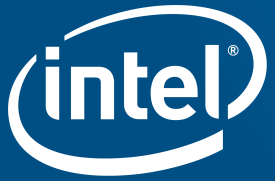

Artificial Intelligence Intel® Xeon Phi™ and Intel® Xeon® Processors

November 13, 2016

Thor Sewell
Director Xeon Phi Marketing
Data Center Group, Intel

Legal Notices and Disclaimers

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL® PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. INTEL PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.

Intel may make changes to specifications and product descriptions at any time, without notice.

All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

Intel, processors, chipsets, and desktop boards 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.

Any code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user.

Intel product plans in this presentation do not constitute Intel plan of record product roadmaps. Please contact your Intel representative to obtain Intel's current plan of record product roadmaps.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests 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 http://www.intel.com/performance

Intel, the Intel logo, Xeon and Xeon logo , Xeon Phi and Xeon Phi logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries

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

This document contains information on products in the design phase of development.

Optimization Notice

Optimization Notice

Intel® compilers, associated libraries and associated development tools may include or utilize options that optimize for instruction sets that are available in both Intel® and non-Intel microprocessors (for example SIMD instruction sets), but do not optimize equally for non-Intel microprocessors. In addition, certain compiler options for Intel compilers, including some that are not specific to Intel micro-architecture, are reserved for Intel microprocessors. For a detailed description of Intel compiler options, including the instruction sets and specific microprocessors they implicate, please refer to the "Intel® Compiler User and Reference Guides" under "Compiler Options." Many library routines that are part of Intel® compiler products are more highly optimized for Intel microprocessors than for other microprocessors. While the compilers and libraries in Intel® compiler products offer optimizations for both Intel and Intel-compatible microprocessors, depending on the options you select, your code and other factors, you likely will get extra performance on Intel microprocessors.

Intel® compilers, associated libraries and associated development tools 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 Intel® Streaming SIMD Extensions 2 (Intel® SSE2), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), and Supplemental Streaming SIMD Extensions 3 (Intel® 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.

While Intel believes our compilers and libraries are excellent choices to assist in obtaining the best performance on Intel® and non-Intel microprocessors, Intel recommends that you evaluate other compilers and libraries to determine which best meet your requirements. We hope to win your business by striving to offer the best performance of any compiler or library; please let us know if you find we do not.

Notice revision #20101101

Intel® Scalable System Framework

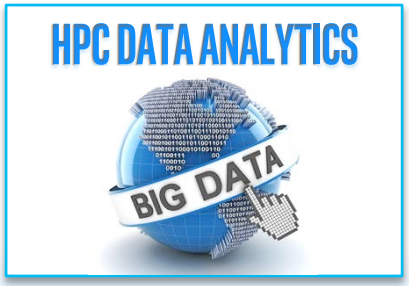

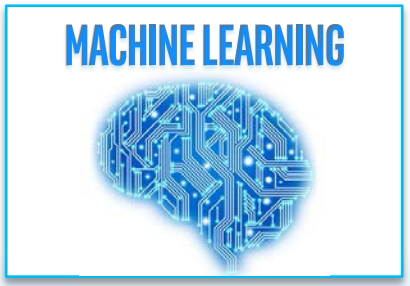

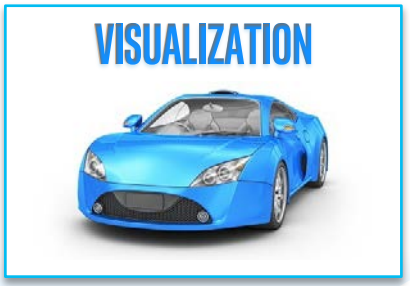

MANY WORKLOADS – ONE FRAMEWORK

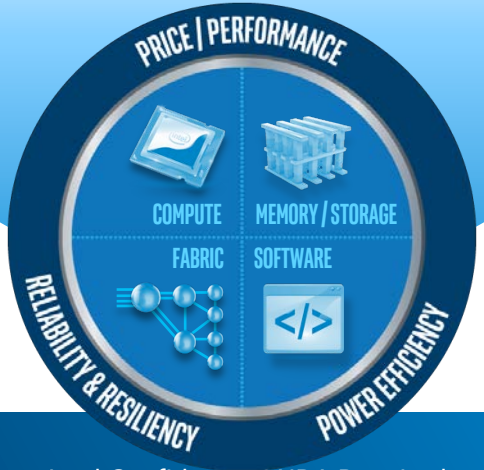

A Flexible Framework for Today & Tomorrow

Enabling Breakthrough System Performance

Agenda

- Artificial Intelligence
  - Intel® Xeon Phi™ Processor
  - Intel® Xeon® Processors
  - Summary

Data + Analytics Creates Unique Opportunities

Companies that use analytics best are...

2X	5X	3X	2X
Make data-driven decisions	Make decisions faster than others	Execute on decisions faster	Have top-quartile financial results

