

The Intel logo, consisting of the word "intel" in a white, lowercase, sans-serif font, is positioned inside a solid blue square. The background of the entire slide is a dark blue, 3D-rendered grid of cubes that recede into the distance, creating a sense of depth and a futuristic, technological atmosphere.

intel[®]

Embedded World 2026

Press Briefing

Michael Masci | Vice President of Product Management



Expanding Edge AI Offerings



Intel® Core™ Series 2 processors with P-cores

Formerly codenamed Bartlett Lake 12P



Intel® Core™ Ultra Series 3 processors

Formerly codenamed Panther Lake



Health & Life Sciences AI Suite

Preview

Edge Transformations

Fixed Function
Embedded Designs

Internet
of Things

Computer
Vision

Modernized Edge Infrastructure
with Software Defined Solutions

Agentic &
Physical AI

40+ Year Journey

Trusted, Proven Dependable

INTEL'S EDGE LEADERSHIP

Empowering Customer Success

100K

Partner
deployments

250M+

Units shipped
in last 10 years

Customer Enhancements (per year)

Linux Kernel Contributions (last 10 years)

500+
customers

10K+
enhancements

Dozens
of platforms

>75K
Intel

>2x than AMD
>5x than NVIDIA
>13x than QCOM

INTRODUCING

Intel® Core™ Series 2 processor with P-cores

(codenamed Bartlett Lake 12P)



Consistent High Perf
& Simplified Compute

Up to 12 P-cores

Up to 1.5x higher multi-threaded performance vs prior gen

New 125W SKUs

For higher thermal headroom

Run mainstream AI

on CPU, pair with dGPU for advanced AI



Industrial Durability
& Flexibility

10-year availability

with LTSC OS support

Time-aware Execution

Intel® TCC and TSN

Socket Compatible

across 12th–14th Gen



Built for Developer
Agility & Efficiency

Windows Server Ready

Optimize licensing

Framework ready

with Intel® OpenVINO™, PyTorch, ONNX Runtime, and WinML

Open Edge Platform

Build, deploy, run, secure, and manage edge solutions



For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex).

Built for Deterministic Edge Performance

Industrial-ready processor

RTOS
Linux RT

intel
CORE

Intel® TCC
(Time Coordinated Computing)

intel
CORE

TSN support

ULTRA



Deterministic control cycle execution

Ensure real-time loops execute within strict timing to maintain reliable system behavior.



Run mixed criticality workloads on the same platform

Temporal isolation of workloads maximizes system efficiency and reducing hardware footprint.



Fluid robot motion & grasp stability

Minimize jitter so robotic movements and gripping actions are smoother and more consistent.



Accelerate development and deployment

Reduce manual configuration effort, enabling faster setup and deployment.

