HLRN User Workshop 3-6 Nov 2020

Cascade Lake Advanced Performance

CLX-AP Overview

Klaus-Dieter Oertel

intel®

Cascade Lake

Building Block for CLX-AP

2-socket+ Intel® Xeon® Roadmap

Thurley Platform		Romley Platform		Grantley Platform		Purley Platform	
Intel® Microarchitecture Codenamed Nehalem		Intel® Microarchitecture Codenamed Sandy Bridge		Intel® Microarchitecture Codenamed Haswell		Intel® Microarchitecture Codenamed Skylake	
Nehalem	Westmere	Sandy Bridge	Ivy Bridge	Haswell	Broadwell	Skylake	Cascade Lake
45nm	32nm	32nm	22nm	22nm	14nm	14nm	14nm
New Micro-architecture		New Micro-architecture		New Micro-architecture		New Micro-architecture	

Brickland Platform is Ivy Bridge-EX, Haswell-EX, and Broadwell-EX

Cascade Lake refreshes Purley platform with higher performance and new capabilities

Cascade Lake Overview on Purley Platform

CPU	**Cascade Lake:** Up to 28Core with Intel HT Technology (Drop in to Purley Platform @ 70W-205W)
New Capabilities	**Frequency and architecture improvements, VNNI (for AI/DL) and Intel® Optane™ DC persistent memory* - module support** on select SKUs**, **Intel® Speed Select Technology** on select SKUs
Socket	Socket P
Scalability	2S, 4S, & glueless 8S (>8S via xNC support)
Memory	6 channels DDR4 R/LRDIMM per CPU/ 12 DIMMs per socket, up to 2666 MT/s 2DPC, **up to 2933 MT/s 1DPC; 16Gb DDR4 based DIMMs support**** **Intel® Optane™ DC persistent memory (up to 512 GB / module)****
UPI **Ultra Path Interconnect**	Up to 3 links per CPU x20, speed: 9.6 and 10.4 GTS
PCIe	PCIe Gen 3: 48 lanes per CPU (bifurcation support: x16, x8, x4)
Host Fabric	Discrete Intel® Omni-Path Architecture adapter (100Gb/s) **[Integrated Fabric SKUs available on Skylake only]**
FPGA	**Support for discrete Intel Arria® 10 FPGA**
PCH – Lewisburg	Intel® QuickAssist Tech (QAT), eSPI, Integrated Intel Ethernet Connection: up to 4x10Gb/1Gb ports, Up to 20 ports PCIe* 3.0 (8 GT/s) Up to 14 SATA 3, Up to 14 USB 2.0, Up to 10 USB 3.0

Purley 2S Configuration

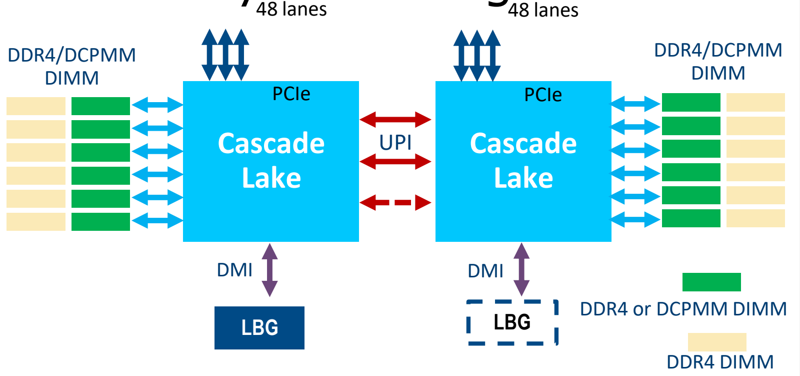

Purley 4S Configuration

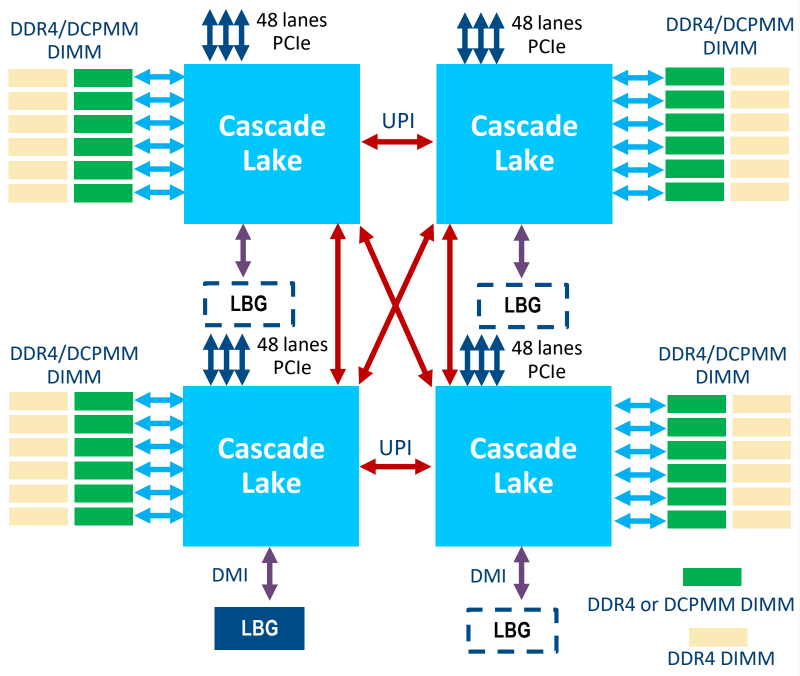

New capabilities/Changes relative to Skylake/Purley in **Bolded Blue**

* Intel® Optane™ persistent memory (DCPMM)

Formerly (codename Apache Pass, AEP)

Intel® Speed Select on Cascade Lake

- Capability to configure the CPU to run at 3 distinct operating points
  - Each operating point defined by core count with a base frequency associated to that core count
    - Higher core count with lower base frequency
    - Lower core count with higher base frequency
  - SKU Stack will include Speed Select specific SKUs
- Static Boot Time Configuration
  - BIOS discovers capability and prompts user to select from core count / base frequency configurations at boot
- Key Value Prop
  - Multiple CPU personalities based on workload/VM Needs
  - Improved server utilization in data center through SKU consolidation
  - Improved guaranteed per-core performance SLAs

For more info see CDI# 597725 'Customer Enabling Deep Dive: Intel® Speed Select for CLX'

