# 

# Intel HW roadmap and Architecture Specifics

OneAPI workshop with FocusCoE, March 02, 2022

Dr. Jean-Laurent Philippe, Intel EMEA HPC Director

## Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at <u>www.Intel.com/PerformanceIndex</u>

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

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

Includes the effect of Intel Thermal Velocity Boost, a feature that opportunistically and automatically increases clock frequency above single-core and multi-core Intel Turbo Boost Technology frequencies based on how much the processor is operating below its maximum temperature and whether turbo power budget is available. The frequency gain and duration is dependent on the workload, capabilities of the processor and the processor cooling solution.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

All product plans and roadmaps are subject to change without notice.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.



- Server CPUs, focus on Ice Lake-SP and Sapphire Rapids processors
- Intel HW Discrete Graphics Accelerators, focus on Ponte Vecchio



- Server CPUs, focus on Ice Lake-SP and Sapphire Rapids processors
- Intel HW Discrete Graphics Accelerators, focus on Ponte Vecchio

**3<sup>rd</sup> Gen** Intel Xeon Scalable processor Codenamed "Ice Lake"



## Meeting the Demands of Real-World Applications

Updated core architecture, balanced design, fast I/O and increased core counts & memory bandwidth

Flexibility of core and frequency configurations on a single CPU to meet differing application ......... performance requirements

Only mainstream CPU with built-in acceleration for both AI and HPC workloads

years of ecosystem enablement and investment means most HPC applications are optimized for Intel®

Xeon<sup>®</sup> Scalable processors

intel. Xeon<sup>\*</sup>

Easy to deploy on-prem or in Cloud, supported by an unmatched portfolio of libraries, tuners and compilers along with Intel validated solutions

Innovative large capacity system memory and
storage architecture for fine grain data access and high performance

High core, high frequency SKU designed for liquidcooled systems

Built-in set of new instructions and enhanced security capabilities that help address current and future privacy and security concerns

# Intel offers flexible performance to change with your HPC needs

6

## 3rd Gen Intel<sup>®</sup> Xeon<sup>®</sup> Scalable processors for HPC in Financial Services



More & better-performing cores

20% IPC boost, Higher L1 & L2 cache

#### Higher memory bandwidth 8 DDR4 channels & 3200MT/s

### Up to 6TB total memory / socket with Intel® Optane™ PMem

**Faster I/O** via PCIe Gen 4 with 64 lanes/socket

#### Built-in acceleration & flexibility AVX-512, DL Boost, Speed Select

# Key 3<sup>rd</sup> Gen Improvements

### Designed to Move Faster, Store More, Process Everything

Feature	2nd Gen Intel® Xeon® Scalable Processor	3rd Gen Intel <sup>®</sup> Xeon <sup>®</sup> Scalable Processors	Notes
Cores per Socket	4-28	8-40	New Sunny Cove architecture
L1/L2/L3 cache per core	32KB/1MB/1.375MB	48KB/1.25MB/1.5MB	Larger caches to enable fast access to data
Memory Channels and DIMM Speed	6 Up to 2933	8 Up to 3200	Huge boost in memory bandwidth & support for Intel® Optane™ PMem 200
Processor Interconnect: UPI links, speed	2 or 3, 10.4 GT/s	2 or 3, 11.2 GT/s	Improved bandwidth between processors
PCIe lanes per socket	PCIe 3.0, 48 Lanes (x16, x8, x4)	PCIe 4.0, 64 lanes (x16, x8,x4)	2x bandwidth and more PCIe lanes to support new Gen 4 SSD, Ethernet and other adjacencies
Workload Acceleration Instructions	AVX-512 VNNI DDIO	AVX-512, VNNI, DDIO vAES, vPCLMULQDQ, VPMADD52, VBMI, PFR, Crypto, SHA extensions, TME, SGX	Enable new capabilities and speedup performance
Platform Adjacencies		Intel® Optane™ PMem 200 series, Intel® Optane™ P5800X SSD, Intel DC P5510 SSD, Intel E810-C ethernet adapter	

