


## News Byte

May 26, 2019

[Contact Intel PR](#)



Ahead of [Intel's COMPUTEX Opening Keynote](#) on May 28, the company today previewed products that deliver a significant boost in real-world workload performance, including a sneak peek at the company's new 10nm mobile processor (code-named "Ice Lake") and the special edition 9th Gen Intel® Core™ i9-9900KS processor – both shipping this year. The company discussed how performance leadership in the new data-centric era of computing will be defined beyond the traditional core count and frequency. Through the power of software, Intel® Architecture is optimized for real-world workload performance leadership that scales for today's and tomorrow's computing experiences.

Intel is driving this performance leadership with [its redefined product innovation model](#), delivering workload-optimized products by combining technical innovations across six pillars: process and packaging, architecture, memory, interconnect, security and software.

### Press Kit: [Intel at 2019 COMPUTEX](#)

"For every order of magnitude performance potential of a new hardware architecture there are two orders of magnitude performance enabled by software. Intel has more than 15,000 software engineers working to optimize workloads and unlock the performance of Intel processors," said Raja Koduri, chief architect and senior vice president of Intel Architecture, Software and Graphics.

Below are examples of the performance boost to real-world workloads for this new data-centric era:

**Ice Lake providing mobile graphics boost:** As unveiled early this month during [Intel's Investor Meeting](#), the company will begin shipping its first volume 10nm processor, a mobile PC product code-named "Ice Lake." Intel's new Gen11 graphics engine in Ice Lake is enabling the industry's first integrated GPU to incorporate variable rate shading capability by applying variable processing power to different areas of the scene to improve rendering performance. In addition, across a number of popular games, including CS:GO\*, Rainbow Six Siege\* and Total War: Three Kingdoms\*, Gen 11 graphics are expected to nearly double<sup>1</sup> the performance compared with Intel Gen 9 graphics, for stunning visual experiences on the go.



Intel announces a new special edition 9th Gen Intel Core i9-9900KS processor, which has been fine-tuned to deliver 5 GHz all-core. (Source: Intel Corporation)

[» Download full-size image](#)

**Heterogeneous computing architectures bring intelligent performance:** Intel is realizing the benefits of heterogeneous computing for both client and data center in current products through its architecture design and I/O innovations. [Ice Lake is a new highly-integrated platform](#) for laptops, combining the new “Sunny Cove” core architecture and the new Gen11 graphics architecture with both Thunderbolt™ 3 and Intel® Wi-Fi 6 (Gig+) integrated for the first time, providing best-in-class connectivity. This will also be Intel’s first processor designed to enable artificial intelligence (AI) for PC — leading with Intel® Deep Learning Boost (DL Boost) on the CPU, as well as AI instructions on the GPU and low power accelerators — to usher in a new era of intelligent performance for PCs. On top of showing Ice Lake accelerating the workloads that people do every day, such as image deblur and stylizing videos, the company also demonstrated how Intel DL Boost can offer up to 8.8 times<sup>2</sup> higher peak AI inference throughput than other comparable products on the market, as measured by AIXPRT.

For data-centric platforms, the [2nd Generation Intel® Xeon® Scalable processors](#) are the only processors with built-in Intel DL Boost AI accelerators, combining vector neural network instructions and deep learning software optimizations. With Intel DL Boost, 2nd Generation Intel Xeon Scalable processors accelerate AI inference workloads including image-recognition, object-detection and image-segmentation by up to 14 times<sup>4</sup> when compared with the previous generation Intel Xeon Scalable processor.

Compared with one of the commonly used GPUs in the market now, 2nd Generation Intel Xeon Scalable processors provide up to 2.4 times performance<sup>3</sup> on a recommendation system, one of the most popular AI workloads in the cloud today that accounts for over 60% of data center inference<sup>5</sup>.

**New special edition desktop gaming processor:** Intel previewed the 9th Gen Intel Core i9-9900KS special edition processor, the first to feature all 8 cores running at a turbo frequency of 5.0 GHz, making the world’s [best gaming desktop processor](#) even better.

Intel also showcased how the company is optimizing ultimate real-world performance on the most popular games running on Intel processors with both hardware and software innovations. Through the years, Intel has optimized hundreds of games by working with hundreds of thousands game developers.

**More leading performance examples at Intel COMPUTEX Industry Opening Keynote:** Gregory Bryant, Intel senior vice president and general manager of the Client Computing Group, will go into more details on Intel’s performance innovations and new experiences during his [COMPUTEX 2019 Industry Opening Keynote](#) on May 28. More details, including a livestream of the keynote, will be available in the Intel [Newsroom](#).