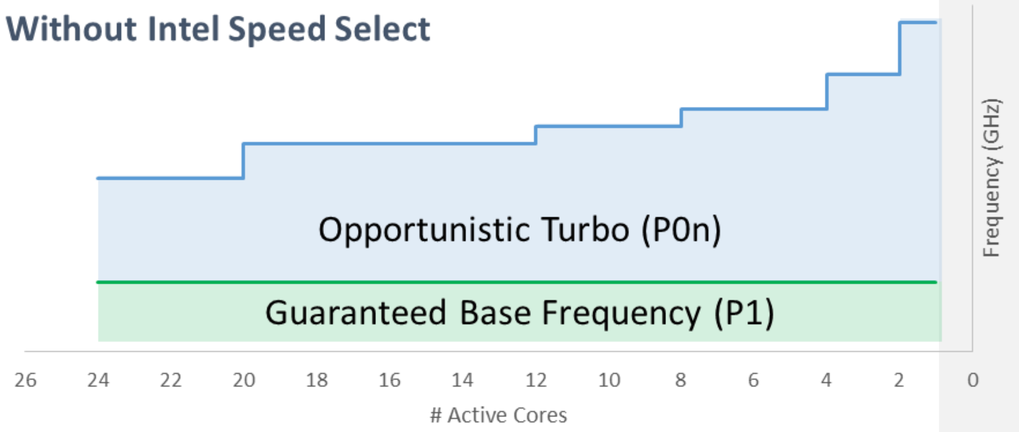

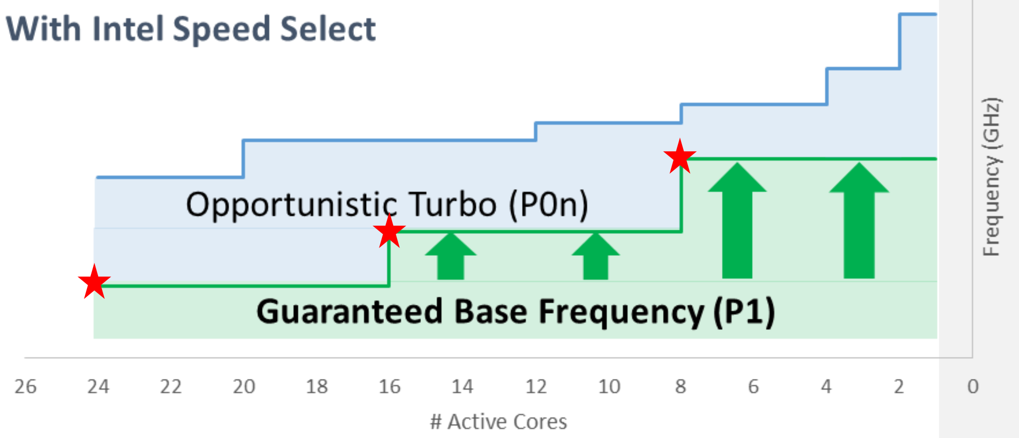

* Frequency and Core Count for Illustration Only

Cascade Lake: Inference Enhancements with Intel® Deep Learning Boost

- ***Intel® Deep Learning Boost (VNNI) on future Intel® Xeon® Scalable processor (codename "Cascade Lake") is designed to deliver significant, more efficient Deep Learning (Inference) acceleration.***

- Intel® DL Boost (VNNI): A new Intel® Advanced Vector Extension (Intel® AVX-512) instruction
  - 8-bit (VPDPBUSD) new instruction, to accelerate Inference performance.
- No hardware changes are required to support Intel® DL Boost on Purley Platform
  - Minimal OS/VMM enabling if Intel® AVX-512F (foundation) state pre-exists
  - SW development support will be enabled through optimizations on popular AI/Deep Learning frameworks (eg: TensorFlow, Caffe & MXNet and libraries (Intel® Math Kernel Library – Deep Neural Networks)
- Intel® DL Boost instruction is available on all CLX-SP XCC B-step, HCC and LCC SKUs

