

Intel® Core™ Ultra 200H & 200U Series Processors for the Edge Overview (Codenamed Arrow Lake H/U)

Jan 2025


intel.

Notices & 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. No product or component can be absolutely secure.
- Results that are based on 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.
- Your costs and results may vary.
- Intel technologies may require enabled hardware, software or service activation.
- Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.
- Intel is committed to respecting human rights and avoiding complicity in human rights abuses. 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 a violation of an internationally recognized human right.
- © 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.



30-Second
**Product
Overview**



Intel® Core™ Ultra 200H & 200U Series Processors Overview (Codenamed Arrow Lake H/U)

Overview

- Successor of Intel® Core™ Ultra processors Series 1
- Next gen Cores plus new GPU architecture
- Faster memory, Thunderbolt™ 5 and integrated Wi-Fi 7 / BT 5.4
- Compatible with boards designed for Meteor Lake
- Up to 5 years availability¹



Value Proposition

Deploy visual inferencing, rich media, and emerging generative AI use cases with greater ease, without the need of entry-level discrete graphics at the edge. With up to **99 total platform TOPS**² Intel® Core™ Ultra processors (Series 2) can scale performance and take AI acceleration to a new level to supercharge edge computing in the era of AI. Accelerate real-time data processing and minimize latency with new CPU cores, a built-in, next-generation Intel® Arc™ GPU³ and a neural processing unit (NPU)—all in a BGA package.

Intel® Core™ Ultra 7 processor 265H vs NVIDIA® Jetson AGX Orin 64GB

- Up to **5.8x** Faster in media performance⁴
- Up to **3.4x** Faster in video analytics end-to-end workload (Media + AI inference) performance⁴
- Up to **8.2x** Better performance per watt per \$⁴




Optimized for Edge

- EDID Correction
- Bezel Compensation
- Pipelock
- HDMI Capture
- Single root I/O virtualization
- Windows 11 IoT Enterprise LTSC 2024
- Ubuntu, Red Hat Enterprise, Wind River, Long-term support (LTS) Linux kernels
- KVM hypervisor, RTS Hypervisor
- Intel® Slim Bootloader, UEFI BIOS

1. 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.
2. Select Intel Core Ultra 200H Series Processors can achieve up to 99 total platform TOPS. Results may vary.
3. Intel® Arc™ GPU only available on 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.
4. Results may vary. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex)



3-Minute
Product
Overview



Intel® Core™ Ultra 200H & 200U Series Processors Overview (Codenamed Arrow Lake H/U)

Overview

- Successor of Intel® Core™ Ultra processors Series 1
- Next gen Cores plus new GPU architecture
- Faster memory, Thunderbolt™ 5 and integrated Wi-Fi 7 / BT 5.4
- Compatible with boards designed for Meteor Lake
- Up to 5 years availability¹



Value Proposition

Deploy visual inferencing, rich media, and emerging generative AI use cases with greater ease, without the need of entry-level discrete graphics at the edge. With up to **99 total platform TOPS**² Intel® Core™ Ultra processors (Series 2) can scale performance and take AI acceleration to a new level to supercharge edge computing in the era of AI. Accelerate real-time data processing and minimize latency with new CPU cores, a built-in, next-generation Intel® Arc™ GPU³ and a neural processing unit (NPU)—all in a BGA package.

Intel® Core™ Ultra 7 processor 265H vs NVIDIA® Jetson AGX Orin 64GB

- Up to **5.8x** Faster in media performance⁴
- Up to **3.4x** Faster in video analytics end-to-end workload (Media + AI inference) performance⁴
- Up to **8.2x** Better performance per watt per \$⁴




Optimized for Edge

- EDID Correction
- Bezel Compensation
- Pipelock
- HDMI Capture
- Single root I/O virtualization
- Windows 11 IoT Enterprise LTSC 2024
- Ubuntu, Red Hat Enterprise, Wind River, Long-term support (LTS) Linux kernels
- KVM hypervisor, RTS Hypervisor
- Intel® Slim Bootloader, UEFI BIOS

1. 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.
2. Select Intel Core Ultra 200H Series Processors can achieve up to 99 total platform TOPS. Results may vary.
3. Intel® Arc™ GPU only available on 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.
4. Results may vary. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex)

Benefits of Intel® Core™ Ultra 200H & 200U Series Processors (Codenamed Arrow Lake H/U)



- Up to 99 total platform TOPS¹, enable/accelerate AI inferencing cost-effectively without discrete accelerator.
- Multiple integrated compute engines for AI — P-cores, E-cores, Intel® Arc™ GPU² with Intel® XM² and Intel® AI Boost, a built-in neural processing unit (NPU) for increased edge AI capabilities at low power.



