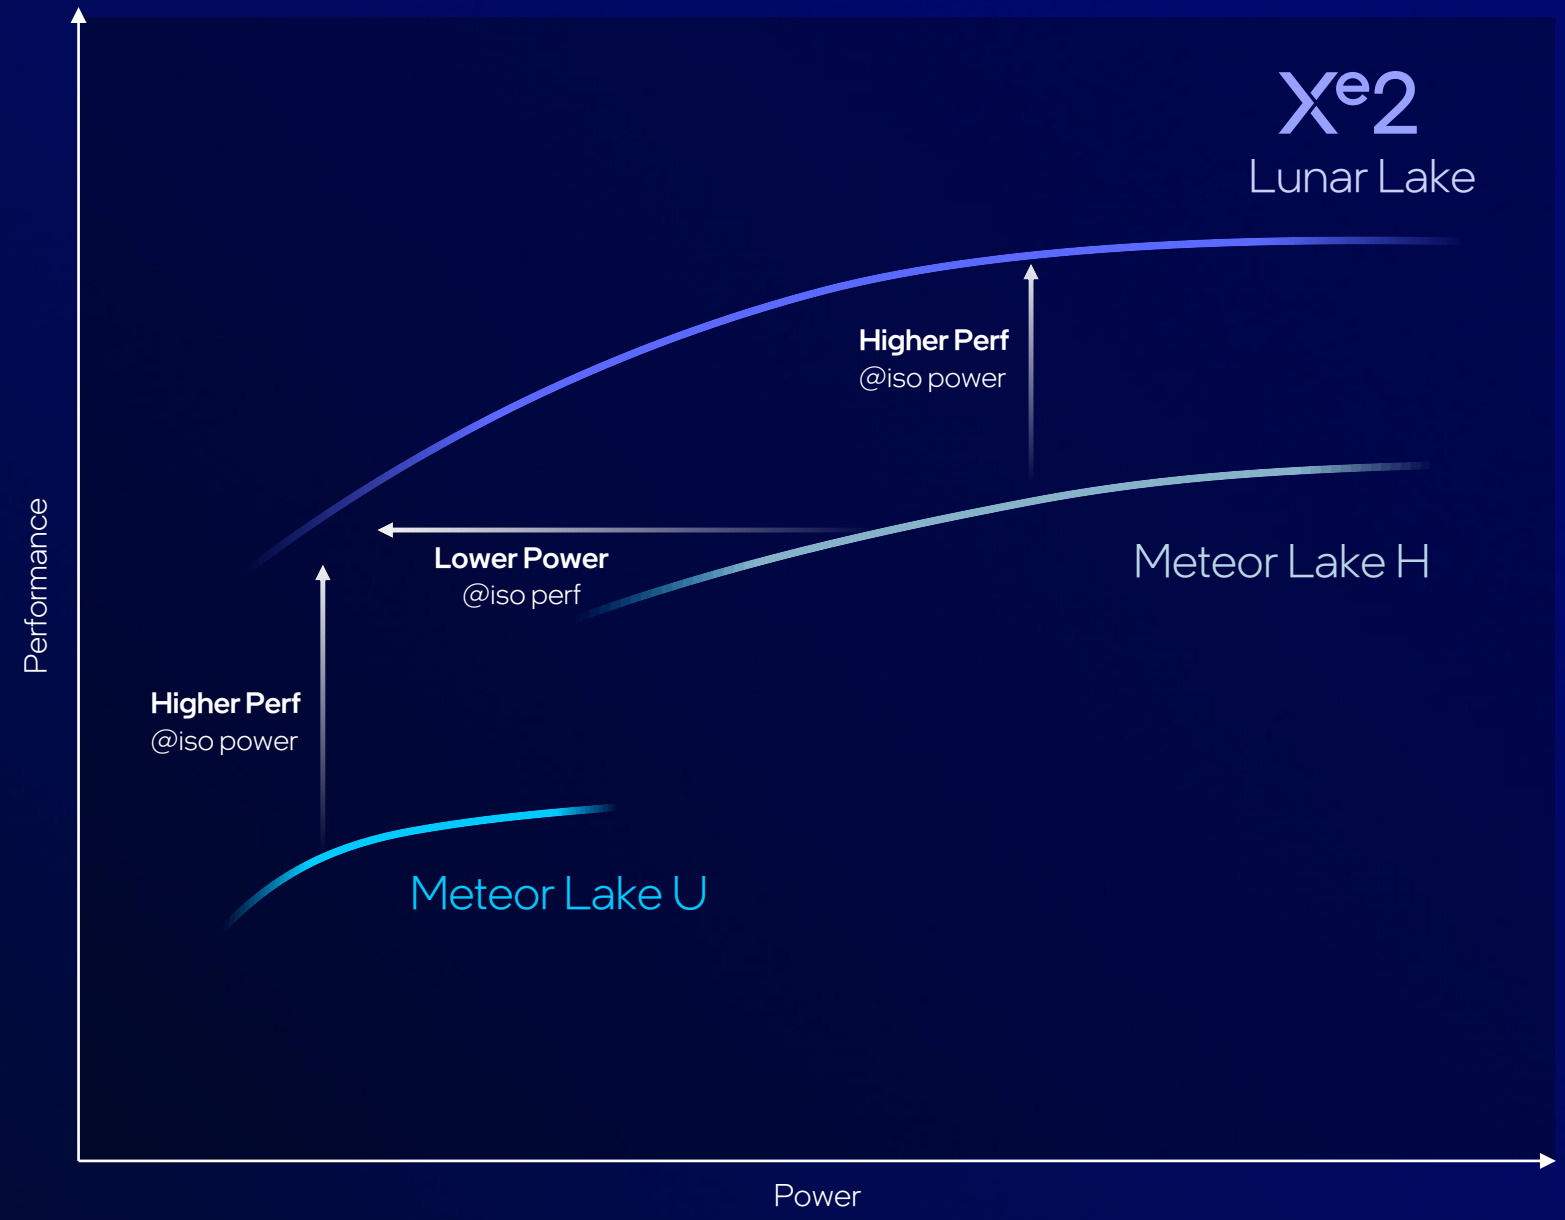

See backup for workloads and configurations. Results may vary.