Growing Gap Between Memory Hierarchy

Limitations to traditional architecture impede unified data management

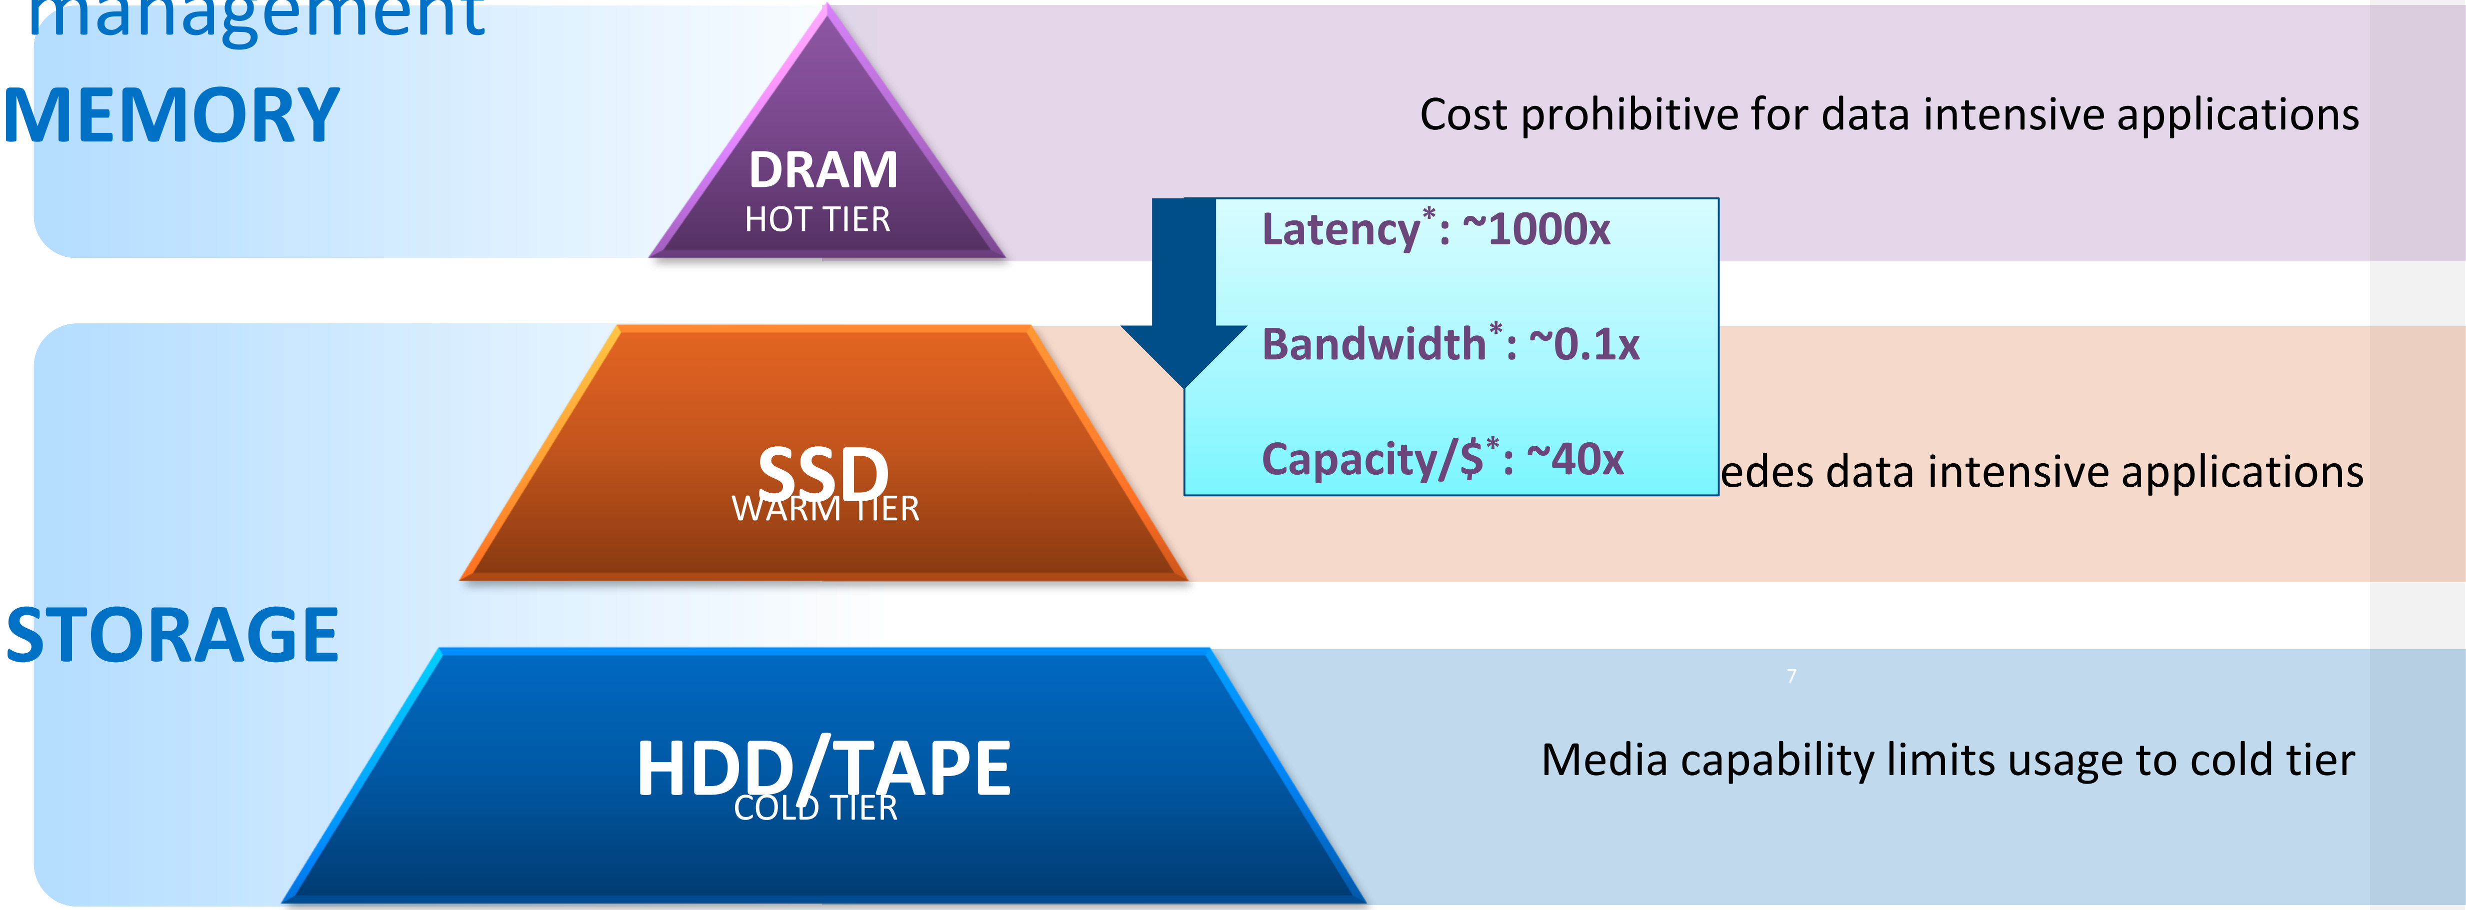

* Actual performance and price may vary

Intel Innovations Address These Gaps

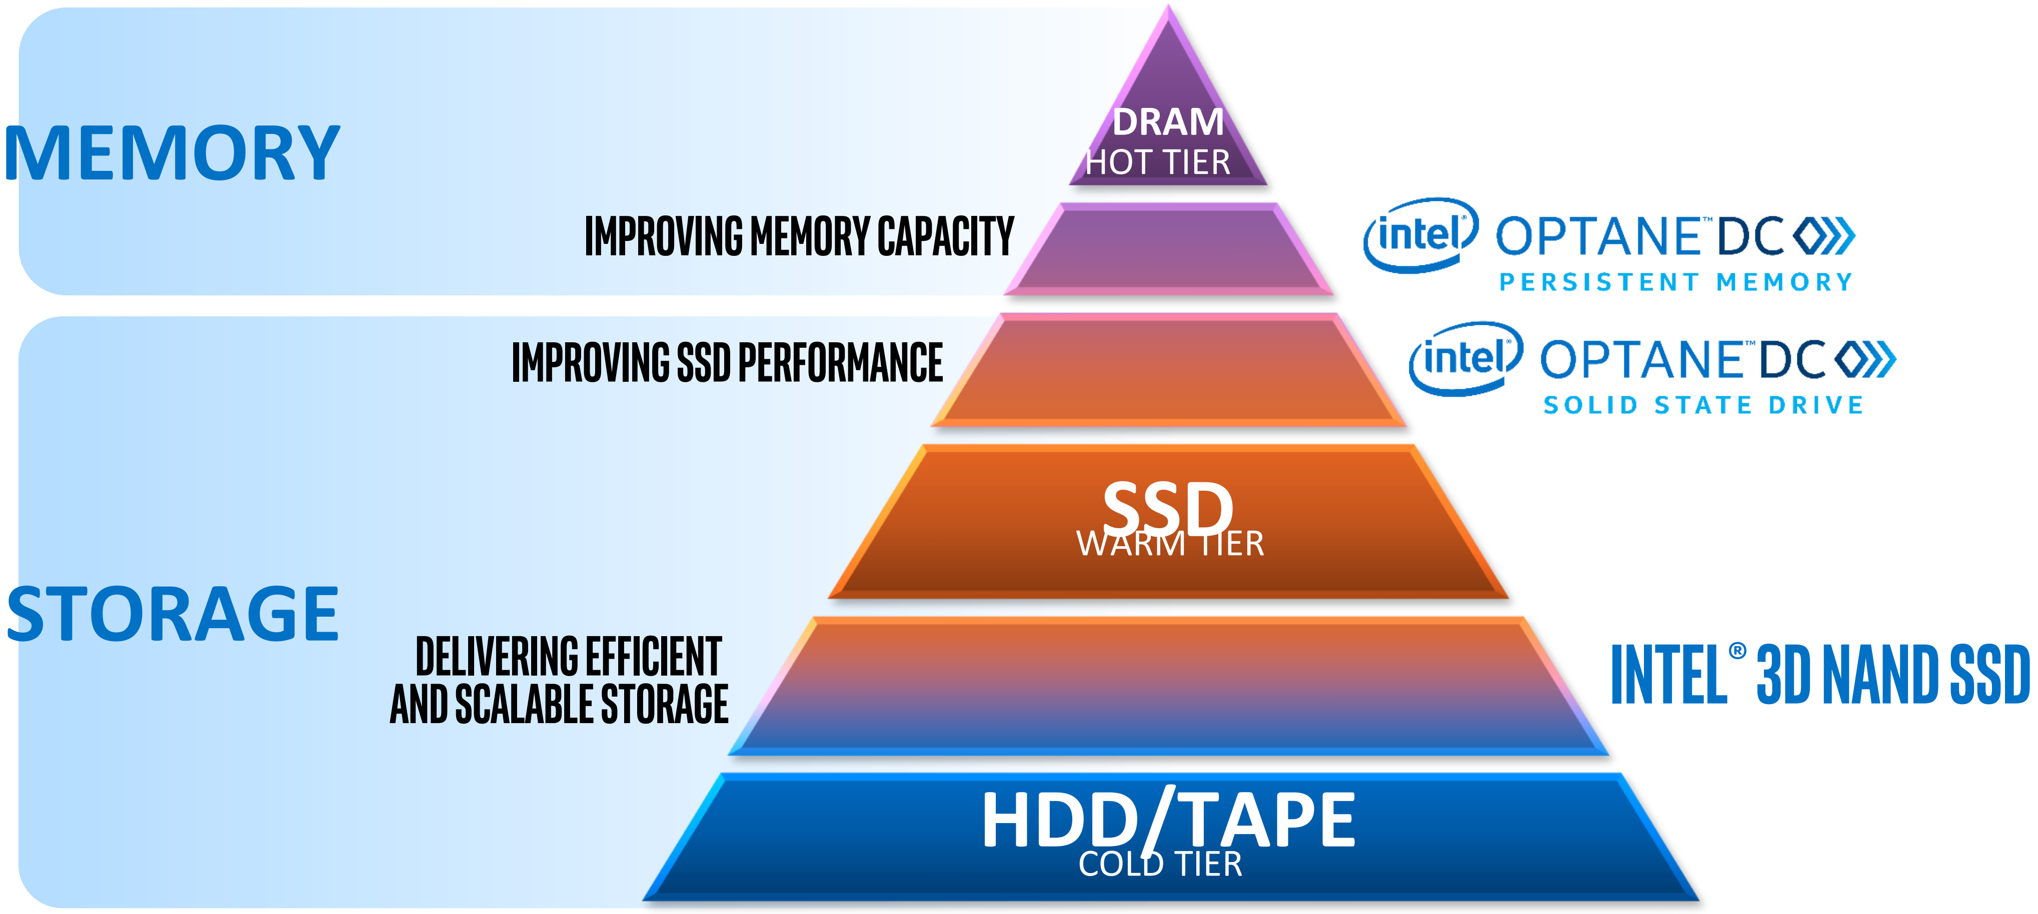

Summary Of CPU Changes

- Core changes
  - Core frequency improvement through speed path fixes and process improvements
  - Core performance updates: LSD, ICCP licensing, AVX512 scatter livelock, PAUSE counter
- Memory controller changes
  - Support for DDR4-2933 1 DPC and 16Gb devices
  - Scheduler improvements to reduce loaded latency
  - Support for Intel® Optane™ DC persistent memory DIMMs
- Uncore changes
  - Latency improvements for select flows in 2S and 4S configurations
  - XPT Prefetch BIOS knob exposed
- Functional improvements and updates
  - Updates to Resource Director Technology (RDT) components – CMT, MBM
  - Updates to PCIe hot-plug (s5353435, s5354002) – LED indicator issue with hot-remove on Linux and surprise hot-add/remove on Win2012 R2

