Unlocking Next-Gen Edge Computing

Intel® Core™ Ultra Processors (Series 3) for the Edge (Codenamed Panther Lake)

June Lim

Dec 2025

The Processor of Choice

for Real-World Edge Success

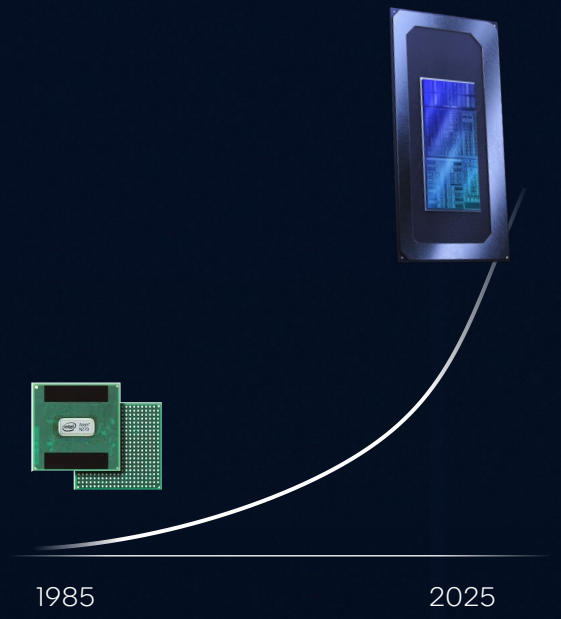

40+ years of edge and embedded mastery

200M+

In 10 years

x86 processors for edge sold

4,000+ edge ecosystem partners

intel

Delivering the Next Frontier of Edge Innovation

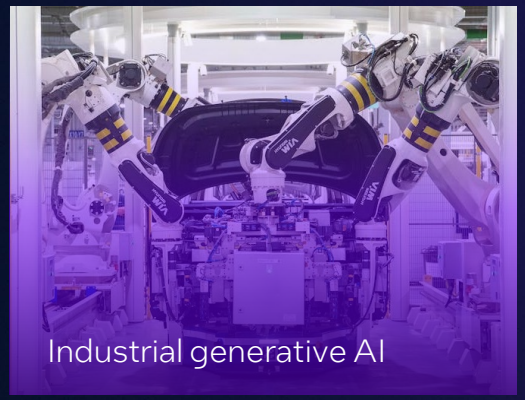

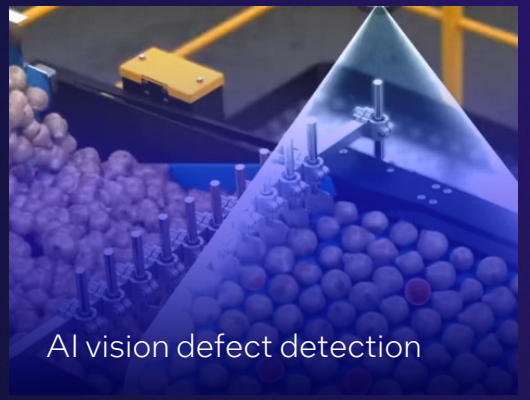

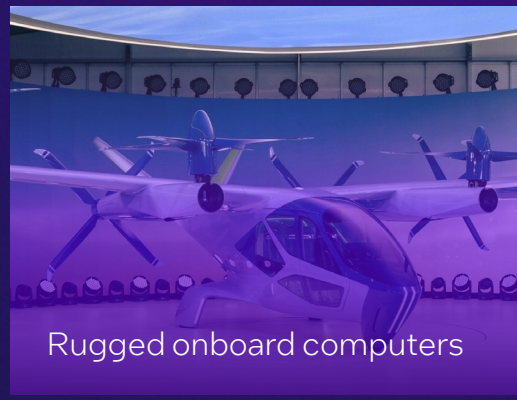

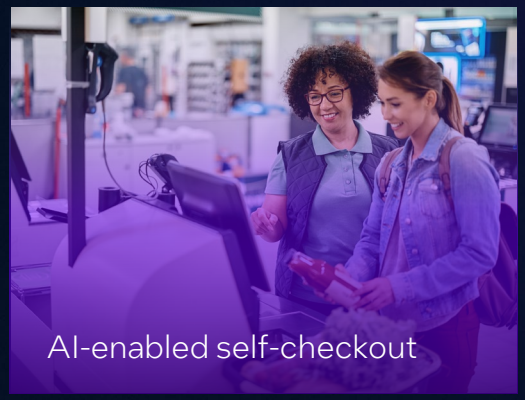

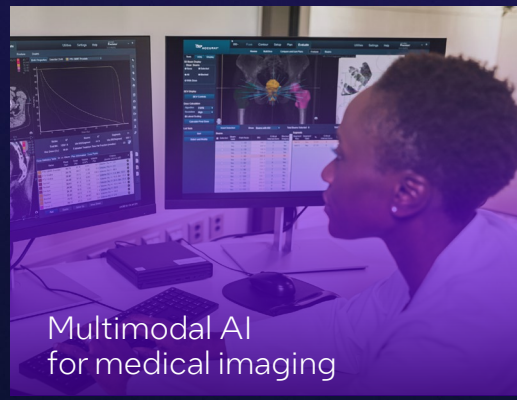

intel

Panther Lake

Proven Ideal for Mission-Critical Edge and Physical AI

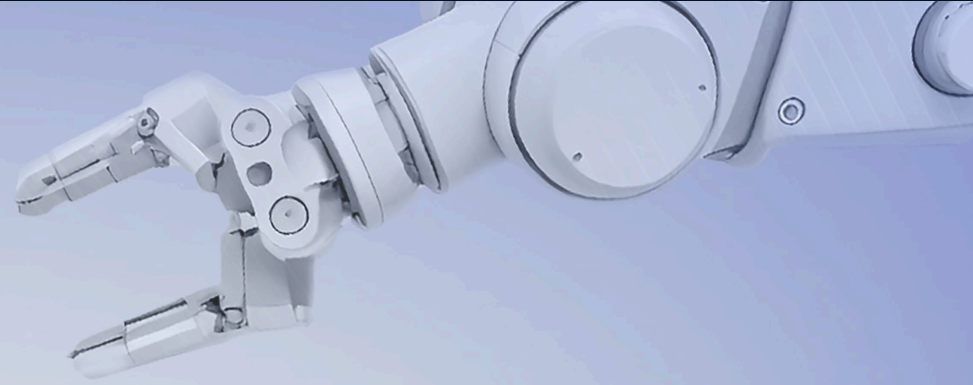