...more likely to

Source: Bain

The Evolution of Artificial Intelligence

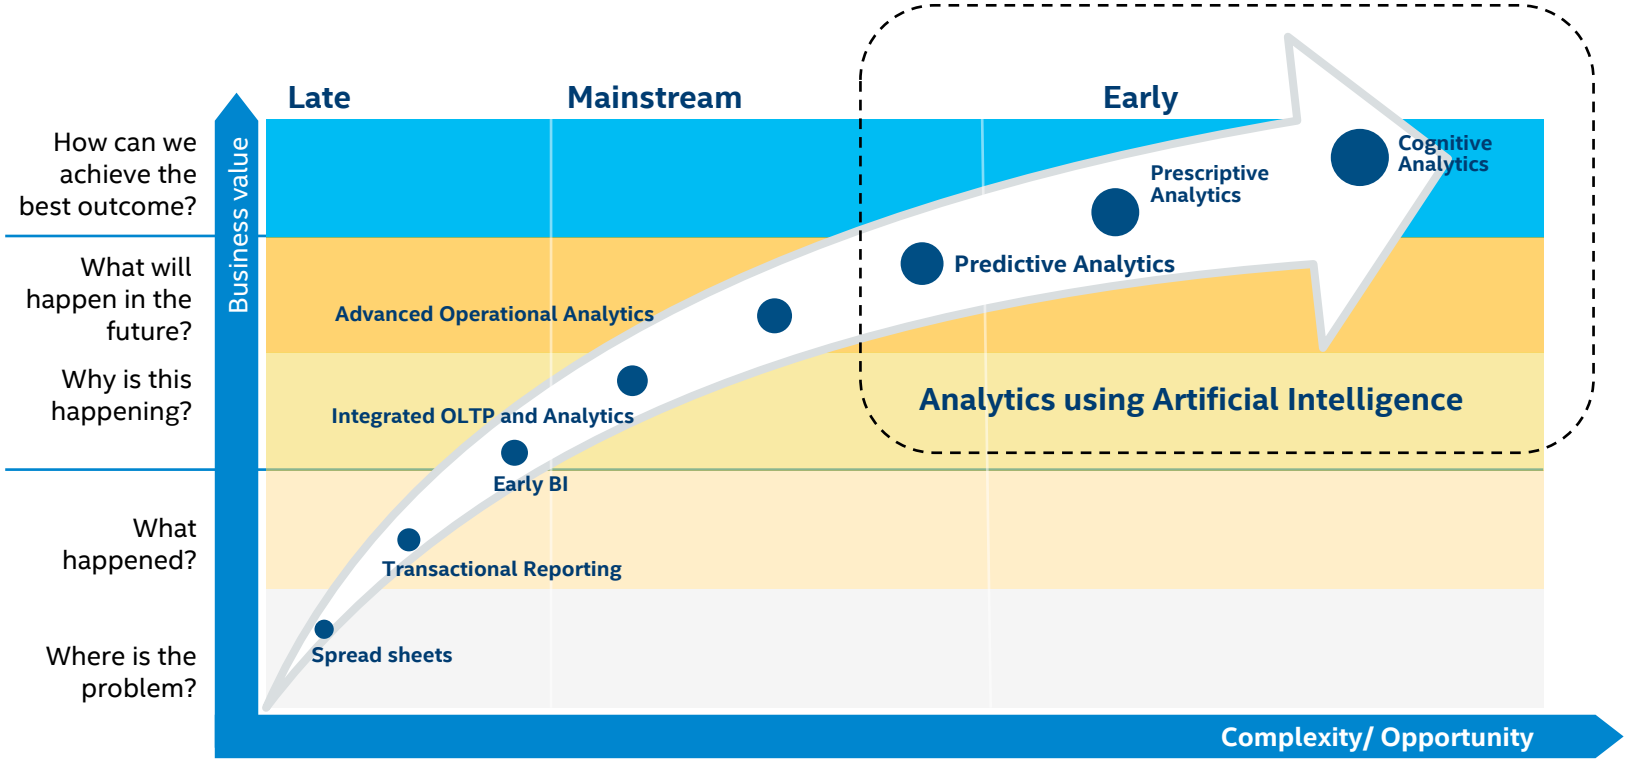

Artificial Intelligence Taxonomy

Artificial Intelligence (AI)

Machines that can sense, reason, act without explicit programming

Machine Learning (ML), a key tool for AI, is the development, and application of algorithms that improve their performance at some task based on experience (previous iterations)

Deep Learning (DL)

Algorithms where abstract ideas are represented by multiple (deep) layers of graphs

- CNN
- RNN
- RBM
- ...

Statistical/Classical Machine Learning

Algorithms based on statistical or other techniques for estimating functions from examples

- Naïve Bayes
- Support Vector Machines
- GA
- Linear Regression

Training: Build a mathematical model based on a data set

Inference: Use trained model to make predictions about new data

Artificial Intelligence Example Use Cases

Cloud Service Providers

Financial Services

Healthcare

Automotive

- Image classification and detection for accurate diagnosis
- Image recognition/ tagging for defect identification
- Natural language recognition (digital assistants)
- Big data pattern detection
- Targeted ads to increase revenue
- Fraud prevention/ face detection
- Gaming, check processing
- Computer server monitoring
- Safe navigation for autonomous vehicles
- Financial forecasting and prediction to avoid risk
- Network intrusion detection

AI Market Opportunity

PRESENT

Server Market (2015)[1]

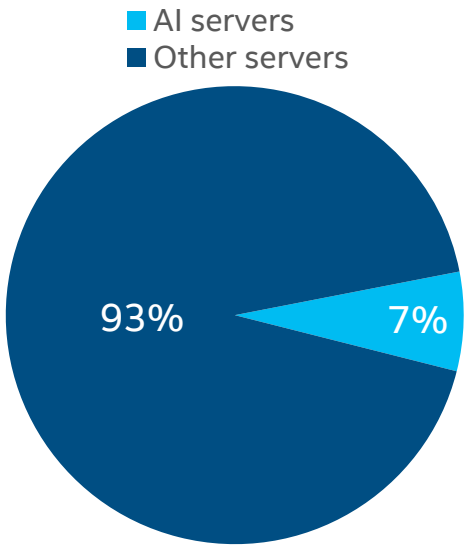

[1]Source: DCG Market Intelligence team

FUTURE

Data is the next disrupter

By 2020, Machine to Machine connections will be 47% of total devices & connections

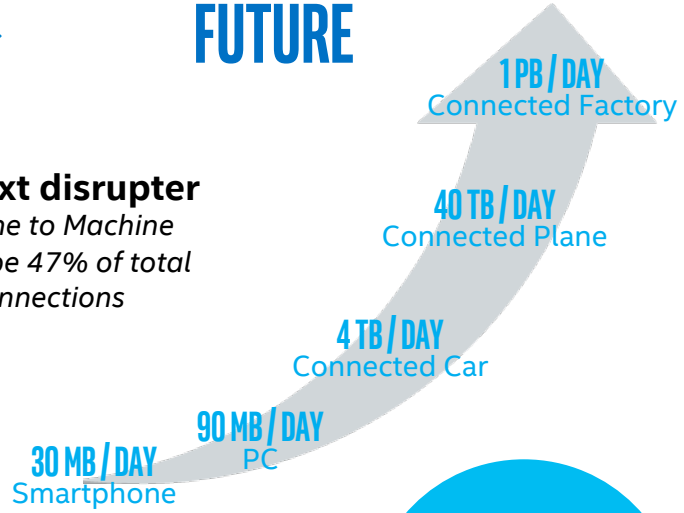

AI-based analytics market[2]

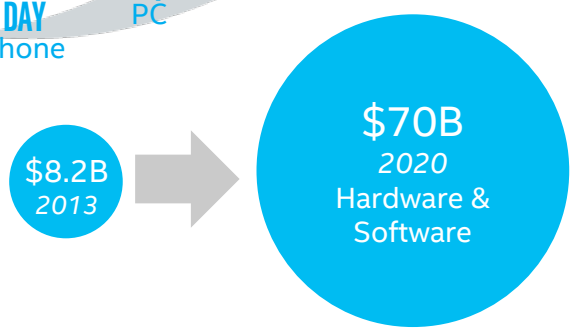

[2]Source: IDC, IOT market related to analytics

Intel AI strategy

Making AI more pervasive by enabling deployment ready AI solutions through a large, open ecosystem

Solution blueprints
for reference across industries

Tools/Platforms
to accelerate deployment of IA solution stack

Optimized Open Frameworks
that scale to multi-node and deliver best performance

Free Libraries/Languages
featuring optimized ML/DL building blocks to enable developers

Best in class hardware
Cross compatible portfolio spanning from data center to edge delivering high perf, perf/TCO, perf/w

