

TRANSFORM YOUR DIGITAL WORLD

From the heart of the data center
To the edge, and back



2nd Generation Intel® Xeon® Scalable Processors


INTRODUCING NEW INTEL® XEON® PLATINUM 9200 PROCESSORS

A new class of advanced performance

Up to 112 cores in a two-socket system

Up to 3.8 GHz with Intel® Turbo Boost Technology


Highest native DDR4 bandwidth




Average of **2x** performance improvement¹



PERFORMANCE TO PROPEL INSIGHTS




Average of 1.33x performance improvement over prior generation²



Built-in agility for evolving workloads with new Intel® Speed Select Technology

Available on select 2nd gen Intel® Xeon® Scalable processors

INFRASTRUCTURE-WIDE AI READINESS



New Intel® Deep Learning Boost with Vector Neural Network Instructions (VNNI)

Up to 30x AI performance improvement compared to prior generation³

BUSINESS RESILIENCE THROUGH MULTI-TENANT DATA PROTECTION



Hardware-enhanced security

New Intel® Security Libraries (Intel® ISeCL-DC)

Featuring new Intel® Threat Detection Technology (Intel® TDT)⁴

GROUNDBREAKING MEMORY AND STORAGE INNOVATION




Delivers support for up to 36 TB system memory (8-socket system with Intel® Optane™ DC persistent memory and DRAM combined)

Supporting reduced cost and enhanced TCO

2x memory capacity
Compared to Intel® Xeon® 8180 processor

Up to 36% more VMs per system⁵

For faster insights



INDUSTRY-LEADING STORAGE SOLUTIONS




New Intel® Optane™ DC D4800X SSD



Performance and resiliency for critical enterprise apps

Dual-port connections for 24x7 data availability and redundancy

New Intel® SSD D5-P4326 E1.L "Ruler"




Up to 1 PB of storage capacity in 1U

Unprecedented data storage capacity and density

INTELLIGENT, FAST NETWORKING

New Intel® Ethernet 800 Series Network Adapters



Up to 100 GbE

Featuring Application Direct Queuing (ADQ)
Prioritizes data flow and workload acceleration

2ND GEN INTEL® XEON® SCALABLE PROCESSORS

The foundation for data-centric innovation



Discover more at www.intel.com/xeonscalable

¹ 2x Average Performance Improvement compared with Intel® Xeon® Platinum 8180 processor. Geo mean of estSPECrate2017_int_base, estSPECrate2017_fp_base, Stream Triad, Intel Distribution of Linpack, server side. Java. Platinum 92xx vs Platinum 8180. 1-node, 2x Intel® Xeon® Platinum 9282 cpun Wolf Pass with 384 GB (12 X 32GB 2933 (2666)) total memory, ucode0x4000013 on RHEL7.6, 3.10.0-957.el7.x86_65, IC18u1, AVX512, HT on all (off Stream, Linpack), Turbo on, result: estintthroughput=526, estfpthroughput=526, Stream Triad=407, Linpack=6411, server side java=332913, test by Intel on 2/16/2019, vs. 1-node, 2x Intel® Xeon® Platinum 8180 cpun Wolf Pass with 384 GB (12 X 32GB 2666) total memory, ucode0x200004D on RHEL7.6, 3.10.0-957.el7.x86_65, IC18u1, AVX512, HT on all (off Stream, Linpack), Turbo on all (off Stream, Linpack), result: estintthroughput=307, estfpthroughput=251, Stream Triad=204, Linpack=3238, server side java=165724, test by Intel on 1/29/2019.

² Up to 1.33x Gen-on-Gen Performance Improvement compared to Intel® Xeon® Gold 5100 Processor. Geomean of estSPECrate2017_int_base, estSPECrate2017_fp_base, Stream Triad, Intel Distribution of Linpack, server side. Java. Gold 5218 vs Gold 5118. 1-node, 2x Intel® Xeon® Gold 5218 cpun Wolf Pass with 384 GB (12 X 32GB 2933 (2666)) total memory, ucode0x4000013 on RHEL7.6, 3.10.0-957.el7.x86_65, IC18u2, AVX2, HT on all (off Stream, Linpack), Turbo on, result: estintthroughput=162, estfpthroughput=172, Stream Triad=185, Linpack=1088, server side java=98333, test by Intel on 12/7/2018, 1-node, 2x Intel® Xeon® Gold 5118 cpun Wolf Pass with 384 GB (12 X 32GB 2666 (2400)) total memory, ucode0x200004D on RHEL7.6, 3.10.0-957.el7.x86_65, IC18u2, AVX2, HT on all (off Stream, Linpack), Turbo on, result: estintthroughput=119, estfpthroughput=134, Stream Triad=148.6, Linpack=822, server side java=67434, test by Intel on 11/12/2018.

