THE FUTURE  UNLEASHED

Intel® Xeon processor Scalable Family: Performance in HPC, AI and other key segments

Dr. Jean-Laurent PHILIPPE, Senior EMEA HPC Technical Specialist
Intel DataCenter Group

NOTICES AND DISCLAIMERS

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as “Spectre” and “Meltdown.” Implementation of these updates may make these results inapplicable to your device or system.

No computer system 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 information go to www.intel.com/benchmarks.

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

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, the Intel logo, the Intel. Experience What’s Inside logo, Intel. Experience What’s Inside, Intel Inside, the Intel Inside logo, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

HPC IS THE FOUNDATION OF INSIGHT

SCIENTIFIC DISCOVERY

Advancing science
and our understanding
of the universe

BUSINESS INNOVATION

High ROI: $672
Average return per $1
of HPC Investment[1]
*30% increase in ROI in past 2 years

ARTIFICIAL INTELLIGENCE

Machine learning
Combines computational
learning theory with HPC

1. Source: IDC HPC and ROI Study Update (September 2017) http://www.hpcuserforum.com/ROI/

INTEL® XEON® SCALABLE PROCESSORS ADDRESS HPC CHALLENGES

Intel® Xeon® Scalable processors

Improve performance

Reduce energy consumption

Lower TCO

Better solutions

Decrease product development time

Improve existing products, reduce cost, add new features

Create new and more competitive products

Smarter strategy

Make better informed decisions, faster

Faster results

Accelerate time to solutions and to market

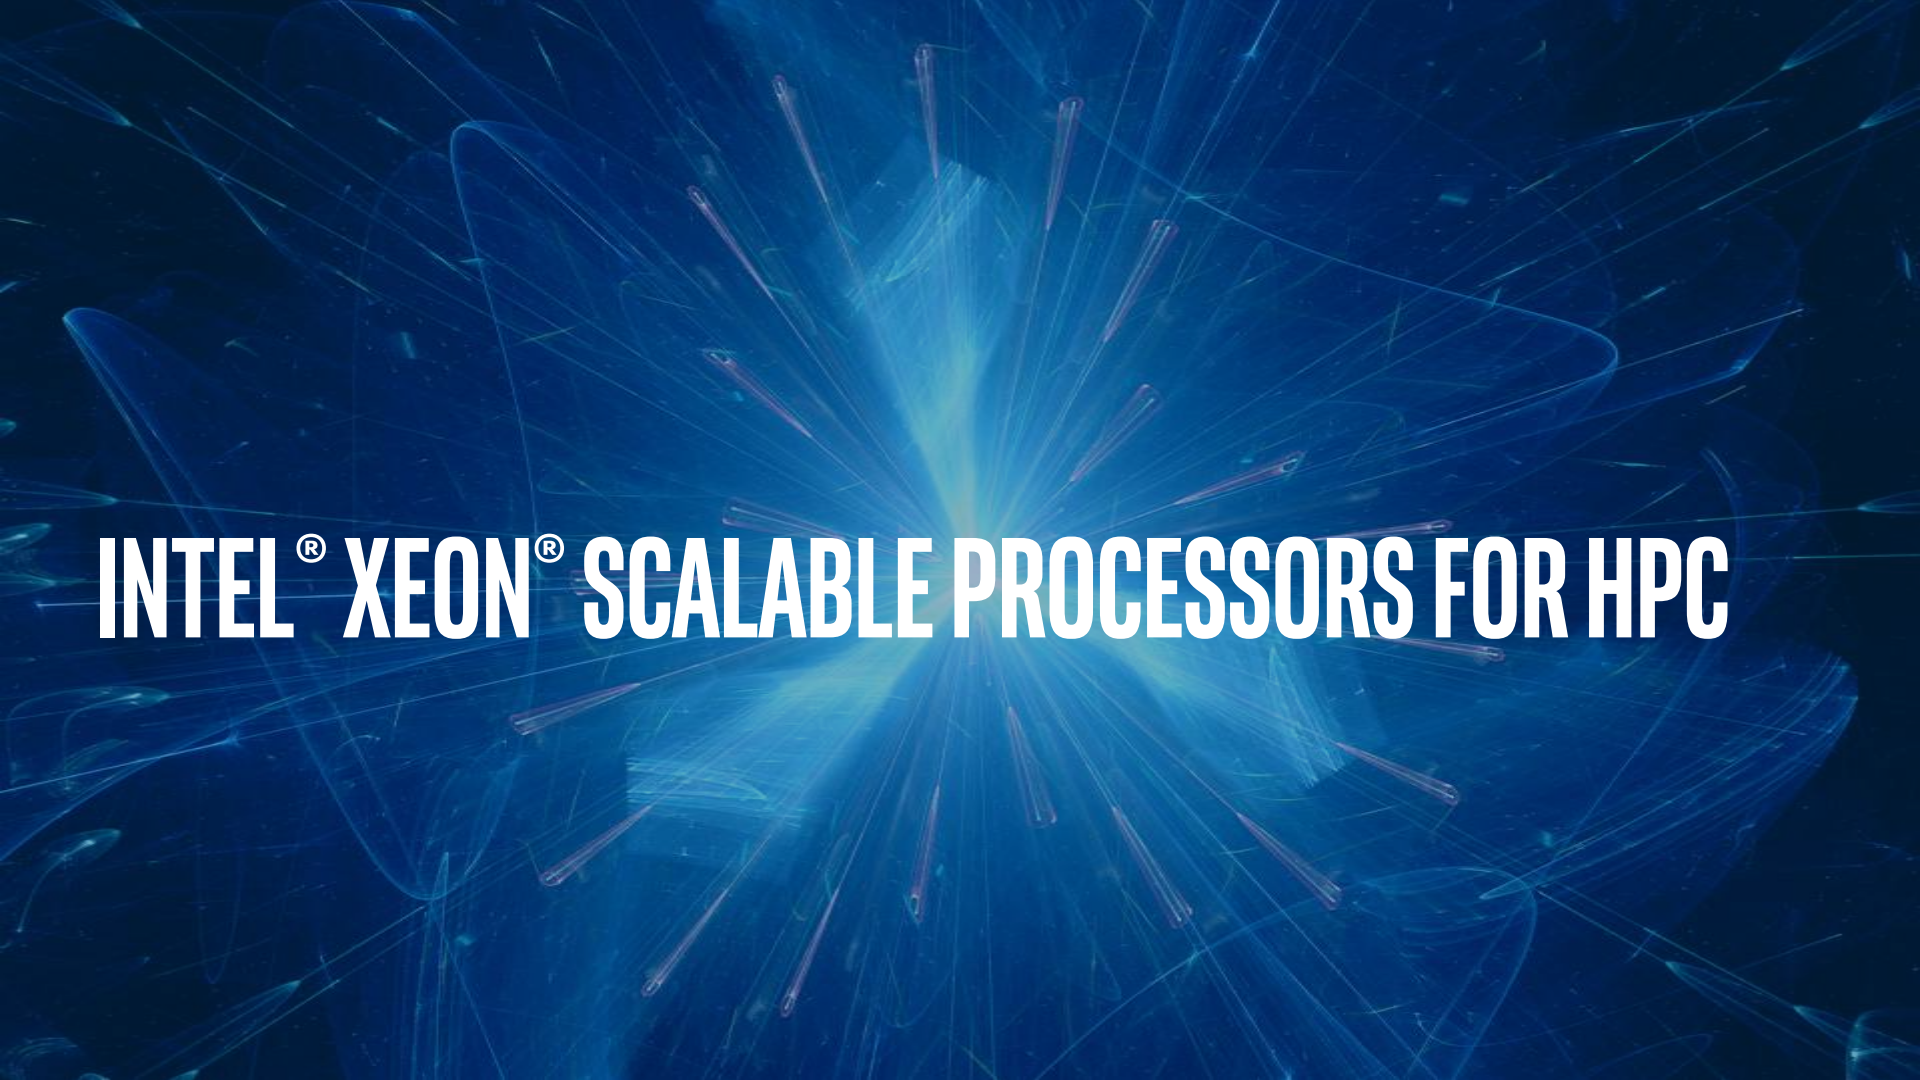
INTEL® XEON® SCALABLE PROCESSORS FOR HPC

INTEL® XEON® SCALABLE PROCESSORS FOR HPC

PERFORMANCE

Delivers HPC results faster

FLEXIBILITY

Ability to run all types of workloads on one infrastructure

ENERGY EFFICIENCY

New design reduces cost of running the HPC cluster

1.63X AVERAGE PERFORMANCE BOOST

over prior generation[1] on HPC applications

1. For disclaimers and configuration, see following slides.
2. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks
3. Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system.

CONFIGURATIONS: TECHNICAL COMPUTE WORKLOADS 1/2

1.) Up to 1.63x gains for HPC workloads: based on Geomean of Weather Research Forecasting - Conus 12Km, HOMME, LSTCLS-DYNA Explicit, INTES PERMAS V16, MILC, GROMACS water 1.5M_pme, VASPSi256, NAMDstmv, LAMMPS, Amber GB Nucleosome, Binomial option pricing, Black-Scholes, Monte Carlo European options

a) **PERMAS by INTES** is an advanced Finite Element software system that offers a complete range of physical models at high performance, quality, and reliability. It plays a mission-critical role in the design process at customers from automotive, ship design, aerospace, and more. E5-2697 v4: 2S Intel® Xeon® processor E5-2697v4, 2.3GHz, 18 cores, turbo on, HT off, NUMA on, BIOS 338.R00, 256 GB total memory (8x 32GB w/ 2400 MT/s, DDR4 LRDIMM), 4x Intel® SSD DC P3600 2 TB in RAID 0 (stripe size 64k). CentOS Linux* release 7.2, kernel 3.10.0-327.13.1.el7.x86_64. Intel® Composer 2015.5.223. INTES PERMAS V16.00. Gold 6148: Intel® Xeon® Gold 6148 processor, 2.4 GHz, 20 cores, turbo on, HT off, NUMA on, BIOS SE5C620.86B.01.00.0412.020920172159, 384 GB total memory (12x 32GB w/ 2400 MT/s, DDR4 LRDIMM), 3x Intel® SSD DC P3600 2 TB in RAID 0 (stripe size 64k), CentOS* Linux* release 7.3, kernel 3.10.0-514.10.2.el7.x86_64. Intel® Composer 2015.7.235. INTES PERMAS V16.00.

b) **LS-DYNA** is the leading product in the crash simulation market. It is used by the automobile, aerospace, construction, military, manufacturing, and bioengineering industries in worldwide. Workload: 2M elements Car2car model with 120ms simulation time. LS-DYNA explicit standard benchmarks tested by Intel, March 2017. E5-2697 V4: 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 18 cores, turbo and HT on, BIOS SE5C610.86B.01.01.0016.033120161139, 128GB total memory, 8 memory channels / 8x16GB / 2400 MT/s / DDR4, Red Hat Enterprise Linux* 7.3 kernel 3.10.0-229.20.1.el6.x86_64.knl2. GOLD 6148: 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 20 cores, turbo and HT on, BIOS version 412, 192GB total memory, 12 memory channels / 12x16GB / 2400 MT/s / DDR4, Red Hat Enterprise Linux* 7.3 kernel 3.10.0-514.el7.x86_64.