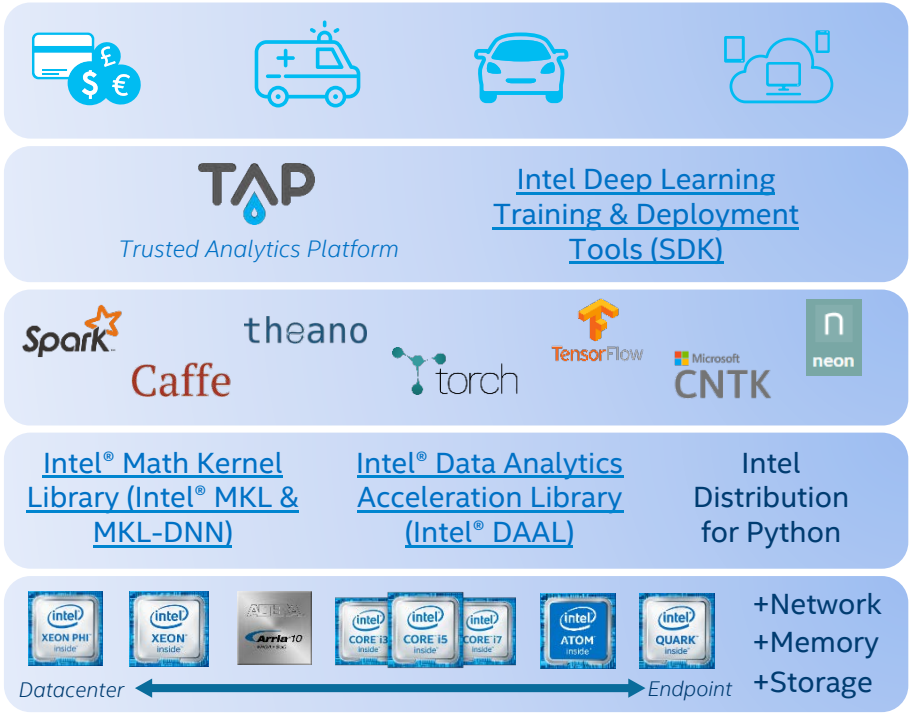

Data Center AI Product Portfolio

Intel® Xeon® Processors

- Most widely deployed machine learning platform
- Optimized for a wide variety of datacenter workloads enabling flexible infrastructure

Intel® Xeon Phi™ Processors

- Higher performance general purpose machine learning solution
- Ideal for HPC and enterprise customers running scale-out, highly parallel, memory intensive applications

Intel® Xeon® Processor + Intel® Nervana Engine

- Best in class neural network performance
- Offers unprecedented compute density with high bandwidth interconnect

Intel® Xeon® Processor + FPGA

- Customizable/Programmable
- Offers low latency, and flexible precision with high perf/w for machine learning inference

General Purpose Infrastructure → **Workload optimized**

Available soon

Deep Learning Tools for End-to-End Workflow

Intel® Deep Learning SDK

Deep Learning Training Tool

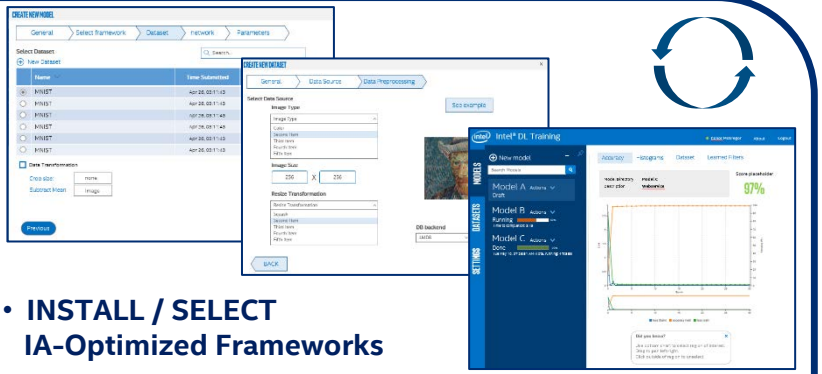

- INSTALL / SELECT IA-Optimized Frameworks
- PREPARE / CREATE Dataset with Ground-truth
- DESIGN / TRAIN Model(s) with IA-Opt. Hyper-Parameters
- MONITOR Training Progress across Candidate Models
- EVALUATE Results and ITERATE

Deep Learning Deployment Tool

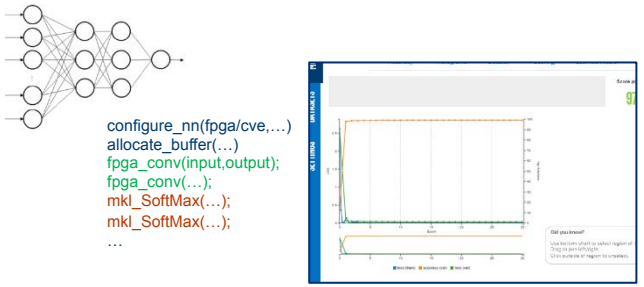

- IMPORT Trained Model (trained on Intel or 3rd Party HW)
- COMPRESS Model for Inference on Target Intel HW
- GENERATE Inference HW-Specific Code (OpenVX, C/C++)
- INTEGRATE with System SW / Application Stack & TUNE
- EVALUATE Results and ITERATE

MKL-DNN Optimized Machine Learning Frameworks

Xeon (local or cloud)

Optimized libraries & run-times (MKL-DNN, OpenVX, OpenCL)
Data acquisition (sensors) and acceleration

Target Inference Hardware Platform (physical or simulated)

Agenda

- Artificial Intelligence
- Intel® Xeon Phi™ Processor
- Intel® Xeon® Processors
- Summary

