



# TOP 5 REASONS TO UPGRADE TO 2ND GEN INTEL® XEON® SCALABLE PROCESSOR-BASED SERVERS

## Designed for the Data-centric Era

As the volume and value of data continues to explode, you need an infrastructure that can move faster, store more and process everything.

Intel's latest processor, memory, storage, network and software innovations deliver the outstanding, balanced performance and capacity you need to gain a competitive edge.

## Accelerate AI and Analytics Insights

It starts with the compute performance of 2nd Gen Intel® Xeon® Scalable processors—with more world-class per-core performance and new features, such as Intel® DL Boost which accelerates AI inference 30x<sup>1</sup> on an Intel® Xeon® Platinum 9200 processor compared to an Intel® Xeon® Platinum 8180 processor (July 2017).

But compute performance only takes you so far. Intel has worked with its vast ecosystem to optimize performance across leading software. A 2nd Gen Intel® Xeon® Scalable processor paired with revolutionary Intel® Optane™ DC persistent memory doubles memory capacity when combined with DRAM (vs. a DRAM-only approach). And, Intel® Optane™ SSDs deliver exceptional storage performance, while Intel networking innovations provide fast I/O between platforms.

## One Infrastructure – Many Demanding Workloads

Why deploy, manage and maintain multiple, disparate, siloed infrastructures to run individual workloads when a single infrastructure can do it all?

2nd Gen Intel® Xeon® Scalable processor-based servers excel at running your most compute-intensive *and* data-demanding workloads—from AI training and inference...to real-time, in-memory analytics...to HPC workloads...to all of your business-critical and everyday workloads.

And that means you can spend your time driving growth, differentiation and innovation instead of managing multiple infrastructures.

## Hardware-hardened Security

Deploying strong security software is critical, but a software-only security approach can only take you so far. That's why Intel designed 2nd Gen Intel® Xeon® Scalable processors with built-in security features.

Silicon-based innovations help to eliminate threats from data center platforms, and additional features greatly accelerate data encryption and decryption for data at rest, in use and in flight.

## Next-level Hybrid Cloud

Run 3.5x more VMs per server versus prior generation<sup>2</sup> (5-year refresh VM density improvement). Take advantage of hardware-based virtualization enhancements. Improve cloud-based resource utilization. And scale server performance and capacity to new highs.

These are just a few of the reasons your hybrid cloud will benefit from 2nd Gen Intel® Xeon® Scalable processor-based servers, delivering new levels of agility, cost-efficiency, scalability and responsiveness.



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 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. Notice Revision #20110804

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 workloads. 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: SE5C620.86B.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), BS=64, No datalayer DummyData: 3x224x224, 56 instance/2 socket, Datatype: INT8 vs Tested by Intel as of July 11th 2017: 2S 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 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® Math Kernel Library (Intel® MKL) small libraries version 2018.0.20170425. Caffe run with "numactl -l".

**2 Up to 3.50X 5-Year Refresh Performance Improvement VM density** compared to Intel® Xeon® E5-2600 v6 processor: 1-node, 2x E5-2697 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\*, 2\*82599 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 VIRT 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\*, 2\*82599 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. \*Other names and brands may be claimed as the property of others.