c) **Binomial option pricing** is a lattice-based approach that uses a discrete-time model of the varying price over time of the underlying financial instrument. This is compute bound, double precision workload. FSI Binomial workload. OS: Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Testing by Intel March 2017. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 slots / 16GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo and HT on, BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

d) **Monte Carlo** is a numerical method that uses statistical sampling techniques to approximate solutions to quantitative problems. In finance, Monte Carlo algorithms are used to evaluate complex instruments, portfolios, and investments. This is compute bound, double precision workload. FSI Monte Carlo workload. OS: Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Testing by Intel March 2017. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 x16GB 2400 MHz DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor@ 2.4GHz, H0QS, 40 cores 150W. QMS1, turbo and HT on, BIOS SE5C620.86B.01.00.0412.020920172159, 192GB total memory, 12 x 16 GB 2666 MHz DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327

e) **Black-Scholes** is a popular mathematical model used in finance for European option valuation. This is a double precision version. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 x16GB 2400 MHz DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor@ 2.4GHz, H0QS, 40 cores 150W. QMS1, turbo and HT on, BIOS SE5C620.86B.01.00.0412.020920172159, 192GB total memory, 12 x 16 GB 2666 MHz DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327

f) **Amber*** is a suite of programs for classical molecular dynamics and statistical analysis. The main MD program is PMEMD (Particle Mesh Ewald Molecular Dynamics) employs two separate algorithms for implicit- and explicit-solvent dynamics. Here performance for explicit solvent (PME) is presented. Amber: Version 16 with all patches applied at December, 2016. Workloads: PME Cellulose NVE(408K atoms), PME stmv(1M atoms), GB Nucleosome (25K), GB Rubisco (75K). No cut-off was used for GB workloads. Compiled with -mic2_spdp –intelmpi - openmp, –DMIC2 * defined. Tests performed on March 2017. E5-2697 v4: Executed with 36 MPI, 2 OpenMP. 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 8x16GB 2400MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Executed with 40 MPI and 2 OpenMP. 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo on, HT on, BIOS 86B.01.00.0412.R00, 12x16GB 2666MHz DDR, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

g) **VASP** CONFIGURATION: The Vienna Ab initio Simulation Package (VASP) is a computer program for atomic scale materials modeling and performs electronic structure calculations and quantum-mechanical molecular dynamics from first principles. VASP provides scientists with fast and precise calculation of materials properties covering wide range of MD methods from DFT, DFT-HF to Random-Phase approximation (GW, ACDFT). Beta VASP, a release candidate for v6.0. Developer branch provided as “Package” included with download: https://github.com/vasp-dev/vasp-knl. AVX512: Intel® Compiler 17.0.1.132, Intel® MPI 2017u1, ELPA 2016.05.004. Optimization Flags: “-O3 -xCORE-AVX512“. AVX2: Intel® Compiler 17.0.1.132, Intel® MPI 2017u1, ELPA 2016.05.004. Optimization Flags: “-O3 -xCORE-AVX2“. E5-2697 v4: 2S Intel® Xeon® processor E5-2697 v4 2.3 GHz , 18 Cores/Socket, 36 Cores, 72 Threads, HT on, turbo off, BIOS 86B0271.R00, 128GB total memory, 2400 MT/s DDR4 RDIMM, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Dual Socket Intel® Xeon® processor Gold 6148 2.4 GHz , 20 Cores/Socket, 40 Cores, 80 Threads, HT on, turbo off, BIOS 86B.01.00.0412, 192GB total memory, 2666 MT/s / DDR4 RDIMM, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

CONFIGURATIONS: TECHNICAL COMPUTE WORKLOADS 2/2

1.) Up to 1.63x gains for HPC workloads: based on Geomean of Weather Research Forecasting - Conus 12Km, HOMME, LSTCLS-DYNA Explicit, INTES PERMAS V16, MILC, GROMACS water 1.5M_pme, VASPSi256, NAMDstmv, LAMMPS, Amber GB Nucleosome, Binomial option pricing, Black-Scholes, Monte Carlo European options

h) **NAMD**: NAMD, recipient of a 2002 Gordon Bell Award, is a parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems. Based on Charm++ parallel objects, NAMD scales to hundreds of cores for typical simulations and beyond 200,000 cores for the largest simulations. Version 2.12 Dec2016. Workloads: apoa1 (92K atoms), stmv(1M atoms). Compiled with –DNAMD_KNL* define. Tests performed on March 2017. E5-2697 v4: 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 8x16GB 2400MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Compiler option "-xCORE-AVX2". Gold 6148: 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo on, HT on, BIOS 86B.01.00.0412.R00, 12x16GB 2666MHz DDR, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Compiler option "-xCORE-AVX512".

i) **LAMMPS**: LAMMPS is a classical molecular dynamics code, and an acronym for Large-scale Atomic/Molecular Massively Parallel Simulator. It is used to simulate the movement of atoms to develop better therapeutics, improve alternative energy devices, develop new materials, and more. E5-2697 v4: 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 36 cores, Intel® Turbo Boost Technology and Intel® Hyperthreading Technology on, BIOS 86B0271.R00, 8x16GB 2400MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, Intel® Turbo Boost Technology and Intel® Hyperthreading Technology on, BIOS 86B.01.00.0412.R00, 12x16GB 2666MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

j) **GROMACS** is a versatile package to perform classical Molecular Dynamics simulations. Heavily optimized for most modern platforms and provides extremely high performance. GROMACS AVX2 CONFIGURATION: Version 2016.3: ftp://ftp.gromacs.org/pub/gromacs/gromacs-2016.3.tar.gz , Intel® Compiler 17.0.1.132, Intel® MPI 2017u1. Optimization Flags: "-O3 -xCORE-AVX2". Cmake options: "-DGMX_FFT_LIBRARY=mkl -DGMX_SIMD=AVX2_256". GROMACS AVX512 CONFIGURATION: Version 2016.3: ftp://ftp.gromacs.org/pub/gromacs/gromacs-2016.3.tar.gz , Intel® Compiler 17.0.1.132, Intel® MPI 2017u1. Optimization Flags: "-O3 -xCORE-AVX512". Cmake options: "-DGMX_FFT_LIBRARY=mkl -DGMX_SIMD=AVX_512". E5-2697 V4: GROMACS AVX2 binary, Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 Cores/Socket, 36 Cores, 72 Threads (HT on, Turbo on), DDR4 128GB, 2400 MHz, Red Hat 7.2. Gold 6148: GROMACS AVX512 binary, Dual Socket Intel® Xeon® processor Gold 6148 2.4 GHz , 20 Cores/Socket, 40 Cores, 80 Threads (HT on, Turbo on), DDR4 192GB, 2666 MT/s DDR4 RDIMMs, Red Hat 7.2.

k) **Weather Research and Forecasting** (WRF) Model is a next-generation mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting needs. It features two dynamical cores, a data assimilation system, and a software architecture facilitating parallel computation and system extensibility. The model serves a wide range of meteorological applications across scales from tens of meters to thousands of kilometers. 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 slots / 16GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA HDD, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Software: WRF version 3.6.1 Compiled using Intel config option with "-O3 -fp-model fast=1 -xCORE-AVX2". Executed with 36 MPI ranks and OMP_NUM_THREADS=1. Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo and HT on, BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Software: WRF version 3.6.1 Compiled using Intel config option with "-O3 -fp-model fast=1 -xCORE-AVX512". Executed with 40 MPI ranks and OMP_NUM_THREADS=1.

l) **HOMME** is the spectral element dynamical core that solves the equations of motion in the CAM-SE atmospheric model, part of the NSF Community Earth System Model (CESM) as well as the related DOE ACME model. CESM is a widely-used Earth system model and an important source of simulations used by the Intergovernmental Panel on Climate Change. HOMME version: https://svn-homme-model.cgd.ucar.edu/branch_tags/dungeon_tags/dungeon016. Compiled with "-O3 -fp-model fast -xCORE-AVX2". Running "perfTestWACCM" benchmark from code repository with size NE=8. Executed with 64 MPI ranks and OMP_NUM_THREADS=1. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 slots / 16GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo and HT on, BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

TECHNOLOGIES DRIVING HPC INNOVATION

More performance, flexibility, and energy efficiency

The **Intel® Xeon® Scalable processor family** is a powerful platform that has been designed to deliver advanced HPC. This platform represents a significant leap forward in the performance and efficiency of cutting-edge HPC systems.

Boost floating point performance

Intel® Advanced Vector Extensions 512 (Intel® AVX-512)

Reduce latency and power

Intel® Mesh Architecture

Intel® Ultra Path Interconnect (Intel® UPI)

Efficient fabric

Intel® Omni-Path Architecture (Intel® OPA)

Host Fabric Interface (HFI)

Accelerate compression

Intel® QuickAssist Technology (Intel® QAT)

RETHINK THE IMPOSSIBLE WITH BREAKTHROUGH HPC INNOVATION

MORE PERFORMANCE, FLEXIBILITY, AND ENERGY EFFICIENCY INTEL® ADVANCED VECTOR EXTENSIONS 512 (INTEL® AVX-512)

- Up to **1.63x** faster HPC[1]
- Up to **2.2x** faster AI/deep learning[2]
- Up to **3.1x** faster cryptographic hashing performance[3]
- Up to **2x** faster data protection[4]

Accelerate data processing

New instructions optimize and accelerate tasks for diverse modern workloads

Learn more at: https://www.intel.com/content/www/us/en/architecture-and-technology/avx-512-animation.html

1. For disclaimers and configuration, see following slides.
2. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks
3. Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system.