Get Measurable Real-Time Advantage

Improve TCO with higher real-time industrial workload density

Higher precision | Lower jitter | Greater reliability

Intel® Core™ 9 processor 273PE vs AMD Ryzen™ 7 9700X

Up to

3.8x

Better deterministic
performance

Up to

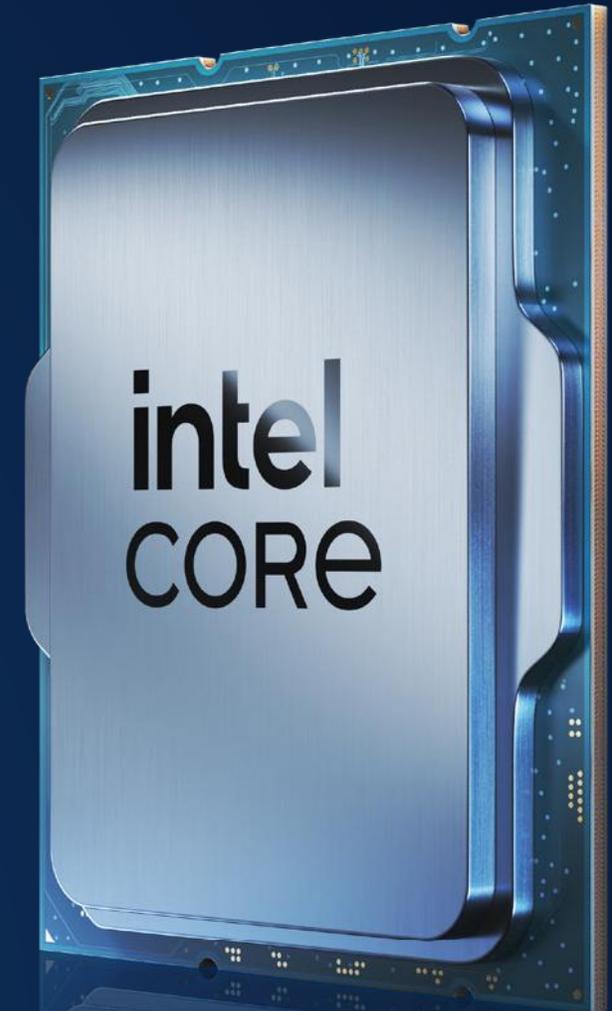
2.5x

More deterministic
response time

Up to

4.4x

Lower max
PCIe latency



As of Feb 2026
For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex).



Elevating the Edge with Partners

NEUROCLE

Quality Inspection

...1.4x reduction in inference latency ...more responsive defect detection, faster image analysis.

Hongsuk Lee - CEO

SAIMOS®

Safety/Security Monitoring

...2.3x gain in efficiency, more cameras per server, lower energy, real-time intelligence for busy public spaces...

Karim Elhanafi - CEO

XYZ Robotics

Warehouse Automation

...5-10% improvement in productivity... more output from the same robots, faster payback.

Peter Yu - CTO

CODESYS

Industrial Control

...1.67x increase in virtual PLC density ...elevates what's possible in a modern control architecture.

Hilmar Panzer - CTO



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Performance varies by use, configuration, and other factors. See [intel.com/performanceindex](https://www.intel.com/performanceindex) for details.



***Intel[®] Core[™]
Ultra Series 3
processors***
at the edge

Intel® Core™ Ultra Series 3 processors

Proven Ideal for Mission-Critical Edge and Physical AI



Power-efficient
AI and graphics

Up to 1.8x more TOPS
vs. prior gen

**Up to 50% more GPU
cores** vs. prior gen

Supercharge agents
Concurrent reasoning with
integrated AI acceleration



Deterministic
performance

Time-aware execution
Intel® TCC and Discrete TSN

Reliable Real-Time & AI
Ideal for mixed-criticality
workloads

AI with real-time response
Ensures low-latency, instant
decision-making



Industrial Durability
& Flexibility

Industrial Grade
Support for extended temp
and ECC memory

Accelerated FuSa dev
with Intel® Silicon Integrity
Tech & FSEDP

Design for 100% operation
with up to 10 years availability
and Windows and Linux long
term service channel