³ Up to 30X AI performance with Intel® DL Boost compared to Intel® Xeon® Platinum 8180 processor (July 2017). Tested by Intel as of 2/26/2019. Platform: Dragon rock 2 socket Intel® Xeon® Platinum 9282(56 cores per socket), HT ON, turbo ON, Total Memory 768 GB (24 slots/ 32 GB/ 2933 MHz), BIOS=ESC620.86B.0D.01.0241.112020180249, Centos 7 Kernel 3.10.0-957.5.1.el7.x86_64, Deep Learning Framework: Intel® Optimization for CaffeVersion: <https://github.com/intel/caffe> d554cf1, ICC 2019.2.187, MKL DNN version: v0.17 (commit hash: 830a10059a018cd2634d94195140cf2d8790a75a), model: [https://github.com/intel/caffe/blob/master/models/intel_optimized_models/m8/resnet50_int8_full_conv_prototxt_B5-64_No_dataLayerDummy_data3x224x224_56_instance/2_socket_DataType_INT8_vs_Tested_by_Intel_as_of_July_11th_2017_25_Intel_Xeon_Platinum_8180_CPU_@_2.50GHz_\(28_cores\)_HT_disabled_turbo_disabled_scaling_governor_set_to_performance_via_intel_ptstatalayerDriver_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_25mm_MLC](https://github.com/intel/caffe/blob/master/models/intel_optimized_models/m8/resnet50_int8_full_conv_prototxt_B5-64_No_dataLayerDummy_data3x224x224_56_instance/2_socket_DataType_INT8_vs_Tested_by_Intel_as_of_July_11th_2017_25_Intel_Xeon_Platinum_8180_CPU_@_2.50GHz_(28_cores)_HT_disabled_turbo_disabled_scaling_governor_set_to_performance_via_intel_ptstatalayerDriver_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_25mm_MLC). Performance measured with: Environment variables KMP_AFFINITY=granularity=full, compact, OMP_NUM_THREADS=56, CPU frequency with cpupowerfrequency-set -d 2.5G -u 3.8G -g performance. Caffe: (<http://github.com/intel/caffe/>), revision 96b759f71b2281835f690af267188b82b150b5c. 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_\(ResNet-50\)](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".

⁴ No product or component can be absolutely secure.

⁵ Up to 36% more VMs per system. Tested by Intel on 1/31/2019. Config: 1 (DRAM only), Intel Reference Platform, 1-Node, 2 sockets, Intel® Xeon® Platinum 8272L processor, HT: on, Turbo: on, BKC version: WW42, Intel Optane DC persistent memory FW version 5253, System DDR Mem Config: slots / cap / run-speed 24 slots / 32GB / 2666, Total Memory/Node (DDR, DCPMM) 768GB, 0, Storage - boot 1x Samsung PM963 M.2 960GB, Storage - application drives 7 x Samsung PM963 M.2 960GB, 4x Intel 5520s 54000 (1.92TB), NIC: Intel X520 SR2 (10Gb), PCIe LBG Q3/PQ - T - B2, OS: Windows Server 2019 RS5-17763, Workload & Version: OLTP Cloud Benchmark, Config: 2 Intel Optane DC Persistent Memory, 1-Node, 2-Sockets, Intel® Xeon® Platinum 8272L processor, Intel Reference Platform, HT: on, Turbo: on, BKC version: WW42, Intel Optane DC persistent memory FW version 5253, System DDR Mem Config: slots / cap / run-speed 12 slots / 16 GB / 2666, System DCPMM Config: slots / cap / run-speed 8 slots / 128GB / 2666, Total Memory Node: 192GB, 1TB, Storage-boot: 1x Samsung PM963 M.2 960GB, Storage-applications: 7x Samsung PM963 M.2 960GB, 4x Intel SSDs 54600 (1.92TB), NIC: 1x Intel X520 SR2 (10Gb), PCIe: LBG OS/PQ - T - B2, OS: Windows Server 2019 RS5-17763, Workload: OLTP Cloud Benchmark.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. For more complete information about performance and benchmark results, visit <http://www.intel.com/benchmarks>.

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 compiled 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 visit <http://www.intel.com/benchmarks>.

Intel® Advanced Vector Extensions (Intel® AVX) provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at or below the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at <http://www.intel.com/go/turbo>.

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 SSE3 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.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

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

©Intel names and brands may be claimed as the property of their owners. ©Intel Corporation.