



Intel Announces New AI Innovations at Computex — Chip to Rackscale AI Solutions Delivered to Customers with the Help of Strategic Industry Partners

TAIPEI, Taiwan--(BUSINESS WIRE)-- Today at Computex 2026, Intel unveiled new innovations that address customers' chip-to-systems-level AI needs with solutions tailored to address their specific industry challenges, including:

- **New rackscale AI infrastructure:** Intel announced rackscale AI infrastructure for customers interested in scaling their inference and agentic workloads based on Intel® Xeon® processors and SambaNova SN-50 Reconfigurable Dataflow Units (RDUs).
- **Agentic Cloud Offering for Disaggregated Inference:** Vector Core Compute, a new purpose-built enterprise inference cloud formed by Vista Equity Partners and Cambium Capital, unveiled fully disaggregated inference running on Intel Xeon processors, SambaNova RDUs, and NVIDIA Blackwell GPUs.
- **Deep industry solutions:** Strategic collaborations with industry leaders, including Foxconn, Siemens, Hitachi, Echo Neurotechnologies, and Greenstone Biosciences focused on delivering integrated vertical customer solutions based on Intel processors and purpose-built silicon.
- **Intel Xeon 6+ processors:** Next-generation data center CPU built on Intel 18A and designed for high-density, scale-out workloads.
- **PC, gaming handheld, and physical AI momentum:** Broad partner support and customer uptake for the Series 3 family of processors.

“For more than five decades, Intel, its ecosystem partners, and Taiwan have brought the world the foundational technologies for the PC, Internet, and now AI eras,” said Lip-Bu Tan, CEO of Intel. “Today, with the rise of inference, agentic, and physical AI, Intel is poised to bring the world new innovations from the chip to systems level that promise to transform industry and society for the better. We are proud to join all our partners in building great products that will delight customers and bring the power of AI to more people as we create a brighter future together.”

Rackscale AI Infrastructure for Inference and Agentic Workloads

As the training of AI models has matured, and more AI applications have moved into production, the industry has witnessed an exponential rise in the demand for cost-effective and power-efficient AI inference. With the emergence of agentic AI, the growing demand for AI inference is changing the balance of power in the data center, returning the CPU to a position of prominence.

According to Creative Strategies CEO and principal analyst Ben Bajarín, while “the training-

era world looked closer to a one-CPU-per-four-GPU relation in AI deployments, agentic inference changes that relationship to roughly a one-CPU-to-one-GPU (or less) ratio.”

Seeking to capitalize on this trend at a systems level, Intel, SambaNova, and Foxconn today announced their intent to build rackscale AI infrastructure for data center, hyperscale, and intelligence center deployments—built on Intel Xeon processors.

The companies are demonstrating production-ready racks that combine Intel Xeon processors with SambaNova SN-50 RDUs, which together are designed to deliver high performance AI inference with improved cost and power efficiency. As part of the collaboration, Foxconn will provide system integration capabilities for the new rackscale AI infrastructure. Foxconn also plans to manufacture a CPU-dense variant of the rackscale infrastructure for workloads that do not require additional acceleration, including cost-optimized inference, data processing, and hybrid AI.

Agentic Cloud Offering for Fully Disaggregated Inference

Vector Core Compute, a new purpose-built enterprise inference cloud formed by Vista Equity Partners and Cambium Capital, unveiled fully disaggregated inference. Running onstage at Computex, Intel, SambaNova, Vista Equity Partners and Cambium Capital showcased the first real-world demonstration of a disaggregated inference system, using Intel Xeon 6 processors for orchestration and execution, SambaNova SN40 RDUs for decode, and NVIDIA Blackwell GPUs for prefill—operating from a Vector Core Compute data center in Los Angeles, California.

Together.ai is the first commercial customer running workloads on Vector Core Compute’s agentic cloud, which delivered the fastest enterprise inference on the MiniMax 2.5 model of any architecture to date. Vista Equity Partners has secured early access to the company’s high-quality, low-cost inference solutions for its 90+ portfolio companies which serve more than 2.5 million enterprise customers and 750 million users worldwide.