Built for
Developer Agility

Accelerate AI inference
with OpenVINO™ across CPU,
GPU and NPU

Fast track development
using Open Edge Platform's
modular open-source AI tools

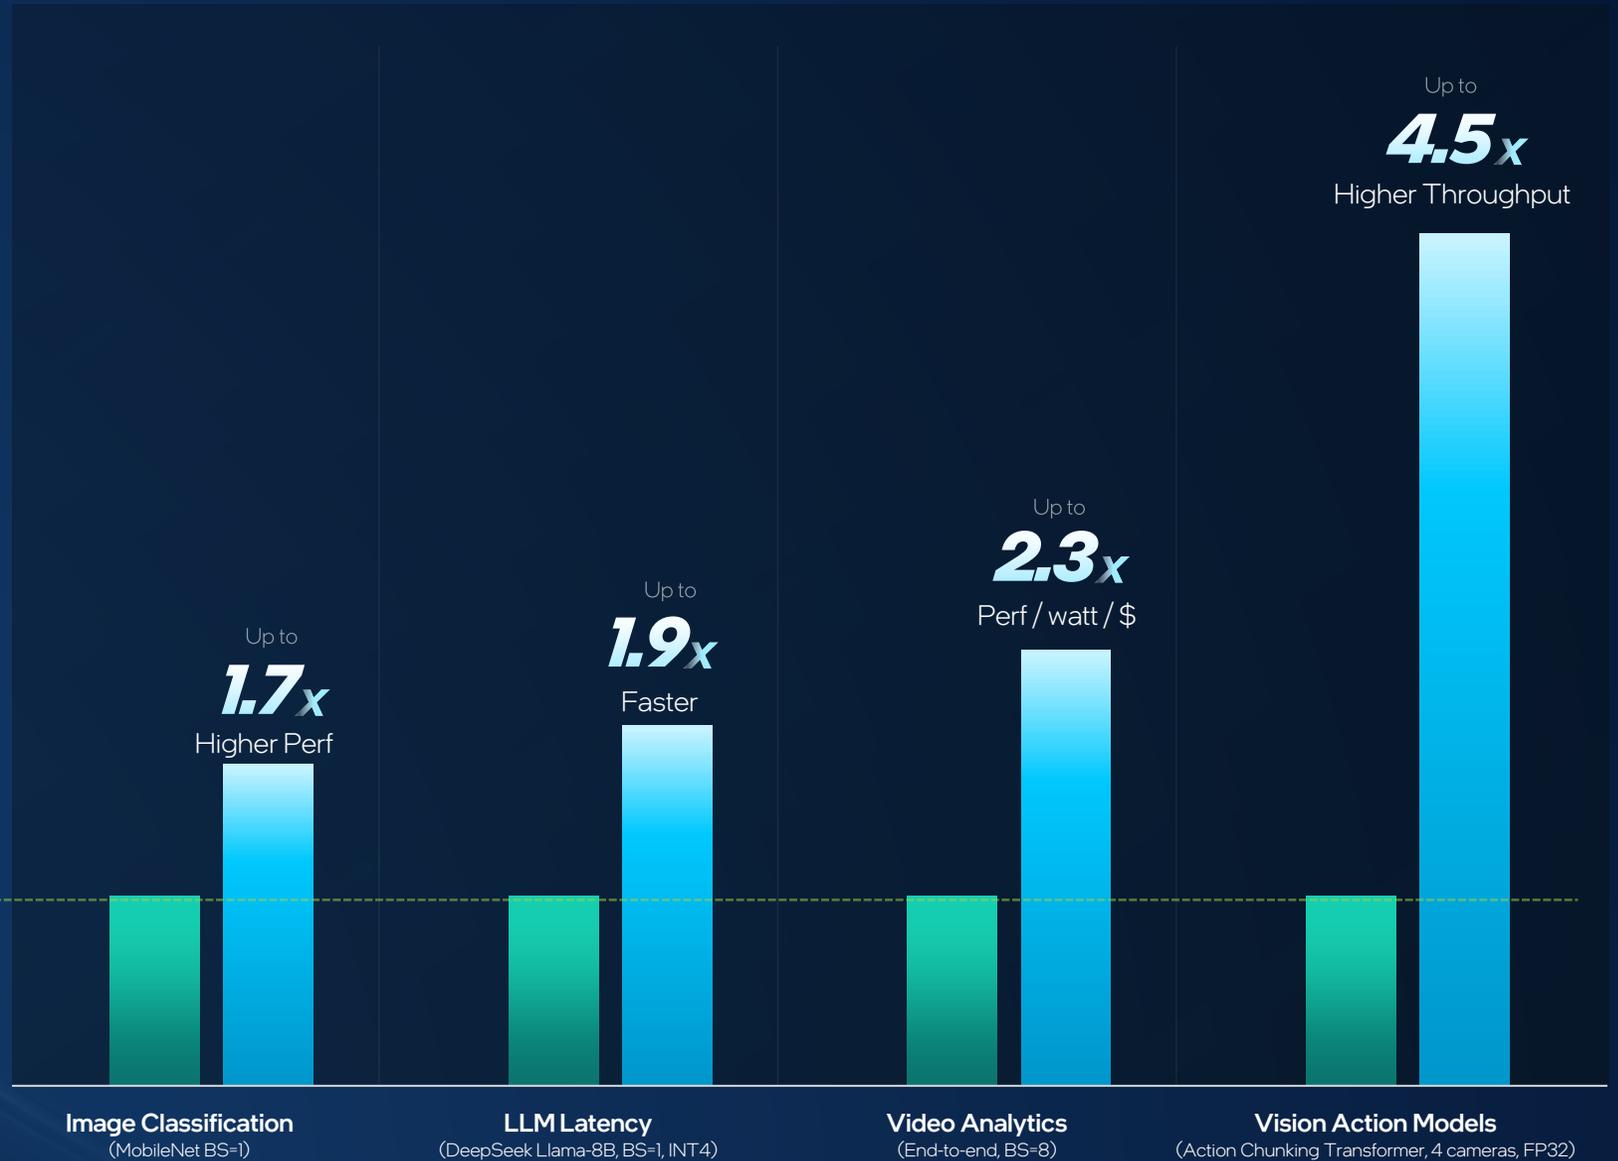
Industry curated AI
Reference apps, verified
systems and benchmarks for
fast time-to-value