Lunar Lake

Media & display engines

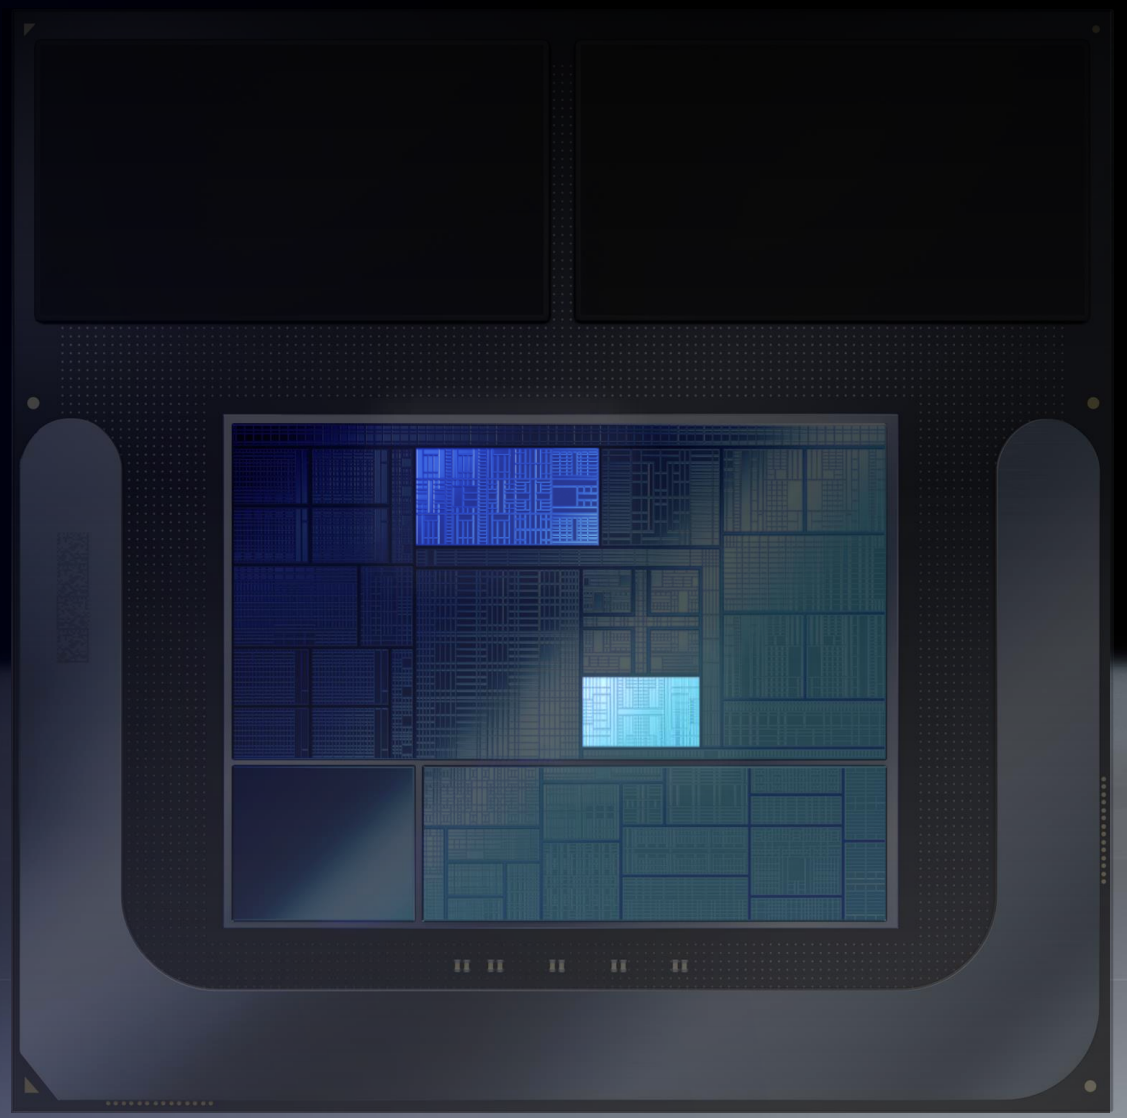

New Media & Display Engines

Media engine

- Encode & decode
- Decode

Display engine

- 1x eDP 1.5
- DisplayPort 2.1
- HDMI 2.1
- 3 display pipes

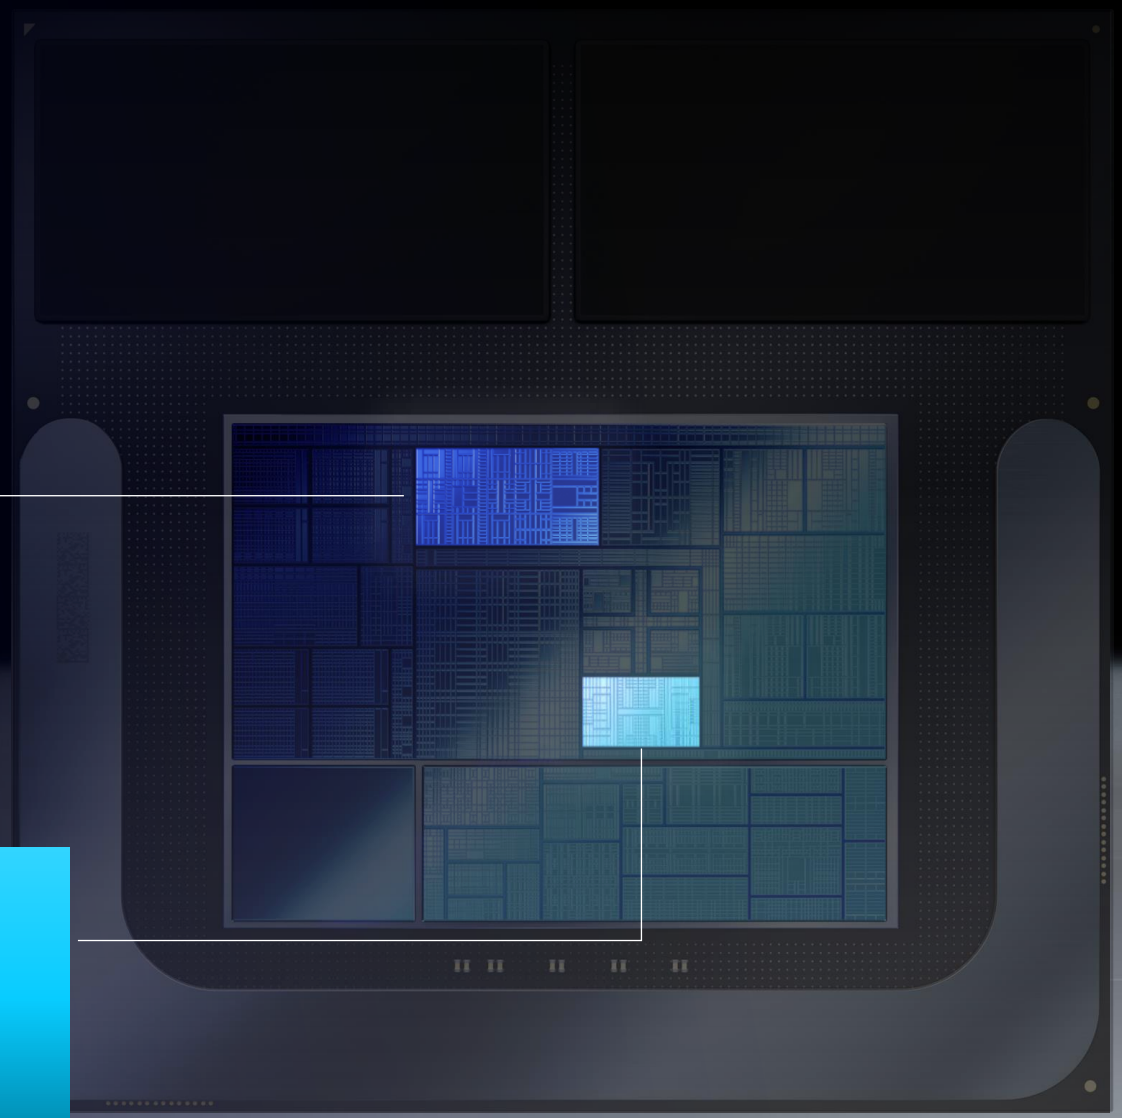

VVC

Significantly reducing bitrate at the same quality

10%

Reduction over AV1

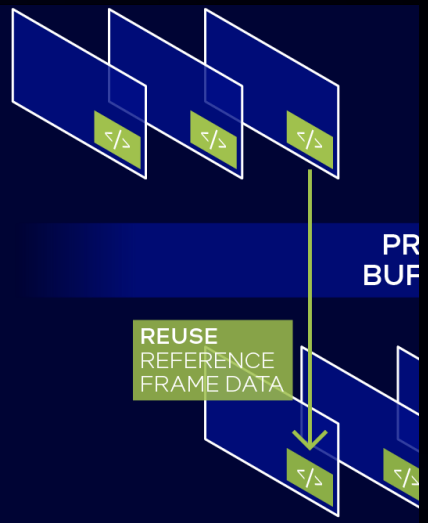

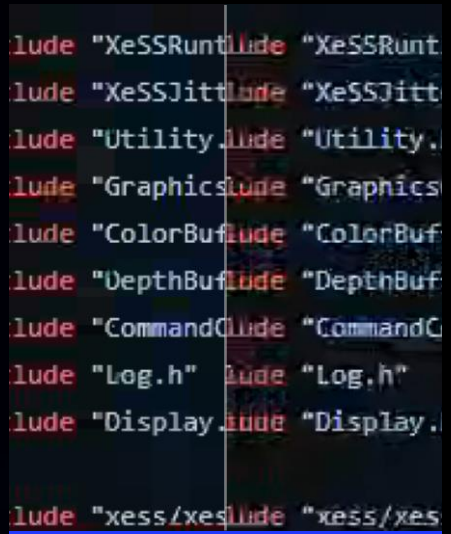

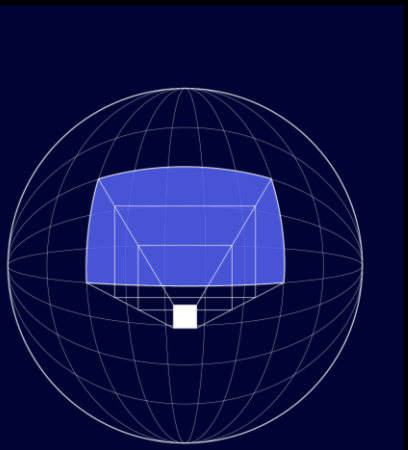

Reduction in file size

Adaptive resolution streaming

Screen content coding

360-degree & panoramic

DisplayPort

eDisplayPort 1.5

Panel Replay

- Evolution of Panel Self Refresh
- Selective update with early transport
- Adaptive sync with panel replay