Industry Specific Solutions Based on Intel Processors and Purpose-built Silicon

It is often stated that AI is transforming every industry. It is also true that the computing needs of specific industries vary widely due to differences in their business environments, processes, workflows, and customers.

Today, Intel announced several strategic partnerships designed to co-develop industry-specific vertical solutions based on Intel processors and purpose-built silicon, including:

- **Foxconn:** The world’s largest electronics manufacturer is working with Intel to provide systems integration capabilities for rackscale AI infrastructure and explore collaboration in design services and custom silicon development.
- **Siemens:** The leading technology company focused on industry, infrastructure, transport, and healthcare and Intel have expanded their existing collaboration. In 2023, Siemens and Intel first joined forces; now the two companies are strengthening their collaboration across the entire value chain from design to manufacturing to chips embedded in Siemens products. Siemens brings its capabilities for the design, manufacturing, and lifecycle management of chips, as well as fab digitalization, automation, and electrification. This collaboration will enable the exploration of use

cases for purpose-built Intel silicon for Siemens' varied compute requirements, which may include edge devices, high-performance computing (HPC), and robotics.

- **Hitachi:** A global leader in digital innovation and sustainable solutions and Intel intend to work together on a range of solutions including foundry tools and quantum computing.
- **Echo Neurotechnologies:** The developer of neuroscience and brain-computer interface solutions and Intel are exploring new neuromorphic technologies to advance neuro-AI, speech neuroscience, brain-computer interfaces, and Intel's future neuromorphic and conventional hardware architectures.
- **Greenstone Biosciences:** The Silicon Valley biotech company plans to use Intel processors, purpose-built silicon, and the Intel Health and Life Sciences AI Suite to accelerate human-centric drug development using stem cells, organoids, genomics, and AI.

Intel Xeon 6+ Processors for Next Generation Data Centers

Extending this week's announcements from data centers and racks down to chip-level innovations, Intel announced the availability of [Intel Xeon 6+ processors](#), which provide greater performance density, power efficiency, and operational scale for cloud-native, agentic AI, and network-intensive workloads.

Built on Intel 18A—its first use in a data center CPU—Xeon 6+ is engineered for sustained performance under real-world power constraints—addressing the orchestration, concurrency, and data movement demands of emerging agentic AI.

Xeon 6+ can be configured for AI rackscale infrastructure purpose-built for hosting agents at maximum density. For example, a single liquid-cooled rack can deliver 36,864 cores using 32U of compute space, which provides the highest agent density available (at approximately 100-kilowatt rack power compute).

Optimized for environments where watts per rack, throughput per core, and latency predictability are critical, Xeon 6+ emphasizes scale-out performance—making room for new AI workloads without requiring disruptive data center redesign.

Series 3 Scale and Momentum

Core Ultra Series 3, built on Intel 18A, continues to experience strong customer uptake for a platform that now powers more than 325 consumer and commercial PC designs. Leveraging the same advanced IP as Ultra, the recently launched [Core processors](#) are enabling a new class of thin, sleek, powerful, and efficient PCs at affordable price points. Series 3 also pushes into the growing market of handheld gaming with the new [Intel Arc G-series processors](#), which will be available starting this month. The expansion of the Series 3 processor family is being accelerated by increased 18A yields and strong customer and partner engagement.

Beyond the PC, Intel has powered edge devices in manufacturing, robotics, retail, and smart cities for decades. For the first time, the latest Series 3 IP scaling in the PC ecosystem will deploy in parallel to thousands of edge customers globally. Over 130 customers have already chosen Series 3 to power [edge AI and robotics designs](#).

About Intel

Intel (Nasdaq: INTC) designs and manufactures advanced semiconductors that connect and power the modern world. Every day, our engineers create new technologies that enhance and shape the future of computing to enable new possibilities for every customer we serve. Learn more at [intel.com](https://www.intel.com).

Vector Core Compute cloud inference performance testing by Artificial Analysis. For details visit [Artificial Analysis](#). Results may vary.

Performance varies by use, configuration, and other factors. Learn more on the [Performance Index site](#). Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

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

© 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.

View source version on businesswire.com:

<https://www.businesswire.com/news/home/20260601425603/en/>

Intel Media Relations
ContactPR@Intel.com

Source: Intel Corporation