Intel® Xeon Phi™ Product Roadmap

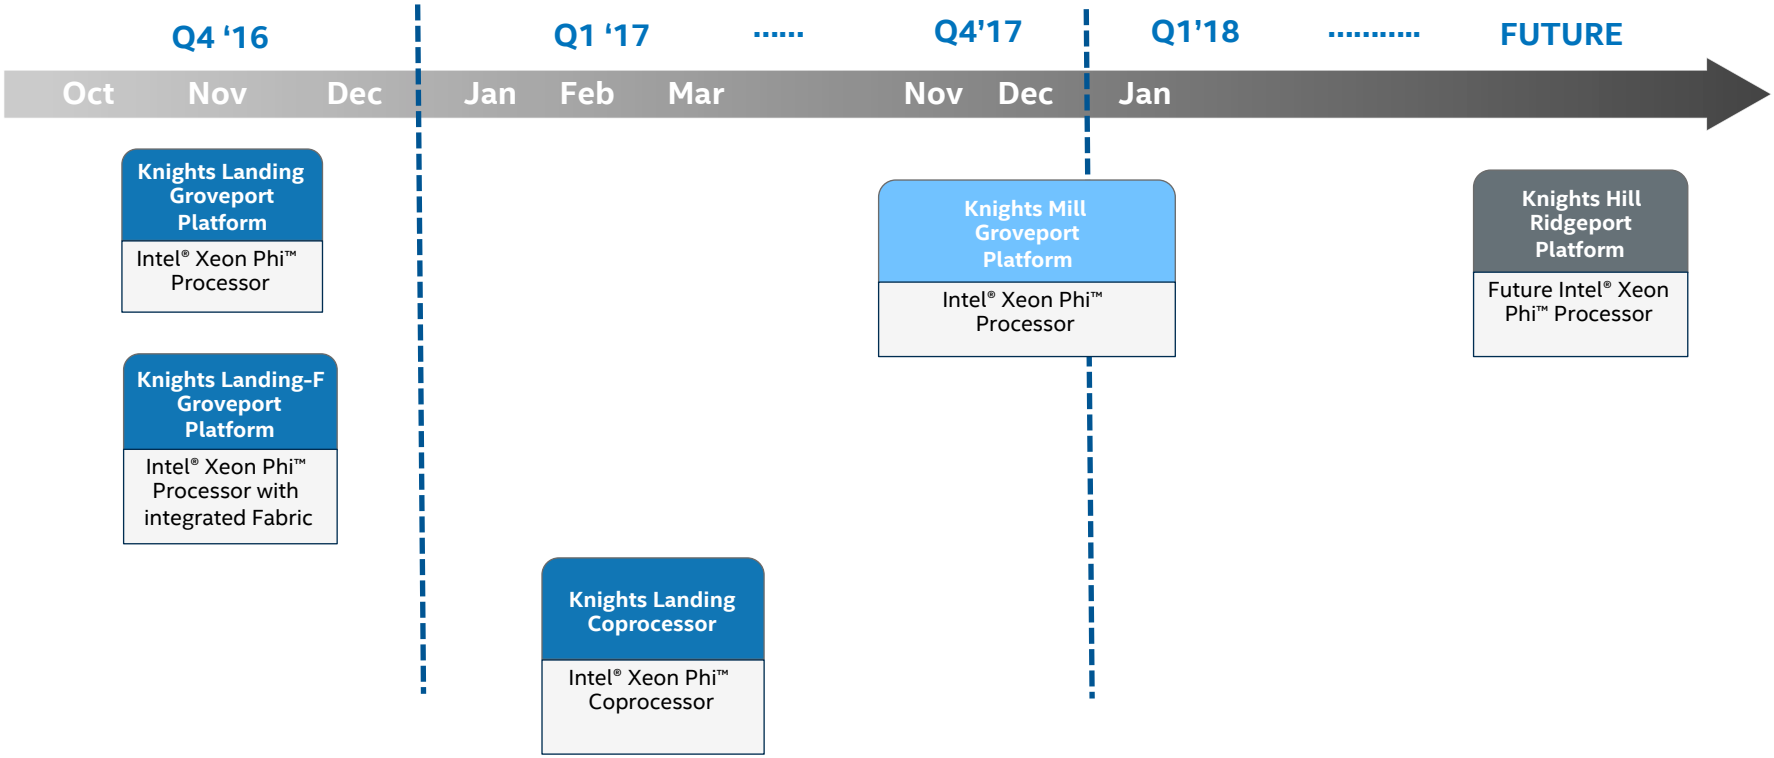

Intel® Xeon Phi™ Processor
(Knights Landing)

Self-Boot Processor
Binary-compatibility with Xeon, 3+ TFLOPS[1] (DP)

On-package memory
16GB, Up to 490 GB/s STREAM TRIAD

Platform Memory
Up to 384GB (6ch DDR4-2400 MHz)

Other Key Features
- ✓ 2D Mesh Architecture
- ✓ Out-of-Order Cores
- ✓ 3X Single-Thread vs. KNC
- ✓ Intel® AVX-512 Instructions
- ✓ Scatter/Gather Engine
- ✓ Integrated Fabric - OPA

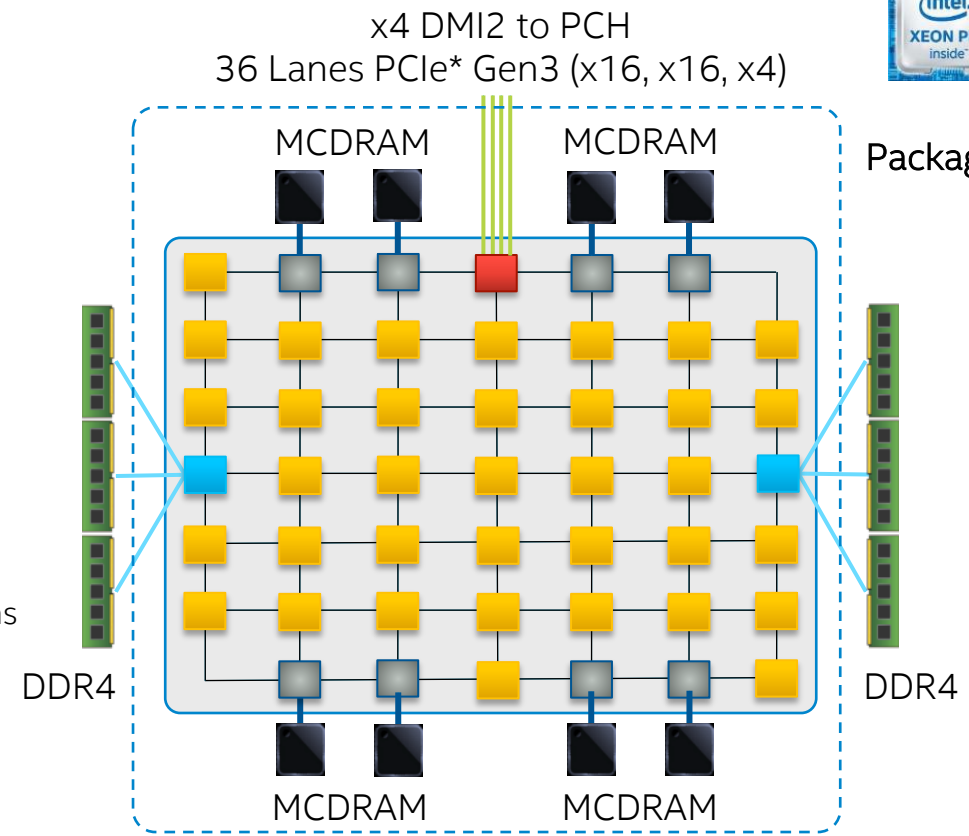

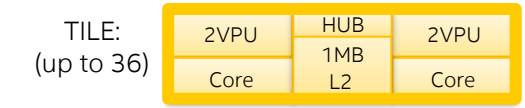

Tile | EDC (Embedded DRAM Controller) | IMC (Integrated Memory Controller) | IIO (Integrated I/O Controller)

[1]Theoretical peak performance

Intel® Xeon Phi™ Top500 Listings – Nov 2016

9 new entrants for 31 systems powered by Intel Xeon Phi

- Intel Xeon Phi Processors (Knights Landing) now power 5 systems in the top 50! (see list below)
- Only 2 Nvidia Pascal* listings