Power-efficient AI and graphics

Up to 1.8x more TOPS vs. Arrow Lake-H

Up to 50% more GPU cores vs. Arrow Lake-H

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

Benefits of Intel® Core™ Ultra Processors (Series 3)

(Codenamed Panther Lake)

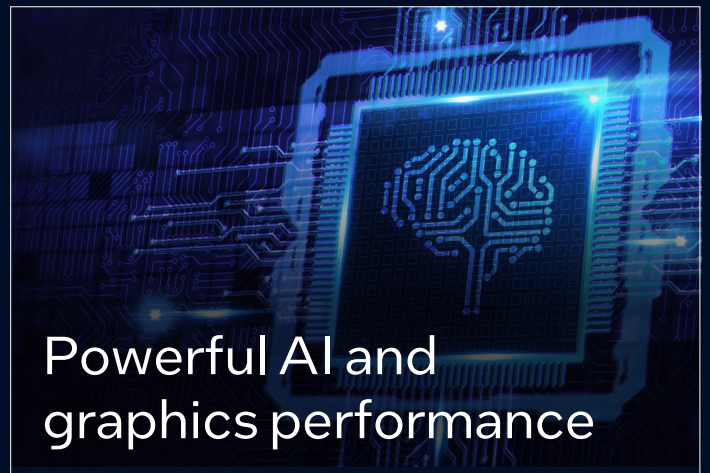

Powerful AI and graphics performance

- Up to 180 total platform TOPS[1], enable/accelerate AI inferencing with built-in NPU & GPU cost-effectively without discrete accelerator.
- Up to 50% more GPU cores (12 Xe cores) than ARL for AI-intensive workloads at the edge.
- Access up to 64GB of LPDDR5 8533 or up to 128GB of DDR5 7200 memory—an advantage over dGPUs with fixed VRAM, ideal for large AI workloads.