#### 3rd Gen Intel<sup>®</sup> Xeon<sup>®</sup> Scalable Processors intel. intel. Xeon Xeon (Two-socket platform) PLATINUM GOLD Performance made flexible Built-in Upto Up to Up to Upto AI&HPC 3.7 GHz 20% 1.53x 40Acceleration increased instructions higher geomean Turbo frequency cores per clock<sup>2</sup> HPC performance per processor Intel<sup>®</sup> Deep Learning Boost single-core & Intel® AVX-512 vs prior gen<sup>1</sup> Advanced Security Technologies Next-Gen Platform Upto Jp to lp to 1001 0110 System Memory DDR4-3200 Lanes Capacity (Per Socket) 2DPC PCI Express 4 Intel Platform Intel Total Memory Memory Capacity Increase vs. Intel Software Guard Intel Crypto (Per Socket) DRAM+PMem (Per Socket) Extensions Firmware Resilience Encryption Acceleration prior gen Scalable, Flexible, Customizable oneAPI Intel ® Optane<sup>™</sup> Persistent Intel Speed Select Optimized Memory 200 series Software Technology

See [108] at www.intel.com/3gen-xeon-config. Results may vary.

9

### Intel® AVX -512 Built-in acceleration and outstanding performance

3rd Gen Intel Xeon Scalable processors are the only data center CPU with 512-bit instruction processing. This wider vectorization speeds computation processes per clock, increasing frequency over the prior generation.

AVX 512 takes advantage of the processors' increased memory bandwidth new core architecture, improved frequency management .Additionally, 2 \*FMA throughput is now available across the Platinum, Gold and Silver skus

The latest version of Intel AVX-512 is easy to enable by using Intel tools, Intel compilers, Intel Math Kernel Library (Intel MKL), and Intel oneAPI Data Analytics Library. All of these have Intel AVX-512 built right in.

Vectorization and can dramatically increase performance for neural network processing, and for other operations that can be parallelized in this way.



### Learn more at intel.com/avx512

<sup>1</sup> Intel<sup>®</sup> AVX 2.0 delivers 16 double precision and 32 single precision floating point operations per second per clock cycle within the 256-bit vectors, with up to two 256-bit fused-multiply add (FMA) units. For workloads and configurations visit www.Intel.com/PerformanceIndex. Results may vary. Check full sku configurations for feature availability.



10

## Intel® Deep Learning Boost A Vector Neural Network Instruction (VNNI) Extends Intel® AVX-512 to accelerate AI/DL Inference

### **Problems Solved**

For workloads and configurations see backup and visit www.Intel.com/PerformanceIndex. Results may vary.

Performance

### End Customer Value



Framework Libraries

Intel Proprietary – OneAPI workshop with FocusCoE, March 02, 2022

Comms Service

Providers

11

Enterprise

**Cloud Service** 

Providers

## Intel® Speed Select Technology



## Flexible Performance

Family of features designed to give you more active and nuanced control over CPU performance

- 3 configurations in a single CPU
- Match your CPU to your workload needs
- Dynamically re-configure processors at runtime
- Enables more efficient systems
- Reduce complexities & costs

## Intel<sup>®</sup> Speed Select Technology Features

Intel® Speed Select Technology is an umbrella term for a collection of features that provide more granular control over CPU performance



\* Power control unit (PCU)

Performance made flexible.

# +49% Geomean

- Still

### **Generational performance leap** for 20 critical HPC workloads

Section of the section of



Intel Proprietary – OneAPI workshop with FocusCoE, March 02, 2022

See backup for configuration details. Results may vary.

# +32% Geomean

# Better performance across workloads that matter

intel XEON Scalable

VS. AMD proce

processors





Codenamed "Ice Lake"







- Server CPUs, focus on Ice Lake-SP and Sapphire Rapids processors
- Intel HW Discrete Graphics Accelerators, focus on Ponte Vecchio

Codenamed "Ice Lake"





## Next Gen Xeon Scalable Processors

Codenamed Sapphire Rapids

### **3<sup>rd</sup> Gen** Xeon Processors

Codenamed Ice Lake





2021



### Ponte Vecchio



### Next Gen Xeon Processors

### **3<sup>rd</sup> Gen** Xeon Processors

### CPU









### Ponte Vecchio



### Next Gen Xeon Processors +**HBM**









**CPU+HBM** 









# Solving the Entire Workload



# Solving the Entire Workload



# Solving the Entire Workload

Next Gen Intel<sup>®</sup> Xeon<sup>®</sup> Scalable Processors codenamed **Sapphire Rapids**  Xe HPC based GPU
Ponte Vecchio



## Next Gen Xeon®

Processors

Built for Supercomputing

----

1999 Carline

Annun .