NERSC	#5	**Cori** (NERSC, USA); Cray XC – 14 Pflops/s
JCAHPC	#6	**Oakforest PACS** (JCAHPC, Japan); Fujitsu CX1640 M1 – 13.5 Pflops/s
CINECA	#12	**Marconi** (CINECA, Italy); Lenovo – 6.2 Pflops/s
Argonne NATIONAL LABORATORY	#18	**Theta** (Argonne National Lab, USA); Cray XC40 – 5.1 Pflops/s
KYOTO UNIVERSITY FOUNDED 1897	#33	**Camphor 2** (ACCMS, Kyoto University, Japan); Cray XC40 – 3.1 Pflops/s

Intel® Xeon Phi™ Processor Family (SKUs)

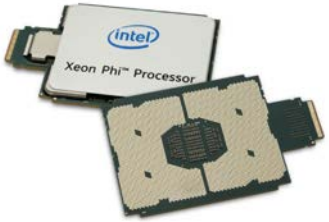

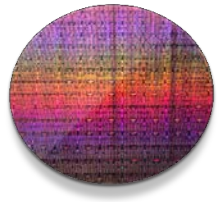

	CORES	GHZ	INTEGRATED MEMORY	INTEGRATED FABRIC*	DDR4	POWER**
7290* Best Performance/Node	72	1.5	16GB 7.2 GT/s	Yes	384GB 2400 MHz	245W
7250 Best Performance/Watt	68	1.4	16GB 7.2 GT/s	Yes	384GB 2400 MHz	215W
7230 Best Memory Bandwidth/Core	64	1.3	16GB 7.2 GT/s	Yes	384GB 2400 MHz	215W
7210 Best Value	64	1.3	16GB 6.4 GT/s	Yes	384GB 2133 MHz	215W

**Add 15 watts for integrated fabric

Intel® Xeon Phi™ Processor Increases Customer Value through More Cores, Wider Vectors, and Memory BW

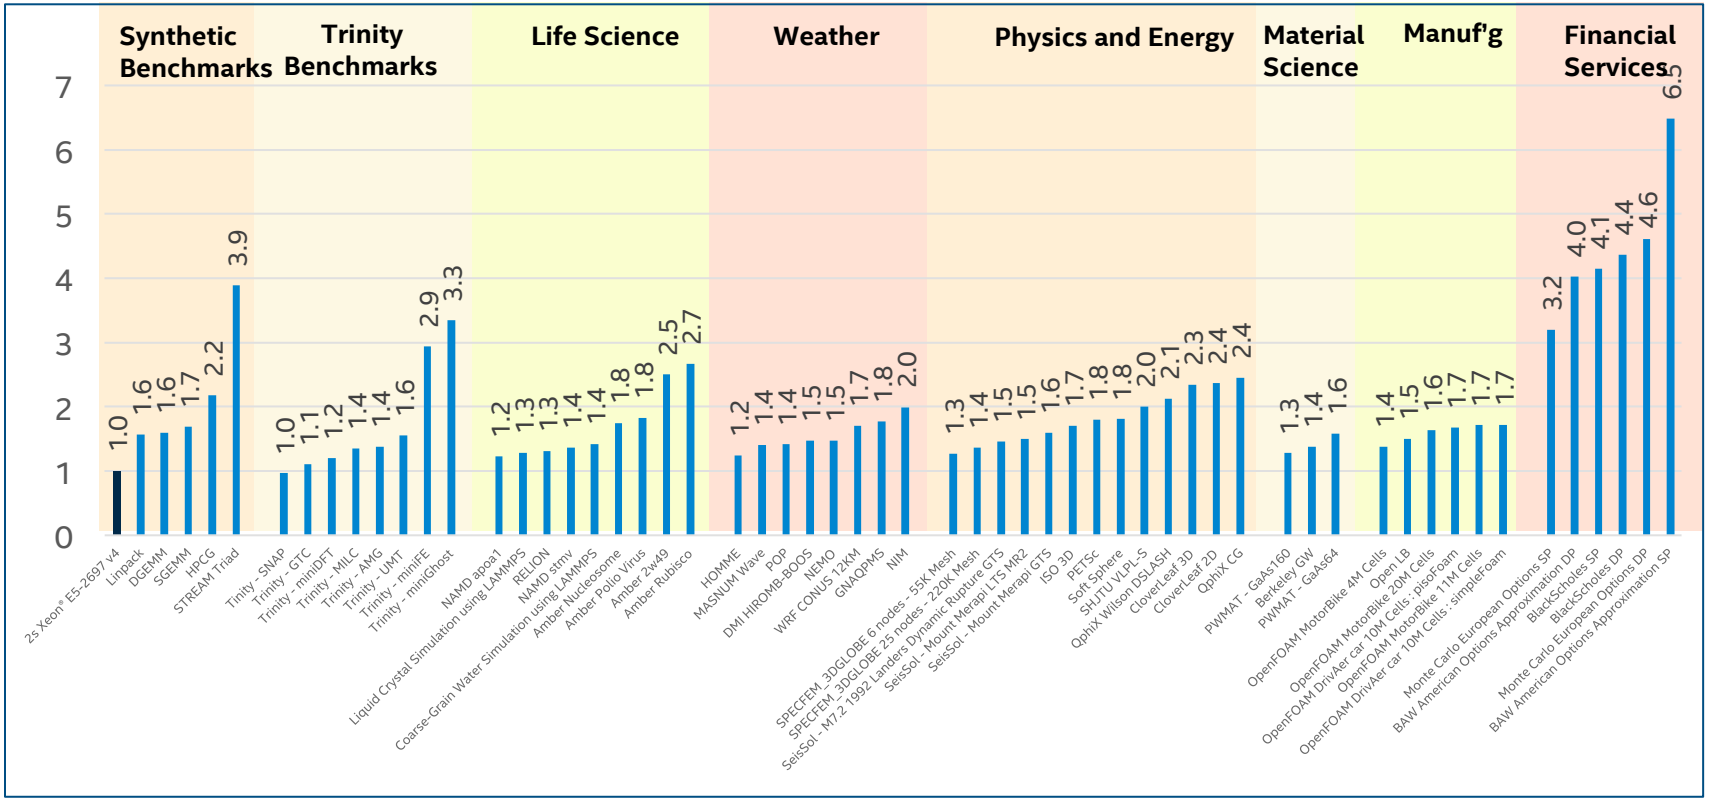

Intel® Xeon Phi ™ Processor Image Classification Training Throughput

Single node: 1s Xeon Phi up to 1.8x better than two Intel® Xeon® processor E5-2697v4

