

## Intel Unpacks Architectural Innovations and Reveals New Transistor Technology at Architecture Day 2020

*Architecture Day 2020 Unveils Willow Cove, Tiger Lake and Xe Architectures, Introduces New Transistor Technology as Intel Advances its Six Pillars of Innovation*

August 13, 2020 – At the Architecture Day 2020 press event, Chief Architect Raja Koduri and Intel fellows and architects provided details on the progress Intel is making on its six pillars of technology innovation strategy. Intel revealed its 10nm SuperFin technology, representing the largest single, intranode enhancement in the company's history and delivering performance improvement comparable to a full-node transition.

The company also unveiled details of its Willow Cove microarchitecture and the Tiger Lake SoC architecture for mobile client and provided a first look at its fully scalable Xe graphics architectures that serve markets ranging from consumer to high-performance computing to gaming usages. Intel's disaggregated design approach, together with advanced packaging technology, XPU offerings and software-centric strategy, the company is focused on developing leading products across its portfolio to customers.

### 10nm SuperFin Technology

- After years of refining the FinFET transistor, Intel is redefining the technology to enable the largest single intranode enhancement in its history, delivering performance improvement comparable to a full-node transition. **10nm SuperFin** technology combines Intel's enhanced FinFET transistors with Super metal insulator metal capacitor. SuperFin technology offers enhanced epitaxial source/drain, improved gate process and additional gate pitch to enable greater performance by:
  - Enhancing epitaxial growth of crystal structures on the source and drain, thus increasing strain and reducing resistance to allow more current through the channel.
  - Improving gate process to drive higher channel mobility, which enables charge carriers to move more quickly.
  - Providing an additional gate pitch option for higher drive current in certain chip functions that require the utmost performance.
  - Using a novel thin barrier to reduce resistance by 30% and enhance interconnect performance.
  - Delivering a 5x increase in capacitance within the same footprint when compared to industry standard, driving a voltage droop reduction that translates to dramatically improved product performance. The technology is enabled by a new class of "Hi-K" dielectric materials stacked in ultra-thin layers just several angstroms thick to form a repeating "super lattice" structure. This is an industry-first technology that is ahead of current capabilities of other manufacturers.
- Intel's next-generation mobile processor, code-named Tiger Lake, is based on 10nm SuperFin technology. Tiger Lake is in production, and shipping to customers with original equipment manufacturer systems is expected for the holiday season.

## Packaging

- **Hybrid bonding test chip** taped out in Q2, 2020. Hybrid bonding is an alternative to the traditional “thermocompression” bonding used in most of today’s packaging technologies. This new technology enables very aggressive bump pitches of 10 microns and below, delivering much higher interconnect density and bandwidth, along with lower power.

## Willow Cove and Tiger Lake CPU Architectures

- **Willow Cove** is Intel’s next-generation CPU microarchitecture. Built on the latest process advancements, 10nm SuperFin technology and the foundation of the Sunny Cove architecture, Willow Cove delivers more than a generational increase in CPU performance with large frequency improvements and increased power efficiency. It also introduces a redesigned caching architecture to a larger non-inclusive 1.25MB MLC and security enhancements with Intel® Control Flow Enforcement Technology.
- **Tiger Lake** will offer intelligent performance and groundbreaking advancements in the key vectors of compute. With optimizations spanning the CPU, AI accelerators and being the first system-on-chip architecture with the new Xe-LP graphics microarchitecture, Tiger Lake will deliver more than a generational increase in CPU performance, massive AI performance improvements and a huge leap in graphics performance with a full set of best-in-class IPs throughout the SoC like the new, integrated Thunderbolt 4. Tiger Lake SoC architecture offers:
  - New Willow Cove CPU core with significant frequency uplift leveraging 10nm SuperFin technology advancements.
  - New Xe graphics with up to 96 execution units (EUs) with significant performance-per-watt efficiency improvements.
  - Power management – autonomous dynamic voltage frequency scaling in coherent fabric, increased fully integrated voltage regulator efficiency.
  - Fabrics and memory – 2x increase in coherent fabric bandwidth, ~86GB/s memory bandwidth, validated LP4x-4267, DDR4-3200; LP5-5400 architecture capability.
  - Gaussian Network Accelerator (GNA) 2.0 dedicated IP for low-power neural inferencing offloading from the CPU. ~20% lower CPU utilization on GNA vs. CPU (running noise suppression workload).
  - IO - Integrated TB4/USB4, integrated PCIe Gen 4 on CPU for low-latency, high-bandwidth device access to memory.
  - Display – up to 64GB/s of isochronous bandwidth to memory for multiple high-resolution displays. Dedicated fabric path to memory to maintain quality of service.
  - IPU6 – up to six sensors with 4K30 video, 27MP image, up to 4K90 and 42MP image architectural capability.

## Hybrid Architecture

- Intel is advancing its hybrid architecture with **Alder Lake, the company’s next-generation client product**. Alder Lake will combine two upcoming architectures: Golden Cove and Gracemont, optimized to offer great performance per watt.

## Xe Graphics Architectures

- Intel detailed the **Xe-LP** (low power) microarchitecture and software optimized to deliver efficient performance for mobile platforms. Xe-LP is Intel’s most efficient architecture for PC and mobile computing platforms with up to 96 EUs, and comes with new architecture designs including asynchronous compute, view instancing, sampler feedback, updated media engine with AV1 and

updated display engine. This will enable new end-user features with instant game tuning, capture, and stream-and-image sharpening. On software optimization, Xe-LP will have driver improvements with a new DX11 path and optimized compiler.