Intel Proprietary – OneAPI workshop with FocusCoE, March 02, 2022

HBM

Accelerators

CHARTER CLARACE CONTRACT

THE PARTY

11

UPI

Memory Controller

Memory Controller

UPI

Accelerators

HBM



- Server CPUs, focus on Ice Lake-SP and Sapphire Rapids processors
- Intel HW Discrete Graphics Accelerators, focus on Ponte Vecchio

## Sapphire Rapids

Next Gen Intel<sup>®</sup> Xeon<sup>®</sup> Scalable Processors

### Ponte Vecchio X<sup>e</sup> HPC based GPU

manne





Ponte Vecchio X<sup>e</sup> HPC based GPU





## Deep Dive Inside the Xe-HPC Micro-Architecture

Built on four hierarchical building blocks



## The Xe Core



## The Xe Slice

### 16 Xe-cores totaling 8MB of L1 Cache

### 16 Ray Tracing Units

that provide fixed-function computation for Ray Traversal, Bounding Box Intersection, and Triangle Intersection

### 1Hardware Context

that enables execution of multiple applications concurrently without expensive software-based context switches

X <sup>e</sup> . slice							
Xecore Main and Main Main and Main and Main Main and Main and Main Main and Main and Main and Main and Main Main and Main and Main and Main and Main and Main and Main Main and Main and Main Main and Main an	X <sup>0</sup> -core Table an an Anne Table an an Anne Table an an Anne Table an Anne Table an Anne Table an Anne Table an Anne Table an Anne	X9-core	X <sup>0</sup> -core <u>100</u> 00 00 000 <u>100</u> 00 00 000 <u>100</u> 00 00 000 <u>100</u> 00 00 000 <u>100</u> 00 00 <u>100</u> 000 <u>100</u> 000 <u>100</u> 000	Non-corre           Nation         Anno         Anno           Nation         Anno         Anno	X <sup>0</sup> -core 1017 00 00 000 1017 000 1000 1000 1000 1000 1000 1000 1000 1000 1000	X <sup>e</sup> -core 2 2 2 2 2 2 2 2 2 2 2 2 2	X°-core
No Taxon Link	Hay Turing Link	Ray Taxang Link	Ray Taning Link	Hay Tuning Unit	Reflected	HarTrangibil	RegTinengthet
Ray Taxing Link	Ray Tracing Line	Ray Teory Line	Ray Tracing Unit	Ray Tocarg Unit	Ray Taking Like	RoTrangilit	ReyTourseyLee
		H LAISAN					

## The Xe Stack

#### Contains:

4 Slices which adds up to 64 Xe Cores, 64 Ray Tracing Units, and 4 hardware contexts

Massive L2 Cache (up to 204MB)

1 Media Engine

 $8 X^{e}$ -Links (90GB/s per link)



Scalable to 2 stacks allowing for multi-stack design

# Ponte Vecchio X<sup>e</sup> Link Tile

Per Tile <b>8 X<sup>e</sup> Links</b>	8 ports Embedded Switch
Built on TSMC N7	Up to <b>90G</b> Serdes



# Ponte Vecchio Execution Progress

AO Silicon Current Status

## >45 TFLOPS FP32 Throughput



## >5 TBps

Memory Fabric Bandwidth

>2TBps

Connectivity Bandwidth





Performance

Install Base

intel. 39





Performance

intel. 40



# Falcon Shores

## XPU

Scalable Architecture for all Super Computing Workloads



#### Performance

oneAPI

For Illustrative purposes only, represents scalable architecture designed to address full install base across a range of performance requirements.







Performance

Performance targets based on estimates relative to current platforms in February 2022. Results may vary.

Thank you

Intel Proprietary – OneAPI workshop with EuroCC/Castiell, February 16-17, 2022

# 

## Configuration Details: Generational Performance, slide 8 (1 of 2)

For the following applications (8280 vs 8380): NAMD, LAMMPS, RELION, GROMACS, Monte Carlo, Black-Scholes, Binomial, LS-Dyna, Fluent, OpenFOAM, RADIOSS, Converge, Numeca, WRF: See [108] at www.intel.com/3gen-xeon-config. Results may vary.

