

December 12, 2018



New Intel Architectures and Technologies Target Expanded Market Opportunities

Intel Demonstrates 10nm-based PCs, Data Center and Networking Systems, Next-Gen 'Sunny Cove' Architecture with AI and Crypto Acceleration, and Industry's First 3D Logic Chip Packaging Technology

SANTA CLARA, Calif.--(BUSINESS WIRE)-- At Intel "Architecture Day," top executives, architects and fellows revealed next-generation technologies and discussed progress on a strategy to power an expanding universe of data-intensive workloads for PCs and other smart consumer devices, high-speed networks, ubiquitous artificial intelligence (AI), specialized cloud data centers and autonomous vehicles.

This press release features multimedia. View the full release here:

<https://www.businesswire.com/news/home/20181212005197/en/>

At Intel Architecture Day on Tuesday, Dec. 11, 2018, Raja Koduri, Intel chief architect, senior vice president of Core and Visual Computing Group and general manager of Edge Computing Solutions, explains how the company is positioned to deliver leadership products across architectures and workloads for an expanding data-centric market. (Credit: Intel Corporation)

Intel demonstrated a range of 10nm-based systems in development for PCs, data centers and networking, and

previewed other technologies targeted at an expanded range of workloads.

More: [New Intel Architectures and Technologies Target Expanded Market Opportunities](#) (Q&A with Intel's Raja Koduri)

The company also shared its technical strategy focused on six engineering segments where significant investments and innovation are being pursued to drive leaps forward in technology and user experience. They include: advanced manufacturing processes and packaging; new architectures to speed-up specialized tasks like AI and graphics; super-fast memory; interconnects; embedded security features; and common software to unify and simplify programming for developers across Intel's compute roadmap.

Together these technologies lay the foundation for a more diverse era of computing in an expanded addressable market opportunity of more than \$300 billion by 2022.¹

Intel Architecture Day Highlights:

Industry-First 3D Stacking of Logic Chips: Intel demonstrated a new 3D packaging technology, called "Foveros," which for the first time brings the benefits of 3D stacking to enable logic-on-logic integration.

Foveros paves the way for devices and systems combining high-performance, high-density

and low-power silicon process technologies. Foveros is expected to extend die stacking beyond traditional passive interposers and stacked memory to high-performance logic, such as CPU, graphics and AI processors for the first time.

The technology provides tremendous flexibility as designers seek to “mix and match” technology IP blocks with various memory and I/O elements in new device form factors. It will allow products to be broken up into smaller “chiplets,” where I/O, SRAM and power delivery circuits can be fabricated in a base die and high-performance logic chiplets are stacked on top.

Intel expects to launch a range of products using Foveros beginning in the second half of 2019. The first Foveros product will combine a high-performance 10nm compute-stacked chiplet with a low-power 22FFL base die. It will enable the combination of world-class performance and power efficiency in a small form factor.

Foveros is the next leap forward following Intel’s breakthrough Embedded Multi-die Interconnect Bridge (EMIB) 2D packaging technology, introduced in 2018.

New Sunny Cove CPU Architecture: Intel introduced Sunny Cove, Intel’s next-generation CPU microarchitecture designed to increase performance per clock and power efficiency for general purpose computing tasks, and includes new features to accelerate special purpose computing tasks like AI and cryptography. Sunny Cove will be the basis for Intel’s next-generation server (Intel® Xeon®) and client (Intel® Core™) processors later next year. Sunny Cove features include:

- Enhanced microarchitecture to execute more operations in parallel.
- New algorithms to reduce latency.
- Increased size of key buffers and caches to optimize data-centric workloads.
- Architectural extensions for specific use cases and algorithms. For example, new performance-boosting instructions for cryptography, such as vector AES and SHA-NI, and other critical use cases like compression and decompression.

Sunny Cove enables reduced latency and high throughput, as well as offers much greater parallelism that is expected to improve experiences from gaming to media to data-centric applications.

Next-Generation Graphics: Intel unveiled new Gen11 integrated graphics with 64 enhanced execution units, more than double previous Intel Gen9 graphics (24 EUs), designed to break the 1 TFLOPS barrier. The new integrated graphics will be delivered in 10nm-based processors beginning in 2019.

The new integrated graphics architecture is expected to double the computing performance-per-clock compared to Intel Gen9 graphics. With >1 TFLOPS performance capability, this architecture is designed to increase game playability. At the event, Intel showed Gen11 graphics nearly doubling the performance of a popular photo recognition application when compared to Intel’s Gen9 graphics. Gen11 graphics is expected to also feature an advanced media encoder and decoder, supporting 4K video streams and 8K content creation in constrained power envelopes. Gen11 will also feature Intel® Adaptive Sync technology

enabling smooth frame rates for gaming.

Intel also reaffirmed its plan to introduce a discrete graphics processor by 2020.