Massive Leap in Graphics

Up to 50% better graphics performance over Meteor Lake

Real-time ray-tracing

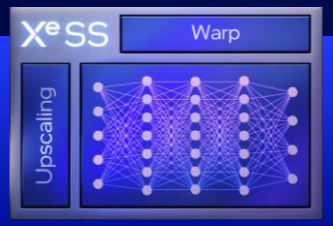

XeSS AI-based upscaling

intel ARC™

Software stack

Up to 67 TOPS

VVC

Decode

HDMI

DisplayPort & HDMI 2.1

Lunar Lake

New NPU 4.0

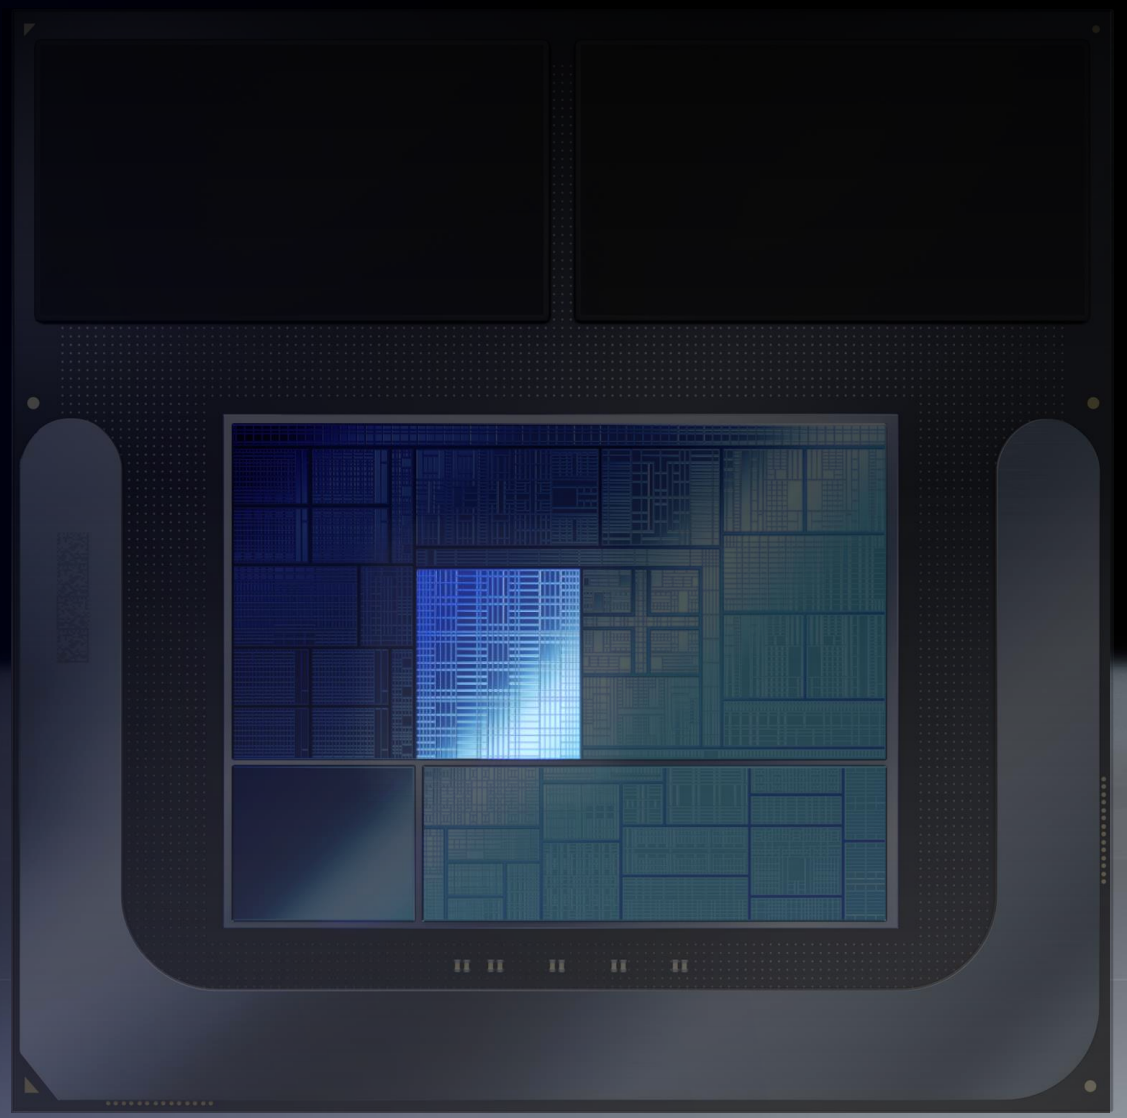

Next Gen NPU 4

Architecture goals

Increase NPU size
To run next gen AI workloads

Increase clock & efficiency
To increase performance and battery life