Ansys Mechanical: Intel 8280; Config Date: 1/21/2021; Platform: Wolf Pass; CPU Details: 2 CPUs per node; Stepping: 6; 28c @ 2.7Ghz; # CPU Sockets: 2; # CPU Cores: 28; CPU Base Frequency: 2.7 GHz; CPU Max Frequency: 4.0 GHz; CPU Base TDP: 205 W; RAM: 192GB, 12x16GB 2933MHz DDR4, Micron 18ASF2G72PDZ-2G9E1; Hard Drive: 480GB Model: ATA INTEL SSDSC2KG48 (scsi); Cluster File System: OPA based Lustre; BIOS: SE5C620.86B.02.01.0012.070720200218; BIOS Settings: HT=on TURBO=ON; Microcode: 0x4002f01; Intel Management: 04.01.04.381; BMC: 2.48. Operating System: CentOS Linux 8.3.2011; Kernel: 4.18.0-240.1.1.el8\_3.crt1.x86\_64; OFED stack: OPA 10.10.3.1.1, Lustre 2.10.8-default.App Version: 2021 R1; Data collected: 2/2/2021; Build notes: Hyper-Threading Enabled, Turbo Enabled, Intel FORTRAN Compiler 19.0.0; Intel C/C++ Compiler 19.0.0; Intel Math Kernel Library 2020.0.0; Intel MPI Library 2018 Update 3; One thread per core, 52ppn. Intel 8380; Config Date: 1/21/2021; Platform: Coyote Pass; CPU Details: 2 CPUs per node; Stepping: 6; 40c @ 2.3Ghz; # CPU Sockets: 2; # CPU Cores: 40; CPU Base Frequency: 2.3 GHz; CPU Max Frequency: 2.3 GHz; CPU Base TDP: 250 W; RAM: 256GB 16\*16GB 3200MT/s DDR4, Hynix HMA82GR7CJR8N-XN; Hard Drive: SSDSC2KG96 960GB; Cluster File System: HDR based Lustre; BIOS: SE5C6200.86B.2021.D40.2103100308; BIOS Settings: HT=on TURBO=ON; Microcode: 0x8d055260; Intel Management: 04.04.03.249; BMC: 2.66; Operating System: CentOS Linux release 8.3.2011; Kernel: 4.18.0-240.1.1.el8\_3.crt1.x86\_64; OFED stack: Mellanox mlnx-5.1-2.5.8.0-default. App Version: 2021 R1; Data collected: 4/8/2021; Build notes: Intel(R) FORTRAN Compiler 19.0.0; Intel(R) C/C++ Compiler 19.0.0; Intel(R) Math Kernel Library 2020.0.0; Intel(R) MPI Library 2018 Update; cores: 80ppn. Tested by Intel as of October 2021.

Simulia Abaqus/Explicit: Intel Xeon Platinum 8380/2/6 (40c) ; 2.3/2.1/1.8 . RAM 256GB 16\*16GB 3200MT/s DDR4, Hynix HMA82GR7CJR8N-XN; SSDSC2KG96 960GB; HDR based Lustre; BIOS SE5C6200.86B.0020.P23.2103261309; HT=ON; TURBO=ON; Microcode 0xd000270; Intel Management Engine: 04.04.04.053; BMC 2.78; OFED stack Mellanox mlnx-5.1-2.5.8.0-default. Intel Xeon Platinum 8280L 2/6 (2.7GHz) 2.7/2.2/1.8. RAM 192GB, 12x16GB 2933MHz DDR4, Hynix HMA82GR7CJR8N-WM; hard drive 480GB Model: ATA INTEL SSDSC2KG48 (scsi); OPA based Lustre; BIOS SE5C620.86B.02.01.0012.070720200218; HT=ON; TURBO=ON; Microcode 0x4002f01. Abaqus2021HF4. MKL 2020 Update 2. Tested by Intel as of October 2021.

Simulia Abaqus/Standard: Intel Xeon Platinum 8380/2/6 (40c) ; 2.3/2.1/1.8 . RAM 256GB 16\*16GB 3200MT/s DDR4, Hynix HMA82GR7CJR8N-XN; SSDSC2KG96 960GB; HDR based Lustre; BIOS SE5C6200.86B.0020.P23.2103261309; HT=ON; TURBO=ON; Microcode 0xd000270; Intel Management Engine: 04.04.04.053; BMC 2.78; OFED stack Mellanox mlnx-5.1-2.5.8.0-default. Intel Xeon Platinum 8280L 2/6 (2.7GHz) 2.7/2.2/1.8. RAM 192GB, 12x16GB 2933MHz DDR4, Hynix HMA82GR7CJR8N-WM; hard drive 480GB Model: ATA INTEL SSDSC2KG48 (scsi); OPA based Lustre; BIOS SE5C620.86B.02.01.0012.070720200218; HT=ON; TURBO=ON; Microcode 0x4002f01. Abaqus2021HF4. MKL 2020 Update 2. Tested by Intel as of October 2021.