Intel, the Intel logo, Intel Core and Intel Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

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 information go to [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

Performance results are based on testing as of date specified in the Configuration Disclosure and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Optimization Notice: Intel’s compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the

availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

<sup>1</sup>Ice Lake gaming performance: Total War: Three Kingdoms 2.08X, Rainbow Six Siege 1.82X, CS:GO 1.72X: Based on gaming performance on those titles with the following settings: Total War:Three Kingdoms: "Battle" benchmark scenario at 1920x1080 resolution – full screen, V-Sync: off, Low Quality Preset, Resolution Scaling: 100%; Rainbow Six: Siege – Y4S1: 5 minutes of gameplay in "Suburban Extraction" situation at 1920x1080 resolution – full screen, Vsync: off, Medium Quality Preset, Measured with: PresentMon, 300Seconds; Counter-Strike: Global Offensive – 1.36.9.5: 5 minutes of gameplay against bots of Dust II map at 1920x1080 resolution – full screen, Medium quality Presets, Multicore Rendering: Enabled, FXAA: Disabled, Texture Filtering Mode: Anisotropic 4X, Vsync: Off. = Configuration: Intel preproduction system, ICL-U, PL1=15W, 4C/8T, Turbo TBA, Intel Gen11 Graphics, preproduction GFX driver, Memory: 8GB LPDDR4X-3733, Storage: Intel SSD Pro 760P 256GB, OS: Microsoft Windows\* 10, RS5 Build 475, preproduction Bios vs. Intel preproduction system, WHL U, Intel® Core™ i7 8565U 1.8 GHz, up to 4.6 GHz Turbo PL1=20W TDP, 4C/8T, Intel UHD Graphics 620, Graphics driver: 26.20.100.6709, Memory: 16GB DDR4-2400, Storage: Intel SSD 760P 512GB, OS: Microsoft Windows\* 10 RS5 Build Version 475, Measured by Intel as of May 2019

<sup>2</sup>Ice Lake AI Performance on IAXPRT: Workload: 7.6X more images per second using IAXPRT Community Preview 2 with Int8 precision on ResNet-50 and 8.8X higher peak AI inference throughput using IAXPRT Community Preview 2 on ResNet-50 Configuration: Intel preproduction system, ICL-U, PL1=15W, 4C/8T, Turbo TBA, Intel Gen11 Graphics, preproduction GFX driver, Memory: 8GB LPDDR4X-3733, Storage: Intel SSD Pro 760P 256GB, OS: Microsoft Windows\* 10, RS5 Build 475, preproduction Bios Vs. Commercially available OEM system with AMD\* Ryzen 7 3700U 2.3 GHz Turbo up to 4 GHz, 4C/8T, AMD\* Radeon\* Vega 10 graphics, Adrenalin 2019 19.4.3 GFX driver, Memory: 8GB DDR4-2400, Storage: SK Hynix BC501 256GB, OS: Microsoft Windows\* 10 RS5 Build 475, Bios: F.07. Measured by Intel as of May 2019

<sup>3</sup>Up to 2.41x performance advantage over Nvidia\* V100 GPUs: 2 socket Intel® Xeon® Platinum 8268 Processor, 24 cores HT On Turbo ON Total Memory 384 GB (12 slots/ 32GB/ 2933 MHz), BIOS: SE5C620.86B.0D.01.0286.011120190816 (ucode:0x4000013), CentOS 7.6, Kernel 4.19.5-1.el7.elrepo.x86\_64, SSD 1x INTEL SSDSC2KG96 960GB, Deep Learning Framework: MXNet <https://github.com/apache/incubator-mxnet.git> commit f1de8e51999ce3acaa95538d21a91fe43a0286ec applying [https://github.com/intel/optimized-models/blob/v1.0.2/mxnet/wide\\_deep\\_criteo/patch.diff](https://github.com/intel/optimized-models/blob/v1.0.2/mxnet/wide_deep_criteo/patch.diff), Compiler: gcc 6.3.1, MKL DNN version: commit: 08bd90cca77683dd5d1c98068cea8b92ed05784, Wide & Deep: [https://github.com/intel/optimized-models/tree/v1.0.2/mxnet/wide\\_deep\\_criteo](https://github.com/intel/optimized-models/tree/v1.0.2/mxnet/wide_deep_criteo) commit: c3e7cbde4209c3657ecb6c9a142f71c3672654a5, Dataset: Criteo Display Advertisement Challenge, Batch Size=512, 2 instance/2 socket, Datatype: FP32; with recommendation results: 678,000 records /seconds. vs. host system: 2 socket Intel® Xeon® Platinum 8180 processor (28 cores), HT ON, Total memory 128 GB (16 slots/8 GB/ 2666 MHz), Ubuntu 18.04.2 LTS Accelerator: Nvidia\* Turing V100 GPU accelerator, 32GB HBM2, 32GB/sec Interconnect BW, System interface x16 PCIe Gen3, Driver Version 410.78, CUDA Version 10.0.130, CUDNN Version 7.5, CUDA CUBLAS 10.0.130 Deep learning workload: MxNet 1.4.0 <https://pypi.org/project/mxnet-cu92/>, DataType:FP32, Batch Size= 512, Running 2 instances Model: Wide & Deep: [https://github.com/intel/optimized-models/blob/master/mxnet/wide\\_deep\\_criteo/model.py](https://github.com/intel/optimized-models/blob/master/mxnet/wide_deep_criteo/model.py) Commit ID for the current state is c3e7cbde4209c3657ecb6c9a142f71c3672654a5 Training dataset (8,000,000 samples): wget [https://storage.googleapis.com/dataset-uploader/criteo-kaggle/large\\_version/train.csv](https://storage.googleapis.com/dataset-uploader/criteo-kaggle/large_version/train.csv) Evaluation dataset (2,000,000 samples): wget [https://storage.googleapis.com/dataset-uploader/criteo-kaggle/large\\_version/eval.csv](https://storage.googleapis.com/dataset-uploader/criteo-kaggle/large_version/eval.csv) python3 inference.py –batch-size \$bs –num-batches 10000 >> \$outdir/bs\$bs-\$runid.2xbgout 2>&1 & python3 inference.py –batch-size \$bs –num-batches 10000 >> \$outdir/bs\$bs-\$runid.2xfout 2>&1. Recommendation results: 281,211 records/second. Tested by Intel as of March 2019