Power-efficient compute and graphics

- As low as 15W for fanless designs and up to 65W.
- Built-in GPU reduces power consumption, lower BOM costs and enables smaller form factor design.
- Faster connectivity with Thunderbolt™ 4 and integrated Wi-Fi 7 R2

Industrial reliability and flexibility

- Maximize reliability with industrial-grade features such as support for Functional Safety, Intel® TCC, and extended temperature support (-40°C to +100°C).
- Ensure long-term durability through select SKUs designed for 100% operation over 10 years in industrial conditions.

1. Select H-SKUs of Intel Core Ultra processors (Series 3) can achieve up to 180 total platform TOPS. Results may vary.
2. Intel® Arc™ GPU only available on select H-SKUs, Intel® Core™ Ultra processor powered systems with at least 16GB of system memory in a dual-channel configuration. OEM enablement required; check with OEM for system configuration details.

intel

What Truly Matters for Edge AI

Intel Vs Nvidia

Intel® Core™ Ultra X9 processor 388H (180 TOPS) VS

NVIDIA® Jetson AGX Orin 64GB (275 TOPS)

Up to

1.9x

Higher LLM performance

Up to

1.7x

Higher image classification performance

Up to

2.3x

Better performance per watt per $ on end-to-end video analytics

Up to

4.5x

Higher throughput on vision language action models

For more complete information about performance and benchmark results, visit intel.com/PerformanceIndex.

TCO: Infer + Fine Tune on a Single SoC

- Fine-tune at ~87% of a discrete GPU performance — for only 17% of the cost, plus better inference.
- More savings in energy cost

5.8X or

$5,549

Estimate savings if you move away from dual-systems to Intel Core Ultra processor (Series 3)

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

Leadership Performance

Across ALL AI Engines

Geekbench AI 1.6 Scores

Series 3 intel CORE ULTRA X9

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	DNR
	FP16	3.1x	DNR
	FP32	2.3x	DNR
NPU	INT8	DNR	1.3x
	FP16	DNR	2.2x
	FP32	DNR	2.6x

See intel.com/performanceindex for details. Results may vary.

Object Recognition and Image Classification Performance vs AMD

YoloV11m and Resnet-50 (BS8) vs competitions

Intel® Core™ Ultra X9 processor 388H
(codenamed Panther Lake H)

VS

AMD Ryzen AI 9 HX370

Up to
4.8x
Higher GPU performance in object detection

Average
8.3x
Higher GPU performance in image classification

Intel® Core™ Ultra X9 processor 388H
(codenamed Panther Lake H)

VS

AMD Ryzen AI Max+ 395

Up to
2x
Higher GPU performance in object detection

Average
2.2x
Higher GPU performance in image classification

For more complete information about performance and benchmark results, visit intel.com/PerformanceIndex.

Gen over Gen Performance Improvements

Intel® Core™ Ultra X9 processor 388H iGPU
(codenamed Panther Lake H)

VS

Intel® Core™ Ultra 9 processor 285H iGPU
(codenamed Arrow Lake H)

Up to
1.8x
Higher performance in object detection

Average
1.6x
Faster 1st token latency on GenAI

Up to
2.7x
Faster 1st token latency on Multimodal AI

Intel® Core™ Ultra X9 processor 388H NPU
(codenamed Panther Lake H)

VS

Intel® Core™ Ultra 9 processor 285H NPU
(codenamed Arrow Lake H)

Up to
9.6x
Higher performance in object detection

Average
3.8x
Faster 1st token latency on GenAI

Up to
6.8x
Higher performance in image classification

For more complete information about performance and benchmark results, visit intel.com/PerformanceIndex.