Brightskies Reverse Time Migration: Intel Xeon Platinum 8280L CPU @ 2.70GHz: Platform/Motherboard: S2600WFD (Wolf Pass) J46732-801; BIOS: SE5C620.86B.0X.02.0001.051420190324; RAM: 12x16GB (192GB) DDR4 2666; OS: Ubuntu 20.04.3 LTS; Kernel: 5.4.0-89-generic; Microcode: 0x5003102. Intel Xeon Platinum 8380 CPU @ 2.30GHz: Platform/Motherboard: M50CYP2SB2U (Coyote Pass) K88091-301; BIOS: SE5C6200.86B.0022.D64.2105220049; RAM: 16x32GB (512GB) DDR4 3200; OS: Ubuntu 20.04.3 LTS; Kernel: 5.4.0-89-generic; Microcode: 0xd0002b1. Application version 4.0.0. Build: Intel® oneAPI 2021.4. CBS = x (best score). Tested by Intel as of October 2021.

## Configuration Details: Generational Performance, slide 8 (2 of 2)

GATK whole genome sequencing:

- Config 1: Test by Intel as of 11/14/2019. 1 application node and 4 compute nodes. Application node configuration: 2x Intel Xeon Gold 6252 processor (24 cores, 2.10 GHz); 1x Intel<sup>®</sup> Server Board S2600WFT; 192 GB (12x 16 GB DDR4 2666 MHz, 1DC); 1x 960 GB Intel<sup>®</sup> SSD D3-S4510 Series (2.5 in SATA 6 Gb/s, 3D2,TLC); 1x 1.6 TB Intel<sup>®</sup> SSD DC P4610 Series (2.5 in PCIe 3.1 x4, 3D2, TLC); Microcode: 0x500002c, BIOS: SE5C620.86B.02.01.0009.092820190230; CentOS Linux Installation ISO (minimal or full) 7.7 build 1910; Intel<sup>®</sup> Cluster Runtimes 2019.4;Intel<sup>®</sup> Cluster Checker 2019.3.5; Intel<sup>®</sup> Select HPC Solution for RPM packages for EL7 2018.0;OpenHPC 1.3.8. 4x compute nodes configuration: 2x Intel Xeon Gold 6252 processor (24 cores, 2.10 GHz); 1x Intel Server Board S2600WFT; 384 GB (12x 32 GB DDR4 2933MHz); 1x 960 GB Intel SSD D3-S4510 Series (2.5 in SATA 6 Gb/s, 3D2, TLC); 1x 1.6 TB Intel SSD DC P4610 Series (2.5 in PCIe 3.1 x4, 3D2, TLC); Network devices: 1x Intel<sup>®</sup> C620 Series Chipset Ethernet Connection; Intel<sup>®</sup> Ethernet Adapter X722 onboard 10 GbE; Microcode: 0x500002c, BIOS: SE5C620.86B.02.01.0009.092820190230; 1x distributed 10 GB Lustre 2.10 ZFS system, 6 OST, 3 OSS, Lnet Router with single 10 GB link for all I/O traffic clients to Lustre servers plus 1X 1.6 TB Intel SSD DC P4610 Series (2.5 in PCIe 3.1 x4, 3D2, TLC).
- Config 2: Test by Intel as of 8/8/2021. 1 front-end node and 4 compute nodes, all using Intel® Server Board M50CYP2SB-003; Front-end node configuration: 2x Intel® Xeon® Gold 6348 processor (28 cores) 2.90 GHz, Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, total memory 256 GB (16 slots/16 GB/3200 MHz); BIOS version: 22D08; BMC 2.66, SDR 0.31, CPLD 3p0; uCode: 0x0b000280; CentOS Linux installation ISO (minimal or full) 8 build 2011; storage boot drive 1x Intel® SSD P4610 1.6 TB (3D NAND PCIe 3.1 x4, 3D1, TLC); high-performance network: 1x Intel® Ethernet Converged Network Adapter X550-T2, model X550T2. Compute node configuration: 2x Intel Xeon Gold 6348 processor (28 cores @2.60 GHz,) Intel Hyper-Threading Technology = ON, Intel Turbo Boost Technology = ON, total memory 512 GB (16 slots/32 GB/3200 MHz); BIOS version: 22D08; BMC 2.66, SDR 0.31, CPLD 3p0; uCode: 0x0b000280; storage Scratch drive: 1x Intel SSD P4610 1.6 TB (3D NAND PCIe 3.1 x4, 3D1, TLC); high-performance network: 1x Intel Ethernet Converged Network Adapter X550-T2 (10 GbE), model X550T2.