CONFIGURATIONS: TECHNICAL COMPUTE WORKLOADS 1/3 (AVX512)

1.) Up to 1.63x gains for HPC workloads: based on Geomean of Weather Research Forecasting - Conus 12Km, HOMME, LSTCLS-DYNA Explicit, INTES PERMAS V16, MILC, GROMACS water 1.5M_pme, VASPSi256, NAMDstmv, LAMMPS, Amber GB Nucleosome, Binomial option pricing, Black-Scholes, Monte Carlo European options

a) **PERMAS by INTES** is an advanced Finite Element software system that offers a complete range of physical models at high performance, quality, and reliability. It plays a mission-critical role in the design process at customers from automotive, ship design, aerospace, and more. E5-2697 v4: 2S Intel® Xeon® processor E5-2697v4, 2.3GHz, 18 cores, turbo on, HT off, NUMA on, BIOS 338.R00, 256 GB total memory (8x 32GB w/ 2400 MT/s, DDR4 LRDIMM), 4x Intel® SSD DC P3600 2 TB in RAID 0 (stripe size 64k). CentOS Linux* release 7.2, kernel 3.10.0-327.13.1.el7.x86_64. Intel® Composer 2015.5.223. INTES PERMAS V16.00. Gold 6148: Intel® Xeon® Gold 6148 processor, 2.4 GHz, 20 cores, turbo on, HT off, NUMA on, BIOS SE5C620.86B.01.00.0412.020920172159, 384 GB total memory (12x 32GB w/ 2400 MT/s, DDR4 LRDIMM), 3x Intel® SSD DC P3600 2 TB in RAID 0 (stripe size 64k), CentOS* Linux* release 7.3, kernel 3.10.0-514.10.2.el7.x86_64. Intel® Composer 2015.7.235. INTES PERMAS V16.00.

b) **LS-DYNA** is the leading product in the crash simulation market. It is used by the automobile, aerospace, construction, military, manufacturing, and bioengineering industries in worldwide. Workload: 2M elements Car2car model with 120ms simulation time. LS-DYNA explicit standard benchmarks tested by Intel, March 2017. E5-2697 V4: 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 18 cores, turbo and HT on, BIOS SE5C610.86B.01.01.0016.033120161139, 128GB total memory, 8 memory channels / 8x16GB / 2400 MT/s / DDR4, Red Hat Enterprise Linux* 7.3 kernel 3.10.0-229.20.1.el6.x86_64.knl2. GOLD 6148: 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 20 cores, turbo and HT on, BIOS version 412, 192GB total memory, 12 memory channels / 12x16GB / 2400 MT/s / DDR4, Red Hat Enterprise Linux* 7.3 kernel 3.10.0-514.el7.x86_64.

c) **Binomial option pricing** is a lattice-based approach that uses a discrete-time model of the varying price over time of the underlying financial instrument. This is compute bound, double precision workload. FSI Binomial workload. OS: Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Testing by Intel March 2017. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 slots / 16GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo and HT on, BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

d) **Monte Carlo** is a numerical method that uses statistical sampling techniques to approximate solutions to quantitative problems. In finance, Monte Carlo algorithms are used to evaluate complex instruments, portfolios, and investments. This is compute bound, double precision workload. FSI Monte Carlo workload. OS: Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Testing by Intel March 2017. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 x16GB 2400 MHz DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor@ 2.4GHz, H0QS, 40 cores 150W. QMS1, turbo and HT on, BIOS SE5C620.86B.01.00.0412.020920172159, 192GB total memory, 12 x 16 GB 2666 MHz DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327

e) **Black-Scholes** is a popular mathematical model used in finance for European option valuation. This is a double precision version. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 x16GB 2400 MHz DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor@ 2.4GHz, H0QS, 40 cores 150W. QMS1, turbo and HT on, BIOS SE5C620.86B.01.00.0412.020920172159, 192GB total memory, 12 x 16 GB 2666 MHz DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327

f) **Amber*** is a suite of programs for classical molecular dynamics and statistical analysis. The main MD program is PMEMD (Particle Mesh Ewald Molecular Dynamics) employs two separate algorithms for implicit- and explicit-solvent dynamics. Here performance for explicit solvent (PME) is presented. Amber: Version 16 with all patches applied at December, 2016. Workloads: PME Cellulose NVE(408K atoms), PME stmv(1M atoms), GB Nucleosome (25K), GB Rubisco (75K). No cut-off was used for GB workloads. Compiled with -mic2_spdp –intelmpi - openmp, –DMIC2 * defined. Tests performed on March 2017. E5-2697 v4: Executed with 36 MPI, 2 OpenMP. 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 8x16GB 2400MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Executed with 40 MPI and 2 OpenMP. 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo on, HT on, BIOS 86B.01.00.0412.R00, 12x16GB 2666MHz DDR, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

g) **VASP** CONFIGURATION: The Vienna Ab initio Simulation Package (VASP) is a computer program for atomic scale materials modeling and performs electronic structure calculations and quantum-mechanical molecular dynamics from first principles. VASP provides scientists with fast and precise calculation of materials properties covering wide range of MD methods from DFT, DFT-HF to Random-Phase approximation (GW, ACDFT). Beta VASP, a release candidate for v6.0. Developer branch provided as “Package” included with download: https://github.com/vasp-dev/vasp-knl. AVX512: Intel® Compiler 17.0.1.132, Intel® MPI 2017u1, ELPA 2016.05.004. Optimization Flags: “-O3 -xCORE-AVX512“. AVX2: Intel® Compiler 17.0.1.132, Intel® MPI 2017u1, ELPA 2016.05.004. Optimization Flags: “-O3 -xCORE-AVX2“. E5-2697 v4: 2S Intel® Xeon® processor E5-2697 v4 2.3 GHz , 18 Cores/Socket, 36 Cores, 72 Threads, HT on, turbo off, BIOS 86B0271.R00, 128GB total memory, 2400 MT/s DDR4 RDIMM, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Dual Socket Intel® Xeon® processor Gold 6148 2.4 GHz , 20 Cores/Socket, 40 Cores, 80 Threads, HT on, turbo off, BIOS 86B.01.00.0412, 192GB total memory, 2666 MT/s / DDR4 RDIMM, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

CONFIGURATIONS: TECHNICAL COMPUTE WORKLOADS 2/3 (AVX512)

1.) Up to 1.63x gains for HPC workloads: based on Geomean of Weather Research Forecasting - Conus 12Km, HOMME, LSTCLS-DYNA Explicit, INTES PERMAS V16, MILC, GROMACS water 1.5M_pme, VASPSi256, NAMDstmv, LAMMPS, Amber GB Nucleosome, Binomial option pricing, Black-Scholes, Monte Carlo European options

h) **NAMD**: NAMD, recipient of a 2002 Gordon Bell Award, is a parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems. Based on Charm++ parallel objects, NAMD scales to hundreds of cores for typical simulations and beyond 200,000 cores for the largest simulations. Version 2.12 Dec2016. Workloads: apoa1(92K atoms), stmv(1M atoms). Compiled with –DNAMD_KNL* define. Tests performed on March 2017. E5-2697 v4: 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 8x16GB 2400MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Compiler option "-xCORE-AVX2". Gold 6148: 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo on, HT on, BIOS 86B.01.00.0412.R00, 12x16GB 2666MHz DDR, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Compiler option "-xCORE-AVX512".

i) **LAMMPS**: LAMMPS is a classical molecular dynamics code, and an acronym for Large-scale Atomic/Molecular Massively Parallel Simulator. It is used to simulate the movement of atoms to develop better therapeutics, improve alternative energy devices, develop new materials, and more. E5-2697 v4: 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 36 cores, Intel® Turbo Boost Technology and Intel® Hyperthreading Technology on, BIOS 86B0271.R00, 8x16GB 2400MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, Intel® Turbo Boost Technology and Intel® Hyperthreading Technology on, BIOS 86B.01.00.0412.R00, 12x16GB 2666MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

j) **GROMACS** is a versatile package to perform classical Molecular Dynamics simulations. Heavily optimized for most modern platforms and provides extremely high performance. GROMACS AVX2 CONFIGURATION: Version 2016.3: ftp://ftp.gromacs.org/pub/gromacs/gromacs-2016.3.tar.gz , Intel® Compiler 17.0.1.132, Intel® MPI 2017u1. Optimization Flags: "-O3 -xCORE-AVX2". Cmake options: "-DGMX_FFT_LIBRARY=mkl -DGMX_SIMD=AVX2_256". GROMACS AVX512 CONFIGURATION: Version 2016.3: ftp://ftp.gromacs.org/pub/gromacs/gromacs-2016.3.tar.gz , Intel® Compiler 17.0.1.132, Intel® MPI 2017u1. Optimization Flags: "-O3 -xCORE-AVX512". Cmake options: "-DGMX_FFT_LIBRARY=mkl -DGMX_SIMD=AVX_512". E5-2697 V4: GROMACS AVX2 binary, Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 Cores/Socket, 36 Cores, 72 Threads (HT on, Turbo on), DDR4 128GB, 2400 MHz, Red Hat 7.2. Gold 6148: GROMACS AVX512 binary, Dual Socket Intel® Xeon® processor Gold 6148 2.4 GHz , 20 Cores/Socket, 40 Cores, 80 Threads (HT on, Turbo on), DDR4 192GB, 2666 MT/s DDR4 RDIMMs, Red Hat 7.2.

k) **Weather Research and Forecasting** (WRF) Model is a next-generation mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting needs. It features two dynamical cores, a data assimilation system, and a software architecture facilitating parallel computation and system extensibility. The model serves a wide range of meteorological applications across scales from tens of meters to thousands of kilometers. 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 slots / 16GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA HDD, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Software: WRF version 3.6.1 Compiled using Intel config option with "-O3 -fp-model fast=1 -xCORE-AVX2". Executed with 36 MPI ranks and OMP_NUM_THREADS=1. Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo and HT on, BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Software: WRF version 3.6.1 Compiled using Intel config option with "-O3 -fp-model fast=1 -xCORE-AVX512". Executed with 40 MPI ranks and OMP_NUM_THREADS=1.

l) **HOMME** is the spectral element dynamical core that solves the equations of motion in the CAM-SE atmospheric model, part of the NSF Community Earth System Model (CESM) as well as the related DOE ACME model. CESM is a widely-used Earth system model and an important source of simulations used by the Intergovernmental Panel on Climate Change. HOMME version: https://svn-homme-model.cgd.ucar.edu/branch_tags/dungeon_tags/dungeon016. Compiled with "-O3 -fp-model fast -xCORE-AVX2". Running "perfTestWACCM" benchmark from code repository with size NE=8. Executed with 64 MPI ranks and OMP_NUM_THREADS=1. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 slots / 16GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo and HT on, BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

CONFIGURATIONS: TECHNICAL COMPUTE WORKLOADS 3/3 (AVX512)

2.) 24X and 2.2X deep learning inference and training performance: Inference throughput batch size 1, Training throughput batch size 256. Source: Intel measured as of June 2017 Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. 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. 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. Deep Learning Frameworks: Neon: ZP/MKL_CHWN branch commit id:52bd02acb947a2adabb8a227166a7da5d9123b6d. Dummy data was used. The main.py script was used for benchmarking , in mkl mode. ICC version used : 17.0.3 20170404, Intel MKL small libraries version 2018.0.20170425. Platform: Platform: 2S Intel® Xeon® CPU E5-2699 v4 @ 2.20GHz (22 cores), HT enabled, turbo disabled, scaling governor set to "performance" via acpi-cpufreq driver, 256GB DDR4-2133 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3500 Series (480GB, 2.5in SATA 6Gb/s, 20nm, MLC). Performance measured with: Environment variables: KMP_AFFINITY='granularity=fine, compact,1,0', OMP_NUM_THREADS=44, CPU Freq set with cpupower frequency-set -d 2.2G -u 2.2G -g performance. Deep Learning Frameworks: Neon: ZP/MKL_CHWN branch commit id:52bd02acb947a2adabb8a227166a7da5d9123b6d. Dummy data was used. The main.py script was used for benchmarking , in mkl mode. ICC version used : 17.0.3 20170404, Intel MKL small libraries version 2018.0.20170425.

3.) and 4.) 3.1x, 1.2x, 2x ISA-L configuration. Intel® Xeon® processor Scalable family: Intel® Xeon® Platinum 8180 processor, 28C, 2.5 GHz, H0, Neon City CRB, 12x16 GB DDR4 2666 MT/s ECC RDIMM, BIOS PLYCRB1.86B.0128.R08.1703242666.

Intel® Xeon® E5-2600v4 series processor, E5-2650 v4, 12C, 2.2 GHz, Aztec City CRB, 4x8 GB DDR4 2400 MT/s ECC RDIMM, BIOS GRRFCRB1.86B.0276.R02.1606020546

Operating System: Redhat Enterprise Linux* 7.3, Kernel 4.2.3, ISA-L 2.18, BIOS Configuration, P-States: Disabled, Turbo: Disabled, Speed Step: Disabled, C-States: Disabled, ENERGY_PERF_BIAS_CFG: PERF.