- Up to 8 Xe-cores for graphics / media-intensive workloads at the edge.
- Built-in GPU reduces power consumption, lower BOM costs and enables smaller form factor design.
- Faster connectivity with Thunderbolt™ 5 and integrated Wi-Fi 7




- As low as 12W in TDP option for fanless design.
- Full performance with 65W
- Compatible with boards designed for MTL

1. Select Intel® Core Ultra 200H Series of Intel Core Ultra processors can achieve up to 99 total platform TOPS. Results may vary.

2. Intel® Arc™ GPU only available on H-SKU, 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.



30-Minute
Product
Overview





Agenda

■ The Edge is the Center of Gravity for AI Inference	10
■ Edge AI Design Challenges	11
■ Edge Workload Examples	12
■ Intel® Core™ Ultra 200H & 200U Series Processors Overview	13
■ Benefits of Intel® Core™ Ultra 200H & 200U Series Processors	14
■ What Truly Matters for Edge AI (Intel Vs Nvidia)	15
■ Intel® Core™ Ultra 9 200H Series Gen over Gen Performance Improvements	16
■ Intel® Core™ Ultra 7 200H Series Gen over Gen Performance Improvements	17
■ Intel® Core™ Ultra 7 200U Series Gen over Gen Performance Improvements	18
■ Technical Advancements Over Previous Generations (H SKU)	19
■ Technical Advancements Over Previous Generations (U SKU)	20
■ Segment Applications and Feature Mapping	21
■ Drive TCO/Sustainability Advantages	22
■ Featured Customer Use Cases	23
■ Additional resources	29

The Edge is the Center of Gravity for AI Inference

Running alongside compute and media

We've been at the edge for years, working with early adopters to digitize operations

Top drivers: data security, operational efficiency, business resilience¹



Increasing use of AI tightly interconnected with existing applications to drive better business outcomes