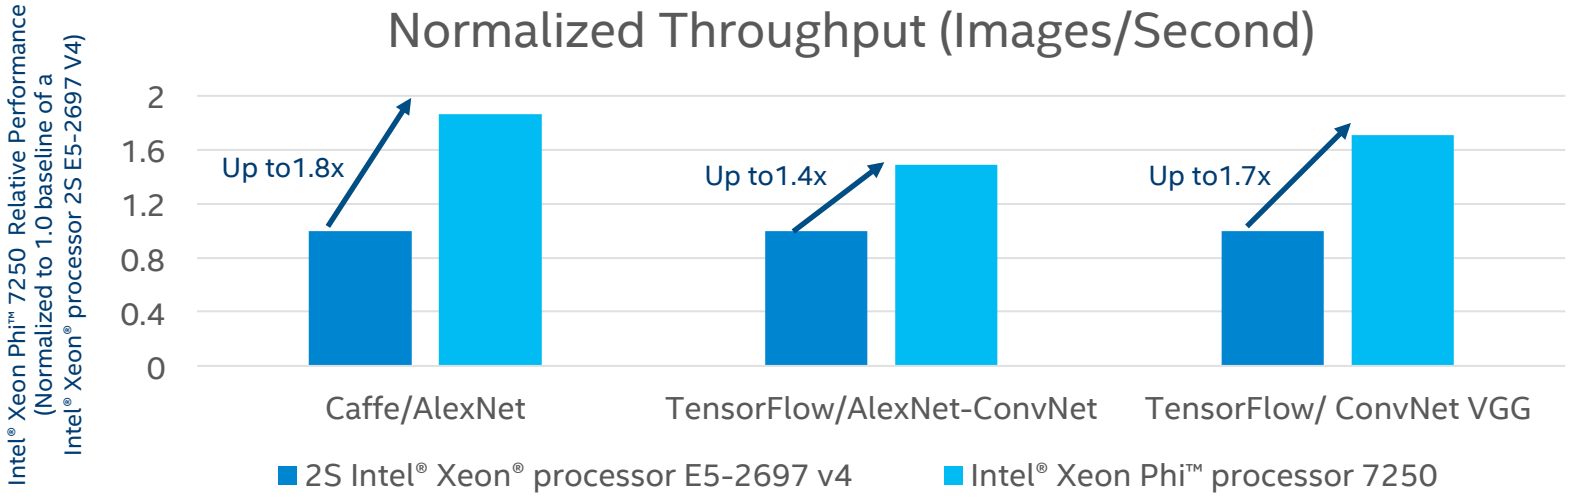

Configuration details on slide: 12

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

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

Time-To-Train Scaling Efficiency up to 97% on 32 nodes of Intel® Xeon Phi™ Processor cluster with Intel® Omni Path Fabric

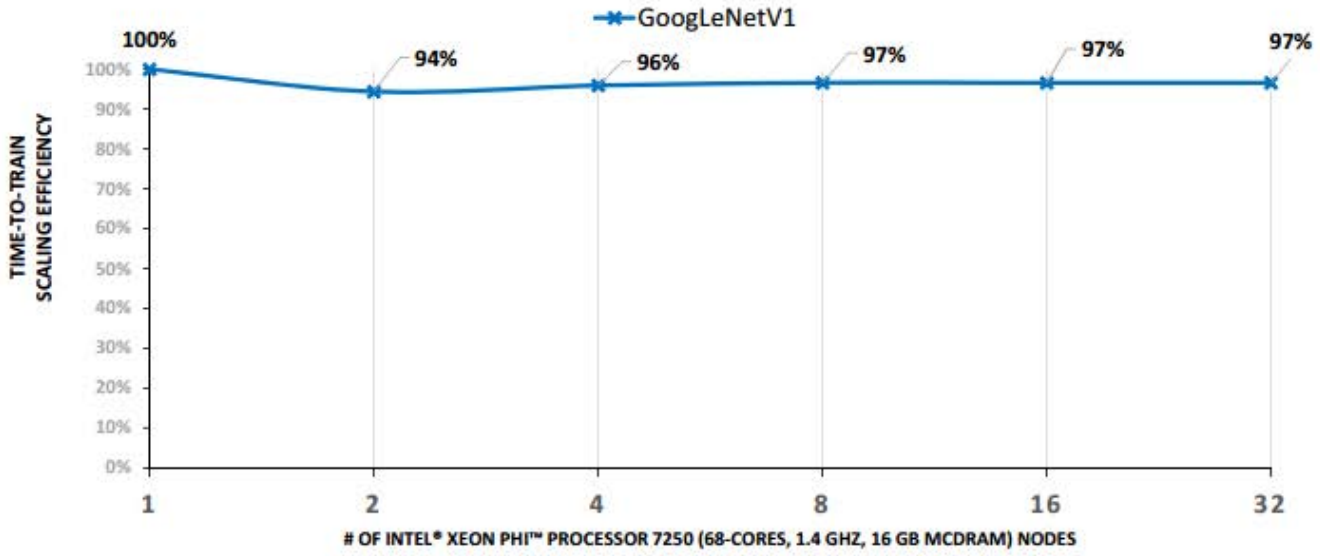

Configuration details on slide: 8

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

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 or estimated

Data Center AI Product Portfolio

Intel® Xeon® Processors

- Most widely deployed machine learning solution

Intel® Xeon Phi™ Processors

- High performance, general purpose machine learning

Intel® Xeon® Processor + Intel® Nervana Engine

- Best in class neural network performance

Intel® Xeon® Processor + FPGA

- Programmable, low latency inference

General Purpose Infrastructure → **Workload optimized**

Knights Mill & Groveport Platform Overview

Trains Machines Faster

- Up to 2.5X* Single Precision performance over Knights Landing for deep learning workloads
- Industry leader variable precision QVNNI up to 4X* faster performance
- Highly distributed multi-node scaling

Memory Flexibility & Bootable Host-CPU

- High memory bandwidth with integrated 16GB MC DRAM and bootable host-CPU reduces offloading & latency challenges
- 384GB 6-channel DDR4 memory capacity for massive AI use cases

Consistent Programming Models

- Common Intel® Xeon® & Intel® Xeon Phi™ programming
- Optimized for industry standard Open Source ML frameworks
- Flexibility to run vast workloads across x86 infrastructure

*NOTE: Performance theoretical wrt KNL7250 SKU based on KNM architectural changes.