“One API” Software: Intel announced the “One API” project to simplify the programming of diverse computing engines across CPU, GPU, FPGA, AI and other accelerators. The project includes a comprehensive and unified portfolio of developer tools for mapping software to the hardware that can best accelerate the code. A public project release is expected to be available in 2019.

Memory and Storage: Intel discussed updates on Intel® Optane™ technology and the products based upon that technology. Intel® Optane™ DC persistent memory is a new product that converges memory-like performance with the data persistence and large capacity of storage. The revolutionary technology brings more data closer to the CPU for faster processing of bigger data sets like those used in AI and large databases. Its large capacity and data persistence reduces the need to make time-consuming trips to storage, which can improve workload performance. Intel Optane DC persistent memory delivers cache line (64B) reads to the CPU. On average, the average idle read latency with Optane persistent memory is expected to be about 350 nanoseconds when applications direct the read operation to Optane persistent memory, or when the requested data is not cached in DRAM. For scale, an Optane DC SSD has an average idle read latency of about 10,000 nanoseconds (10 microseconds), a remarkable improvement.² In cases where requested data is in DRAM, either cached by the CPU’s memory controller or directed by the application, memory sub-system responsiveness is expected to be identical to DRAM (<100 nanoseconds).

The company also showed how SSDs based on Intel’s 1 Terabit QLC NAND die move more bulk data from HDDs to SSDs, allowing faster access to that data.

The combination of Intel Optane SSDs with QLC NAND SSDs will enable lower latency access to data used most frequently. Taken together, these platform and memory advances complete the memory and storage hierarchy providing the right set of choices for systems and applications.

Deep Learning Reference Stack: Intel is releasing the Deep Learning Reference Stack, an integrated, highly-performant open source stack optimized for Intel® Xeon® Scalable platforms. This open source community release is part of our effort to ensure AI developers have easy access to all of the features and functionality of the Intel platforms. The Deep Learning Reference Stack is highly-tuned and built for cloud native environments. With this release, Intel is enabling developers to quickly prototype by reducing the complexity associated with integrating multiple software components, while still giving users the flexibility to customize their solutions.

- **Operating System:** Clear Linux* OS is customizable to individual development needs, tuned for Intel platforms and specific use cases like deep learning;
- **Orchestration:** Kubernetes* manages and orchestrates containerized applications for multi-node clusters with Intel platform awareness;
- **Containers:** Docker* containers and Kata* containers utilize Intel® Virtualization Technology to help secure container;

- **Libraries:** Intel® Math Kernel Library for Deep Neural Networks (MKL DNN) is Intel's highly optimized math library for mathematical function performance;
- **Runtimes:** Python* providing application and service execution runtime support is highly tuned and optimized for Intel architecture;
- **Frameworks:** TensorFlow* is a leading deep learning and machine learning framework;
- **Deployment:** KubeFlow* is an open-source industry-driven deployment tool that provides a fast experience on Intel architecture, ease of installation and simple use.

About Intel

Intel (NASDAQ: INTC), a leader in the semiconductor industry, is shaping the data-centric future with computing and communications technology that is the foundation of the world's innovations. The company's engineering expertise is helping address the world's greatest challenges as well as helping secure, power and connect billions of devices and the infrastructure of the smart, connected world – from the cloud to the network to the edge and everything in between. Find more information about Intel at newsroom.intel.com and intel.com.

¹Intel calculated 2022 total addressable market opportunity derived from industry analyst reports and internal estimates.

²Average idle read latency is the mean time for read data to return to a requesting processor. This is an average; some latencies will be longer. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Forward-Looking Statements

Statements in this news summary that refer to future plans and expectations, including with respect to Intel's future products and the expected availability and benefits of such products, are forward-looking statements that involve a number of risks and uncertainties. Words such as "anticipates," "expects," "intends," "goals," "plans," "believes," "seeks," "estimates," "continues," "may," "will," "would," "should," "could," and variations of such words and similar expressions are intended to identify such forward-looking statements. Statements that refer to or are based on estimates, forecasts, projections, uncertain events or assumptions, including statements relating to total addressable market (TAM) or market opportunity and anticipated trends in our businesses or the markets relevant to them, also identify forward-looking statements. Such statements are based on the company's current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in these forward-looking statements. Important factors that could cause actual results to differ materially from the company's expectations are set forth in Intel's earnings release dated October 25, 2018, which is included as an exhibit to Intel's Form 8-K furnished to the SEC on such date. Additional information regarding these and other factors that could affect Intel's results is included in Intel's SEC filings, including the company's most recent reports on Forms 10-K and 10-Q. Copies of Intel's Form 10-K, 10-Q

and 8-K reports may be obtained by visiting our Investor Relations website at www.intc.com or the SEC's website at www.sec.gov.

Intel and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

*Other names and brands may be claimed as the property of others.

View source version on businesswire.com:

<https://www.businesswire.com/news/home/20181212005197/en/>

Agnes Kwan
408-398-2573
agnes.ck.kwan@intel.com

Source: Intel Corporation