Customer Testimonials

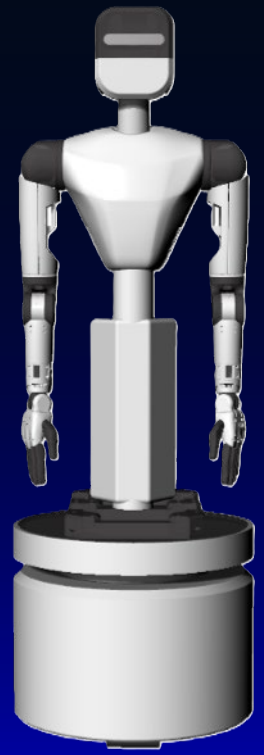

circulus

"...3.7X faster TTS, 3.9X higher VLM throughput, and 5.4× 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

Baidu 百度

"...1.4X lower latency and 2.7X higher power efficiency compared to an NVIDIA RTX 4060 GPU... enables real-time document parsing and text recognition on compact edge devices, helping businesses automate workflows..."

Alex Zhang
Senior Product Manager

TROSSEN

"...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... sustained a 30 Hz control rate in the LeRobot data-collection pipeline and maintained 30 Hz and 200 Hz control rates...

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.

Customer Testimonials

activeloop

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

Davit Buniatyan
CEO

NANOXAI

"...2.5× faster overall... 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. This near–real-time performance accelerates clinical decision-making while keeping sensitive cardiac data securely on-device."

Sharon Saban
General Manager

ASRock Industrial

"...run the entire AI stack on a single platform...nearly 2x the 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

- 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.

Customer Testimonials

SENSORY Ai™

"...delivered nearly 2X faster time-to-first-token while operating at roughly 22 watts versus close to 90 watts on a discrete RTX 5080 GPU, enabling responsive, stable physical AI systems without the cost and complexity of a separate accelerator."

Keith Tan
CEO

CARTKEN

"...shows great promise in this regard, as we saw a 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/watt exceeding 80%... translates to more powerful neural networks & algorithms that allow our robots to act smarter..."

Jonas Witt
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. For more complete information about performance and benchmark results, visit intel.com/PerformanceIndex.

Proven TCO for Edge AI

Driven by a more powerful and integrated GPU and NPU

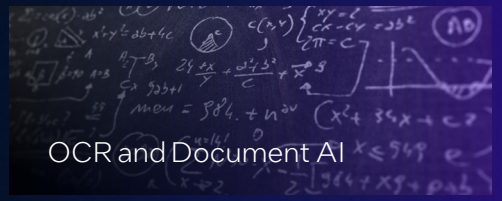

Displaced NVIDIA RTX 4060 in text recognition and document parsing

Up to 39% or $745
savings per system over 5 years

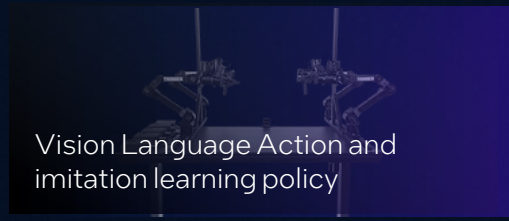

Displaced NVIDIA Jetson AGX Orin 64GB in robotics AI workloads

Up to 42% or $849
savings per system

Displaced NVIDIA Jetson Thor in humanoid AI workloads

Up to 67% or $2,447
savings per system over 5 years

Performance varies by use, configuration, and other factors. For more complete information about performance and benchmark results, visit intel.com/PerformanceIndex.