EDA front end and back end (7 workloads): Xeon 6248R (24 cores): 3.0GHz/4.0GHz/3.6GHz; HDD 2x 1.2TB SAS; BIOS 3.3-IN001, 02/12/2020; OS SLES12 SP5; kernel 4.12.14-122.7-default. Turbo enabled, Hyperthreading disabled. Xeon 6342 (24 cores): 2.8GHz/3.5GHz/3.3GHz; HDD 2x 2.4TB SAS; BIOS 1.0-IN001, 03/12/2021; OS SLES12 SP5; kernel 4.12.14-122.60-default. Turbo enabled, Hyperthreading disabled. Tested by Intel as of April 2021.

### Footnotes for page 16

- 1. Xeon 8380 vs Xeon 8280
- 2. Xeon 6348 vs Xeon 6252
- 3. Xeon 6342 vs 6248R

## Configuration Details: Competitive Performance, slide 9 (1 of 2)

Binomial Options: 2S Xeon® Platinum 8380 (40C, 2.3GHz, 270W): Intel "Coyote Pass"; 256GB, 16x16GB 3200MHz DDR4; CentOS Linux 8.3.2011, 4.18.0-240.1.1.el8\_3.crt1.x86\_64, SE5C6200.86B.0021.D40.2101090208, 0x261. App Version: v1.0; Build notes: Tools: Intel C Compiler 2020u4, Intel Threading Building Blocks ; threads/core: 2; Turbo: used; Build knobs: - O3 -xCORE-AVX512 -qopt-zmm-usage=high -fimf-domain-exclusion=31 -fimf-accuracy-bits=11 -no-prec-div -no-prec-sqrt. 2S AMD EPYC 7763 (64C, 2.45GHz, 280W): GIGABYTE R282-Z92 server; 512GB, 16x32GB 3200MHz DDR4; NPS4, CTDP=280W, Determinism=Power; Red Hat Enterprise Linux 8.3, 4.18, ucode 0xa001114. App Version: v1.0; Build notes: Tools: Intel C Compiler 2020u4, Intel Threading Building Blocks ; threads/core: 2; Turbo: used; Build knobs: -O3 -march=core-avx2 -fimf-domain-exclusion=31 -fimf-accuracy-bits=11 -no-prec-div -no-prec-sqrt. Workloads tested by Intel and results as of May 2021.

Monte Carlo: See [37] at www.intel.com/3gen-xeon-config. Results may vary.

MatLogica Provided Vector Accelerator Library XVA pricing benchmark: 2S Intel® Xeon® Platinum 8380 @ 2.30GHz 40 cores on Intel platform with 512 GB DDR4 memory (8(1DPC)/32GB/SK Hynix/1.2v/3200 MT/s), HT on, Turbo on, CentOS Linux 8.4, internal Bios SE5C6200.86B.0020.P23.2103261309 Release Date: 03/16/2021, CentOS Linux release 8.4.2105, Metalogical Libraries AADC-demo-2021-10-01-cd0737f-M9s6, Run Instructions: taskset -c \$EXEFILE \$CONFIG \$INPUT file (for Intel® XEON® avx512 CONFIG=512 used), export OMP\_NUM\_THREAD=1, 2 thread/core, optimized with AVX512 and Metalogical AADC libraries, test by Intel on 10/14/2021. 2S AMD EPYC™ 7763 @ 2.45GHz 64-Core Processor on GIGABYTE R282-Z92-00 with 512 GB DDR4 memory (8(1DPC)/32GB/SK Hynix/1.2v/3200 MT/s), HT on, Turbo on, Bios: GIGABYTE M07 Release Date: 09/03/2021, CentOS Linux release 8.4.2105, Metalogical Libraries AADC-demo-2021-10-01-cd0737f-M9s6, Run Instructions: taskset -c \$EXEFILE \$CONFIG \$INPUT file (for AMD EPYC ™ 7763 @ 2.45GHz 64-Core Processor on GIGABYTE R282-Z92-00 with 512 GB DDR4 memory (8(1DPC)/32GB/SK Hynix/1.2v/3200 MT/s), HT on, Turbo on, Bios: GIGABYTE M07 Release Date: 09/03/2021, CentOS Linux release 8.4.2105, Metalogical Libraries AADC-demo-2021-10-01-cd0737f-M9s6, Run Instructions: taskset -c \$EXEFILE \$CONFIG \$INPUT file (for AMD EPYC avx2 CONFIG=256 used), export OMP\_NUM\_THREAD=1, 2 threads/core, optimized with AVX2 and Metalogical AADC libraries, test by Intel on 10/14/2021.