Up to 3.1x gains for Cryptographic Hashing **using SHA Algorithms and Up to 2x Data Protection using Reed Solomon Erasure Code:** Performance on single core with frequency obfuscation comparing Intel® Xeon® Platinum 8180 Processor vs Intel® Xeon® Processor E5-2650v4.

Configurations: Intel Xeon® Processor Scalable Family: Platinum 8180 Processor, 28C, 2.5 GHz, H0, Neon City CRB, 12x16 GB DDR4 2666 MT/s ECC RDIMM, BIOS PLYCRB1.86B.0128.R08.1703242666 vs. Intel® Xeon® E5-2600v4 Series Processor, E5-2650 v4, 12C, 2.2 GHz, Aztec City CRB, 4x8 GB DDR4 2400 MT/s ECC RDIMM, BIOS GRRFCRB1.86B.0276.R02.1606020546. Operating System: Redhat Enterprise Linux 7.3, Kernel 4.2.3, ISA-L 2.18, BIOS Configuration, P-States: Disabled, Turbo: Disabled, Speed Step: Disabled, C-States: Disabled, ENERGY_PERF_BIAS_CFG: PERF

INTEL® XEON® SCALABLE PROCESSOR

Re-architected from the Ground Up

- Skylake core microarchitecture, with data center specific enhancements
- Intel® AVX-512 with 32 DP flops per core
- Data center optimized cache hierarchy – 1MB L2 per core, non-inclusive L3
- New mesh interconnect architecture
- Enhanced memory subsystem
- Modular IO with integrated devices
- New Intel® Ultra Path Interconnect (Intel® UPI)
- Intel® Speed Shift Technology
- Security & Virtualization enhancements (MBE, PPK, MPX)
- Optional Integrated Intel® Omni-Path Fabric (Intel® OPA)

Features	Intel® Xeon® Processor E5-2600 v4	Intel® Xeon® Scalable Processor
Cores Per Socket	Up to 22	**Up to 28**
Threads Per Socket	Up to 44 threads	**Up to 56 threads**
Last-level Cache (LLC)	Up to 55 MB	**Up to 38.5 MB (non-inclusive)**
QPI/UPI Speed (GT/s)	2x QPI channels @ 9.6 GT/s	**Up to 3x UPI @ 10.4 GT/s**
PCIe* Lanes/ Controllers/Speed(GT/s)	40 / 10 / PCIe* 3.0 (2.5, 5, 8 GT/s)	**48 / 12 / PCIe 3.0 (2.5, 5, 8 GT/s)**
Memory Population	4 channels of up to 3 RDIMMs, LRDIMMs, or 3DS LRDIMMs	**6 channels of up to 2 RDIMMs, LRDIMMs, or 3DS LRDIMMs**
Max Memory Speed	Up to 2400	**Up to 2666**
TDP (W)	55W-145W	**70W-205W**

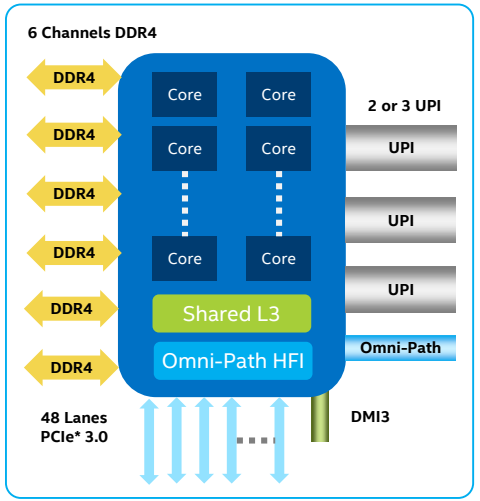

NEW MESH INTERCONNECT ARCHITECTURE

Intel® Xeon® Processor E7 family (24-core die)

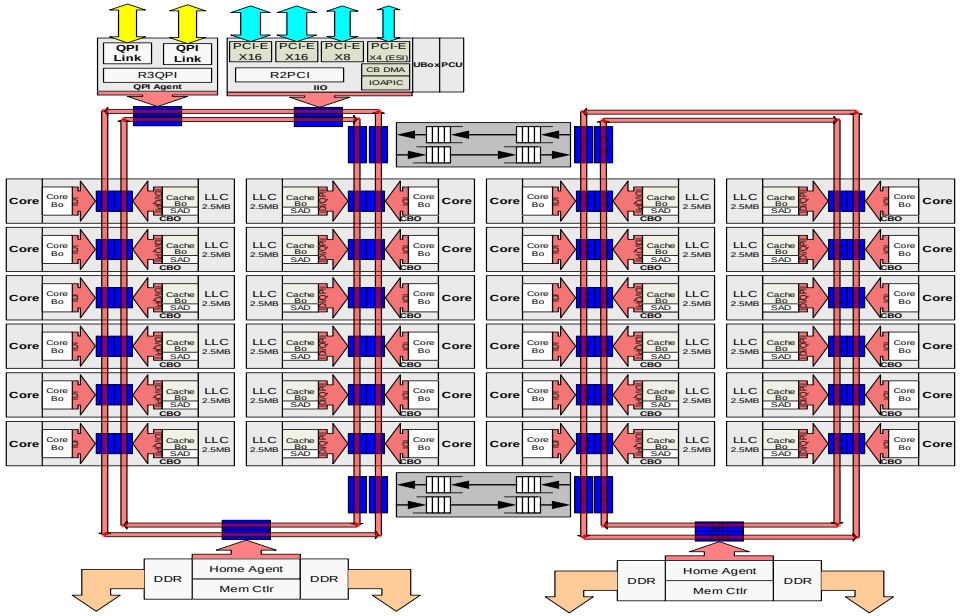

Intel® Xeon® Scalable Processor (28-core die)

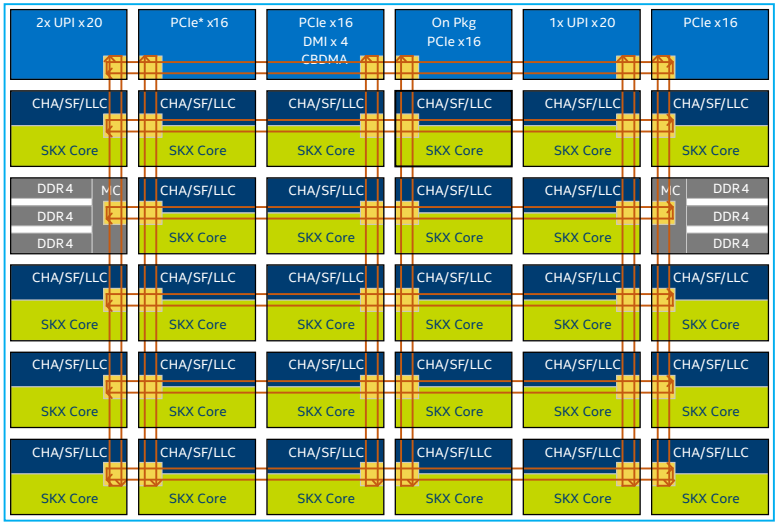

CHA – Caching and Home Agent ; SF – Snoop Filter; LLC – Last Level Cache;
SKX Core – Skylake Server Core; UPI – Intel® UltraPath Interconnect

MESH IMPROVES SCALABILITY WITH HIGHER BANDWIDTH AND REDUCED LATENCIES

DISTRIBUTED CACHING AND HOME AGENT (CHA)

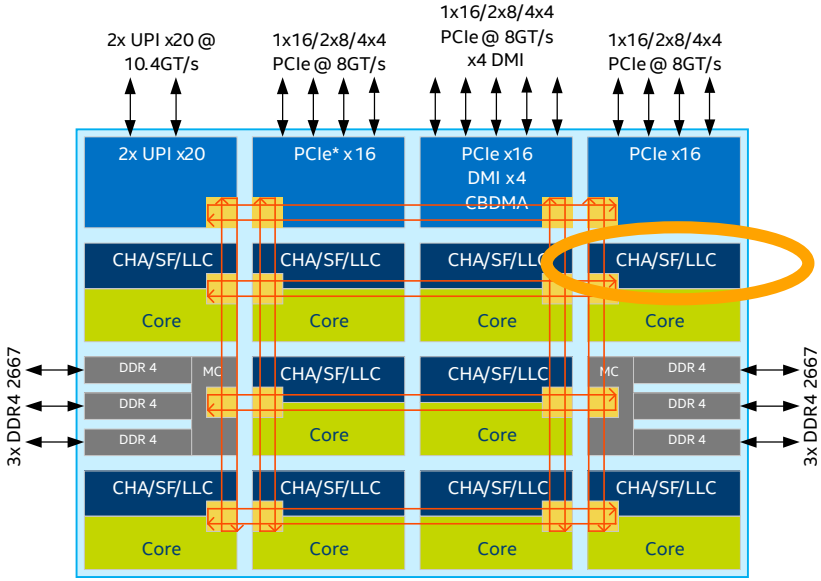

- Intel® UPI caching and home agents are distributed with each LLC bank
  - Prior generation had a small number of QPI home agents
- Distributed CHA benefits
  - Eliminates large tracker structures at memory controllers, allowing more requests in flight and processes them concurrently
  - Reduces traffic on mesh by eliminating home agent to LLC interaction
  - Reduces latency by launching snoops earlier and obviates need for different snoop modes