Technical Advancements Over Previous Generations

Processor Family	Intel® Core™ Ultra processors (Series 1) MTL H	Intel® Core™ Ultra Processors (Series 2) ARL H	Intel® Core™ Ultra Processors (Series 3) PTL
Core	Up to 16 cores (6P+8E+2LPE)	Up to 16 cores (6P+8E+2LPE)	Up to 16 cores (4P+8E+4LPE)
Integrated AI	Up to 34 total platform TOPS Intel® Arc™ GPU (up to 18 TOPS) Integrated NPU (up to 11 TOPS) Intel Deep Learning Boost	Up to 99 total platform TOPS[1] Intel® Arc™ GPU w/ Intel® XMX (up to 77 TOPS) Integrated NPU (up to 13 TOPS) Intel Deep Learning Boost	~180 total platform TOPS Intel® Arc™ GPU w/ Intel® XMX (up to 120 TOPS) Integrated NPU (up to 50 TOPS) Intel® Deep Learning Boost
Graphics	Intel® Arc™ GPU with up to 8 Xᵉ cores	Intel® Arc™ GPU with up to 8 2nd Gen Xᵉ cores	Intel® Arc™ GPU with up to 12 3rd Gen Xᵉ cores Intel® Graphics with up to 4 3rd Gen Xᵉ cores
Memory	LPDDR5 6400 MT/s LPDDR5x 7467 MT/s (T4 board) DDR5 5600 MT/s	LPDDR5x 8400 MT/s DDR5 6400 MT/s	Up to LPDDR5x (CAMM & Down) 8533 MT/s (T3 board) or 9600 MT/s (T4 board)[3] Up to DDR5-7200[3] MT/s
I/O Connectivity	Up to 8x PCIe 5 and 20x PCIe 4 4x Integrated Thunderbolt 4 Integrated Wi-Fi 6E, Bluetooth 5.3	Up to 8x PCIe 5 and 20x PCIe 4 4x Integrated Thunderbolt 4 Discrete Thunderbolt 5 Integrated Wi-Fi 7, Bluetooth 5.4	Up to 12x PCIe 5 and 8x PCIe 4 4x Integrated Thunderbolt™ 4 Discrete Thunderbolt™ 5 Integrated Wi-Fi 7 R2, Bluetooth 6
Embedded/Indu use conditions	Yes, via PTR	Yes, via PTR	Yes, via PTR FuSa[3], Extended Temp[3], TCC, IBECC

1. Select H-SKUs of Intel Core Ultra processors (Series 3) can achieve up to 180 total platform TOPS. Results may vary.
2. Intel® Arc™ GPU only available on select H-SKUs, Intel® Core™ Ultra processor powered systems with at least 16GB of system memory in a dual-channel configuration. OEM enablement required; check with OEM for system configuration details.
3. Select SKU

intel

High Level SKUs Comparison

Processor family	Intel Core™ Ultra processors (Series 3) (15W-45W), codenamed Panther Lake	Intel Core™ Ultra processors (Series 3) (15W-65W), codenamed Panther Lake	Intel Core™ Ultra processors (Series 3) (15W-65W), codenamed Panther Lake
Core	Up to 8 cores (4P+0E+4LPE)	Up to 16 cores (4P+8E+4LPE)	Up to 16 cores (4P+8E+4LPE)
Integrated AI	~100 total platform TOPS: NPU (50 TOPS) Intel® Graphics, up to 4 Xᵉ cores (40 TOPS) Intel® XMX, Intel® DL Boost	~100 total platform TOPS: NPU (50 TOPS) Intel® Graphics, up to 4 Xᵉ cores (40 TOPS) Intel® XMX, Intel® DL Boost	~180 total platform TOPS: NPU (50 TOPS) Intel® Arc™ GPU,[1] up to 12 Xᵉ cores (120 TOPS) Intel® XMX, Intel® DL Boost
Graphics	Intel® Graphics with up to 4 Xᵉ cores	Intel® Graphics with up to 4 Xᵉ cores	Intel® Arc GPU[1] with up to 12 Xᵉ cores
I/O capacity	Up to 12 PCIe lanes (x4 Gen 5, x8 Gen 4) 4x Integrated Thunderbolt 4 Discrete Thunderbolt™ 5 Integrated Intel® Wi-Fi 7 R2, Bluetooth 6	Up to 20 PCIe lanes (x12 Gen 5, x8 Gen 4) 4x Integrated Thunderbolt 4 Discrete Thunderbolt 5 Integrated Intel® Wi-Fi 7 R2, Bluetooth 6	Up to 12 PCIe lanes (x4 Gen 5, x8 Gen 4) 4x Integrated Thunderbolt 4 Discrete Thunderbolt 5 Integrated Intel® Wi-Fi 7 R2, Bluetooth 6
Memory	LPDDR5x up to 7467 MT/s; 96GB max capacity DRR5 up to 6400 MT/s; 128GB max capacity	LPDDR5x up to 8533 MT/s (T3 board), up to 9600 MT/s (T4 board); 96GB max capacity DDR5 up to 7200 MT/s; 128GB max capacity	LPDDR5x up to 8533 MT/s (T3 board), up to 9600 MT/s (T4 board); 96GB max capacity DDR5 up to 7200 MT/s; 128GB max capacity
SKU	Intel Core™ Ultra Processor 5, 7, and 9	Intel Core™ Ultra Processor 5, 7, and 9	Intel Core™ Ultra Processor 5, 7, and 9
Embedded/Indu use conditions and availability[2]	Yes, via PTR Industrial SKUs with extended temp Up to 10 years	Yes, via PTR Industrial SKUs with extended temp Up to 10 years	Yes, via PTR Up to 10 years