Ultra Performance at the Edge

Intel® Core™ Ultra X9 388H vs. NVIDIA® Jetson AGX Orin 64GB

NVIDIA Jetson AGX Orin 64G

- Intel® Core™ Ultra X9 388H
- NVIDIA® Jetson AGX Orin 64GB



Relative performance (higher is better)



As of Dec 2025
Details at intel.com/performanceindex. Results may vary.

Leadership Performance Across ALL AI Engines

Geekbench AI 1.6 Scores



Intel® Core™ Ultra X9 388H

		vs. AMD HX 370	vs. Qualcomm 84-100
CPU	INT8	1.2x	1.5x
	FP16	2.3x	1.7x
	FP32	1.2x	2.7x
GPU	INT8	9.0x	Does Not Run ❌
	FP16	3.1x	Does Not Run ❌
	FP32	2.3x	Does Not Run ❌
NPU	INT8	Does Not Run ❌	1.3x
	FP16	Does Not Run ❌	2.2x
	FP32	Does Not Run ❌	2.6x



As of Dec 2025
See [intel.com/performanceindex](https://www.intel.com/performanceindex) for details. Results may vary.

Proven TCO savings for Edge AI

Intel Core Ultra Series 3 displaces competition across multiple verticals and customers



Document Parsing

OCR and Document AI

Up to

39%

Savings per system over 5 years

Displaced NV RTX 4060

The background of this panel shows a chalkboard filled with mathematical equations and diagrams, including $D(x) = a$, $x^2 + y^2 = ab + 4c$, and $A \cdot T = B$.

Robotics AI

Vision Language Action and imitation learning policy

Up to

42%

Savings per system

Displaced NV Jetson AGX Orin 64GB

The background of this panel features a humanoid robot with a white head and torso, set against a blurred background of a modern interior.

Humanoid Robotics

Speech synthesis, VLM and multi-task vision

Up to

67%

Savings per system over 5 years

Displaced NV Jetson Thor

The background of this panel shows a close-up profile of a humanoid robot's head, highlighting its eye and ear sensors.

TCO

Infer + Fine Tune on a Single SoC

\$5,549

or **5.8x**

Estimated savings
if you move away from dual-systems
to Intel® Core™ Ultra Series 3

Fine-tune at ~87%
of a discrete GPU¹ performance —
for only 17% of the cost, plus better inference

More savings
in energy cost



As of Dec 2025
Performance varies by use, configuration, and other factors. For more complete information, visit intel.com/PerformanceIndex.



Pushing the Edge with Partners



Medical Imaging

...localization up to 152X faster on the NPU, ...accelerates decision-making while keeping sensitive data secure.

Sharon Saban - General Manager



Humanoid

...3.9X higher VLM throughput, 5.4x faster multi-task vision against NVIDIA Jetson AGX Orin...

Park Jong Gun - CEO



Barista Robot

...nearly 2x faster time-to-first-token at roughly 22 watts, responsive, stable Physical AI without cost and complexity...

Keith Tan - CEO & Founder



Safety/Security Monitoring

...1.2x faster with 1.7x lower latency ...smoother and responsive pipeline from ingestion to output.