Quantifi Credit Option Pricing Al Inference: 2S Intel Xeon 8380 CPU @ 2.30 GHz (40 cores/processor), Turbo ON, HT ON. 512GB DDR4-3200. CentOS Linux Version 8. BIOS Version: SE5C6200.86B.0022.D08.2103221623. BIOS Release Date: 03/22/2021. 2S AMD EPYC 7763 @ 2.45 GHz (64 cores/processor). Turbo ON, HT ON. 512GB DDR4-3200. CentOS Linux Version 8. BIOS Version: M06. BIOS Release Date: 07/10/2021. Tested by Intel. Libraries: python 3.8.11, intel-tensorflow 2.6.0. Turbo=ON. All platforms: python Quantifi\_Inference.py. 2 threads/core (all platforms). Tested by Intel as of 10-12-2021.

Quantum ESPRESSO AUSURF112, PSIWAT: 2S Xeon® Platinum 8380 (40C, 2.3GHz, 270W): Intel "Coyote Pass"; 256GB, 16x16GB 3200MHz DDR4; CentOS Linux 8.3.2011, 4.18.0-240.1.1.el8\_3.crt1.x86\_64, SE5C6200.86B.0021.D40.2101090208, 0x261. App Version: 6.3; Build notes: Tools: Intel Fortran Compiler 2021.2, Intel MKL 2021.2, Intel MPI 2021.2, ELPA 2020.11; threads/core: 2; Turbo: used; Build knobs: -O2 -align array64byte -threads -heap-arrays 4096 -xCORE-AVX512 -qopt-zmm-usage=high -fp-model fast=2 -complex-limited-range assume byterecl –qopenmp. 2S AMD EPYC 7763 (64C, 2.45GHz, 280W): GIGABYTE R282-Z92 server; 512GB, 16x32GB 3200MHz DDR4; NPS4, CTDP=280W, Determinism=Power; Red Hat Enterprise Linux 8.3, 4.18, ucode 0xa001114.. App Version: 6.3; Build notes: Tools: Intel Fortran Compiler 2021.2, Intel MKL 2021.2, ELPA 2020.11; threads/core: 2; Turbo: used; Build knobs: -O2 -align array64byte -threads -heap-arrays 4096 -march=core-avx2 -fp-model fast=2 -complex-limited-range -assume byterecl –qopenmp. Workloads tested by Intel and results as of May 2021.

BWA-MEM2 on Intel Xeon vs AMD EPYC: 2S Intel<sup>®</sup> Xeon<sup>®</sup> Platinum 8380 (40C, 2.3GHz, 270W). Intel "Coyote Pass" - 256GB, 16x16GB 3200MHz DDR4; CentOS Linux release 8.4.2105, 4.18.0-240.22.1.el8\_3.crt4.x86\_64, ucode 0xd000270. 2S AMD EPYC 7763 (64C, 2.45GHz, 280W): GIGABYTE R282-Z92 server - 512GB; 16x32GB 3200MHz DDR4; CentOS Linux release 8.4.2105, 4.18.0-240.22.1.el8\_3.crt4.x86\_64, ucode 0xa00111d; NPS=1 yields best performance for EPYC. Baseline code: BWA-MEM, v0.7.17; optimized code (OpenOmics): BWA-MEM2, v2.2.1. Datasets used: Reference sequence: version GRCh38; Read dataset: 10M reads from SRR7733443 (151bp each) and HG004 (250bp each). Tested by Intel as of September 2021.