Cascade Lake AP

2 CLX dies per CPU

Intel® Xeon® Platinum 9200 Processor Overview

- Intel® Xeon® Platinum 9200 Processors consist of two die in a BGA package
  - Multi-chip processor with single hop latency from any die to memory in a 2S system
- Key IO/mem features
  - 12 ch DDR4 2933 MT/s per CPU (6 per die)
  - 4 UPI x20 wide at 10.4GTs per CPU
  - x80 PCIe G3 lanes per 2S Node in Intel® Server Systems

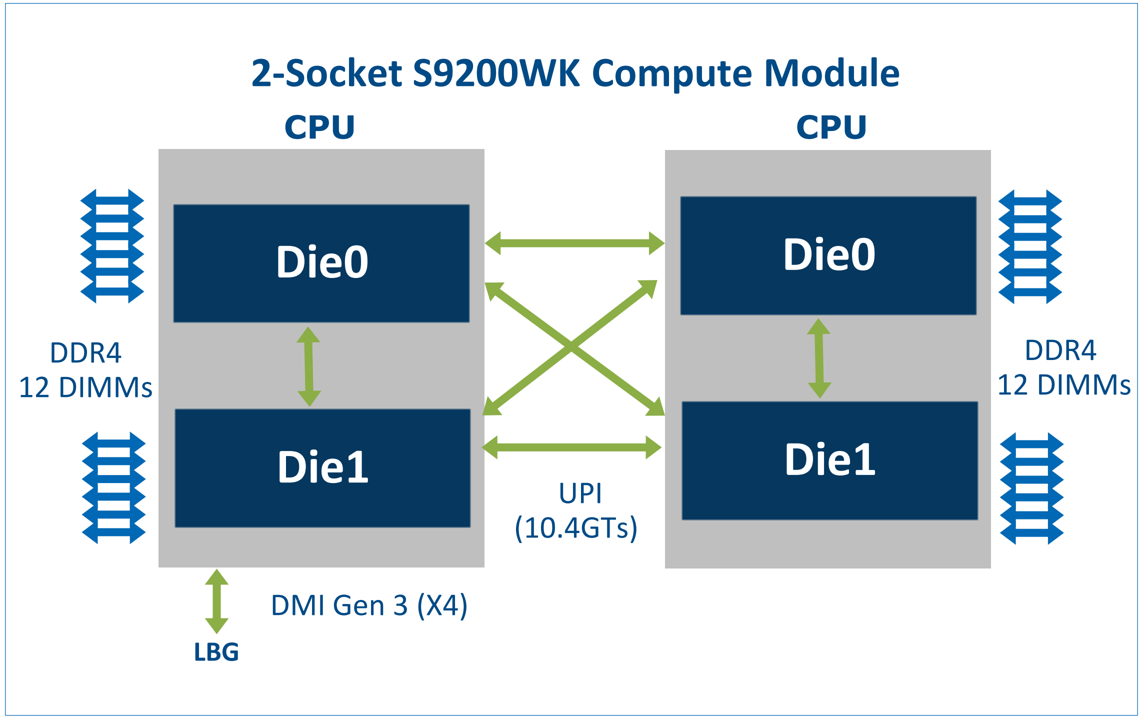

Delivering 4S Performance in 2S Form Factor

Intel® Xeon® Platinum 9200 Processors SKU Stack

SKU Description	Active Cores	Cooling	TDP (W)	Cache (MB)	(P1) TDP Freq (GHz)	All Core Turbo Freq (GHz)	AVX512 Freq (GHz)	AVX512 All Core Turbo Freq (GHz)	DDR4 1DPC (MHz)
9282	56	Liquid	400	77	2.6	3.4	1.6	2.6	2933
9242	48	Liquid/Air	350	71.5	2.3	3.1	1.6	2.3	2933
9222	32	Liquid	250	71.5	2.3	3.0	1.5	2.8	2933
9221	32	Liquid/Air	250	71.5	2.3	3.0	1.5	2.7	2933

12

Intel® Server System S9200WK for HPC & AI

▪ Providing Increased Performance

- Intel® Xeon® Platinum 9200 processors
  - 56, 48, or 32 cores, 12 memory channels
- Intel® Server System S9200WK Data Center Block
  - 2U/4N or 2U/2N, liquid or air cooled, configurable options

LEADERSHIP PERFORMANCE
FOR HPC & AI

"… average of **31% higher performance** than 2S AMD EPYC "Rome" 7742 (64C)"[1]

▪ Simplifying Solutions

- Fully validated, unbranded server systems include Intel's latest data center technology
- Data Center Block can be configured to support a wide range of memory, storage, I/O, and cooling options

[1] See backup for configuration details. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.
(2nd Gen Intel® Xeon® Scalable Processors - claim #31).