Eugene Beytenbrod - Global Director of Engineering



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Performance varies by use, configuration, and other factors. See intel.com/performanceindex for details.

INTRODUCING the

Health & Life Sciences AI Suite

Accelerate the next generation of healthcare AI on Intel

Bring multimodal patient monitoring to life

Multi-sensor input for a 360-degree, non-invasive view powered by Intel® Core™ Ultra Processors

Enable real-time insights at the point of care

Deliver context-aware accurate diagnostics, reduce false positives

Enhance data privacy and reliable performance

Keep patient data private and deliver low-latency, real-time health insights



Multi-parameter monitoring simulation



3D human pose estimation



AI ECG arrhythmia detection



Camera-based physiological sensing



Benchmarking tools



Optimized deployment on Intel®





3 Things to Remember

Intel® Core™ Series 2 processor with P-cores

Intel® Core™ Series 2 processor with up to 12 P-cores delivers consistent, reliable, real-time compute at the edge, for flexibility and long-term ROI.

Intel® Core™ Ultra Series 3 processor

Intel Core Ultra Series 3 processor brings integrated AI acceleration and deterministic behavior to VLMs, VLAs, and GenAI with low latency and powerful media processing.

Edge AI Suites for Health and Life Sciences

The new Health & Life Sciences AI Suite shows how real-time, multimodal AI can move directly to the device for faster, more private, more responsive and predictive insights.



Appendix

Compute

Built on

Intel
7

CPU

Up to
12
P-cores

Up to
24
Threads

Up to
5.9
GHz

Base Power

125W

65W

45W

Memory

Up to
5600
DDR5

Up to
192 GB
max capacity

Support for
ECC
memory

I/O

Up to
16x
PCIe Gen 5

Up to
4x
PCIe Gen 4

Up to
4x
4K displays

Technologies

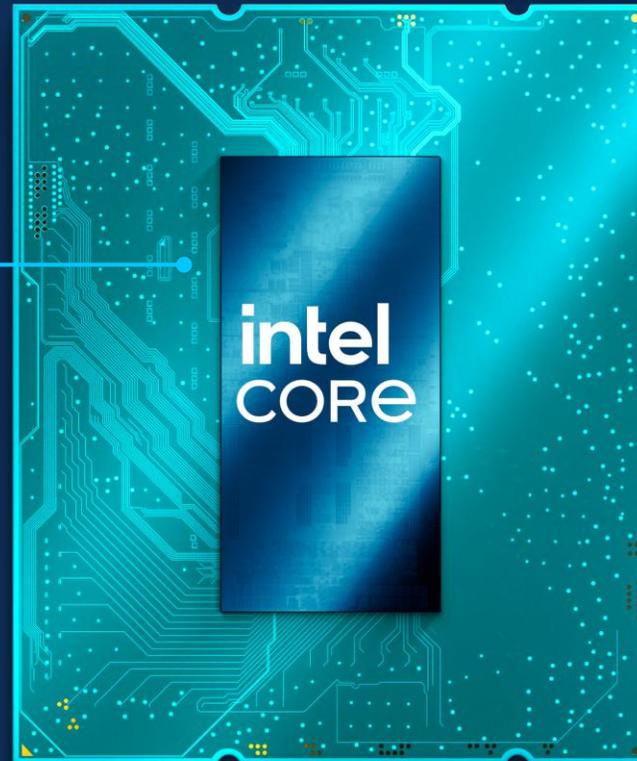
Intel®
TCC

intel
graphics

More...