Groveport Platform

Bootable Host CPU

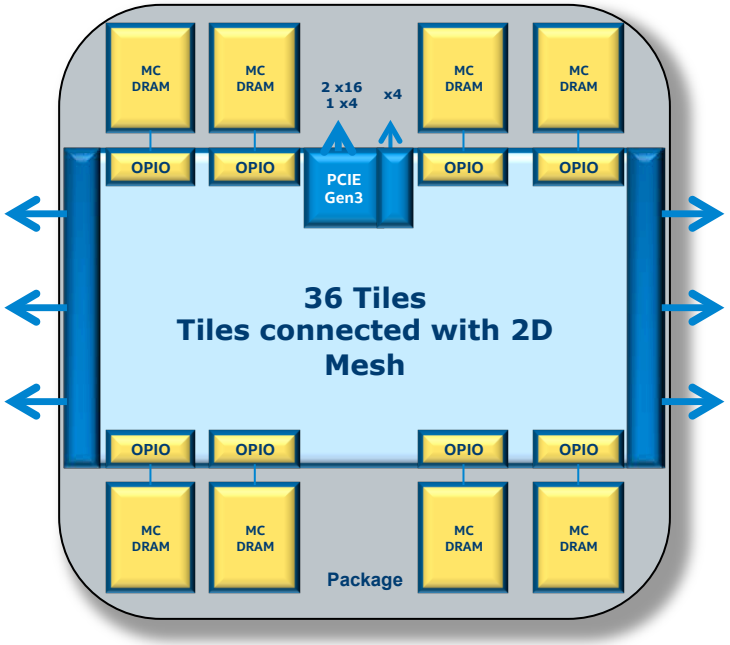

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks. Performance estimate wrt KNL 7250 SKU SGEMM. Performance Calculation= AVX freq X Cores X Flops per Core X Efficiency

Knights Mill QFMA for Faster Performance

Enhanced ISA QFMA instructions in Knights Mill delivers:

- ✓ Higher Peak Flops for CNN, RNN, DNN, LSTM
- ✓ Higher Efficiency (One Quad FMA executed in two cycles)
- ✓ 2X FP operations per cycle

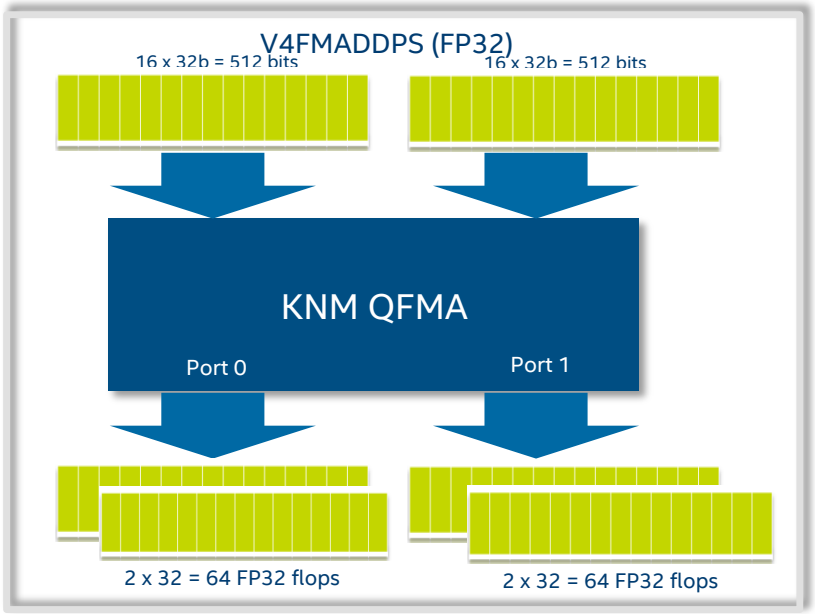

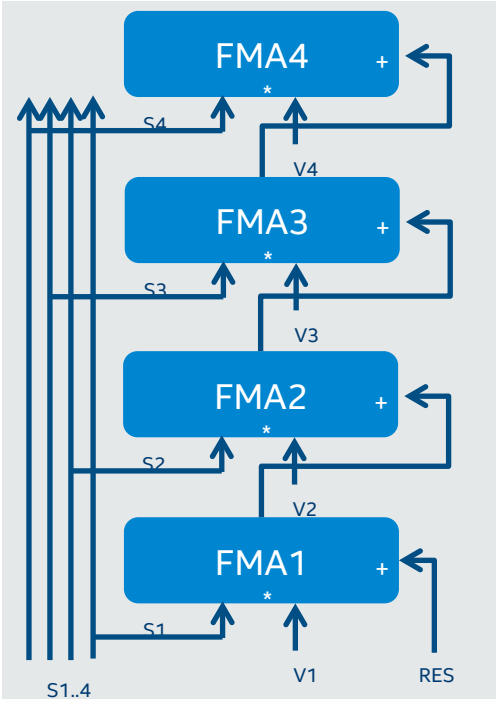

QMADD packs 4 IEEE FMA ops in a single instruction

*2X faster than KNL SP

Knights Mill Variable Precision Performance

Enabling Faster Throughput for Machine Learning Training

VNNI

- 2x the flops by using INT16 inputs
- Similar accuracy as SP by using INT32 accumulated output