<sup>4</sup><https://arxiv.org/ftp/arxiv/papers/1704/1704.04760.pdf>, <https://research.fb.com/wp-content/uploads/2017/12/hpca-2018-facebook.pdf>

<sup>5</sup>Up to 14x AI Performance Improvement with Intel® DL Boost compared to Intel® Xeon® Platinum 8180 Processor when launched (July 2017). Tested by Intel as of 2/20/2019. 2 socket Intel® Xeon® Platinum 8280 Processor, 28 cores HT On Turbo ON Total Memory 384 GB (12 slots/ 32GB/ 2933 MHz), BIOS: SE5C620.86B.0D.01.0271.120720180605 (ucode: 0x200004d), Ubuntu 18.04.1 LTS, kernel 4.15.0-45-generic, SSD 1x sda INTEL SSDSC2BA80 SSD 745.2GB, nvme1n1 INTEL SSDPE2KX040T7 SSD 3.7TB, Deep Learning Framework: Intel® Optimization for Caffe version: 1.1.3 (commit hash: 7010334f159da247db3fe3a9d96a3116ca06b09a), ICC version 18.0.1, MKL DNN version: v0.17 (commit hash: 830a10059a018cd2634d94195140cf2d8790a75a, model:[https://github.com/intel/caffe/blob/master/models/intel\\_optimized\\_models/int8/resnet50\\_int8\\_full\\_conv.prototxt](https://github.com/intel/caffe/blob/master/models/intel_optimized_models/int8/resnet50_int8_full_conv.prototxt), BS=64, DummyData, 4 instance/2 socket, Datatype: INT8 vs Tested by Intel as of July 11th2017: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50 GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel\_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86\_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC). Performance measured with: Environment variables: KMP\_AFFINITY='granularity=fine, compact', OMP\_NUM\_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (<http://github.com/intel/caffe/>), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time – forward\_only" command, training measured with "caffe time" command. For "ConvNet" topologies, dummy dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from [https://github.com/intel/caffe/tree/master/models/intel\\_optimized\\_models](https://github.com/intel/caffe/tree/master/models/intel_optimized_models)(ResNet-50), Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

Tags: [2019 Computex](#), [9th Gen Intel Core](#), [Ice Lake](#), [Intel Xeon](#), [PC innovation](#), [Taiwan News](#)

## Other News



April 8, 2021

[SD Supercomputer Center Selects Habana, Intel for Efficient AI](#)

April 6, 2021

[Intel Xeon Advances Nasdaq's Homomorphic Encryption R&D](#)

April 6, 2021

[3rd Gen Intel Xeon Scalable Launch](#)

### 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](https://newsroom.intel.com) and [intel.com](https://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.

## Latest News: Artificial Intelligence



April 8, 2021

[SD Supercomputer Center Selects Habana, Intel for Efficient AI](#)



April 6, 2021

[Intel Launches Its Most Advanced Performance Data Center Platform](#)



April 1, 2021

[At John Deere, 'Hard Iron Meets Artificial Intelligence'](#)

[Read More](#)