Optimize for modern AI
For efficiently running LLMs and transformers

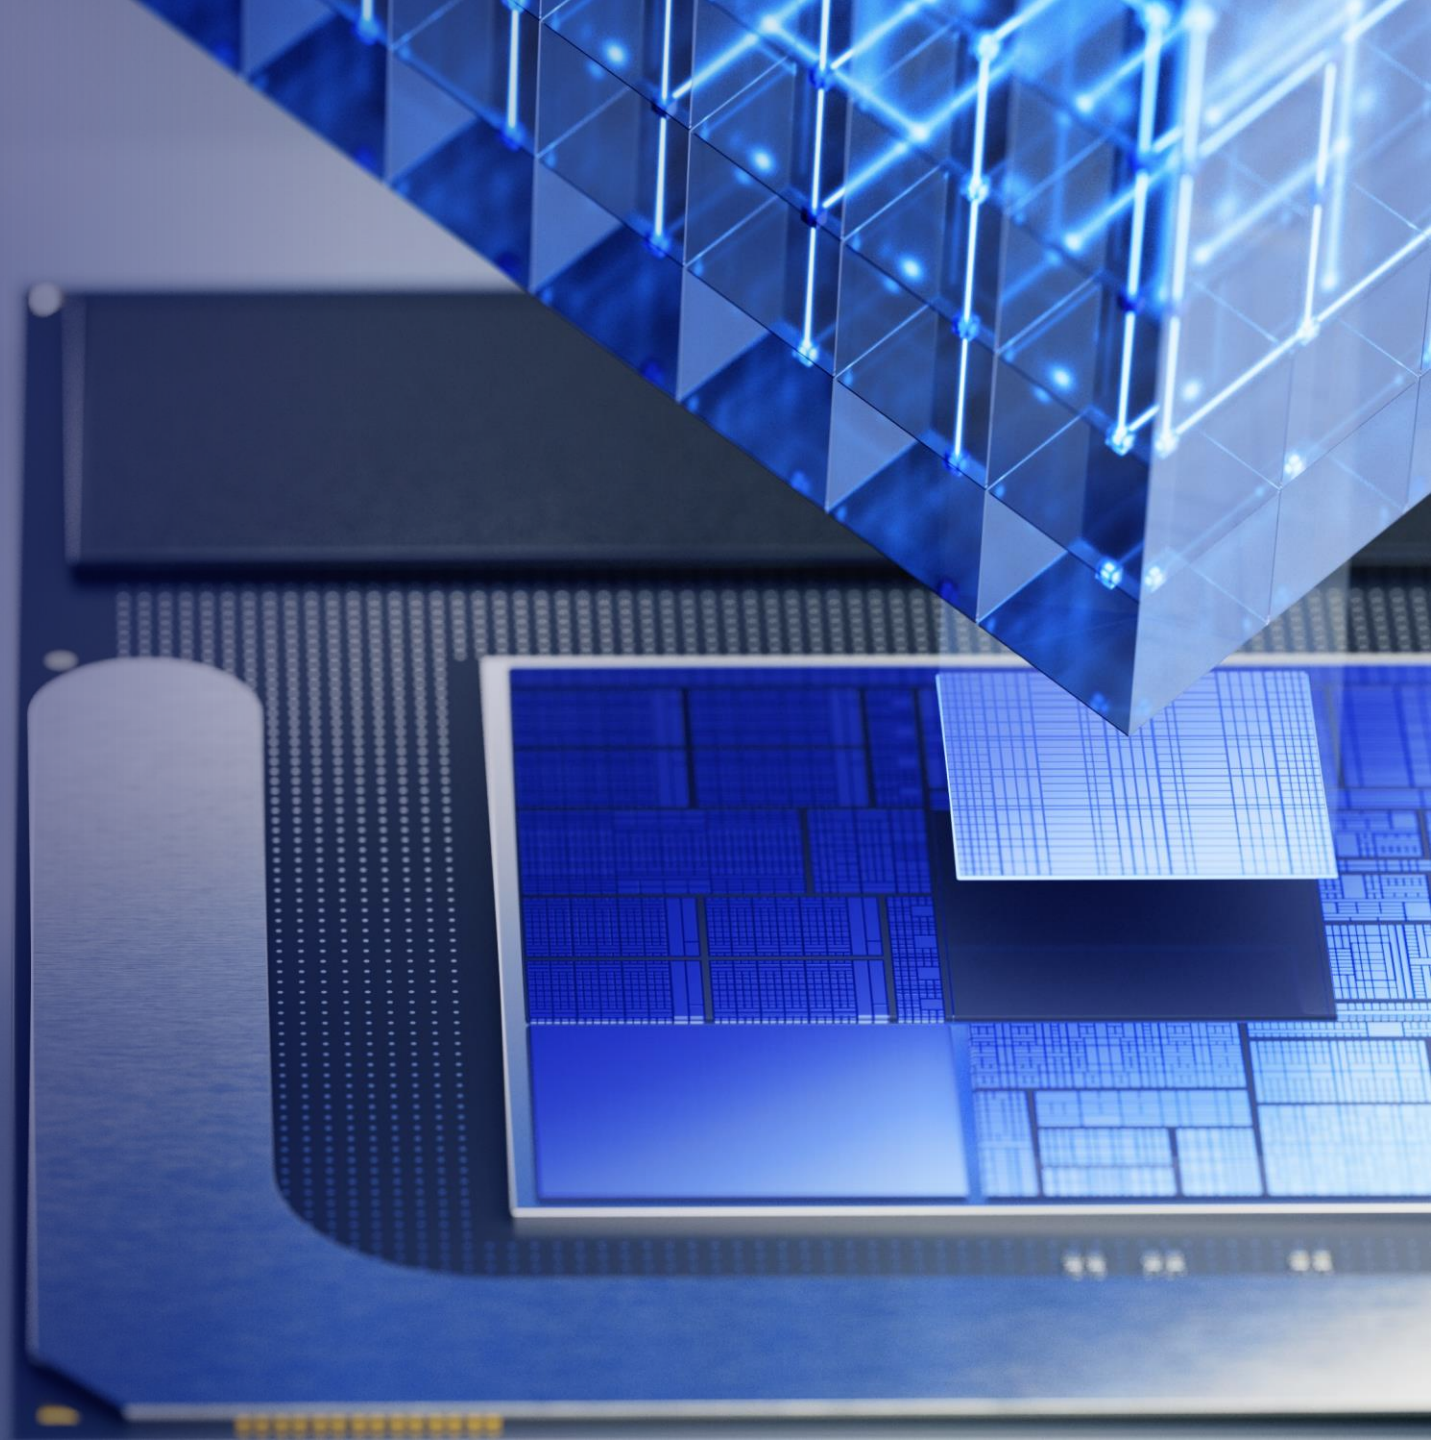

Next Gen NPU 4

Largest integrated and dedicated AI accelerator for the AI PC

Native activation function & data conversion support

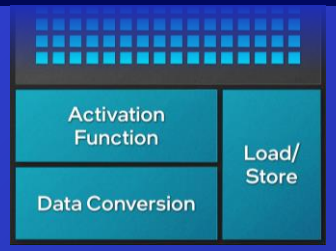

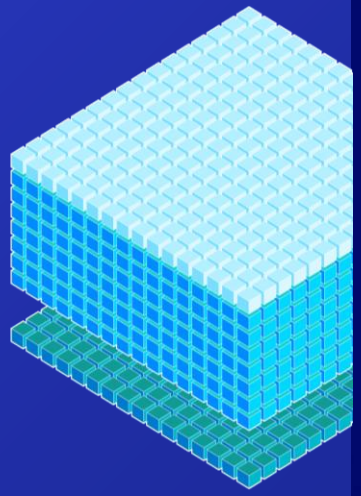

Efficiency optimized MAC array

up to **48** TOPS

2x Bandwidth

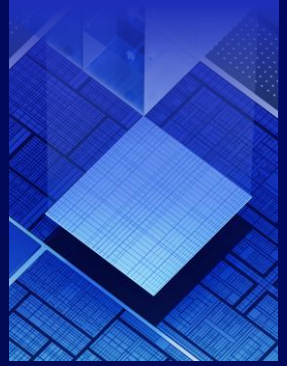

DMA