SAMSUNG MEDISON

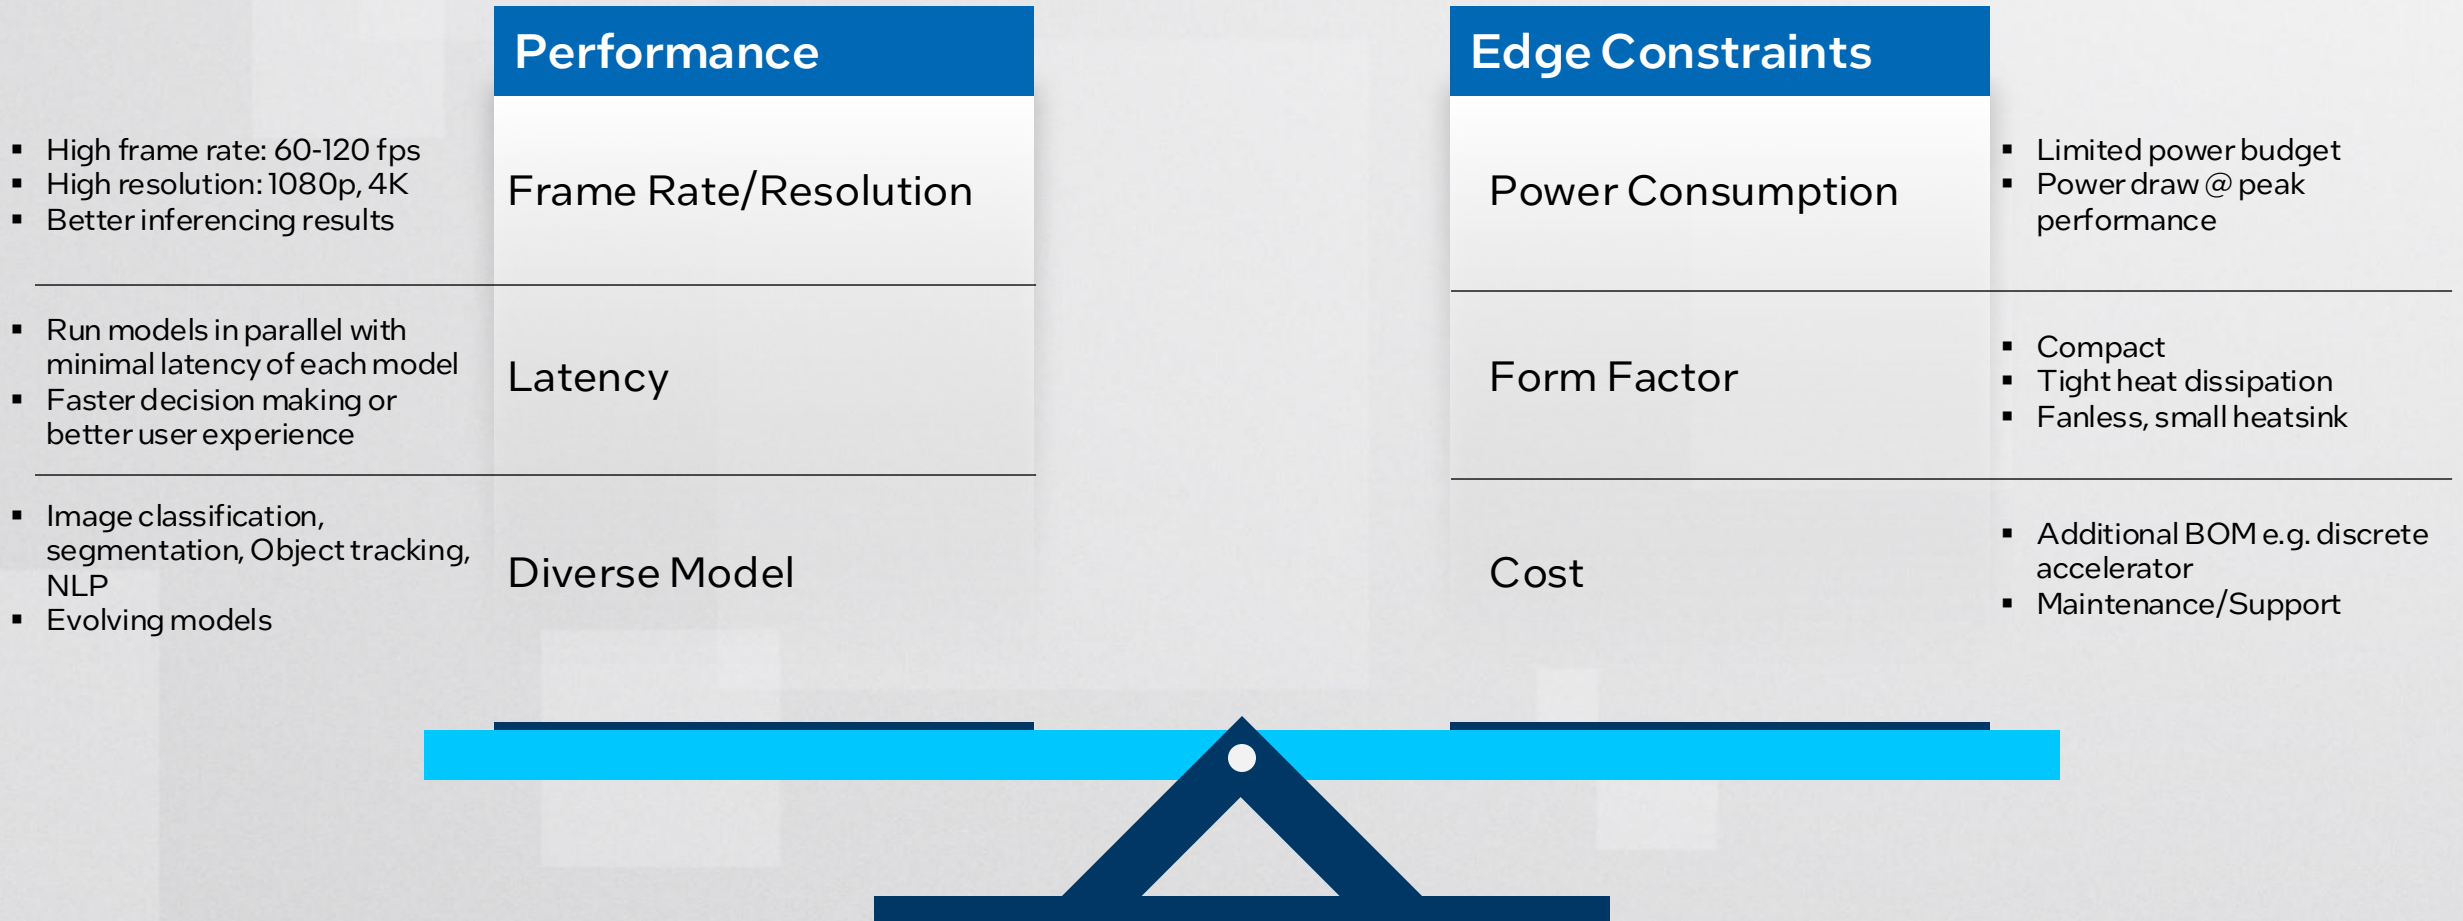
Edge AI requires unique solutions built from expertise