CLX-AP Performance

Intel® Xeon® Platinum 9200 Processors Performance

2S Linpack

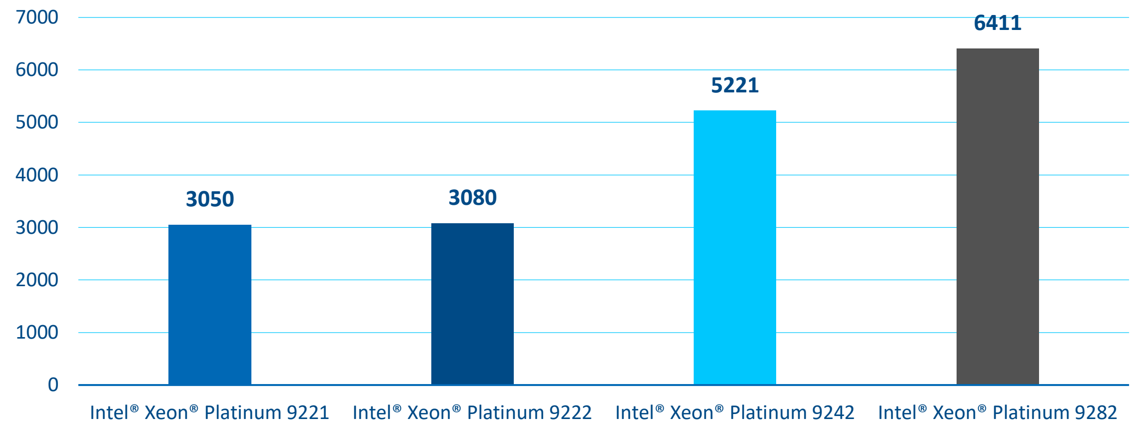

2S Memory Bandwidth

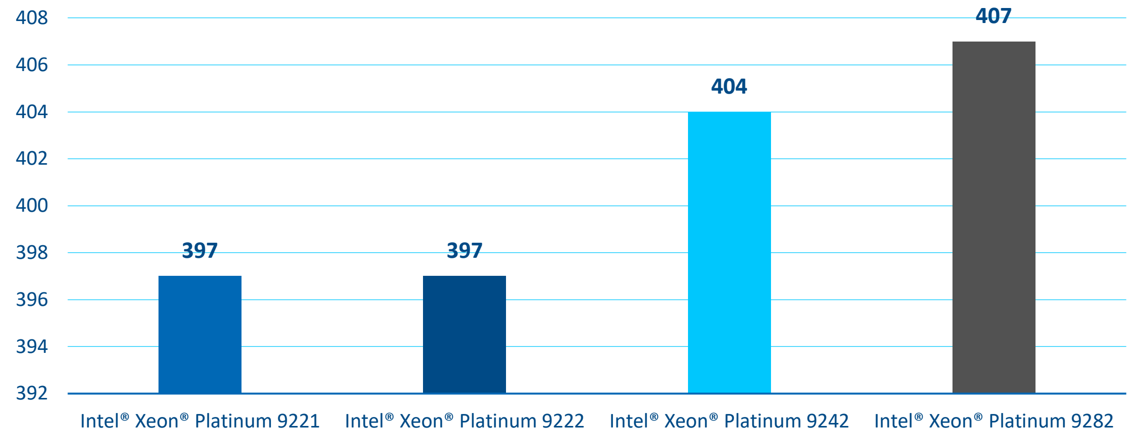

2S Integer Throughput

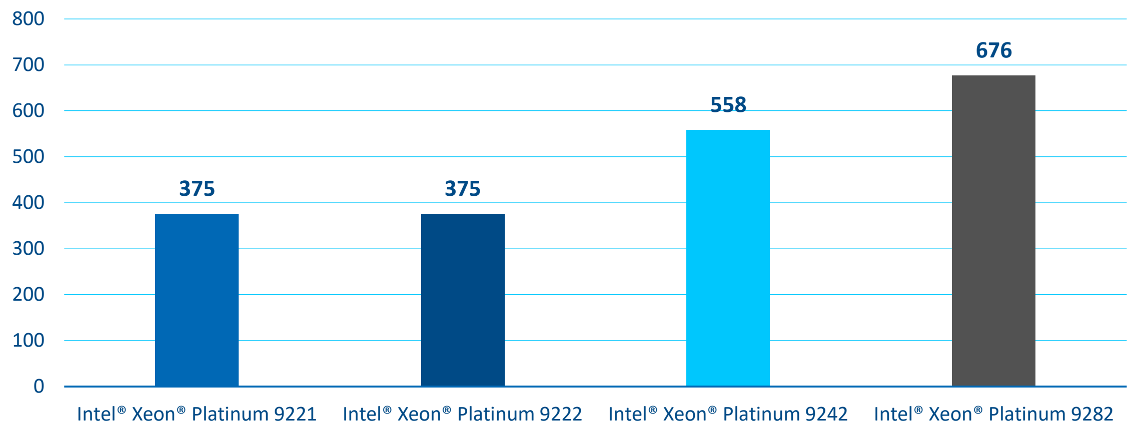

2S Floating Point Throughput

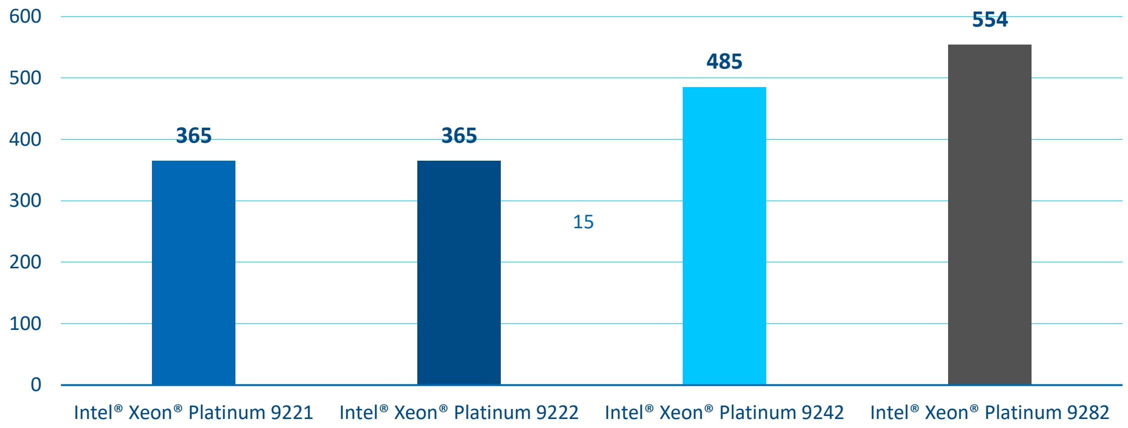

See backup for configuration details. – For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

9222 & 9221 performance values are estimated. See config on slide 23 for 9242 and 9282 configuration details

HPC Workload Description & Sensitivities

√ signifies approximate 'first order' sensitivity

Application	Application Description	Benefits from AVX-512	Compute Bound	Memory Bound
STREAM (Triad)	Benchmark that measures sustainable memory bandwidth (in MB/s)	√		√
HPCG	High Performance Conjugate Gradients (HPCG) Benchmark			√
HPL	High Performance Linpack (HPL) Benchmark	√	√	
OpenFOAM	Open source software for Computational Fluid Dynamics (CFD)			√
ANSYS® Fluent®	General purpose Computational Fluid Dynamics (CFD) and Multiphysics solver			√
LS-Dyna	General-purpose finite element program		√	
VASP	Vienna Ab initio Simulation Package (VASP) used for atomic scale materials modeling	√		√
NAMD	Parallel molecular dynamics code for simulation of large biomolecular systems	√	√	
GROMACS	GROningen MAchine for Chemical Simulations used for Molecular Dynamics	√	√	
LAMMPS	Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS)	√ 16	√	
WRF	The Weather Research and Forecasting (WRF) Model is a weather prediction system			√
FSI Binomial	Binomial Options Pricing models the price of an European call option	√	√	
FSI Black Scholes	Black Scholes is a mathematical model used for European option valuation	√	√	
FSI Monte Carlo	Monte Carlo algorithms are used to calculate the value of an option	√	√	