DISTRIBUTED CHA ARCHITECTURE SUSTAINS HIGHER BANDWIDTH AND LOWERS LATENCY

RE-ARCHITECTED L2 & L3 CACHE HIERARCHY

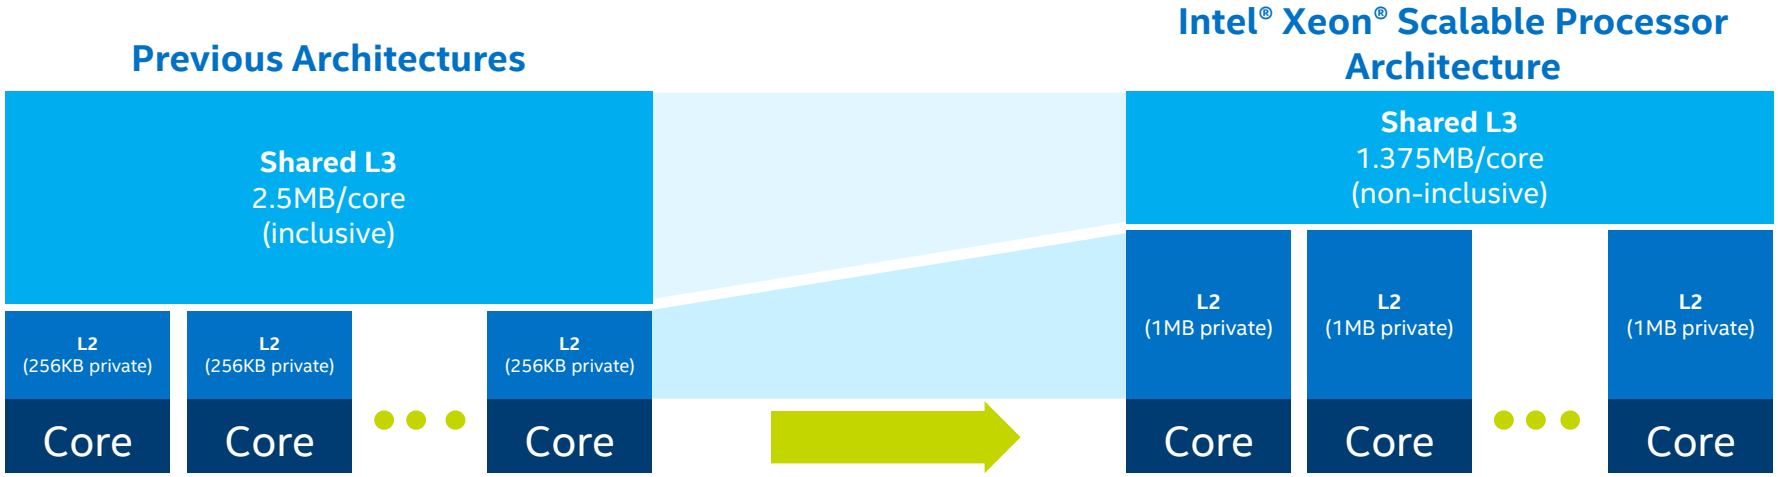

- On-chip cache balance shifted from shared-distributed (prior architectures) to private-local (Skylake architecture):
  - Shared-distributed ➔ shared-distributed L3 is primary cache
  - Private-local ➔ private L2 becomes primary cache with shared L3 used as overflow cache
- Shared L3 changed from inclusive to non-inclusive:
  - Inclusive (prior architectures) ➔ L3 has copies of all lines in L2
  - Non-inclusive (Skylake architecture) ➔ lines in L2 ***may not*** exist in L3

SKYLAKE-SP CACHE HIERARCHY ARCHITECTED SPECIFICALLY FOR DATA CENTER USE CASE

INTEL® ADVANCED VECTOR EXTENSIONS-512 (AVX-512)

End Customer Value: Workload-optimized performance, throughput increases, and H/W-enhanced security improvements for familiar analytics, HPC, video transcode, cryptography, and compression software.

Problems Solved:

1. Achieve more work per cycle (doubles width of data registers)

2. Minimize latency & overhead (doubles the number of registers) with ultra-wide (512-bit) vector processing capabilities (that that 2x FMA processing engines are available on Intel® Xeon® Platinum and Intel® Xeon® Gold Processors)

VALUE PILLARS

PERFORMANCE

SECURITY

SEGMENTS

PROOF POINTS

Up to 2X FLOPS/clock cycle[1]

Up to 4X greater throughput[2]

Accelerates performance for your most demanding computational tasks

* FLOPs = Floating Point Operations

1 Peak performance vs. Intel® AVX2. As measured by Intel® Xeon® Processor Scalable Family with Intel® AVX-512 compared to an Intel® Xeon® E5 v4 with Intel® AVX2

2 Vectorized floating-point throughput. As measured by Intel® Xeon® Processor Scalable Family with Intel® AVX-512 compared to an Intel® Xeon® E5 v4 with Intel® AVX2

INTEL® ADVANCED VECTOR EXTENSIONS-512 (AVX-512)

- 512-bit wide vectors
- 32 operand registers
- 8 64b mask registers
- Embedded broadcast
- Embedded rounding

Microarchitecture	Instruction Set	SP FLOPs / cycle	DP FLOPs / cycle
Skylake	Intel® AVX-512 & FMA	64	32
Haswell / Broadwell	Intel AVX2 & FMA	32	16
Sandybridge	Intel AVX (256b)	16	8
Nehalem	SSE (128b)	8	4

Intel AVX-512 Instruction Types	
AVX-512-F	AVX-512 Foundation Instructions
AVX-512-VL	Vector Length Orthogonality : ability to operate on sub-512 vector sizes
AVX-512-BW	512-bit Byte/Word support
AVX-512-DQ	Additional D/Q/SP/DP instructions (converts, transcendental support, etc.)
AVX-512-CD	Conflict Detect : used in vectorizing loops with potential address conflicts

POWERFUL INSTRUCTION SET FOR DATA-PARALLEL COMPUTATION

PERFORMANCE AND EFFICIENCY WITH INTEL® AVX-512

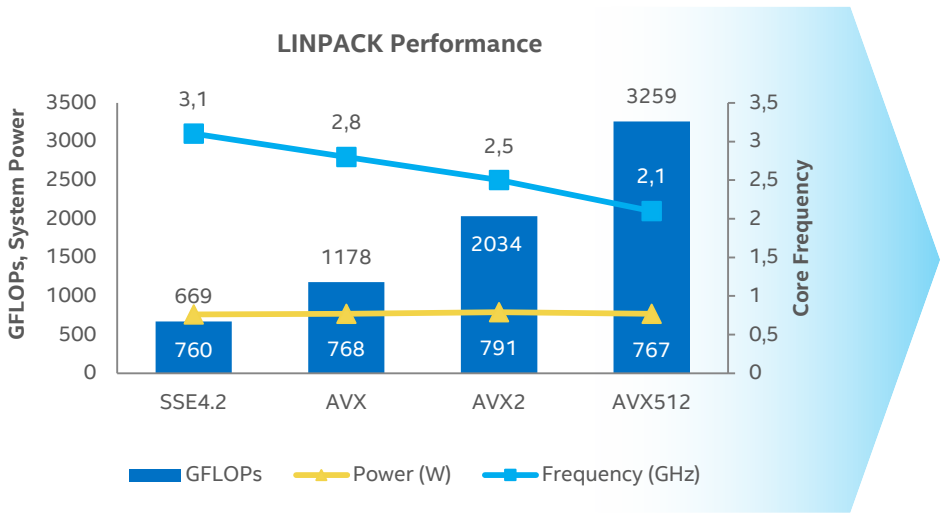

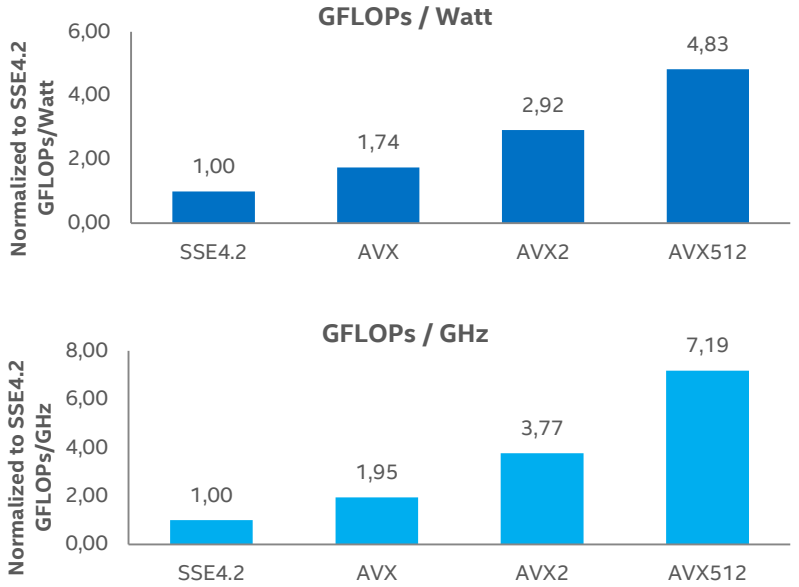

INTEL® AVX-512 DELIVERS SIGNIFICANT PERFORMANCE AND EFFICIENCY GAINS

Source as of June 2017: Intel internal measurements on platform with Xeon Platinum 8180, Turbo enabled, UPI=10.4, SNC1, 6x32GB DDR4-2666 per CPU, 1 DPC. 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.

INTEL® OPTANE™ SSD DC P4800X FOR STORAGE BUILDERS

SPDK Performance: Platform Comparison

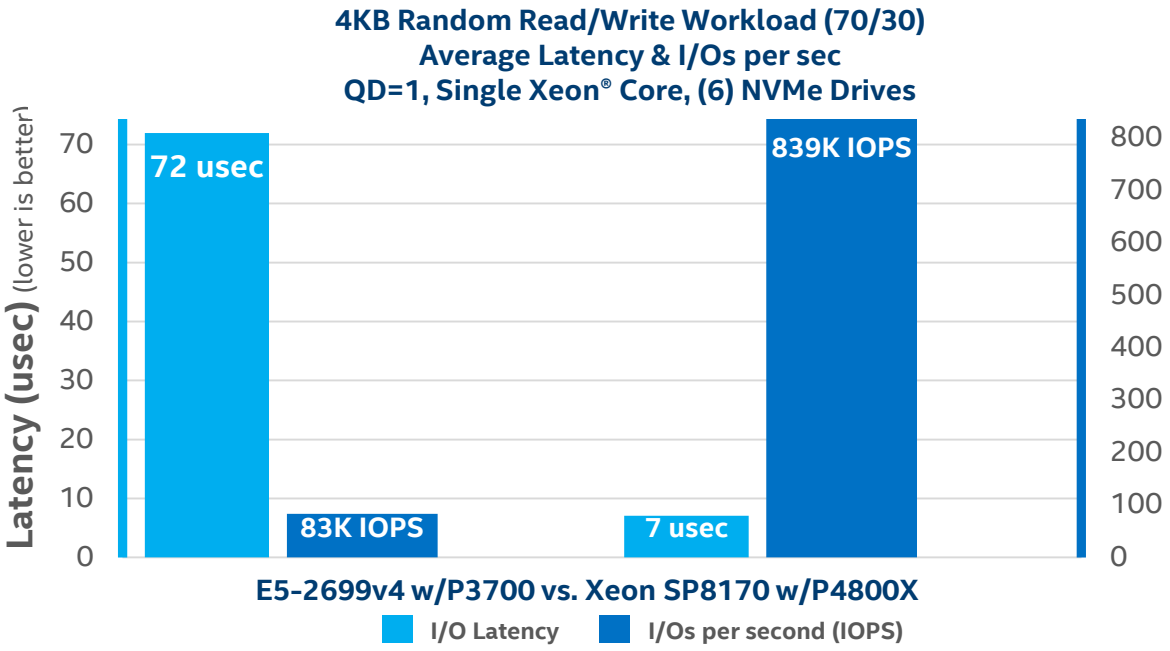

Intel® Xeon® Scalable Processor Platinum Family + Intel® Optane™

- 10X higher throughput
- 10X lower latency
- Up to 27 cores remaining for:
  - Virtual Machines
  - Big Data/Analytics
  - Machine Learning
  - Storage services like erasure coding, de-duplication, compression, or encryption.
- Platform offers RDMA
  - Enables NVMe over Fabrics
  - No more trapped I/O capacity

See notices, configurations, disclaimers

INTEL® OMNI-PATH ARCHITECTURE (INTEL® OPA)

End-to-end HPC fabric solution that cost-effectively scales from entry clusters to the largest supercomputers

Builds on the industry's best technologies

- Highly leverage existing Aries and Intel® True Scale Fabric
- Adds innovative new features and capabilities to improve performance, reliability, and QoS

Robust product offerings and ecosystem

- >100 OEM designs[1]
- Strong ecosystem with 80+ Fabric Builders members

HFI Adapters
Single port
x8 and x16

x16 Adapter (100 Gb/s)

x8 Adapter (58 Gb/s)

Edge Switches
1U Form Factor
24 and 48 port

48-port Edge Switch

24-port Edge Switch

Director Switches
QSFP-based
192 and 768 port

768-port Director Switch (20U chassis)

192-port Director Switch (7U chassis)

Silicon
OEM custom designs
HFI and Switch ASICs

HFI silicon
Up to 2 ports
(50 GB/s total b/w)

Switch silicon
up to 48 ports
(1200 GB/s total b/w

Software
Open Source
Host Software and Fabric Manager

Cables
Third Party Vendors
Passive Copper
Active Optical

MAXIMIZE PRICE/PERFORMANCE, FREEING UP CLUSTER BUDGETS FOR INCREASED COMPUTE AND STORAGE CAPABILITY

1. Intel internal information. Design win count based on OEM and HPC storage vendors who are planning to offer either Intel-branded or custom switch products, along with the total number of OEM platforms that are currently planned to support custom and/or standard Intel® OPA adapters. Design win count as of November 1, 2015 and subject to change without notice based on vendor product plans.