1. Source: IDC InfoBrief, sponsored by Intel, Breaking Boundaries: Edge-Native Infrastructure Powers AI Advancements, doc #US52123724, June 2024

According to Gartner®: By 2026, at least **50%** of edge computing deployments will involve ML.²

2. Gartner®, Hyperscalers Stretching to the Digital Edge, By Thomas Bittman, 24 July 2023. GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All right reserved. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Edge AI Design Challenges



Effective edge AI solution needs to satisfy power, performance, and cost.

Edge Workloads Examples



Data Collection and Preprocessing

Edge devices often collect and preprocess raw data from sensors and **other** sources before transmitting it to the cloud. This can involve tasks such as data filtering, noise reduction, and data aggregation.




Image and Video Processing

Edge devices in surveillance, industrial monitoring, and IoT cameras often process images and videos locally. This can include tasks like object detection, tracking, facial recognition, and image compression.



Sensor Data Analysis

Edge hardware can analyze data from various sensors, such as temperature sensors, accelerometers, and gyroscopes. These analyses might involve anomaly detection, pattern recognition, or predicting system failures.



Real-Time Analytics

Edge devices might perform real-time analytics on incoming data to extract insights and trigger immediate actions. Examples include monitoring machinery for predictive maintenance or analyzing customer behavior in retail settings.



Natural Language Processing (NLP)

Edge devices with speech recognition and NLP capabilities can enable voice-controlled interfaces and language processing without relying on cloud services. This is common in smart signage/kiosk and voice assistants.




Local AI Inference

AI models are often deployed on edge devices for tasks like image recognition, natural language understanding, and sentiment analysis. These models make predictions locally, reducing latency and dependence on cloud resources.



Autonomous Systems

Edge devices in robotics and autonomous vehicles process sensor data to make decisions and navigate in dynamic environments.



Security and Surveillance

Edge devices can process video feeds for intrusion detection, access control, and identifying security threats.

Intel® Core™ Ultra 200H & 200U Series Processors Overview (Codenamed Arrow Lake H/U)

Overview

- Successor of Intel® Core™ Ultra processors Series 1
- Next gen Cores plus new GPU architecture
- Faster memory, Thunderbolt™ 5 and integrated Wi-Fi 7 / BT 5.4
- Compatible with boards designed for Meteor Lake
- Up to 5 years availability¹



Value Proposition

Deploy visual inferencing, rich media, and emerging generative AI use cases with greater ease, without the need of entry-level discrete graphics at the edge. With up to **99 total platform TOPS**² Intel® Core™ Ultra processors (Series 2) can scale performance and take AI acceleration to a new level to supercharge edge computing in the era of AI. Accelerate real-time data processing and minimize latency with new CPU cores, a built-in, next-generation Intel® Arc™ GPU³ and a neural processing unit (NPU)—all in a BGA package.