Intel® Core™ Series 2 processors with P-cores

Die Package Configuration & PCH



Platform Control Hub

I/O

Up to
12x
PCIe Gen 4

Up to
16x
PCIe Gen 3

Up to
8x
DMI Gen 4

Up to
14x
USB 3.2

Connectivity

Discrete
Intel® Wi-Fi 7 &
Bluetooth® 5.4

Integrated
2.5GbE MAC with
TSN support

Intel® Core™ Series 2 Processors with P-cores

Proc #	Brand String	TDP	Functional Core Config	P Core Base Freq (PI)	Cache	P Core Max Turbo Freq	P Core All Core Turbo Freq	Gfx Eus	Max DDR5 Freq	PCIe Config
273PQE	Intel® Core™ 9 273PQE	125	12+0	3.400	36	5.9	5.300	32	5600	1x16+4, 2x8+4
253PQE	Intel® Core™ 7 253PQE	125	10+0	3.500	33	5.7	5.300	32	5600	1x16+4, 2x8+4
223PQE	Intel® Core™ 5 223PQE	125	8+0	4.000	24	5.5	5.300	32	5600	1x16+4, 2x8+4
273PE	Intel® Core™ 9 273PE	65	12+0	2.300	36	5.7	5.200	32	5600	1x16+4, 2x8+4
253PE	Intel® Core™ 7 253PE	65	10+0	2.500	33	5.5	5.100	32	5600	1x16+4, 2x8+4
223PE	Intel® Core™ 5 223PE	65	8+0	2.900	24	5.4	4.800	32	5600	1x16+4, 2x8+4
213PE	Intel® Core™ 5 213PE	65	8+0	2.700	24	5.2	4.600	24	4800	1x16+4, 2x8+4
273PTE	Intel® Core™ 9 273PTE	45	12+0	1.400	36	5.5	4.600	32	5600	1x16+4, 2x8+4
253PTE	Intel® Core™ 7 253PTE	45	10+0	1.800	33	5.4	4.600	32	5600	1x16+4, 2x8+4
223PTE	Intel® Core™ 5 223PTE	45	8+0	2.300	24	5.4	4.800	32	5600	1x16+4, 2x8+4
213PTE	Intel® Core™ 5 213PTE	45	8+0	2.100	24	5.2	4.600	24	4800	1x16+4, 2x8+4

Modern Edge Workload Mastery

Intel® Core™ Ultra X9 388H vs. AMD Ryzen AI 9 and Ryzen AI Max+

Object Detection

YoloV11m | BS=8

Up to
4.8x
Higher Perf

Image Classification

Resnet-50 | BS=8

Up to
8.3x
Higher Perf

- Intel® Core™ Ultra X9 388H
- AMD Ryzen AI 9 HX370
- AMD Ryzen AI Max+ 395

AMD Ryzen AI 9



Details at intel.com/performanceindex. Results may vary.

Modern Edge Workload Mastery

Intel® Core™ Ultra X9 388H vs. Intel® Core™ Ultra 9 285H

- Intel® Core™ Ultra X9 388H
- Intel® Core™ Ultra 9 285H

Object Detection

YoloV11m | BS=8

Up to
9.6x
Higher Perf

288H

NPU

GPU

Image Classification

Resnet-50 | BS=8

Up to
6.8x
Faster

NPU

Time to First Token

Llama-3.2-3B, DeepSeek-Qwen-1.5B,
Phi-4-Mini-3.8B, Mistral-7B v0.1, and
DeepSeek-Llama-8B

Up to
3.8x
Faster

NPU

GPU



Streamlining

healthcare and industrial workflows at the Edge

... orchestrate sophisticated models across the heterogeneous CPU, GPU, and NPU engines, we achieve exceptional performance while maintaining minimal power consumption... ensures our robots deliver the ultra-low latency and strict data privacy...

Vincenzo Latino
R&D Engineering Manager



Oversonic humanoid solution



Speech to text Whisper model on NPU
Text to speech Piper model on CPU
Gemma 3 4B model on iGPU



Software layer underpinning trained model



Hardware layer with built-in AI accelerators



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Performance varies by use, configuration, and other factors. Read the full list of customer testimonials on the Intel Press Portal.



NANOXAI

Empowering *medical imaging at the Edge*

...localization runs up to 152X faster on the NPU, heart segmentation up to 25X faster on the CPU, and calcium segmentation nearly 27X faster on the integrated Arc GPU compared to our cloud setup. This near-real-time performance **accelerates clinical decision-making while keeping sensitive cardiac data securely on-device.**

Sharon Saban
General Manager



Nanox digital imaging solution



Heart segmentation model on CPU
Localization model on NPU
Calcium segmentation model on iGPU



Software layer underpinning trained model



Hardware layer with built-in AI accelerators



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

Performance varies by use, configuration, and other factors. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex). Read the full list of customer testimonials on the Intel Press Portal.