MORE PERFORMANCE , FLEXIBILITY, AND ENERGY EFFICIENCY INTEL® OMNI-PATH ARCHITECTURE (INTEL® OPA)

CPU/Fabric Integration

- Lower power and cost, higher density
- Increased bandwidth to each socket
- Improved latency to each socket

Optimized Host Implementation

- High MPI message rate
- Maintains low end-to-end latency at scale
- Enhancements for file system traffic
- IB software compatibility

Enhanced Fabric Architecture

- High traffic throughput, low port-to-port latency
- Latency efficient error detection & correction
- Deterministic latency and QoS features

SKYLAKE PERFORMANCE

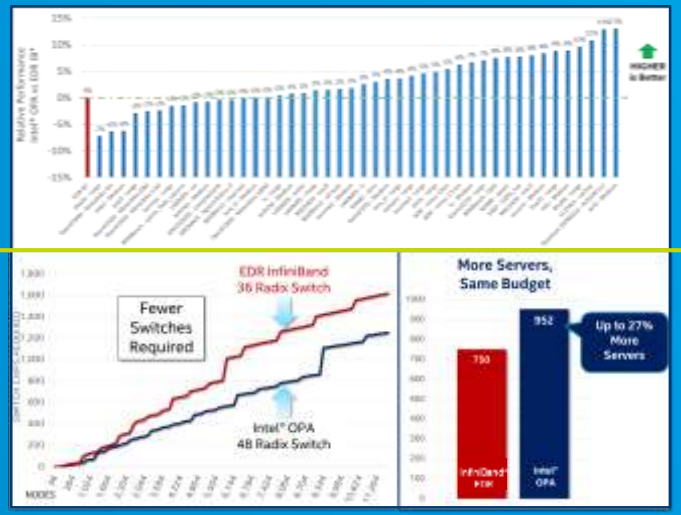

1. Configuration assumes a 750-node cluster, and number of switch chips required is based on a full bisectional bandwidth (FBB) Fat-Tree configuration. Intel® OPA uses one fully-populated 768-port director switch, and Mellanox EDR solution uses a combination of 648-port director switches and 36-port edge switches. Mellanox component pricing from www.kernelsoftware.com, with prices as of April 4, 2017. Compute node pricing based on Dell PowerEdge R730 server from www.dell.com, with prices as of April 4, 2017. Intel® OPA pricing based on pricing from www.kernelsoftware.com as of August 15, 2017
2. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks
3. Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system.
4. See Performance/Configuration slide for configuration information

CONFIGURATION FOR APPLICATION PERFORMANCE (OPA)

INTEL® XEON® PLATINUM 8170 PROCESSORS 1/2

Common configuration for all applications listed below (Unless otherwise specified for the application)

Internal Intel testing with dual socket Intel® Xeon® Platinum 8170 processor nodes. Unless otherwise specified, one MPI rank per physical CPU core is used. 192 GB 2666 MHz DDR4 memory on first 8 nodes, 64 GB 2666 MHz DDR4 memory on second group of 8 nodes. This is due to current limited memory supply. Red Hat Enterprise Linux* Server release 7.3 (Maipo), 3.10.0-514.16.1.el7.x86_64 kernel. Intel® Turbo Boost and Hyperthreading Technology enabled. All compute nodes connected to one edge switch device. All data written over NFSv3 with 1GbE to Intel SSDSC2BB48 drive storage. Intel MPI 2017 Update 3, Open MPI 2.1.1 or 1.10.4-hfi as packaged with IFS for Intel OPA. Intel OPA: Intel Fabric Suite 10.3.1.0.22. Intel Corporation Device 24f0 – Series 100 HFI. OPA Switch: Series 100 Edge Switch – 48 port. Non-default HFI parameters: krcvqs=4, eager_buffer_size=8388608, max_mtu=10240. The better performance shown using either I_MPI_FABRICS tmi or shm:tmi EDR IB* based on internal testing: MLNX_OFED_LINUX-4.0-2.0.0.1 (OFED-4.0-2.0.0). Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR InfiniBand switch. Forward Error Correction is automatically disabled because <=2M copper IB* cables were used in the testing.. Unless otherwise noted, I_MPI_FABRICS=shm:dapl for Intel MPI. For Open MPI, MXM/FCA as packaged with MOFED, flags used: "-x LD_PRELOAD=/opt/mellanox/hcoll/lib/libhcoll.so -mca coll_hcoll_enable 1 -x HCOLL_MAIN_IB=mlx5_0:1 -x HCOLL_ENABLE_MCAST_ALL=0".

LAMMPS (Large-scale Atomic/Molecular Massively Parallel Simulator) Feb 16, 2016 stable version release. Official Git Mirror for LAMMPS (http://lammps.sandia.gov/download.html) ls, rhodo, sw, and water benchmark. Intel MPI 2017 Update 3 with 52 ranks per node and 2 OMP threads per rank. Common parameters: I_MPI_PIN_DOMAIN=core Run detail: Number of time steps=100, warm up time steps=10 (not timed) Number of copies of the simulation box in each dimension: 8x8x4 and problem size: 8x8x4x32k = 8,192k atoms Build parameters: Modules: yes-asphere yes-class2 yes-kspace yes-manybody yes-misc yes-molecule yes-mpiio yes-opt yes-replica yes-rigid yes-user-omp yes-user-intel. Binary to be built: lmp_intel_cpu. . Runtime lammps parameters: -pk intel 0 -sf intel -v n 1

NAMD version 2.10b2, stmv and apoa1 benchmark. Intel MPI 2017 Update 3 Build detail: CHARM 6.6.1. FFTW 3.3.4. Relevant build flags: ./config Linux-x86_64-icc --charm-arch mpi-linux-x86_64-ifort-smp-mpicxx --cxx icpc --cc icc --with-fftw3. 52 MPI ranks per node.

NWCHEM release 6.6. Binary: nwchem_armci-mpi_intel-mpi_mkl with MPI-PR run over MPI-1. Workload: siosi3 and siosi5. http://www.nwchem-sw.org/index.php/Main_Page. Intel OPA: -genv PSM2_SDMA=0. EDR parameters: I_MPI_FABRICS=shm:ofa. 2 ranks per node, 1 rank for computation and 1 rank for communication. -genv CSP_VERBOSE 1 -genv CSP_NG 1 -genv LD_PRELOAD libcasper.so

BSMBench - An HPC Benchmark for BSM Lattice Physics Version 1.0. 32 ranks per node. Parameters: global size is 64x32x32x32, proc grid is 8x4x4x4. Machine config build file: cluster.cfg

GROMACS version 2016.2. http://www.prace-ri.eu/UEABS/GROMACS/1.2/GROMACS_TestCaseB.tar.gz lignocellulose-rf benchmark. Build detail: -xCORE-AVX2 (Open MPI) and AVX512 for Intel MPI, -g -static-intel. CC=mpicc CXX=mpicxx -DBUILD_SHARED_LIBS=OFF -DGMX_FFT_LIBRARY=mkl -DGMX_MPI=ON -DGMX_OPENMP=ON -DGMX_CYCLE_SUBCOUNTERS=ON -DGMX_GPU=OFF -DGMX_BUILD_HELP=OFF -DGMX_HWLOC=OFF -DGMX_SIMD=AVX2_256 (OpenMPI) or AVX512 for Intel MPI. GMX_OPENMP_MAX_THREADS=256. Run detail: gmx_mpi mdrun -s run.tpr -gcom 20 -resethway -noconfout

Spec MPI2007, https://www.spec.org/mpi/. *Intel Internal measurements marked estimates until published. Applications listed with "-Large" or "-Medium" in the name were part of the spec MPI suite. OPA parameters: better of I_MPI_FABRICS shm:tmi and tmi. EDR parameters: better of I_MPI_FABRICS shm:dapl and shm:ofa. Intel Parallel Studio 2017 Update 4, Intel compiler 17.0.4. -O3 -xCORE-AVX2 -no-prec-div. Intel MPI: mpiicc, mpiifort, mpiicpc. Open MPI: mpicc, mpifort, mpicxx. Run detail: mref and lref suites, 3 iterations. 121.pop2: CPORTABILITY=-DSPEC_MPI_CASE_FLAG. 126.lammps: CXXPORTABILITY = -DMPICH_IGNORE_CXX_SEEK. 127.wrf2: CPORTABILITY = -DSPEC_MPI_CASE_FLAG -DSPEC_MPI_LINUX. 129.tera_tf=default=default=default: srcalt=add_rank_support 130.socorro=default=default=default: srcalt=nullify_ptrs FPORTABILITY = -assume nostd_intent_in CPORTABILITY = -DSPEC_EIGHT_BYTE_LONG CPORTABILITY = -DSPEC_SINGLE_UNDERSCORE. Intel® OPA: 32 MPI ranks per node for 115.fds4 benchmark

CONFIGURATION FOR APPLICATION PERFORMANCE

INTEL® XEON® PLATINUM 8170 PROCESSORS 2/2

Quantum ESPRESSO is an integrated suite of Open-Source computer codes for electronic-structure calculations and materials modeling at the nanoscale. It is based on density-functional theory, plane waves, and pseudopotentials. http://www.quantum-espresso.org/ Build detail: MKLROOT=/opt/intel/compilers_and_libraries_2017.4.196/linux/mkl export FC=mpifort export F90=$FC export F77=$FC export MPIF90=$FC export FCFLAGS="-O3 -xCORE-AVX2 -fno-alias -ansi-alias -g -mkl -I$MKLROOT/include/fftw -I${MKLROOT}/include/intel64/ilp64 -I${MKLROOT}/include -qopenmp -static-intel" export FFLAGS=$FCFLAGS export CC=mpicc export CPP="icc -E" export CFLAGS=$FCFLAGS export AR=xiar export BLAS_LIBS="-L$MKLROOT/lib/intel64 -lmkl_blas95_lp64" export LAPACK_LIBS="-L$MKLROOT/lib/intel64_lin -lmkl_blacs_openmpi_lp64" export SCALAPACK_LIBS="-L$MKLROOT/lib/intel64_lin -lmkl_scalapack_lp64 -lmkl_blacs_openmpi_lp64" export FFT_LIBS="-L$MKLROOT/intel64" ./configure --enable-openmp --enable-parallel. BLAS_LIBS= -lmkl_intel_lp64 -lmkl_intel_thread -lmkl_core ELPA_LIBS_SWITCH = enabled SCALAPACK_LIBS = $(TOPDIR)/ELPA/libelpa.a -lmkl_scalapack_lp64 -lmkl_blacs_openmpi_lp64 DFLAGS= -D__INTEL -D__FFTW -D__MPI -D__PARA -D__SCALAPACK -D__ELPA -D__OPENMP $(MANUAL_DFLAGS) AUSURF112 benchmark, all default options

LS-DYNA, A Program for Nonlinear Dynamic Analysis of Structures in Three Dimensions Version : mpp s R9.1.0 Revision: 113698, single precision (I4R4) OPA parameters: better of I_MPI_FABRICS shm:tmi and tmi EDR parameters: better of I_MPI_FABRICS shm:dapl and shm:ofa. Example pfile: gen { nodump nobeamout dboutonly } dir { global one_global_dir local /tmp/3cars }. 32 MPI ranks per node used for OPA with 3cars benchmark.