HPC Performance Leadership

31% Higher Performance with 2S Intel Xeon-AP vs 2S AMD "Rome" 7742

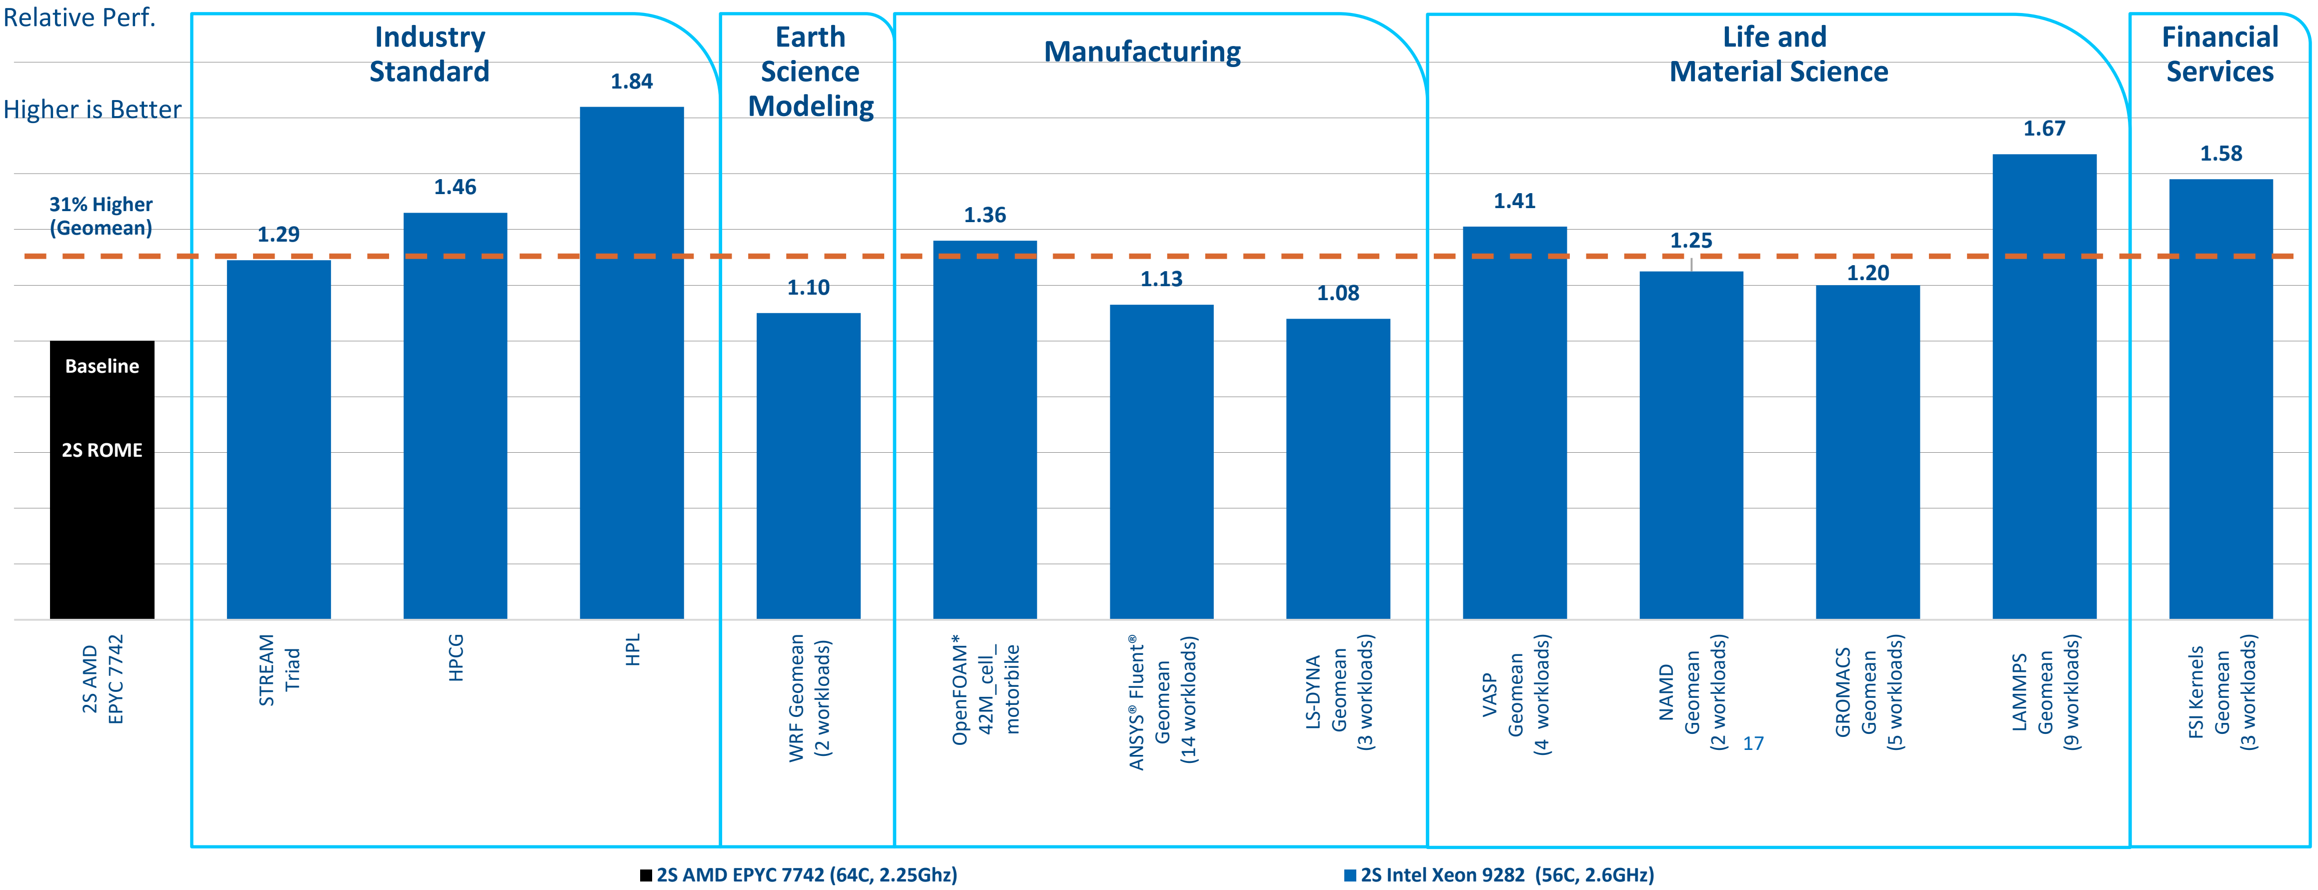

See backup for configuration details. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.
(2nd Intel® Xeon® Scalable Processors - claim #31).
* This offering is not approved or endorsed by OpenCFD Limited, producer and distributor of the OpenFOAM software via www.openfoam.com, and owner of the OPENFOAM® and OpenCFD® trade mark

Continued Innovation Driving DL Gains on Intel® Xeon® Processors

Intel® Optimizations for Caffe ResNet-50 Inference Throughput Performance

BASE

SKX launch

5.7x[1]
vs. BASE

14.5x[1]
vs. BASE

30x[1]
vs. BASE

2S Intel® Xeon® Platinum 8180 processor (28 cores/S)

2S Intel® Xeon® Platinum 8280 processor (28 cores/S)

2S Intel® Xeon® Platinum 9282 processor (56 cores/S)

Intel® Xeon® Scalable Processor (Sky Lake)

2nd Generation Intel® Xeon® Scalable Processor (Cascade Lake)

[1] Based on Intel internal testing: 1X,5.7x,14.5x and 30x performance improvement based on Intel® Optimization for Café Resnet-50 inference throughput performance on Intel® Xeon® Scalable Processor. See Configuration Details 3

Performance results are based on testing as of 7/11/2017(1x) ,11/8/2018 (5.7x), 2/20/2019 (14.5x) and 2/26/2019 (30x) and may not reflect all publically available security updates. No product can be absolutely secure. See configuration disclosure for details. ,

Performance Considerations on CLX-AP

- Will my MPI or OpenMP or MPI/OpenMP hybrid application scale?
  - Consider communication vs. computation ratio
  - Hybrid: use more threading and less MPI
- NUMA effects for OpenMP or hybrid applications?
  Possible solutions:
  - Hybrid: 1 (or more) MPI ranks per die (not per socket!)
  - Pure OpenMP: Distributed initialization of shared memory (Linux first touch)

Tools

Checking the node layout

Basic tools

- Numactl
  - Control NUMA policy for processes or shared memory
  - `numactl --hardware` Display inventory of available nodes on the system
- Hwloc
  - Obtain the hierarchical map of key computing elements, such as: NUMA memory nodes, shared caches, processor sockets, processor cores, and processor "threads"
  - `hwloc-distances` Displays distance matrices
  - `hwloc-info` Show the principle topology of the system
  - `hwloc-ls` Show the detailed topology of the system
- Cpuinfo
  - System overview from Intel MPI

Intel Memory Latency Checker (MLC)

- MLC measures memory latencies and bandwidth
  - Establish a baseline for the system and for performance analysis
  - Default measurements (command line arguments for fine grained control):
    1. Matrix of idle memory latencies from NUMA node to NUMA node
    2. Peak memory b/w measured with varying amounts of reads and writes
    3. Matrix of memory b/w values from NUMA node to NUMA node
    4. Latencies at different b/w points
    5. Cache-to-cache data transfer latencies
- https://software.intel.com/content/www/us/en/develop/articles/intelr-memory-latency-checker.html

Demo

Using the tools to check the node layout

Notices & Disclaimers

Intel technologies may require enabled hardware, software or service activation. Learn more at intel.com or from the OEM or retailer.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

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. https://software.intel.com/en-us/articles/optimization-notice

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. See backup for configuration details. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for details. No product or component can be absolutely secure.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Configs: 2nd Gen Intel® Xeon® Scalable Family
Delivering Performance and Customer Choice Across the Stack

- 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. 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 visit www.intel.com/benchmarks.