Consolidating

AI workloads for scalable autonomous system

...run the entire AI stack on a single platform...**nearly 2x better energy efficiency and costs about half as much as before**... simplifies our robotics compute design and gives us a much more scalable foundation moving forward.

Kenny Chang
Chief Operating Officer



ASRock robotics solution



TTS on CPU, STT on NPU
Llava 7B VLM on iGPU
GPT-OSS-20B reasoning AI on iGPU



Software layer underpinning trained model



Hardware layer with built-in AI accelerators



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Performance varies by use, configuration, and other factors. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex). Read the full list of customer testimonials on the Intel Press Portal.



Automating

**outdoor-indoor
material flow**

...**20% absolute performance gain**
for the built-in GPU running our
neural networks vs. our current
NVIDIA Orin AGX and even more
importantly a boost in **performance
per watt exceeding 80%**.

Jonas Witt
CTO



Cartken mobile robot solutions



Panoptic Model on iGPU
Segmentation Model on iGPU



Software layer underpinning trained model



Hardware layer with built-in AI accelerators



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Performance varies by use, configuration, and other factors. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex). Read the full list of customer testimonials on the Intel Press Portal.



Enabling Humanoid without dGPU

...3.7X faster TTS, 3.9X higher VLM throughput, and 5.4x faster multi-task vision at just 25W against NVIDIA Jetson AGX Orin... **smoother motion, clearer awareness, and more natural interaction in our humanoids...**

Park Jong Gun
CEO



Circulus humanoid robotics solution



Text to speech on CPU
Gemma 12B on iGPU
Yolo11x-seg vision models on NPU



Software layer underpinning trained model



Hardware layer with built-in AI accelerators



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Performance varies by use, configuration, and other factors. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex). Read the full list of customer testimonials on the Intel Press Portal.

TROSSEN

Making

*robotics research &
development more accessible*

... substantial improvement in policy throughput —
**4.9× faster on our lighter model, 8.3× faster on our
mid-range model, and 8.6× faster on our higher-
load model than Jetson AGX Orin...** stronger
foundation to keep improving the experience without
increasing complexity or power demands...

Luke Schmitt
Lead Software Engineer



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Performance varies by use, configuration, and other factors. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex). Read the full list of customer testimonials on the Intel Press Portal.



Empowering Last-Mile Delivery at the Edge

...achieved nearly 9X faster throughput compared to NVIDIA® Jetson Nano™ ... giving delivery robots the context they need to evaluate higher-stakes decisions, ... enables fast perception with deep reasoning, **making possible, VLM-driven intelligence for last-mile delivery robots.**

Davit Buniatyan
CEO



Autonomous delivery robot by Pinkbot



Perceptron/Isaac 3B vision-language on iGPU



Deep Lake vector database for AI



Intel Extension for PyTorch



Hardware layer with built-in AI accelerators



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Performance varies by use, configuration, and other factors. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex). Read the full list of customer testimonials on the Intel Press Portal.

SENSORY AI™

Empowering *hospitality at the Edge*

...In our benchmarks, Intel delivered nearly 2× faster time-to-first-token while operating at roughly 22 watts versus close to 90 watts on a discrete GPU, enabling responsive, stable Physical AI systems without the cost and complexity of a separate accelerator.

Keith Tan
CEO & Founder



Modular suite of embodied AI robots



Mistral 7B on iGPU



Software layer underpinning trained model



Hardware layer with built-in AI accelerators



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Performance varies by use, configuration, and other factors. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex). Read the full list of customer testimonials on the Intel Press Portal.

Notices and Disclaimers

For notices, disclaimers, and details about certain performance claims, visit www.intel.com/PerformanceIndex or scan the QR code:



© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Notices and Disclaimers

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

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

No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

The image features the Intel logo centered within a dense field of 3D blue cubes. The cubes vary in size and are arranged in a perspective that creates a sense of depth, receding into the background. The lighting is soft, highlighting the edges of the cubes and the logo. The overall color palette is a range of blues, from light to dark.

intel®