Embedding tokenization used for LLMs

12 **Enhanced SHAVE DSPs**

Accelerating LLM & transformer operations

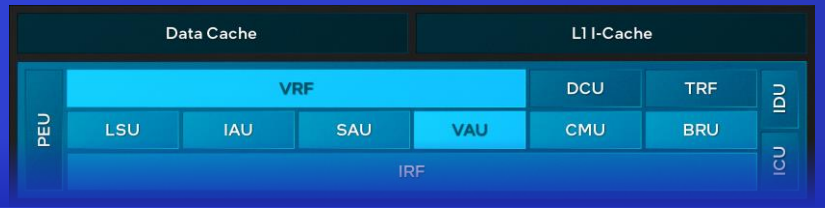

6 Neural compute engines

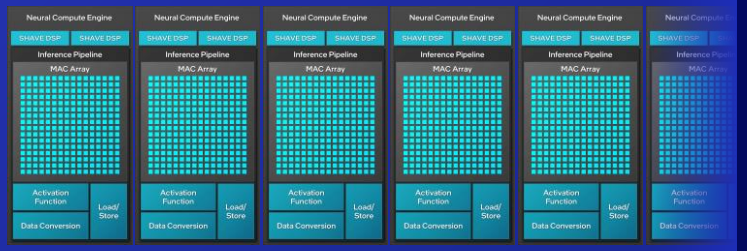

intel

Next Gen NPU 4

Scaling AI performance and efficiency at AI pace

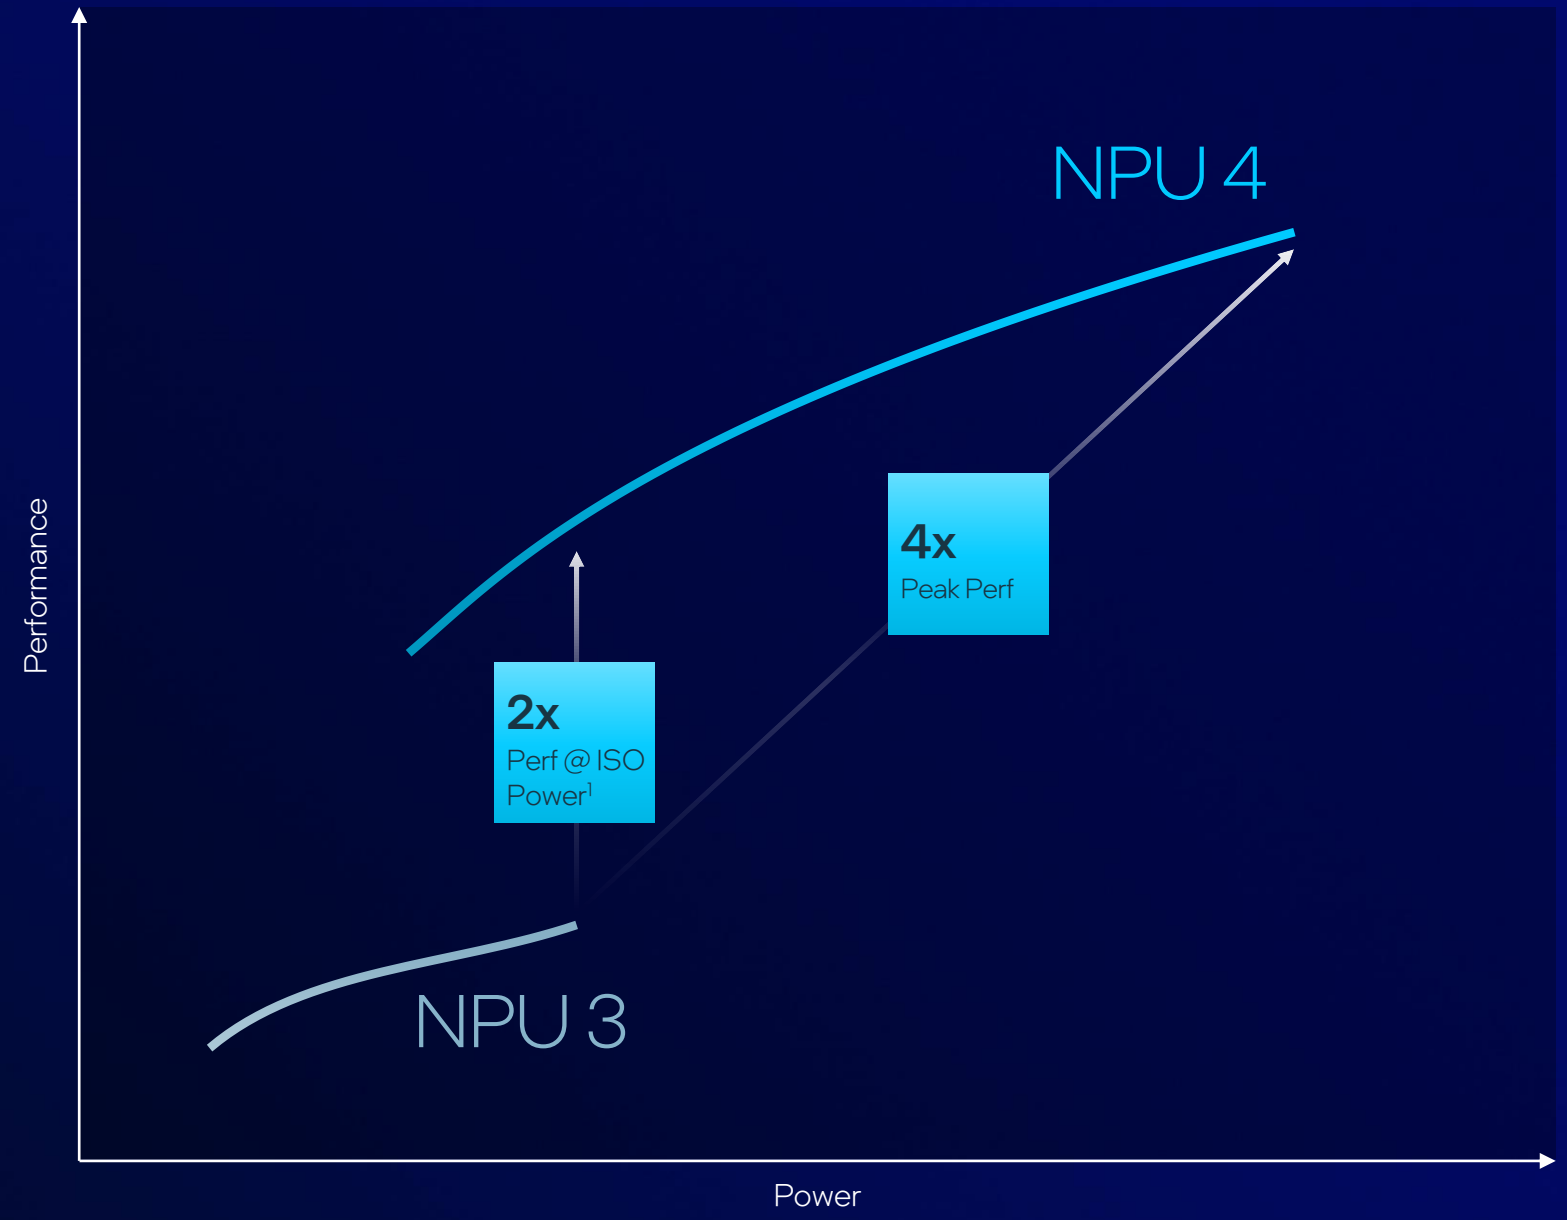

NPU 4 – in Lunar Lake

NPU 3 – in Meteor Lake

[1] Based on pre-production simulation data of a real network. See backup for details.