## Configuration Details: Competitive Performance, slide 9 (2 of 2)

LAMMPS (Polyethylene, Stillinger-Weber, Tersoff, Water): ): 2S Xeon® Platinum 8380 (40C, 2.3GHz, 270W): Intel "Coyote Pass"; 256GB, 16x16GB 3200MHz DDR4; CentOS Linux 8.3.2011, 4.18.0-240.1.1.el8\_3.crt1.x86\_64, SE5C6200.86B.0021.D40.2101090208, 0x261. App Version: v2020-10-29; Build notes: Tools: Intel MKL 2020u4, Intel C Compiler 2020u4, Intel Threading Building Blocks 2020u4, Intel MPI 2019u8; threads/core: 2; Turbo: used; Build knobs: -O3 -ip -xCORE-AVX512 -qopt-zmm-usage=high. 2S AMD EPYC 7763 (64C, 2.45GHz, 280W): GIGABYTE R282-Z92 server; 512GB, 16x32GB 3200MHz DDR4; NPS4, CTDP=280W, Determinism=Power; Red Hat Enterprise Linux 8.3, 4.18, ucode 0xa001114. App Version: v2020-10-29; Build notes: Tools: Intel MKL 2020u4, Intel C Compiler 2020u4, Intel Threading Building Blocks 2020u4, Intel MPI 2019u8; threads/core: 2; Turbo: used; Build knobs: -O3 -ip -xCORE-AVX512 -qopt-zmm-usage=high. 2S AMD EPYC 7763 (64C, 2.45GHz, 280W): GIGABYTE R282-Z92 server; 512GB, 16x32GB 3200MHz DDR4; NPS4, CTDP=280W, Determinism=Power; Red Hat Enterprise Linux 8.3, 4.18, ucode 0xa001114. App Version: v2020-10-29; Build notes: Tools: Intel MKL 2020u4, Intel C Compiler 2020u4, Intel Threading Building Blocks 2020u4, Intel MPI 2019u8; threads/core: 2; Turbo: used; Build knobs: -O3 -ip -march=core-avx2. Workloads tested by Intel and results as of May 2021.

NAMD: See [36] at www.intel.com/3gen-xeon-config. Results may vary.

RELION: See [38] at www.intel.com/3gen-xeon-config. Results may vary.

Ansys Mechanical 2021 R2 on Endeavor (HDR Fabric): 21% improvement across all SP cases (V21sp-1, V21sp-2, V21sp-3, V21sp-4). Intel® Xeon® Platinum 8358 Processor (32 cores, 48M Cache, 2.60 GHz). RAM 256 GB, 16\*16GB 3200 MT/s DDR4, Hynix HMA82GR7CJR8N-XN. HDR based Lustre. BIOS SE5C6200.86B.0020.P23.2103261309. BIOS settings HT=on, Turbo Boost ON, SNC 2. Microcode 0xd000270. Intel Management Engine 04.04.04.053. BMC 2.78. Cent OS 8.4.2105. Kernel 4.18.0-240.22.1.el8\_3.crt2.x86\_64. OFED stack Mellanox mlnx-5.1-2.5.8.0-default. AMD EPYC<sup>™</sup> 7543 (32 cores, Total L3 Cache 256MB, 2.8 GHz base clock). RAM 256 GB, 16\*16GB 3200 MT/s DDR4, SK Hynix HMA82GR7CJR8N-XN. 240GB ATA Intel SSDSC2KB24. HDR based Lustre. BIOS Ver 2.1 Rev 5.22. Default settings, Multi-threading on, Boost on. NPS4 . Microcode 0xa00111d. BMC 3.1. Cent OS 8.4.2105. Kernel 4.18.0-240.22.1.el8\_3.crt2.x86\_64. OFED stack Mellanox mlnx-5.1-2.5.8.0-default. Intel MPI 2019u9 used throughout. 2018u3 used where needed. Intel MKL used for Intel platforms: Intel(R) Math Kernel Library Version 2020.0.0 Product Build 20191122. AMD BLIS used for AMD platforms. Intel Compiler Version: 2019.0.0. Intel(R) FORTRAN Compiler Version 19.0.0 (Build: 20190206). Tested by Intel as of October 2021.