SPECFEM3D_GLOBE simulates the three-dimensional global and regional seismic wave propagation based upon the spectral-element method (SEM). It is a time-step algorithm which simulates the propagation of earth waves given the initial conditions, mesh coordinates/ details of the earth crust. small_benchmark_run_to_test_more_complex_Earth benchmark, default input settings. specfem3d_globe-7.0.0. Intel Parallel Studio XE 2017 Update 4. FC=mpiifort CC=mpiicc MPIFC=mpiifort FCFLAGS=-g -xCORE_AVX2 CFLAGS=-g -O2 -xCORE_AVX2. run_this_example.sh and run_mesher_solver.sh, NCHUNKS=6, NEX_XI=NEX_ETA=80, NPROC_XI=NPROC_ETA=10. 600 cores used, 52 cores per node.

VASP v6 beta, May 9, 2017, https://github.com/vasp-dev/vasp-knl.git Intel Parallel Studio XE 2017 Update 4, CPP_OPTIONS= -DMPI -DHOST=\"IFC17_impi\" \ -DCACHE_SIZE=12000 -Davoidalloc -DMPI_BLOCK=8000 -DscaLAPACK -Duse_collective -DnoAugXCmeta -Duse_bse_te -Duse_shmem -Dtbdyn -Dvasp6 -D_OPENMP -DPROFILING -Dshmem_bcast_buffer -Dshmem_rproj -Dmemalign64 -DELPA -DVTUNE_PROFILING ARCH=-xCORE-AVX512 Testing performed on 8 nodes with 192 GB 2666 MHz DDR4 memory

WRF - Weather Research & Forecasting Model (http://www.wrf-model.org/index.php) version 3.5.1. OPA parameters: better of I_MPI_FABRICS shm:tmi and tmi. EDR parameters: I_MPI_FABRICS shm:dapl and shm:ofa. Intel Parallel Studio XE 2017 Update 4, ifort, icc, mpif90, mpicc. -xCORE_AVX2 -O3 . Net CDF 4.4.1.1 built with icc. Net CDF-fortran version 4.4.4 built with icc.

OpenFOAM is a free, open source CFD software package developed primarily by [OpenCFD](http://www.openfoam.com) since 2004 and is currently distributed by [ESI-OpenCFD](http://www.openfoam.com) and the [OpenFOAM Foundation](http://openfoam.org). It has a large user base across most areas of engineering and science, from both commercial and academic organisations. OpenFOAM has an extensive range of features to solve anything from complex fluid flows involving chemical reactions, turbulence and heat transfer, to acoustics, solid mechanics and electromagnetics. (http://www.openfoam.com/documentation). Version v1606+ used for Intel MPI and Open MPI with OPA. Version v1612+ used for Open MPI with EDR IB*. Gcc version 4.8.5 for Intel MPI, Icc version 17.0.4 used for Open MPI. All default make options.

INTEL® XEON® SCALABLE PROCESSORS

Performance and energy benchmarks

1.63X AVERAGE GAINS ON HIGH PERFORMANCE COMPUTE APPS

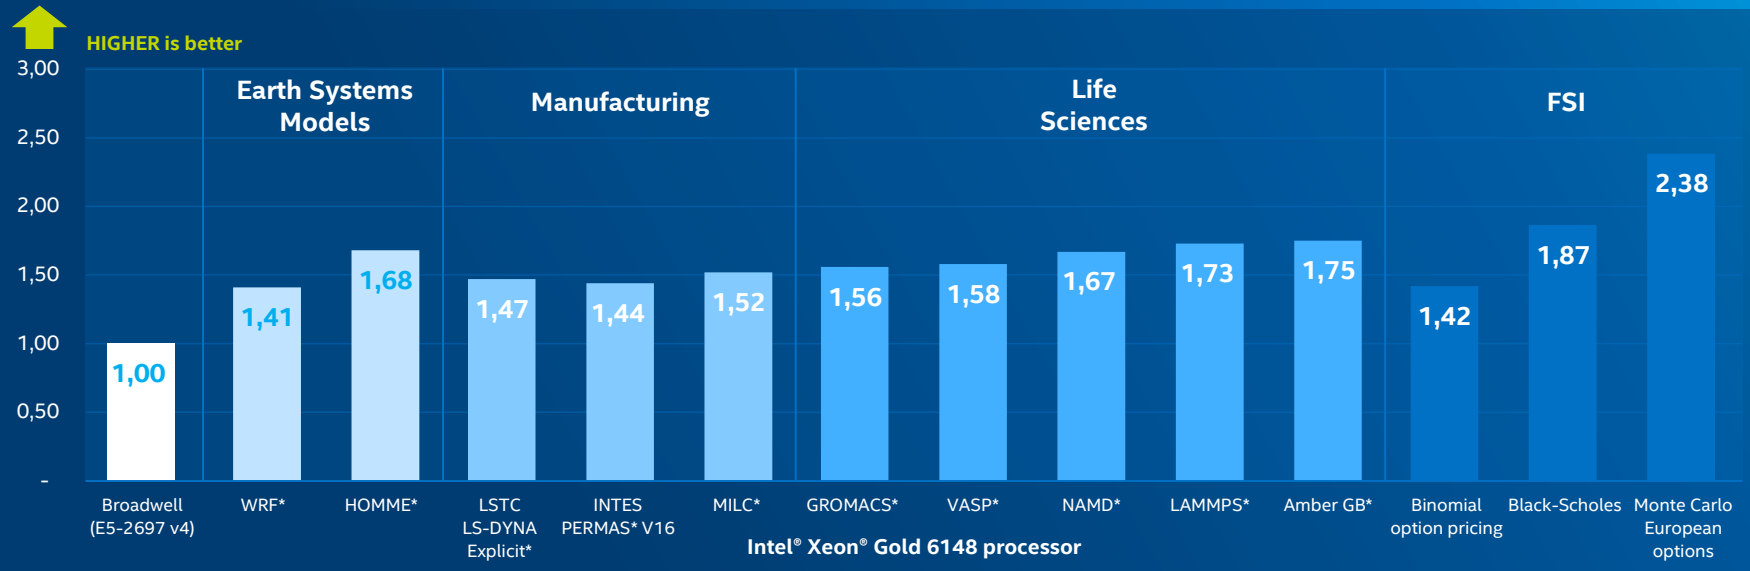

EVEN HIGHER PERFORMANCE WITH INTEL® XEON® PLATINUM 8180 PROCESSOR ON HPC APPS

1. For disclaimers and configuration, see following slides.
2. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks
3. Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system.

CONFIGURATIONS: TECHNICAL COMPUTE WORKLOADS 1/2

1.) Up to 1.63x gains for HPC workloads: based on Geomean of Weather Research Forecasting - Conus 12Km, HOMME, LSTCLS-DYNA Explicit, INTES PERMAS V16, MILC, GROMACS water 1.5M_pme, VASPSi256, NAMDstmv, LAMMPS, Amber GB Nucleosome, Binomial option pricing, Black-Scholes, Monte Carlo European options

a) **PERMAS by INTES** is an advanced Finite Element software system that offers a complete range of physical models at high performance, quality, and reliability. It plays a mission-critical role in the design process at customers from automotive, ship design, aerospace, and more. E5-2697 v4: 2S Intel® Xeon® processor E5-2697v4, 2.3GHz, 18 cores, turbo on, HT off, NUMA on, BIOS 338.R00, 256 GB total memory (8x 32GB w/ 2400 MT/s, DDR4 LRDIMM), 4x Intel® SSD DC P3600 2 TB in RAID 0 (stripe size 64k). CentOS Linux* release 7.2, kernel 3.10.0-327.13.1.el7.x86_64. Intel® Composer 2015.5.223. INTES PERMAS V16.00. Gold 6148: Intel® Xeon® Gold 6148 processor, 2.4 GHz, 20 cores, turbo on, HT off, NUMA on, BIOS SE5C620.86B.01.00.0412.020920172159, 384 GB total memory (12x 32GB w/ 2400 MT/s, DDR4 LRDIMM), 3x Intel® SSD DC P3600 2 TB in RAID 0 (stripe size 64k), CentOS* Linux* release 7.3, kernel 3.10.0-514.10.2.el7.x86_64. Intel® Composer 2015.7.235. INTES PERMAS V16.00.

b) **LS-DYNA** is the leading product in the crash simulation market. It is used by the automobile, aerospace, construction, military, manufacturing, and bioengineering industries in worldwide. Workload: 2M elements Car2car model with 120ms simulation time. LS-DYNA explicit standard benchmarks tested by Intel, March 2017. E5-2697 V4: 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 18 cores, turbo and HT on, BIOS SE5C610.86B.01.01.0016.033120161139, 128GB total memory, 8 memory channels / 8x16GB / 2400 MT/s / DDR4, Red Hat Enterprise Linux* 7.3 kernel 3.10.0-229.20.1.el6.x86_64.knl2. GOLD 6148: 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 20 cores, turbo and HT on, BIOS version 412, 192GB total memory, 12 memory channels / 12x16GB / 2400 MT/s / DDR4, Red Hat Enterprise Linux* 7.3 kernel 3.10.0-514.el7.x86_64.