Unmatched AI Compute

Up to **120** platform TOPS

GPU	Up to **67** TOPS	XMX & DP4a	Gaming & creator AI
NPU	Up to **48** TOPS	Dense vector & matrix math	AI assistants & creation
CPU	Up to **5** TOPS	VNNI & AVX	Light AI workloads

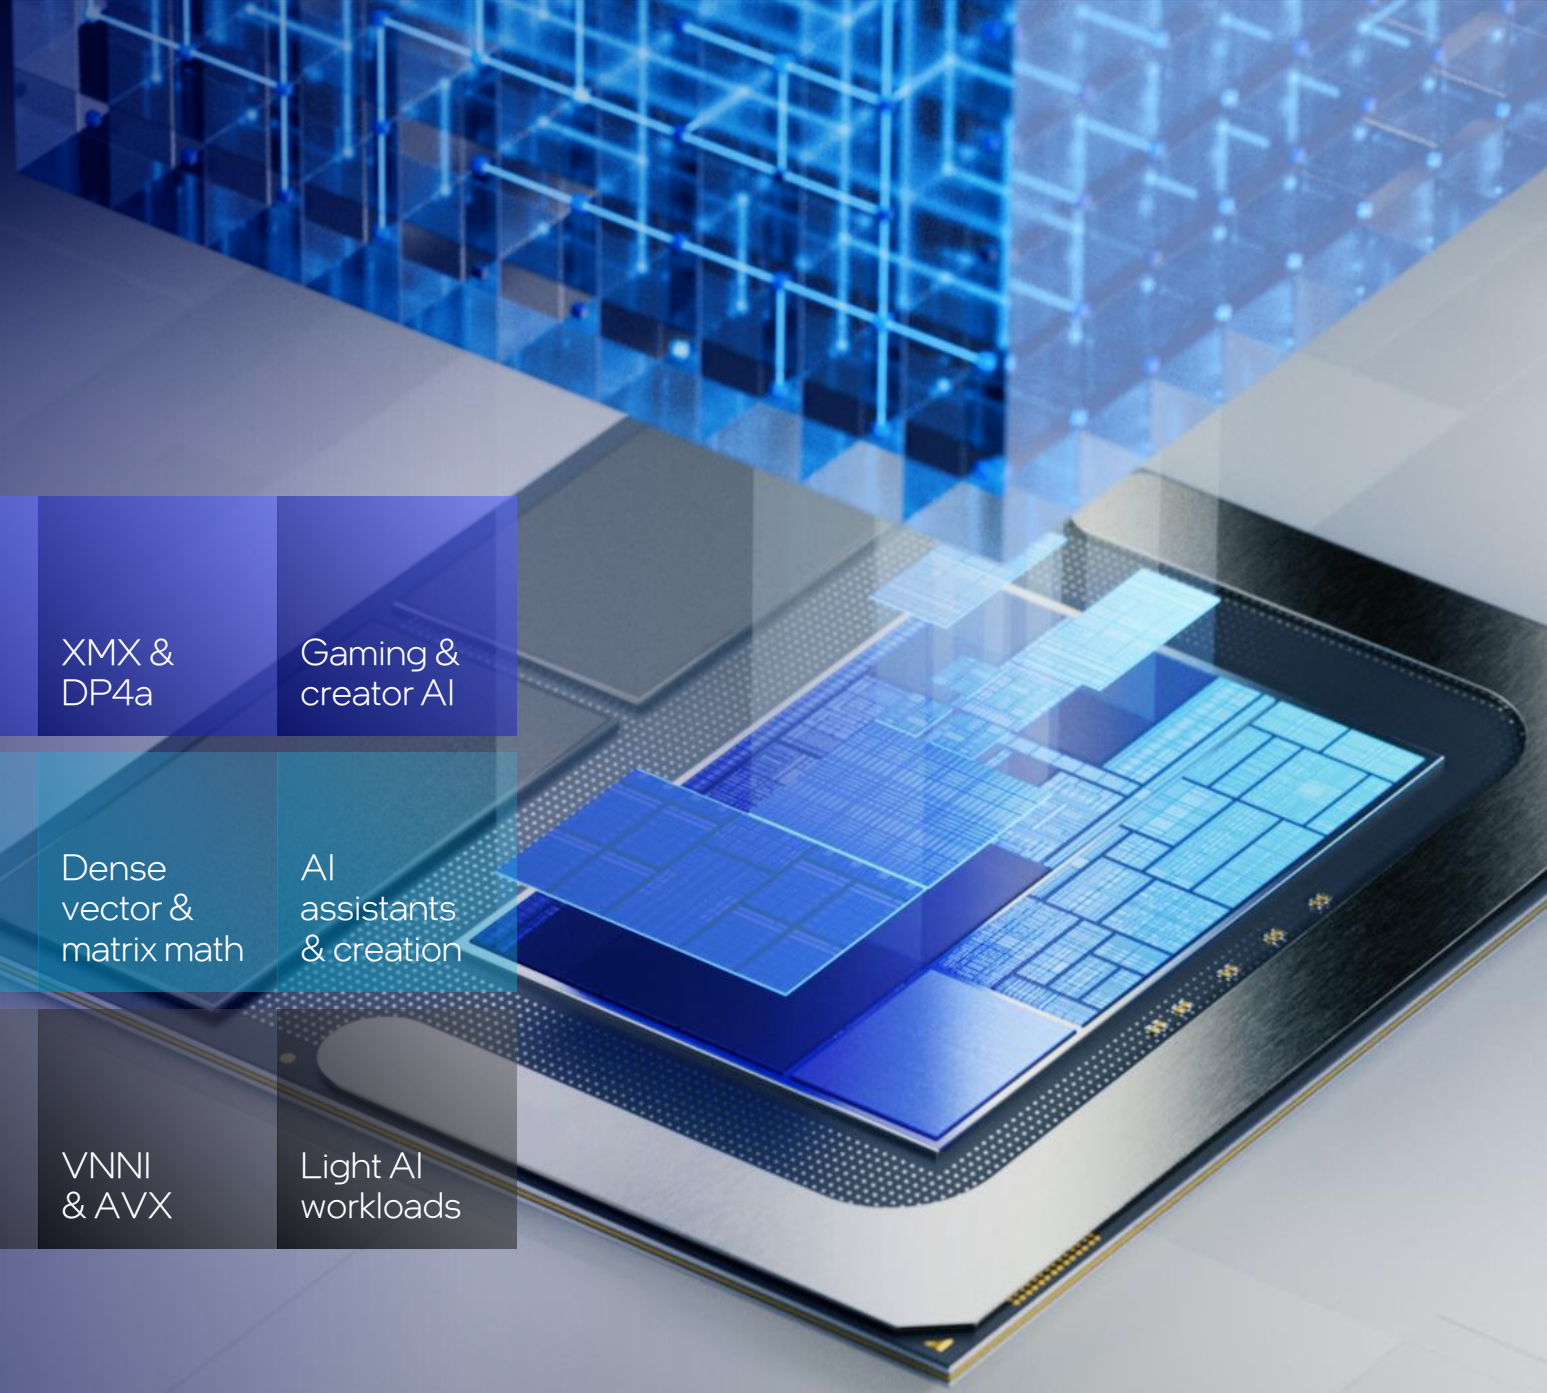

Lunar Lake

Connectivity

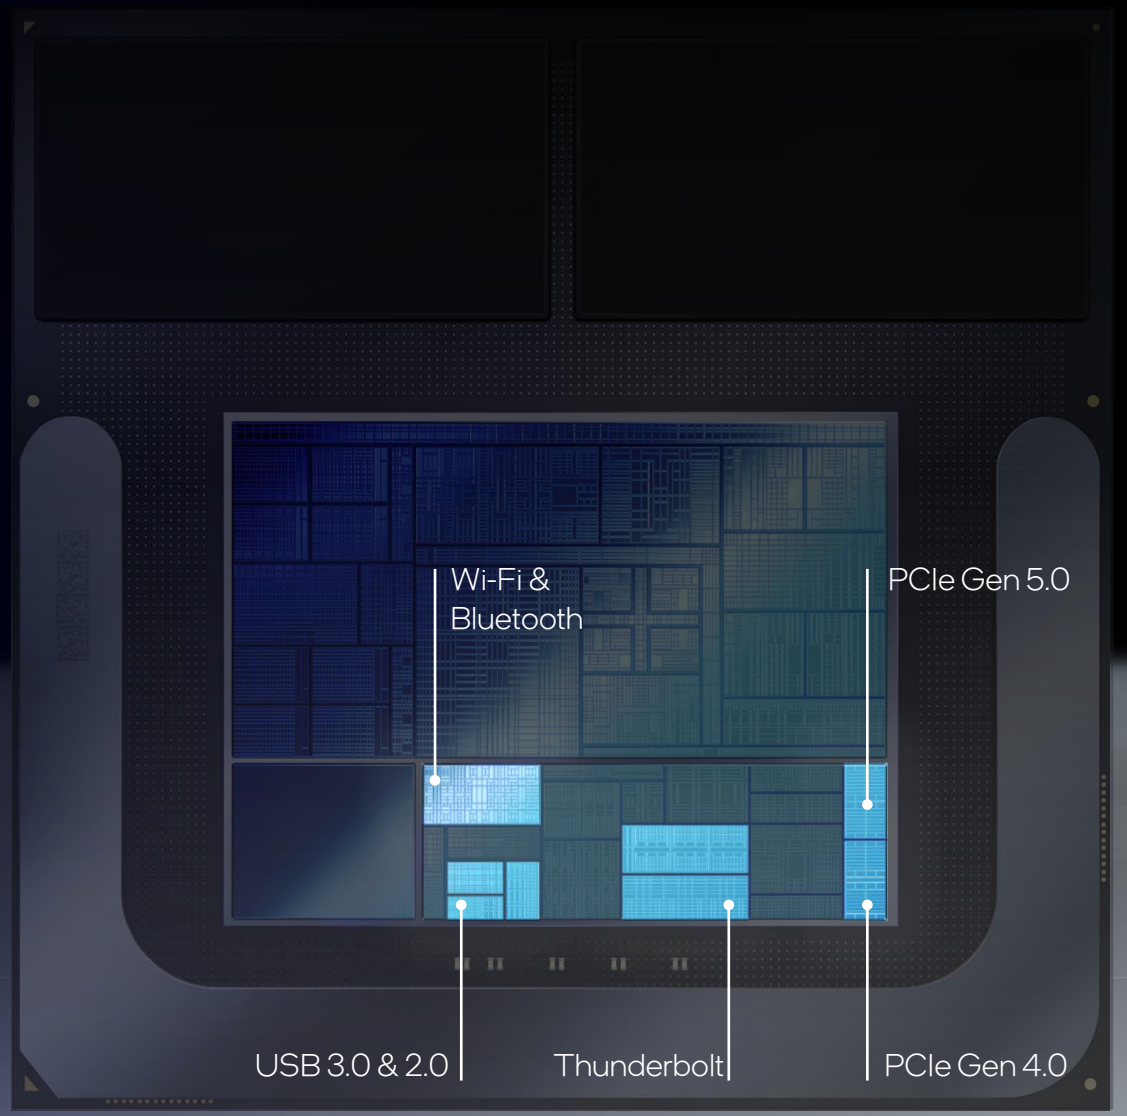

Leadership Connectivity

Integrated right onto the package

Intel Unison

New multi-device experiences

Tablet control

Swift connect

Universal hotspot*

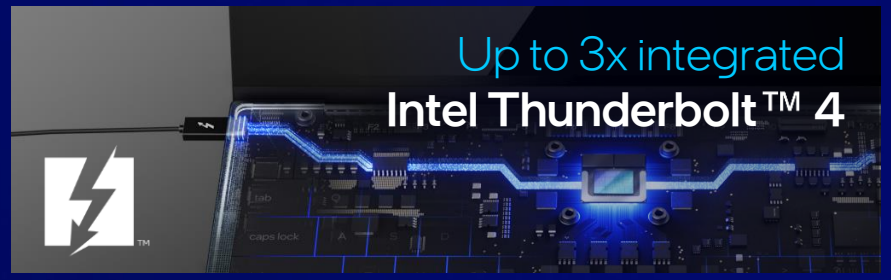