Intel® Core™ Ultra 7 processor 265H vs NVIDIA® Jetson AGX Orin 64GB

- Up to **5.8x** Faster in media performance⁴
- Up to **3.4x** Faster in video analytics end-to-end workload (Media + AI inference) performance⁴
- Up to **8.2x** Better performance per watt per \$⁴




Optimized for Edge

- EDID Correction
- Bezel Compensation
- Pipelock
- HDMI Capture
- Single root I/O virtualization
- Windows 11 IoT Enterprise LTSC 2024
- Ubuntu, Red Hat Enterprise, Wind River, Long-term support (LTS) Linux kernels
- KVM hypervisor, RTS Hypervisor
- Intel® Slim Bootloader, UEFI BIOS

1. 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.
2. Select Intel Core Ultra 200H Series Processors can achieve up to 99 total platform TOPS. Results may vary.
3. Intel® Arc™ GPU only available on 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.
4. Results may vary. For more complete information about performance and benchmark results, visit [intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex)

Benefits of Intel® Core™ Ultra 200H & 200U Series Processors (Codenamed Arrow Lake H/U)



- Up to 99 total platform TOPS¹, enable/accelerate AI inferencing cost-effectively without discrete accelerator.
- Multiple integrated compute engines for AI — P-cores, E-cores, Intel® Arc™ GPU² with Intel® XM² and Intel® AI Boost, a built-in neural processing unit (NPU) for increased edge AI capabilities at low power.



- Up to 8 Xe-cores for graphics / media-intensive workloads at the edge.
- Built-in GPU reduces power consumption, lower BOM costs and enables smaller form factor design.
- Faster connectivity with Thunderbolt™ 5 and integrated Wi-Fi 7



- As low as 12W in TDP option for fanless design.
- Full performance with 65W
- Compatible with boards designed for MTL

1. Select Intel Core Ultra 200H Series of Intel Core Ultra processors can achieve up to 99 total platform TOPS. Results may vary.

2. Intel® Arc™ GPU only available on H-SKU, 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.

What Truly Matters for Edge AI (Intel Vs Nvidia)

Intel® Core™ Ultra 7 processor 265H (97 TOPS) VS NVIDIA® Jetson AGX Orin 64GB (275 TOPS)



Up to
5.8x
Faster in media
performance

Up to
3.4x
Faster in video analytics
end-to-end workload Media
+ AI inference performance

Up to
8.2x
Better performance
per watt per \$


See complete performance information at [intel.com/processorclaims](https://www.intel.com/processorclaims): Intel® Core™ Ultra processors. Results may vary.

Intel® Core™ Ultra 9 200H Series Processor Gen over Gen Performance Improvements



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

VS



**Intel® Core™ Ultra 9
processor 185H**
(codenamed Meteor Lake H)

Up to

2.2x

Higher performance in
Procyon AI Computer
Vision

Up to

3.3x

Higher performance in
Llama 3 8B

Up to

2.3x

Higher performance in
Stable Diffusion 1.5

Up to

1.24x

Faster graphic
performance

Up to

1.23x

(est.) Faster multi-
threaded performance

Up to

1.15x

(est.) Faster single-
threaded performance

Intel® Core™ Ultra 7 200H Series Processor Gen over Gen Performance Improvements



**Intel® Core™ Ultra 7
processor 265H**
(codenamed Arrow Lake H)

VS



**Intel® Core™ Ultra 7
processor 165H**
(codenamed Meteor Lake H)

Up to

1.66x

Higher performance in
Procyon AI Computer
Vision

Up to

1.73x

Higher performance in
Llama 3 8B

Up to

2.85x

Higher performance in
Stable Diffusion 1.5

Up to

1.26x

Faster graphic
performance

Up to

1.23x

(est.) Faster multi-
threaded performance

Up to

1.15x

(est.) Faster single-
threaded performance

Intel® Core™ Ultra 7 200U Series Processor Gen over Gen Performance Improvements



**Intel® Core™ Ultra 7
processor 265U**
(codenamed Arrow Lake U)

Up to

1.17x

Higher performance in
latency-optimized
ResNet50

Up to

1.11x

Higher performance in
Llama 3 8B

Up to

1.08x

Higher performance in
Stable Diffusion 1.5

VS



**Intel® Core™ Ultra 7
processor 165U**
(codenamed Meteor Lake U)

Up to

1.21x

Faster graphic
performance

Up to

1.12x

(est.) Faster multi-
threaded performance

Up to

1.11x

(est.) Faster single-
threaded performance

Technical Advancements Over Previous Generations (200H Series)

