


# ACCELERATE DATA-CENTRIC BUSINESS TRANSFORMATION



2nd Gen Intel® Xeon® Scalable processors deliver performance, agility and scalability



Experience up to **30x AI performance**<sup>1</sup> versus prior generation with new, performance-optimized processor and platform features that let you discover more insights, faster.



## SCALE PERFORMANCE

Platform-wide performance improvements with:

- Up to **112 processor cores** in a 2-socket system
- Up to **2x memory bandwidth**<sup>2</sup> versus prior generation
- Up to **3.5x more VMs**<sup>3</sup> per server (5-year refresh improvement)
- Supports revolutionary **Intel® Optane™ DC Persistent Memory**



## IMPROVE SECURITY

Built-in data protection features helps better secure platforms and data. Provides trusted service delivery and accelerated encryption.

Learn more by  
contacting your  
Intel Technology  
Provider

ITP  
LOGO

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 dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product can be absolutely secure. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. 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 SSE4.2 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. Notice revision 202110804. The benchmark results may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workload. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations. 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.

**1 Up to 30x AI performance with Intel® Deep Learning Boost (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: SES620.86B.0D.01.0241.112020180249, Centos® 7 Kernel 3.10.0-957.5.1.el7.x86\_64, Deep Learning Framework: Intel® Optimization for Caffe® version: <https://github.com/intel/caffe> d554cbf1, ICC 2019.2.187, MKL DNN version: v0.17 (commit hash: 830a10059a018cd-2634d94195140cf2d8790a75a), 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), B5=64, No datalayer DummyData: 3x224x224, 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\_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: (<https://github.com/intel/caffe>), revision f663759f71b218135f69a0af267158b82b150b5c. 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® Math Kernel Library (Intel® MKL) small libraries version 2018.0.20170425. Caffe run with "numactl -s".

**2 Double the memory bandwidth** with 12 memory channels per CPU and 24 memory channels per compute module, compared against CLX-SP product family with 6 memory channels per CPU.

**3 Up to 3.50X 5-Year Refresh Performance Improvement VM density** compared to Intel® Xeon® ES-2600 v6 processor: 1-node, 2x ES-2697 v2 on Canon Pass with 256 GB (16 slots / 16GB / 1600) total memory, ucode 0x42c on RHEL7.6, 3.10.0-957.el7.x86\_65, 1x Intel 400GB SSD OS Drive, 2x P4500 4TB PCIe\*, 2x82599 dual port Ethernet, Virtualization Benchmark, VM kernel 4.19, HT on, Turbo on, score: VM density=74, test by Intel on 1/15/2019. vs. 1-node, 2x 8280 on Wolf Pass with 768 GB (24 slots / 32GB / 2666) total memory, ucode 0x2000056 on RHEL7.6, 3.10.0-957.el7.x86\_65, 1x Intel 400GB SSD OS Drive, 2x P4500 4TB PCIe\*, 2x82599 dual port Ethernet, Virtualization Benchmark, VM kernel 4.19, HT on, Turbo on, score: VM density=21, test by Intel on 1/15/2019.

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