c) **Binomial option pricing** is a lattice-based approach that uses a discrete-time model of the varying price over time of the underlying financial instrument. This is compute bound, double precision workload. FSI Binomial workload. OS: Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Testing by Intel March 2017. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 slots / 16GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo and HT on, BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

d) **Monte Carlo** is a numerical method that uses statistical sampling techniques to approximate solutions to quantitative problems. In finance, Monte Carlo algorithms are used to evaluate complex instruments, portfolios, and investments. This is compute bound, double precision workload. FSI Monte Carlo workload. OS: Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Testing by Intel March 2017. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 x16GB 2400 MHz DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor@ 2.4GHz, H0QS, 40 cores 150W. QMS1, turbo and HT on, BIOS SE5C620.86B.01.00.0412.020920172159, 192GB total memory, 12 x 16 GB 2666 MHz DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327

e) **Black-Scholes** is a popular mathematical model used in finance for European option valuation. This is a double precision version. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 x16GB 2400 MHz DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor@ 2.4GHz, H0QS, 40 cores 150W. QMS1, turbo and HT on, BIOS SE5C620.86B.01.00.0412.020920172159, 192GB total memory, 12 x 16 GB 2666 MHz DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327

f) **Amber*** is a suite of programs for classical molecular dynamics and statistical analysis. The main MD program is PMEMD (Particle Mesh Ewald Molecular Dynamics) employs two separate algorithms for implicit- and explicit-solvent dynamics. Here performance for explicit solvent (PME) is presented. Amber: Version 16 with all patches applied at December, 2016. Workloads: PME Cellulose NVE(408K atoms), PME stmv(1M atoms), GB Nucleosome (25K), GB Rubisco (75K). No cut-off was used for GB workloads. Compiled with -mic2_spdp –intelmpi - openmp, –DMIC2 * defined. Tests performed on March 2017. E5-2697 v4: Executed with 36 MPI, 2 OpenMP. 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 8x16GB 2400MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Executed with 40 MPI and 2 OpenMP. 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo on, HT on, BIOS 86B.01.00.0412.R00, 12x16GB 2666MHz DDR, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

g) **VASP CONFIGURATION:** The Vienna Ab initio Simulation Package (VASP) is a computer program for atomic scale materials modeling and performs electronic structure calculations and quantum-mechanical molecular dynamics from first principles. VASP provides scientists with fast and precise calculation of materials properties covering wide range of MD methods from DFT, DFT-HF to Random-Phase approximation (GW, ACDFT). Beta VASP, a release candidate for v6.0. Developer branch provided as “Package” included with download: https://github.com/vasp-dev/vasp-knl. AVX512: Intel® Compiler 17.0.1.132, Intel® MPI 2017u1, ELPA 2016.05.004. Optimization Flags: “-O3 -xCORE-AVX512“. AVX2: Intel® Compiler 17.0.1.132, Intel® MPI 2017u1, ELPA 2016.05.004. Optimization Flags: “-O3 -xCORE-AVX2“. E5-2697 v4: 2S Intel® Xeon® processor E5-2697 v4 2.3 GHz , 18 Cores/Socket, 36 Cores, 72 Threads, HT on, turbo off, BIOS 86B0271.R00, 128GB total memory, 2400 MT/s DDR4 RDIMM, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Dual Socket Intel® Xeon® processor Gold 6148 2.4 GHz , 20 Cores/Socket, 40 Cores, 80 Threads, HT on, turbo off, BIOS 86B.01.00.0412, 192GB total memory, 2666 MT/s / DDR4 RDIMM, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

CONFIGURATIONS: TECHNICAL COMPUTE WORKLOADS 2/2

1.) Up to 1.63x gains for HPC workloads: based on Geomean of Weather Research Forecasting - Conus 12Km, HOMME, LSTCLS-DYNA Explicit, INTES PERMAS V16, MILC, GROMACS water 1.5M_pme, VASPSi256, NAMDstmv, LAMMPS, Amber GB Nucleosome, Binomial option pricing, Black-Scholes, Monte Carlo European options

h) NAMD: NAMD, recipient of a 2002 Gordon Bell Award, is a parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems. Based on Charm++ parallel objects, NAMD scales to hundreds of cores for typical simulations and beyond 200,000 cores for the largest simulations. Version 2.12 Dec2016. Workloads: apoa1(92K atoms), stmv(1M atoms). Compiled with –DNAMD_KNL* define. Tests performed on March 2017. E5-2697 v4: 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 8x16GB 2400MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Compiler option "-xCORE-AVX2". Gold 6148: 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo on, HT on, BIOS 86B.01.00.0412.R00, 12x16GB 2666MHz DDR, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Compiler option "-xCORE-AVX512".

i) LAMMPS: LAMMPS is a classical molecular dynamics code, and an acronym for Large-scale Atomic/Molecular Massively Parallel Simulator. It is used to simulate the movement of atoms to develop better therapeutics, improve alternative energy devices, develop new materials, and more. E5-2697 v4: 2S Intel® Xeon® processor E5-2697 v4, 2.3GHz, 36 cores, Intel® Turbo Boost Technology and Intel® Hyperthreading Technology on, BIOS 86B0271.R00, 8x16GB 2400MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: 2S Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, Intel® Turbo Boost Technology and Intel® Hyperthreading Technology on, BIOS 86B.01.00.0412.R00, 12x16GB 2666MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

j) GROMACS is a versatile package to perform classical Molecular Dynamics simulations. Heavily optimized for most modern platforms and provides extremely high performance. GROMACS AVX2 CONFIGURATION: Version 2016.3: ftp://ftp.gromacs.org/pub/gromacs/gromacs-2016.3.tar.gz , Intel® Compiler 17.0.1.132, Intel® MPI 2017u1. Optimization Flags: "-O3 -xCORE-AVX2". Cmake options: "-DGMX_FFT_LIBRARY=mkl -DGMX_SIMD=AVX2_256". GROMACS AVX512 CONFIGURATION: Version 2016.3: ftp://ftp.gromacs.org/pub/gromacs/gromacs-2016.3.tar.gz , Intel® Compiler 17.0.1.132, Intel® MPI 2017u1. Optimization Flags: "-O3 -xCORE-AVX512". Cmake options: "-DGMX_FFT_LIBRARY=mkl -DGMX_SIMD=AVX_512". E5-2697 V4: GROMACS AVX2 binary, Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 Cores/Socket, 36 Cores, 72 Threads (HT on, Turbo on), DDR4 128GB, 2400 MHz, Red Hat 7.2. Gold 6148: GROMACS AVX512 binary, Dual Socket Intel® Xeon® processor Gold 6148 2.4 GHz , 20 Cores/Socket, 40 Cores, 80 Threads (HT on, Turbo on), DDR4 192GB, 2666 MT/s DDR4 RDIMMs, Red Hat 7.2.

k) Weather Research and Forecasting (WRF) Model is a next-generation mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting needs. It features two dynamical cores, a data assimilation system, and a software architecture facilitating parallel computation and system extensibility. The model serves a wide range of meteorological applications across scales from tens of meters to thousands of kilometers. 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 slots / 16GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA HDD, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Software: WRF version 3.6.1 Compiled using Intel config option with "-O3 -fp-model fast=1 -xCORE-AVX2". Executed with 36 MPI ranks and OMP_NUM_THREADS=1. Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo and HT on, BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Software: WRF version 3.6.1 Compiled using Intel config option with "-O3 -fp-model fast=1 -xCORE-AVX512". Executed with 40 MPI ranks and OMP_NUM_THREADS=1.

l) HOMME is the spectral element dynamical core that solves the equations of motion in the CAM-SE atmospheric model, part of the NSF Community Earth System Model (CESM) as well as the related DOE ACME model. CESM is a widely-used Earth system model and an important source of simulations used by the Intergovernmental Panel on Climate Change. HOMME version: https://svn-homme-model.cgd.ucar.edu/branch_tags/dungeon_tags/dungeon016. Compiled with "-O3 -fp-model fast -xCORE-AVX2". Running "perfTestWACCM" benchmark from code repository with size NE=8. Executed with 64 MPI ranks and OMP_NUM_THREADS=1. E5-2697 v4: 2S Intel® Xeon® processor CPU E5-2697 v4 , 2.3GHz, 36 cores, turbo and HT on, BIOS 86B0271.R00, 128GB total memory, 8 slots / 16GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. Gold 6148: Intel® Xeon® Gold 6148 processor, 2.4GHz, 40 cores, turbo and HT on, BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800GB INTEL SSD SC2BA80, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

INTEL® XEON® SCALABLE PROCESSORS

Artificial Intelligence

INTEL® XEON® PROCESSOR PLATFORM PERFORMANCE ON AI

Hardware plus optimized software

INFERENCE THROUGHPUT

Up to **198x**

Intel® Xeon® Platinum 8180 Processor
higher Intel optimized Caffe GoogleNet v1 with Intel® MKL
inference throughput compared to
Intel® Xeon® Processor E5-2699 v3 with BVLC-Caffe

TRAINING THROUGHPUT

Up to **127x**

Intel® Xeon® Platinum 8180 Processor
higher Intel Optimized Caffe AlexNet with Intel® MKL
training throughput compared to
Intel® Xeon® Processor E5-2699 v3 with BVLC-Caffe

Inference and training throughput uses FP32 instructions

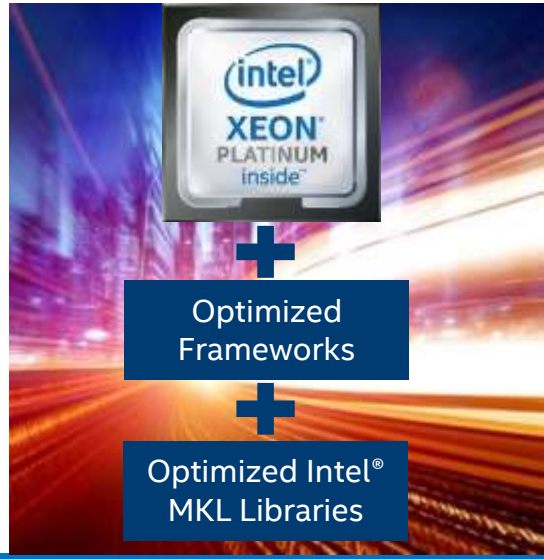

Deliver significant AI performance with hardware and software optimizations on Intel® Xeon® Scalable Family

Up to 191X Intel® Xeon® Platinum 8180 Processor higher Intel optimized Caffe Resnet50 with Intel® MKL inference throughput compared to Intel® Xeon® Processor E5-2699 v3 with BVLC-Caffe
Up to 93X Intel® Xeon® Platinum 8180 Processor higher Intel optimized Caffe Resnet50 with Intel® MKL training throughput compared to Intel® Xeon® Processor E5-2699 v3 with BVLC-Caffe

Performance estimates were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system.

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: http://www.intel.com/performance Source: Intel measured as of June 2017. Configurations: See the last slide in this presentation. *Other names and brands may be claimed as the property of others.

INTEL® XEON® PROCESSOR SCALABLE FAMILY

Scalable performance for widest variety of AI & other datacenter workloads – including deep learning

THE FOUNDATION FOR AI

BUILT-IN ROI

Begin your AI journey today using existing, familiar infrastructure

POTENT PERFORMANCE

DL training in ~~days~~ HOURS with up to 198X[1] perf vs. prior gen

PRODUCTION-READY

Robust support for full range of AI deployments

[1]Configuration details on last slide in presentation

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: http://www.intel.com/performance Source: Intel measured as of November 2016

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

LEARN MORE

Intel Xeon Scalable processors: intel.com/xeonscalable

Intel Omni-Path Architecture: intel.com/omnipath

Intel Cluster Checker: https://software.intel.com/intel-cluster-checker

Intel Parallel Studio XE: https://software.intel.com/parallel-studio-xe

Intel Scalable System Framework: intel.com/ssf

Intel Select Solutions: intel.com/selectsolutions

Intel Software Solutions Showcase: intel.com/softwaresolutions


THANK YOU!