Processor Family	13th Gen Intel® Core™ processors RPLH	Intel® Core™ Ultra processors (Series 1) MTL H	Intel® Core™ Ultra 200H Series Processors ARL H
Core/Thread	Up to 14 cores (6P+8E)	Up to 16 cores (6P+8E+2e)	Up to 16 cores (6P+8E+2e)
AI	Integrated GPU Intel Deep Learning Boost	Built-in Intel® Arc™ GPU (up to 18 TOPS) Integrated NPU (up to 11 TOPS) Intel Deep Learning Boost	Up to 99 total platform TOPS ¹ Built-in Intel® Arc™ GPU w/ Intel® XMX (up to 77 TOPS) Integrated NPU (up to 13 TOPS) Intel Deep Learning Boost
Graphics	Integrated GPU with up to 96EU	Intel® Arc™ GPU with up to 8 X ^e -cores Half rate ray-tracing	Intel® Arc™ GPU with up to 8 X ^e -cores Full rate ray-tracing
Memory	LPDDR4x 4267 LPDDR5/x 6400 DDR5 5200 / DDR4 3200	LPDDR5 6400 LPDDR5x 7467 (Type 4 board) DDR5 5600 (UH package SKUs)	LPDDR5x 8400 DDR5 6400
I/O Connectivity	Up to 8x PCIe 5, 8x PCIe 4, 12x PCIe 3 4x Integrated Thunderbolt 4 Integrated Wi-Fi 5	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

1. Select Intel Core Ultra 200H Series of Intel Core Ultra processors can achieve up to 99 total platform TOPS. Results may vary.

2. Intel® Arc™ GPU only available on H-SKU, 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.

Technical Advancements Over Previous Generations (200U Series)

Processor Family	13th Gen Intel® Core™ processors RPLU	Intel® Core™ Ultra processors (Series 1) MTL U	Intel® Core™ Ultra 200U Series Processors ARLU
Core/Thread	Up to 10 cores (2P+8E)	Up to 12 cores (2P+8E+2e)	Up to 12 cores (2P+8E+2e)
AI	Intel® Graphics Intel Deep Learning Boost	Intel® Graphics (up to 8 TOPs) Integrated NPU (up to 11 TOPs) Intel Deep Learning Boost	Up to 24 total platform TOPs ¹ Intel® Graphics (up to 8 TOPs) Integrated NPU (up to 13 TOPs) Intel Deep Learning Boost
Graphics	Intel® Graphics with up to 96EU	Intel® Graphics with up to 4 Xe-cores	Intel® Graphics with up to 4 Xe-cores
Memory	LPDDR5/x 6400 DDR5 4800	LPDDR5x 7467 DDR5 5600	LPDDR5x 8400 DDR5 6400
I/O Connectivity	Up to 8x PCIe 4, 12x PCIe 3 4x Integrated Thunderbolt 4 Integrated Wi-Fi 6E, Bluetooth 5.3	Up to 20x PCIe 4 4x Integrated Thunderbolt 4 Integrated Wi-Fi 6E, Bluetooth 5.3	Up to 20x PCIe 4 4x Integrated Thunderbolt 4 Integrated Wi-Fi 7, Bluetooth 5.4

¹ Select Intel Core Ultra 200U Series of Intel Core Ultra processors can achieve up to 24 total platform TOPs. Results may vary.

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

Segments Applications



Retail

- AI-Enhanced Checkout
- Self-service Kiosk
- Loss Prevention
- Return Management



Education

- Interactive Whiteboard
- Remote Classroom
- Video Conference



Hospitality

- Digital Menu
- Self-service Kiosk
- In-Store Analytics



Gaming

- Slot Machine
- Electronic Table Game
- Lottery Ticket Kiosks
- Digital Safety



Healthcare

- Lab Diagnostic Equipment
- Workstation on Wheels
- Nurse Stations
- Genomic Sequencers



Smart Cities

- License Plate Recognition
- Traffic Management
- Network Video Recorder



Industrial

- Mobile or Stationary Robots
- Machine Vision based Quality Control
- AI-Augmented Process Control

Key Features

Display: 4 concurrent 4K displays, Pipelock, EDID, Bezel Compensation
 Media: Integrated HDMI capture
 GPU virtualization with SR-IOV

AI-Capable: Up to 99 total platform TOPS¹ inferencing with CPU, NPU and iGPU w/ Intel® XMN

Intel® Core™ Ultra Processors: Up to 16 Cores, 8 Xe-Cores, 8 lanes PCIe 5.0, 20 lanes PCIe 4, LPDDR5 8400 memory

1. Select 200H SKUs of Intel Core Ultra processors can achieve up to 99 total platform TOPS. Results may vary.

Drive TCO/Sustainability Advantages


Quividi



"...improves AI inference throughput by 1.27X in object detection in average compared to previous generation... all AI processing run on CPU only, challenging the perception that running AI requires discrete GPU."