Note: Platform benchmarks, benefits, and features will vary by SKU. Not all features are available on every SKU. Consult the product lineup for additional details.

1. Built-in Intel® Arc™ GPU [integrated Intel® Arc™ graphics] only available on select Intel Core Ultra processor (Series 3) -powered systems. Other system configurations feature Intel® Graphics.
2. Intel does not commit or guarantee product availability or software support by way of road map guidance. Intel reserves the right to change road maps or discontinue products, software, and software support services through standard EOL/PDN processes. Contact your Intel account rep for additional information.

Target Segments

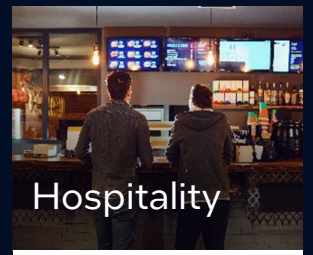

Key Features

Extended Temperature (-40°C to 100°C)
Industrial Use Conditions
FuSa, IBECC
TCC, TSN

4 concurrent 6K displays
SW Genlock, Pipelock
Bezel Correction
EDID

Extended Temperature Industrial Use Conditions

AI-Capable: Inferencing with CPU, NPU and iGPU with Intel® XMX

Up to 16 Cores, 12 Xᵉ-Cores, 20 PCIe lanes, **96**GB LPDDR5 or 128GB DDR memory

Up to 10 Years Availability

intel

Benefits for Industrial & Energy

Optimized Performance for Automation/Robotics

With a power as low as 15W and BGA packaging, this solution enables fanless designs in space-constrained environments. The hybrid architecture provides the power and efficiency needed for real-time processing and autonomous navigation in mobile robotics.

Enhanced AI for Industrial Co-Pilot

With next gen GPU and NPU, these processors support industrial AI workloads by facilitating advanced machine learning and data analysis, improving human-machine collaboration and operational safety.

Precision in Computer Vision

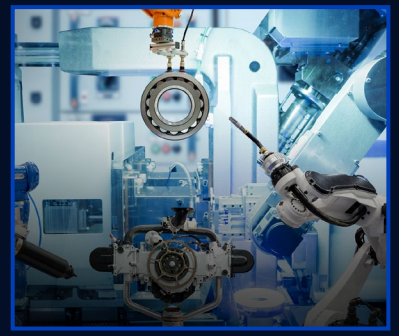

The processors' robust performance and AI capabilities enhance classic computer vision applications, such as defect detection, by enabling high-speed image processing and accurate anomaly identification.

Built-In Safety Features

Integrated functional safety technologies ensure reliable operation and compliance with stringent industrial standards, providing critical safeguards for AI-powered autonomous systems.

Reliability and Long-Term Support

- Withstand harsh conditions with extended temperature and dynamic temperature range (DTR) up to 140°C on select SKUs.
- Up to 10-year availability, s reduce the need for frequent upgrades.

Benefits for Retail & Hospitality

Enhanced AI & Media

Implement AI-driven self-checkout systems, loss prevention solutions, and interactive smart kiosks powered by the latest Intel Core Ultra processor, featuring an integrated NPU and Intel® ARC™ GPU, to efficiently manage multiple camera streams.

Powerful Built-In GPU

Deploy Smart signage solutions including complete video walls with up to four concurrent 6K displays, Genlock and Pipelock synchronization, extended display identification (EDID), and bezel correction.

Efficient Performance

Help reduce infrastructure costs with built-in Intel® Arc™ Graphics that allow you to skip discrete GPU.