- 1. Platinum 92xx vs Platinum 8180: 1-node, 2x Intel® Xeon® Platinum 9282 cpu on Walker Pass with 768 GB (24x 32GB 2933) total memory, ucode 0x400000A on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on all (off Stream, Linpack), result: est int throughput=635, est fp throughput=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 cpu on Wolf Pass with 384 GB (12 X 32GB 2666) total memory, ucode 0x200004D on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on all (off Stream, Linpack), result: est int throughput=307, est fp throughput=251, Stream Triad=204, Linpack=3238, server side java=165724, test by Intel on 1/29/2019. 1-node, 2x Intel® Xeon® Platinum 9242 cpu on Walker Pass with 288 GB (18 X 16GB 2933) total memory, ucode 0x400000A on Linux-4.15.0-45-generic-x86_64-with-Ubuntu-18.04-bionic, , gcc (Ubuntu 7.3.0-27ubuntu1~18.04) 7.3.0, HT on, Turbo on, AIXPRT, OpenVINO R5 (DLDTK Version :1.0.19154) on ResNet 50 using int8 throughput, BS 2, 96 instances, score=4930, test by Intel on 2/20/2019. 1-node, 2x Intel® Xeon® Platinum 8180 cpu on Wolf Pass with 192 GB (12 X 16GB 2666) total memory, ucode 0x200004D on Linux-4.15.0-29-generic-x86_64-with-Ubuntu-18.04-bionic, , gcc (Ubuntu 7.3.0-27ubuntu1~18.04) 7.3.0, HT on, Turbo on, result: AIXPRT, OpenVINO R5 (DLDTK Version :1.0.19154) on ResNet 50 using int8 throughput, BS 4, 28 instances=1419, test by Intel on 2/21/2019.

Configs: 2nd Gen Intel® Xeon® Scalable Family
Delivering Performance and Customer Choice Across the Stack

- 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. 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 visit www.intel.com/benchmarks.
- **1. >2x average generational gains on 2-socket servers with 2nd Gen Intel® Xeon® Platinum 9200 processor**. Geomean of est SPECrate2017_int_base, est SPECrate2017_fp_base, Stream Triad, Intel Distribution of Linpack, server side Java, est AIXPRT OpenVINO/ResNet 50. **Platinum 92xx vs Platinum 8180**: 1-node, 2x Intel® Xeon® Platinum 9282 cpu on Walker Pass with 768 GB (24x 32GB 2933) total memory, ucode 0x400000A on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on all (off Stream, Linpack), result: est int throughput=635, est fp throughput=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 cpu on Wolf Pass with 384 GB (12 X 32GB 2666) total memory, ucode 0x200004D on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on all (off Stream, Linpack), result: est int throughput=307, est fp throughput=251, Stream Triad=204, Linpack=3238, server side java=165724, test by Intel on 1/29/2019. 1-node, 2x Intel® Xeon® Platinum 9242 cpu on Walker Pass with 288 GB (18 X 16GB 2933) total memory, ucode 0x400000A on Linux-4.15.0-45-generic-x86_64-with-Ubuntu-18.04-bionic, , gcc (Ubuntu 7.3.0-27ubuntu1~18.04) 7.3.0, HT on, Turbo on, AIXPRT, OpenVINO R5 (DLDTK Version :1.0.19154) on ResNet 50 using int8 throughput, BS 2, 96 instances, score=4930, test by Intel on 2/20/2019. 1-node, 2x Intel® Xeon® Platinum 8180 cpu on Wolf Pass with 192 GB (12 X 16GB 2666) total memory, ucode 0x200004D on Linux-4.15.0-29-generic-x86_64-with-Ubuntu-18.04-bionic, , gcc (Ubuntu 7.3.0-27ubuntu1~18.04) 7.3.0, HT on, Turbo on, result: AIXPRT, OpenVINO R5 (DLDTK Version :1.0.19154) on ResNet 50 using int8 throughput, BS 4, 28 instances=1419, test by Intel on 2/21/2019.
- **2. Platinum 8280 vs Platinum 8180**: 1-node, 2x Intel® Xeon® Platinum 8280M cpu on Wolf Pass with 384 GB (12 X 32GB 2933) total memory, ucode 0x400000A on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on, result: est int throughput=317, est fp throughput=264, Stream Triad=217, Linpack=3462, server side java=177561, test by Intel on 1/30/2019. 1-node, 2x Intel® Xeon® Platinum 8180 cpu on Wolf Pass with 384 GB (12 X 32GB 2666) total memory, ucode 0x200004D on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on, result: est int throughput=307, est fp throughput=251, Stream Triad=204, Linpack=3238, server side java=165724, test by Intel on 1/29/2019. 1-node, 2x Intel® Xeon® Platinum 8280M cpu on Wolf Pass with 384 GB (12 X 32GB 2933) total memory, ucode 0x400000A on Linux-4.15.0-43-generic-x86_64-with-debian-buster-sid, , gcc (Ubuntu 7.3.0-27ubuntu1~18.04) 7.3.0, HT on, Turbo on, result: AIXPRT, OpenVINO R5 (DLDTK Version :1.0.19154) on ResNet 50 using int8 throughput, BS 4, 32 instances=3266, test by Intel on 2/21/2019. 1-node, 2x Intel® Xeon® Platinum 8180 cpu on Wolf Pass with 192 GB (12 X 16GB 2666) total memory, ucode 0x200004D on Linux-4.15.0-29-generic-x86_64-with-Ubuntu-18.04-bionic, gcc (Ubuntu 7.3.0-27ubuntu1~18.04) 7.3.0, HT on, Turbo on, result: AIXPRT, OpenVINO R5 (DLDTK Version :1.0.19154) on ResNet 50 using int8 throughput, BS 4, 32 instances=1419, test by Intel on 2/21/2019.
- 1-node, 2x **Intel® Xeon® Platinum 9242** cpu on Walker Pass with 384 GB (24 slots / 16GB / 2933) total memory, ucode 0x4000010 on Red Hat 7.6, 3.10.0-957.el7.x86_64, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on, result: est int throughput=543, est fp throughput=483, Stream Triad=404, Linpack=5221, server side java=292765, test by Intel on 1/24/2019. 1-node, 2x Intel® Xeon® Platinum 9242 cpu on Walker Pass with 288 GB (18 X 16GB 2933) total memory, ucode 0x400000A on Linux-4.15.0-45-generic-x86_64-with-Ubuntu-18.04-bionic, , gcc (Ubuntu 7.3.0-27ubuntu1~18.04) 7.3.0, HT on, Turbo on, AIXPRT, OpenVINO R5 (DLDTK Version :1.0.19154) on ResNet 50 using int8 throughput, BS 2, 96 instances, score=4930, test by Intel on 2/20/2019.