Paolo Prandoni
Chief Scientific Officer

CAI CONSTEMS-AI AI and Immersive Solutions



"...40x faster in end-to-end AI pipeline performance than previous gen wth Nvidia RTX A2000, while consuming less power... ensure higher quality control, reduce downtime, and increase overall efficiency."

Amit Srivastava
Co-Founder & Global CTO

SAMSUNG MEDISON



"...over 2.5X more powerful than mainstream discrete GPU, while the NPU slightly outperforms by 5%...eliminating the need for a discrete GPU. making cutting-edge imaging technology more accessible and cost-effective."

SungShik Baik
Principal Engineer


Critical Links



"...boosts GenAI inference throughput by up to 2.3X and cuts latency by up to 4.8X compared to previous generation. These improvements are a game changer for education, especially in remote areas."

Helder Pereira
VP of Engineering

n^x Network Optix



"... 1.35X faster throughput and lower latency compared to previous generation... enable us to seamlessly integrate context-aware monitoring features into our intelligent video solutions... meeting the demands of modern security challenges..."

Robin van Emden
Senior Director of Data Science

sodaclick



"...delivers 1.9x faster time to first token, 1.3x higher tokens per second throughput, all while using 1.4x less memory than the prior generation in Llama 3.2 (3B)... offers unmatched flexibility, performance, and efficiency ..."

Sam Jan
Co-Founder and COO

Empowering Brands at the Edge

Quividi

"...delivers an average 1.25x faster object detection inference than the previous generation... AI performance runs entirely on the CPU, freeing the GPU for signage tasks and ensuring flawless video playback without glitches during AI workload spikes."

Paolo Prandoni
Chief Scientific Officer

High fidelity AI detection providing faster, more accurate insights that are private




Audience measurement platform



Custom object detection multi model



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. See configuration details in the back up. Read the full list of customer testimonials on the Intel Press Portal.