Improve Customer Service with AI Avatars

Robust GenAI capabilities support AI-enabled avatar customer service applications, delivering personalized and efficient customer interactions that enhance satisfaction and loyalty.

Benefits for Smart Cities

Advanced Video Analytics for Safety & Security

With next gen GPU and NPU, these processors provide powerful AI capabilities for security surveillance and video analytics applications, enabling real-time monitoring and analysis to enhance public safety and security in urban environments.

Efficient Performance for Smart Transportation

The hybrid architecture ensures efficient and responsive operation of in-vehicle computers, supporting real-time data processing for navigation, fleet management, and passenger information systems.

Industrial Use Conditions and Extended Temperature Support

Designed to withstand harsh environments, these processors offer extended temperature support and industrial-grade reliability, making them ideal for in-vehicle computing in diverse and challenging conditions.

Reliability and Long-Term Support

Up to 10 years of availability powers a stable baseline throughout your product's lifecycle, reducing total cost of ownership.

Benefits for Healthcare

Powerful Built-In GPU

Removes the need for a separate accelerator, shrinking device size, power draw, and validation effort—so new imaging appliances reach the ward sooner and cost less to maintain.

On-Device Multimodal AI

Sub-second AI summaries and preliminary reads keep patient data in-house, cut cloud latency, and let clinicians act while the patient is still present.

Quad 4K Display Support

Radiologists view original scans, AI overlays, priors, and reports side-by-side—reducing reading time and improving diagnostic confidence.

10-Year Product Availability

Aligns with long certification cycles, protecting capital investments and avoiding mid-life hardware re-qualification costs.