Workload and Configuration Details

- **31% Higher Performance with 2S Intel Xeon-AP vs 2S AMD* EPYC* "Rome" 7742:** Intel measured as of October 8, 2019 using geomean of STREAM Triad, HPCG, HPL, WRF (2 workloads), OpenFOAM 42M_cell_motorbike, ANSYS® (14 workloads), LS-DYNA (3 workloads), VASP (4 workloads), NAMD (2 workloads), GROMACS (9 workloads), LAMMPS (9 workloads), FSI Kernels (3 workloads).
- Intel Xeon 9282 Processor configuration: Intel "Walker Pass" S9200WKL platform with 2-socket 9282 Intel Xeon processors (2.6GHz, 56C), 24x16GB DDR4-2933, 1 SSD, BIOS: SE5C620.86B.2X.01.0053, Microcode: 0x5000029, Red Hat Enterprise Linux 7.7, kernel 3.10.0-1062.1.1
- AMD EPYC 7742 Processor configuration: Supermicro AS-2023-TR4 (HD11DSU-iN) with 2-socket 7742 AMD EPYC "Rome" processors (2.25GHz, 64C), 16x32GB DDR4-3200, 1 SSD, BIOS: 2.0 CPLD 02.B1.01, Microcode: 830101C, CentOS Linux release 7.7.1908, kernel 3.10.0-1062.1.1.el7.crt1.x86_64
- STREAM OMP 5.1 Triad: Intel Xeon 9282: Intel Compiler 2019u5, BIOS: HT ON , Turbo ON, SNC ON, 1 thread/core; AMD EPYC 7242: Intel Complier 2019u5, BIOS: SMT ON, Boost ON, NPS 4, 1 thread/core
- HPCG Intel optimized version: Intel Xeon 9282: Intel Compiler 2019u4, Intel MKL 2019u4, Intel MPI 2019u4, BIOS: HT ON, Turbo OFF, SNC OFF, 1 thread/core; AMD EPYC 7742: Intel Compiler 2019u4, Intel MKL 2019u4, Intel MPI 2019u4, BIOS: SMT ON, Boost ON OFF, NPS 4, 1 thread/core
- HPL v2.3: Intel Xeon 9282: Intel Optimized Linpack Benchmark, Intel Distribution for LINPACK Benchmark, Compiler: Intel MPI 2018u1N=80000, NB=384, P=2, Q=1, BIOS: HT ON, Turbo ON, SNC OFF, 1 thread/core; AMD EPYC 7742: AMD official HPL binary https://developer.amd.com/amd-aocl/blas-library/, Compiler: Netlib HPL + BLIS, OpenMPI3 N=16000, NB=192, P=2, Q=4; BIOS: SMT ON, Boost ON, NPS 4, 1 thread/core
- WRF 3.9.1.1: Geomean (2 workloads: conus-12km, conus-2.5km): Intel Xeon 9282: Intel Compiler 2018u3, Intel MPI 2018u3, AVX2 build, BIOS: HT ON, Turbo ON, SNC OFF, 1 thread/core; AMD EPYC 7742: Intel Compiler 2018u3, Intel MPI 2018u3, AVX2 build, BIOS: SMT ON, Boost ON, NPS 4, 1 thread/core
- OpenFOAM v6.0 42M_cell_motorbike: Intel Xeon 9282: Intel Compiler 2019u3, Intel MPI 2019u3, BIOS: HT ON, Turbo ON, SNC OFF, 1 thread/core; AMD EPYC 7742: Intel Compiler 2019u3, Intel MPI 2019u3, BIOS: SMT ON, Boost ON, NPS 4, 1 thread/core.
- ANSYS® Fluent® 2019R1: Geomean (14 workloads: aircraft_wing_14m, aircraft_wing_2m, combustor_12m, combustor_16m, combustor_71m, exhaust_system_33m, f1_racecar_140m, fluidized_bed_2m, ice_2m, landing_gear_15m, oil_rig_7m, pump_2m, rotor_3m, sedan_4m): Intel Xeon 9282: Intel Compiler 2017u3, Intel MPI 2018u3, BIOS: HT ON, Turbo ON, SNC ON, 1 thread/core; AMD EPYC 7742: Intel Compiler 2017u3, Intel MPI 2018u3, BIOS: SMT ON, Boost ON, NPS 4, 1 thread/core
- LS-DYNA v9.3: Geomean (3 workloads: 3cars/150ms, car2car/120ms, ODB_10M/30ms): Intel Xeon 9282: Intel Compiler 2016u3, Intel MPI 2018u3, AVX2 build, BIOS: HT OFF, Turbo ON, SNC ON, 1 thread per core; AMD EPYC 7742: Intel Compiler 2016u3, Intel MPI 2018u3, AVX2 build, BIOS: SMT OFF, Boost ON, NPS 4, 1 thread/core
- VASP, developer branch based on v5.4.4: Geomean (4 workloads: CuC, PdO4, PdO4_K221, Si): Intel Xeon 9282: Intel Compiler 2019u4, Intel MKL 2019u4, Intel MPI 2019u4, BIOS: HT ON, Turbo OFF, SNC OFF, 1 thread per core; AMD EPYC 7742: Intel Compiler 2019u4, Intel MKL 2019u4, Intel MPI 2019u4, BIOS: SMT ON, Boost ON, NPS 4, 1 thread per core
- NAMD v2.13: Geomean (2 workloads: Apoa1, STMV): Intel Xeon 9282: Intel Compiler 2019u4, Intel MPI 2019u4, BIOS: HT ON, Turbo ON, SNC OFF, 2 threads per core; AMD EPYC 7742: Compiler: AOCC 2.0, Intel MPI 2019u4, BIOS: SMT ON, Boost ON, NPS 4, 2 threads/core
- GROMACS 2019.3: Geomean (5 workloads: archer2_small, ion_channel_pme, lignocellulose_rf, water_pme, water_rf): Intel Xeon 9282: Intel Compiler 2019u4, Intel MKL 2019u4, Intel MPI 2019u4, AVX-512 build, BIOS: HT ON, Turbo OFF, SNC OFF, 2 threads per core; AMD EPYC 7742: Intel Compiler 2019u4, Intel MKL 2019u4, Intel MPI 2019u4, AVX2 build, BIOS: SMT ON, Boost ON, NPS 4, 1 threads per core
- LAMMPS v2019: Geomean (9 workloads: Atomic Fluid, Copper, DPD, Liquid Crystal, Polyethylene, Protein, Stillinger-Weber, Tersoff, Water): Intel Xeon 9282: Intel Compiler 2019u5, BIOS: HT ON, Turbo ON, SNC ON, 2 threads/core; AMD EPYC 7742: Compiler: AOCC 2.0, Intel MPI 2019u5, BIOS: SMT ON, Boost ON, NPS 4, 2 threads/core
- FSI Kernels v2.0: Geomean (3 workloads: Binomial Options, Black Scholes, Monte Carlo): Intel Xeon 9282: Intel Compiler 2019u5, Intel MKL 2019u5, BIOS: HT ON, Turbo ON, SNC OFF, 2 threads/core, HT OFF, Turbo ON, SNC OFF, 1 threads/core, HT ON, Turbo ON, SNC OFF, 2 threads/core; AMD EPYC 7742: Intel Compiler 2019u5, Intel MKL 2019u5, BIOS: SMT ON, Boost ON, NPS 4, 2 threads/core, SMT OFF, Boost ON, NPS 4, 1 thread/core, SMT ON, Boost ON, NPS 4, 2 threads/core

Configurations: Continued Innovation Driving DL Gains on Xeon®

- **1x inference throughput improvement in July 2017 (baseline):** Tested by Intel as of July 11th 2017: Platform: 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: (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 (ResNet-50),and https://github.com/soumith/convnet-benchmarks/tree/master/caffe/imagenet_winners (ConvNet benchmarks; files were updated to use newer Caffe prototxt format but are functionally equivalent). Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with “numactl -l“.
- **5.7x inference throughput improvement in December 2018 vs baseline:** Tested by Intel as of November 11th 2018 :2 socket Intel® Xeon® Platinum 8180 CPU @ 2.50GHz / 28 cores HT ON , Turbo ON Total Memory 376.46GB (12slots / 32 GB / 2666 MHz). CentOS Linux-7.3.1611-Core, kernel: 3.10.0-862.3.3.el7.x86_64, SSD sda RS3WC080 HDD 744.1GB,sdb RS3WC080 HDD 1.5TB,sdc RS3WC080 HDD 5.5TB , Deep Learning Framework Intel® Optimization for caffe version: 551a53d63a6183c233abaa1a19458a25b672ad41 Topology::ResNet_50_v1 BIOS:SE5C620.86B.00.01.0014.070920180847 MKLDNN: 4e333787e0d66a1dca1218e99a891d493dbc8ef1 instances: 2 instances socket:2 (Results on Intel® Xeon® Scalable Processor were measured running multiple instances of the framework. Methodology described here: https://software.intel.com/en-us/articles/boosting-deep-learning-training-inference-performance-on-xeon-and-xeon-phi) NoDataLayer. Datatype: INT8 Batchsize=64 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: (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 (ResNet-50). Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with “numactl -l“.
- **14.5x inference throughput improvement vs baseline:** 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, BS=64, DummyData, 4 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: (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 (ResNet-50),. Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with “numactl -l“.
- **30x inference throughput improvement with CascadeLake-AP vs baseline:** 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: 830a10059a018cd2634d94195140cf2d8790a75a), model: 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: (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 (ResNet-50),. Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with “numactl -l“.

intel®