- The first **Xe-HP** chip has been powered on and back from the labs. Xe-HP is the industry's first multi-tiled, highly scalable, high performance architecture, providing data center-class, rack-level media performance, GPU scalability and AI optimization. It covers a dynamic range of compute from one tile, to two and four tiles, functioning like a multicore GPU. At Architecture Day, Intel demonstrated Xe-HP transcoding 10 full streams of high-quality 4K video at 60 frames per second on a single tile. Another demo showed the compute scalability of Xe-HP across multiple tiles. Intel is now sampling Xe-HP with key customers and plans to enable it in Intel® DevCloud for developers. Xe-HP will be available next year.
- Intel introduced a new Xe microarchitecture variant, **Xe-HPG**, a gaming-optimized microarchitecture, combining good performance-per-watt building blocks from Xe-LP, leveraging the scale from Xe-HP for a bigger configuration and compute frequency optimization from Xe-HPC. A new memory subsystem based on GDDR6 is added to improve performance per dollar and Xe-HPG will have accelerated ray tracing support. Xe-HPG is expected to start shipping in 2021.
- The **Intel® Server GPU (SG1)** is Intel's first discrete GPU based on Xe architecture for the data center. SG1 brings performance from four DG1s in a small form factor to the data center and is targeted for low-latency, high-density Android cloud gaming and video streaming. SG1 will ship later this year and will be in production soon.
- Intel's first Xe-based discrete GPU, code-named **DG1**, is in production and on track to start shipping in 2020. DG1 is now available within the Intel DevCloud, accessible to early access users. Announced at CES, it is Intel's first discrete GPU for PCs based on Xe-LP microarchitecture.
- New features were introduced to the Intel® Graphics Command Center, including **instant game tuning** and **game sharpening**.
  - Instant game tuning is a game-specific driver. Fixes and optimizations can be pushed to users faster than before and without a full driver download and install. It will only require a single opt-in from the user per game.
  - Game sharpening uses perceptual adaptive sharpening – a compute shader-based adaptive sharpening algorithm that boosts image clarity in games. This feature is particularly useful for titles that use resolution scaling to balance performance and image quality and is an opt-in feature within IGCC.

## Data Center Architectures

- **Ice Lake**, the first 10nm-based Intel® Xeon® Scalable processor targeted for the end of 2020, will deliver significant performance in both throughput and responsiveness across workloads. It will bring a set of technologies, including total memory encryption, PCIe Gen 4 and eight memory channels with instruction-set architecture to speed up crypto processing. Variants for network storage and internet of things will also be introduced as part of the Ice Lake family.
- **Sapphire Rapids** is Intel's next-generation Xeon Scalable processor based on enhanced SuperFin technology and will offer leading industry-standard technologies including DDR5, PCIe Gen 5 and Compute Express Link 1.1. Sapphire Rapids will be the CPU used in the Aurora Exascale supercomputer system at Argonne National Lab. It will continue our strategy of built-in AI acceleration with a new accelerator called Advanced Matrix Extensions. Sapphire Rapids is expected to start initial production shipments in the second half of 2021.

- Demonstrating our continued innovation for advancing field programmable gate array (FPGA) technologies and third consecutive generation of transceiver leadership, Intel now has the world's **first next-gen 224G-PAM4 TX transceiver**.

## Software

- The **oneAPI Gold** release will be available later this year, providing developers with production quality and performance across scalar, vector, matrix and spatial architectures. Intel released its eighth iteration of the oneAPI beta in July, delivering new features and enhancements for distributed data analytics, rendering performance, profiling, and the video and threading library. The **DG1 discrete GPU** is currently available to developers in the **Intel DevCloud**, allowing them early access to libraries and toolkits that enable them to write software using oneAPI before they have hardware in hand.

These news disclosures represent the progress of Intel's six pillars of technology innovation strategy. Intel is taking full advantage of its unique position to deliver a mix of scalar, vector, matrix and spatial architectures deployed in CPUs, GPUs, accelerators and FPGAs – unified by an open, industry-standard programming model, [oneAPI](#), to simplify application development.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information about performance and benchmark results, visit <http://www.intel.com/benchmarks>.

Refer to <https://software.intel.com/articles/optimization-notice> for more information regarding performance and optimization choices in Intel software products.

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.

Your costs and results may vary.

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

All product plans and roadmaps are subject to change without notice.

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 not intended to function as trademarks.

## Forward-Looking Statements

Statements in this article that refer to future plans and expectations 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 future products and technology and the expected availability and benefits of such products and technology, market opportunity, and anticipated trends in our businesses or the markets relevant to them, also identify forward-looking statements. Such statements are based on management'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 July 23, 2020, which is included as an exhibit to Intel's Form 8-K furnished to the SEC on such date, and 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](http://www.intc.com) or the SEC's website at [www.sec.gov](http://www.sec.gov). Intel does not undertake, and expressly disclaims any duty, to update any statement made in this article, whether as a result of new information, new developments or otherwise, except to the extent that disclosure may be required by law.

## **About Intel**

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to [newsroom.intel.com](http://newsroom.intel.com) and [intel.com](http://intel.com).

© 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.