Introducing the **Robotics Reference Board**

Faster time to market

Shorten onboarding with optimized reference solution

Lower research and development cost

Scalable & ready-to-use

Modular design for scalability & reuse

Built in EtherCAT, CAN-HD, GMSL and USB extends functionality

Open ecosystem collaboration

Driving industry form-factor standards

Broad collaboration with top industry partners

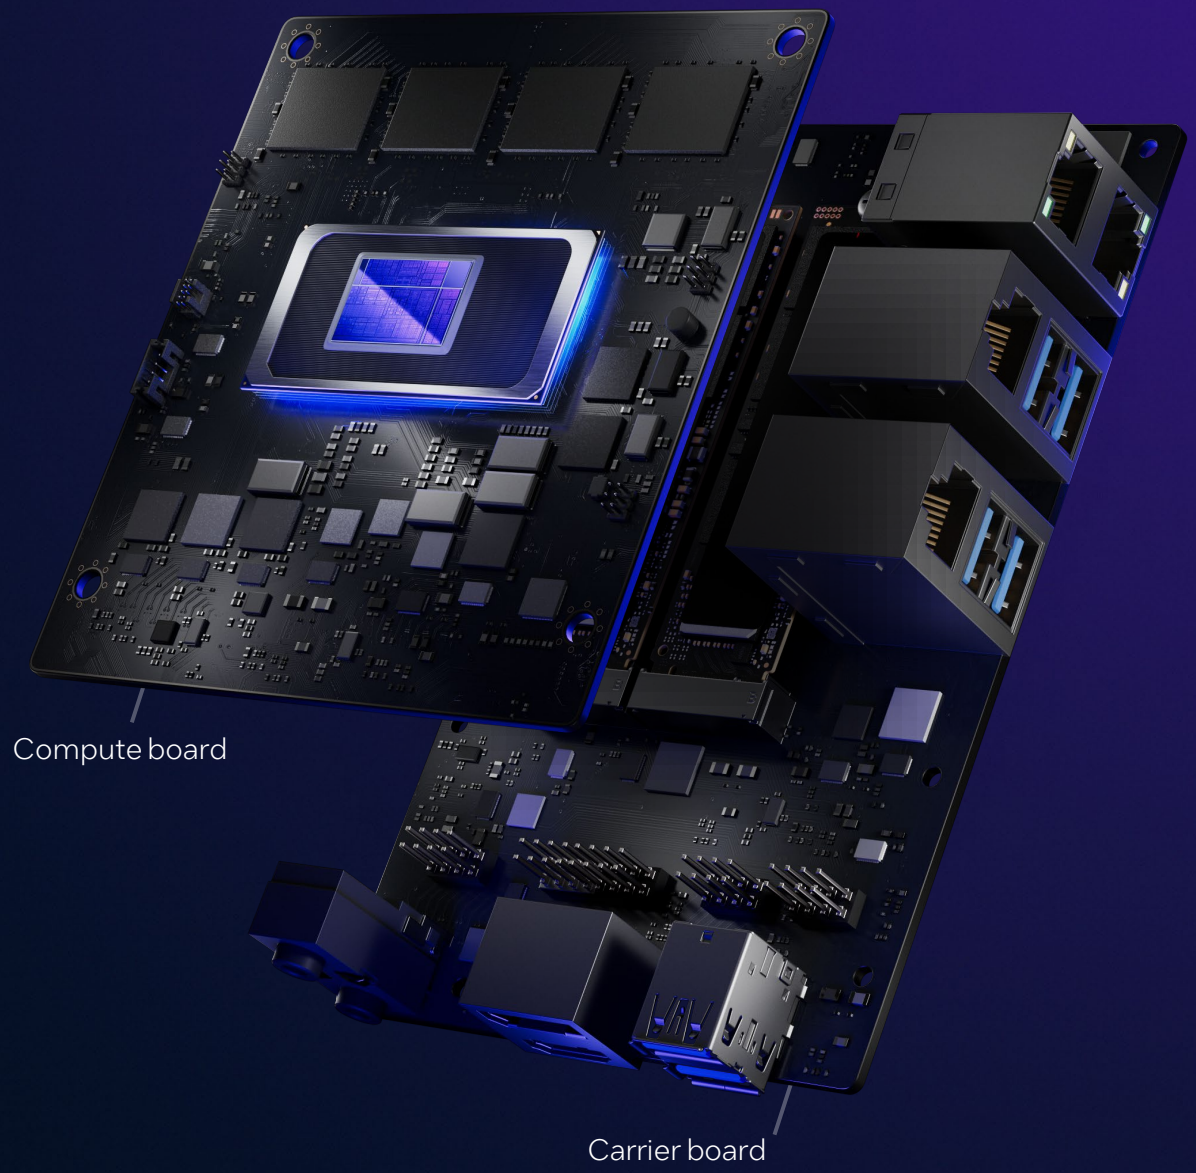

intel

Introducing the **Robotics AI Suite**

Adapt and scale physical AI on Intel® for fast time to value

Accelerate development for all kinds of robotics

Industrial, service, mobile and humanoid robotics
Locomotion, perception, and imitation learning
Run vision AI, media analytics, and gen AI workloads with real-time controls

Innovate and scale fast with an open ecosystem

Validated systems, SDK, and vision AI use cases for accelerate time to value
Near zero-touch deployment and remote device management
Rely on Intel's proven long-standing robotics expertise

Reduce complexity, improve TCO

Latest generation processors for best possible performance
Integrated CPUs, GPUs, NPUs for robotics performance across form factors
Drive real-time controls & AI/perception from one platform

Robotic AI Suite content

- Vision libraries
- Real-time control frameworks
- AI inference engines
- Orchestration-ready modules
- Hardware-aware tuning
- ROS 2 & open standards

intel

Additional Resources

CNDA

- Panther Lake Platform Gold Deck

Public

- N/A

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 backup for configuration details. No product or component can be absolutely secure.

Intel® processors of the same SKU may vary in frequency or power as a result of natural variability in the production process.

All product plans and road maps are subject to change without notice.

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 SEC filings at intc.com.

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and are not intended to function as trademarks.

Not all features are available on all SKUs.

Not all features are supported in every operating system.

Intel may change availability of products and support at any time without notice. All product plans are subject to change without notice.

Your costs and results may vary.

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

Intel Statement on Product Usage: Intel is committed to respecting human rights and avoiding causing or contributing to adverse impacts on human rights. See Intel's Global Human Rights Principles. Intel's products and software are intended only to be used in applications that do not cause or contribute to adverse impacts on human rights.


intel
TECH
tour.US
Thank
you