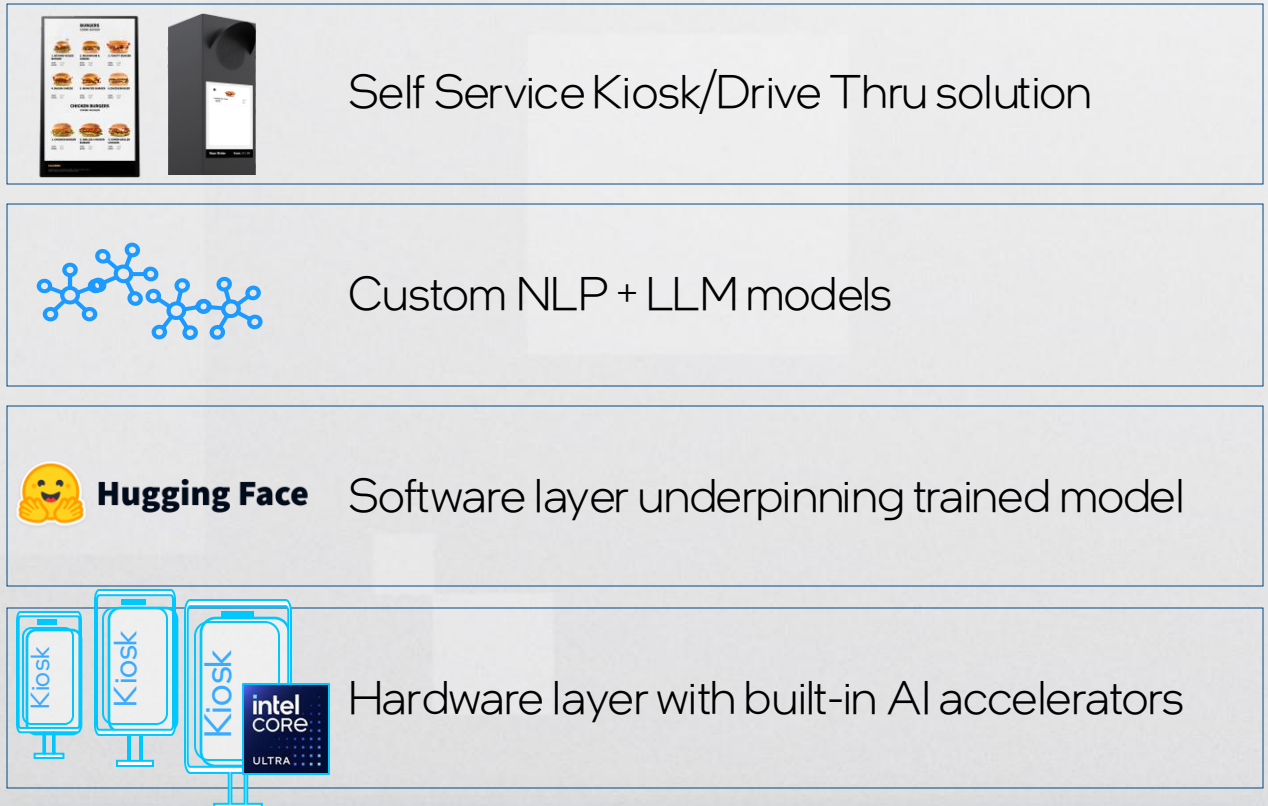
Empowering Hospitality at the Edge



“...delivers 1.9x faster time to first token (TTFT), 1.3x higher tokens per second throughput, all while using 1.4x less memory than the prior generation in Llama 3.2 (3B). With the CPU handling natural language processing (NLP) and the GPU focusing on generative AI... offers unmatched flexibility, performance, and efficiency.”

Sam Jan
Co-Founder and COO

Voice AI based self-service solution providing faster, better experience for quick service restaurant



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 configuration details in the back up. Read the full list of customer testimonials on the Intel Press Portal.

Empowering Teachers and Students at the Edge

Gen AI running locally, doesn't require internet and offers a private and safe option

Critical Links

"...boosts Gen AI inference throughput by 2.3X and cuts latency by 4.8X compared to previous generation. These improvements are a game changer for education, especially in remote areas."

Helder Pereira
VP of Engineering



E-learning and Gen AI applications



Model layer – Llama 3.1 8B



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. See configuration details in the back up. Read the full list of customer testimonials on the Intel Press Portal.


Enhancing Manufacturing at the Edge



“...40x faster in end-to-end AI pipeline performance including media decode, object classification and detection than... previous generation desktop processor paired with Nvidia RTX A2000, while consuming less power ... ensure higher quality control, reduce downtime, and increase overall efficiency.”

Amit Srivastava
Co-Founder & Global CTO

High-definition AI detection providing faster, more accurate insights without the need for dGPU



AI-powered defect detection solution



Pruned version of Resnet-50 & Faster-RCNN



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. See configuration details in the back up. Read the full list of customer testimonials on the Intel Press Portal.

Empowering Healthcare at the Edge

SAMSUNG MEDISON

“...over 2.5X more powerful than mainstream discrete GPU, while the NPU outperforms by up to 5%, all while consuming less than one-third of the power... eliminating the need for a discrete GPU, making cutting-edge imaging technology more accessible and cost-effective.”

SungShik Baik
Principal 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. See configuration details in the back up. Read the full list of customer testimonials on the Intel Press Portal.

Enabling on-device data processing to ensure user privacy and accessibility without the need for dGPU



Live View Assist Ultrasound Imaging



Custom object classification/detection multi model



Software layer underpinning trained model



Hardware layer with built-in AI accelerators

Enhancing Security Monitoring at the Edge



“...1.35X faster throughput and lower latency compared to previous generation... enable us to seamlessly integrate context-aware monitoring features into our intelligent video solutions... meeting the demands of modern security challenges.”


Robin van Emden
Senior Director of Data Science

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 configuration details in the back up. Read the full list of customer testimonials on the Intel Press Portal.

Dynamic and context-aware monitoring without the need for predefined rules



Intelligent Video Solutions



CLIP (Contrastive Language–Image Pre-training)



Software layer underpinning trained model



Hardware layer with built-in AI accelerators

Additional Resources

CNDA

- [NEX Arrow Lake U/H Platform Gold Deck](#)

Public

- N/A

The Intel logo is centered on a dark blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small, bright blue square is positioned above the letter "i". To the right of the word "intel" is a registered trademark symbol (®). The background is a solid dark blue with several faint, semi-transparent squares of varying shades of blue scattered across it, creating a subtle geometric pattern.

intel®