Thunderbolt Share

Share between PCs at Thunderbolt speed

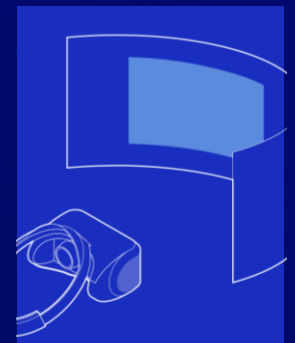

Enhanced VR

With Wi-Fi 7 & Intel Killer Wi-Fi

Intel® Bluetooth® 5.4

For efficient & HD audio

Integrated Intel®Wi-Fi 7 (5 Gig)

5.8Gb/s

Wi-Fi 7 speed

40Gb/s

TBT4 speed

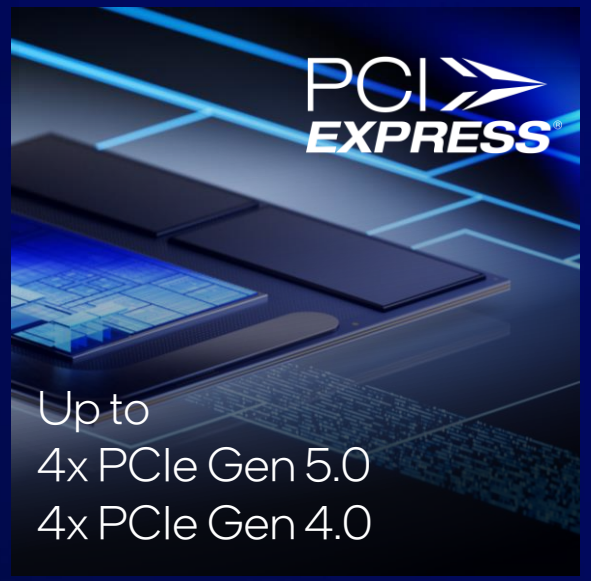

intel.

* Available post LNL launch

Energy Efficiency

With innovations and deep integration across the entire platform

Memory on package

Up to 40% data movement power reduction (PHY)

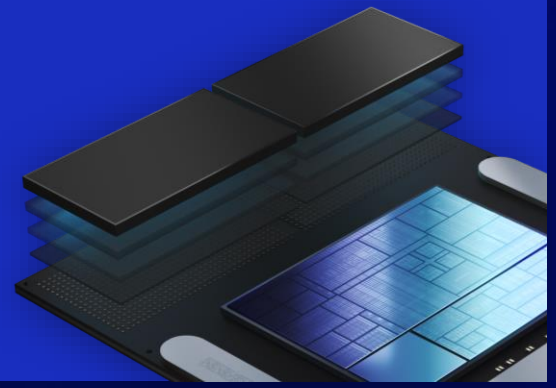

Power management enhancement & integration

Improved **E-core cluster**

New
8MB
Memory side cache

Efficiently “feeding” memory hungry IPs

New power delivery architecture
for better control, enhancing power utilization

Up to
40% lower SOC power*

intel

* Over previous gen

Lunar Lake

Built-in ray-tracing

Bigger NPU 4.0

Integrated Wi-Fi 7

Enhanced Intel® Thread Director

Foveros packaging

4P + 4E core design

Up to **120** total platform **AI TOPS**

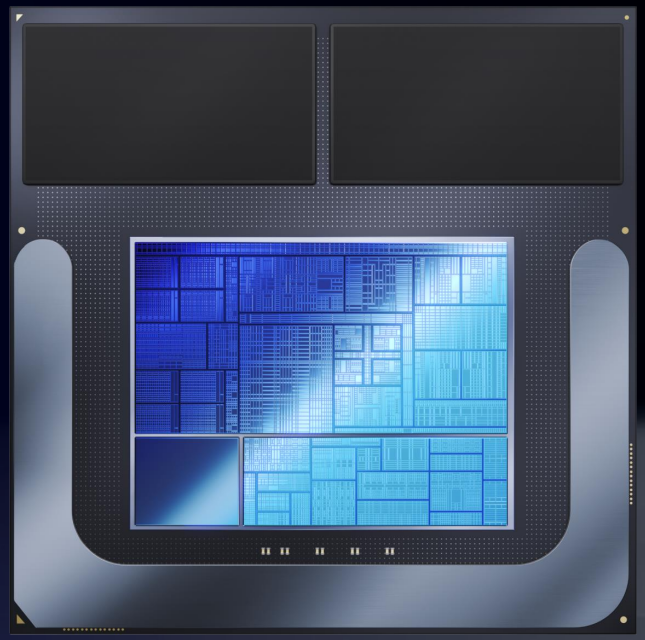

New **Xe2** GPU

Up to **32GB** memory on package

VVC Decode

P-core **Lion Cove**

E-core **Skymont**

Integrated Thunderbolt™ 4

New Intel® Partner Security Engine

New PMIC power delivery architecture

intel


intel
TECH
tour.TW
Thank
You

Notices & Disclaimers

The preceding presentation contains product features that are currently under development. Information shown through the presentation is based on current expectations and subject to change without notice.

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.

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

AI features may require software purchase, subscription or enablement by a software or platform provider, or may have specific configuration or compatibility requirements. Details at www.intel.com/AIPC.

No product or component can be absolutely secure. Intel technologies may require enabled hardware, software or service activation.

All product plans and roadmaps are subject to change without notice.

Performance hybrid architecture combines two core microarchitectures, Performance-cores (P-cores) and Efficient-cores (E-cores), on a single processor die first introduced on 12th Gen Intel® Core™ processors. Select 12th Gen and newer Intel® Core™ processors do not have performance hybrid architecture, only P-cores or E-cores, and may have the same cache size. See ark.intel.com for SKU details, including cache size and core frequency.

Built-in Intel® Arc™ GPU only available on select Intel® Core™ Ultra processor-powered systems; OEM enablement required.

Some images may have been altered or simulated and are for illustrative purposes only.

Built into the hardware, Intel® Thread Director is provided only in performance hybrid architecture configurations of 12th Gen or newer Intel® Core™ processors; OS enablement is required. Available features and functionality vary by OS.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

While Wi-Fi 7 is backward compatible with previous generations, new Wi-Fi 7 features require PCs configured with Intel Wi-Fi 7 solutions, PC OEM enabling, operating system support, and use with appropriate Wi-Fi 7 routers/APs/gateways. 6 GHz Wi-Fi 7 may not be available in all regions. Performance varies by use, configuration, and other factors. For details on performance claims, learn more at www.Intel.com/performance-wireless.

APPENDIX

Claim # & Statement	Slide # & Title/Details
	SLIDE 2: Flagship SoC for the next gen of AI PCs
Up to 40% lower SoC power	Testing by Intel as of May 2024. Data based on Lunar Lake reference validation platform measurement vs Meteor Lake reference validation platform as measured by YouTube 4K 30 AV1.
Similar ST perf at half the power	Testing by Intel as of May 2024. Data based on Lunar Lake reference validation platform as measured CBR24 ST vs. prior generation.
Up to 1.5X better graphics	Testing by Intel as of May 2024. Data based on Lunar Lake reference validation platform measurement vs Meteor Lake reference validation platform as measured by 3DM Time Spy. 3DMark*
	SLIDE 9: Lion Cove P-core
Lion Cove core delivers 14% better IPC vs. Redwood Cove core	IIso frequency benefit estimate across: components of SPECrate2017_int_base and SPECrate2017_fp_base (both estimated) running 1 copy, Cinebench R23 Single Core, Cinebench 2024 Single Core, Geekbench 5.4.5 Single-Core, Geekbench 6.2.1 Single-Core, WebXPRT 4, Speedometer
Lion Cove Performance at different power levels vs. Redwood Cove	Results are based on SPECrate2017_int_base (estimated) running n copies. Based on measurement on an Intel internal reference validation platforms at a fixed PL1 power setting.

APPENDIX

Claim # & Statement	Slide # & Title/Details
	SLIDE 13: Skymont E-core
Skymont IPC on Lunar Lake Low Power Island: 1.38x integer and 1.68x floating point vs. Meteor Lake LP E-core (Crestmont)	Results are based on Intel's internal projections/estimates as of 5.13.2024(+/- 10% Margin of Error) on SPEC CPU 2017 Rate est, GCC12.1-O2 Linux at Fixed Frequency (ISO).
Skymont Power & Performance on Lunar Lake Low Power Island: up to 2x peak ST performance or 1/3 the power at similar ST performance and up to 4x peak MT performance or 1/3 of the power at similar MT performance	Results are based on Intel's internal projections/estimates as of 5.13.2024(+/- 10% Margin of Error) on SPEC CPU 2017_int_base est, GCC12.1-O2 Linux (ISO). Comparing a Skymont E-core cluster (4 Skymont cores) vs. Meteor Lake LP E-core cluster (2 Crestmont cores) to showcase workload coverage increase for the Lunar Lake Low Power Island.
	SLIDE 15: Breakthrough x86 Power Efficiency
>50% peak performance >20-80% per/watt	Illustration of the relative Lunar Lake P-core and E-core performance across the SoC power range.
	SLIDE 21: Improved Experience
35% power reduction when containment & power management optimization are enabled	As of May 2024, based performance estimated with measurements on Lunar Lake reference platform with power optimizations enabled vs. power optimizations disabled.

APPENDIX

Claim # & Statement	Slide # & Title/Details
	SLIDE 26: Next Gen Xe2 GPU
1.5x better vs. Meteor Lake GPU	Testing by Intel as of May2024. Data based on Lunar Lake reference validation platform measurement vs Meteor Lake reference validation platform as measured by 3DM Time Spy. 3DMark*
	SLIDE 28: Next Gen Xe2 GPU
1.5x graphics performance over Meteor Lake	Testing by Intel as of May2024. Data based on Lunar Lake reference validation platform measurement vs Meteor Lake reference validation platform as measured by 3DM Time Spy. 3DMark*
	SLIDE 37: Next Gen NPU 4
2x performance at ISO power vs. Meteor Lake	Testing by Intel as of January 2024. Based on VPU-EM simulation. Power data is generated from the simulation tool based on power data that has been extracted from circuit simulation tools. This simulation, which is a ~100% utilization int8 network, is expected to correlate well with silicon.
4x peak performance	4x peak performance is based on TOPS increase from MTL (11 TOPS) to LNL (48 TOPS).
	SLIDE 41: Energy Efficiency
Up to 40% lower SoC power vs. Meteor Lake	Testing by Intel as of May 2024. Data based on Lunar Lake reference validation platform measurement vs Meteor Lake reference validation platform as measured by YouTube 4K 30 AV1.


intel