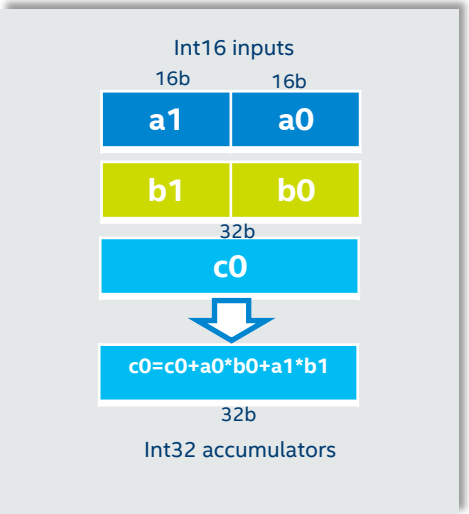

QVNNI

- 2x VNNI operations per port
- 4x* ML performance than regular AVX512-SP

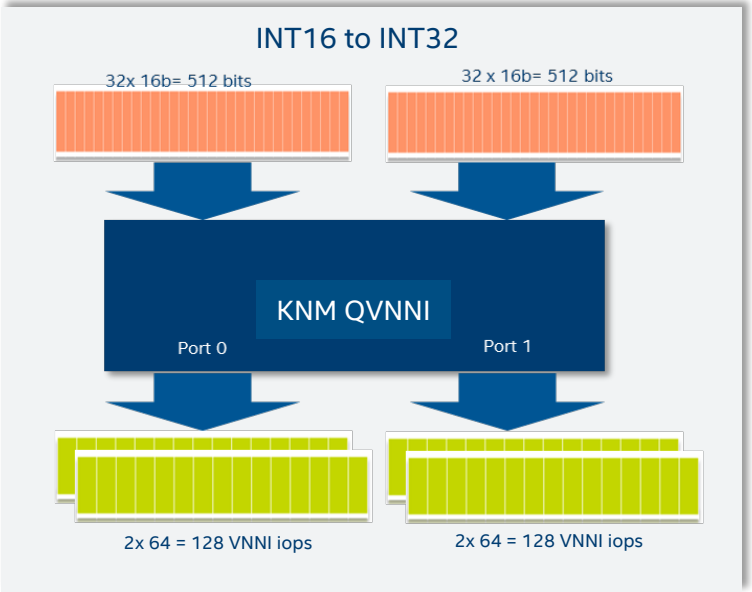

*4x faster than KNL SP

Knights Mill QVNNI Advantages over FP16

FP16

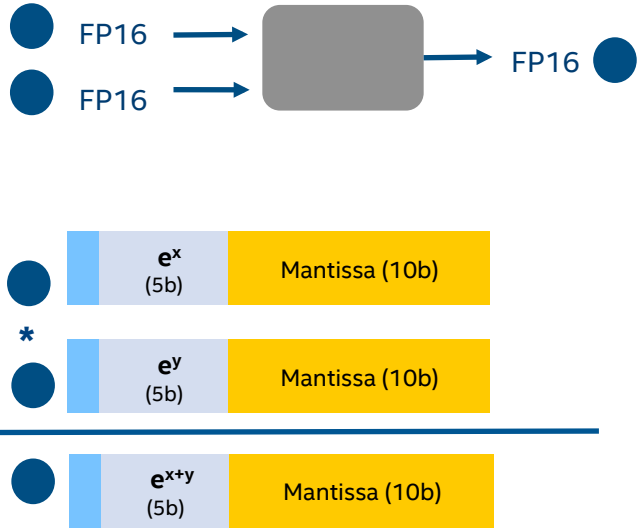

Intel® QVNNI

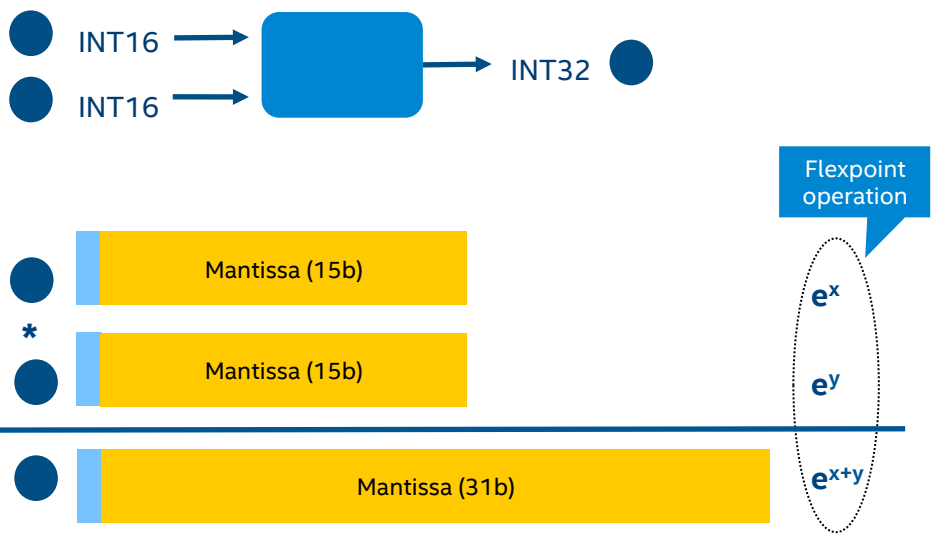

QVNNI for Higher Accuracy and Faster Operations

Knights Mill Highly Efficient Cluster Scale-Out

- Competition cross-system Fabric Bandwidth = 6.25GBs per GPU
- Knights Mill cross-system Fabric Bandwidth = 50GBs per CPU

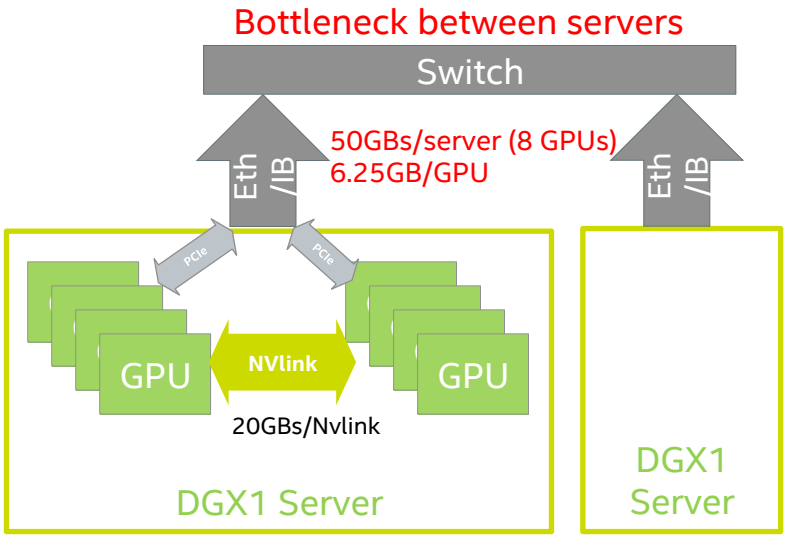

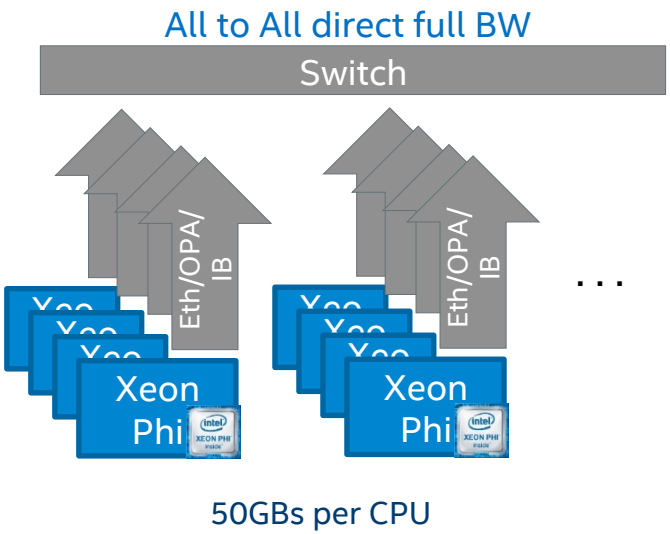

Knights Hill Processor Developments

CPU – fabric integration

- Direct access to KNH CPU resources
- Improved fabric latency
- Lower cost and improved density opportunity

Enhanced performance

- Huge leap in Dual Precision vector performance
- Dramatic leap in Single & 16-bit performance
- High density system options
- Improved Intel® OPA fabric bandwidth

Reduced costs

- TCO via Performance/Watt
- Faster time to solution
- Higher radix Omni-Path switches

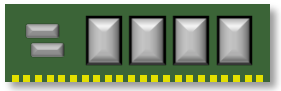

Memory

- Higher capacity and bandwidth in package memory
- Innovations in 3D XPoint™ technology support

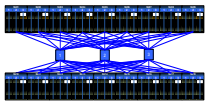

Scaling

- Emphasis on reliability and resiliency
- Storm Lake 2 scaling support for 100K nodes

Agenda

- Artificial Intelligence
- Intel® Xeon Phi™ Processor
- Intel® Xeon® Processors
- Summary

How can I get higher performance & TCO for my apps?

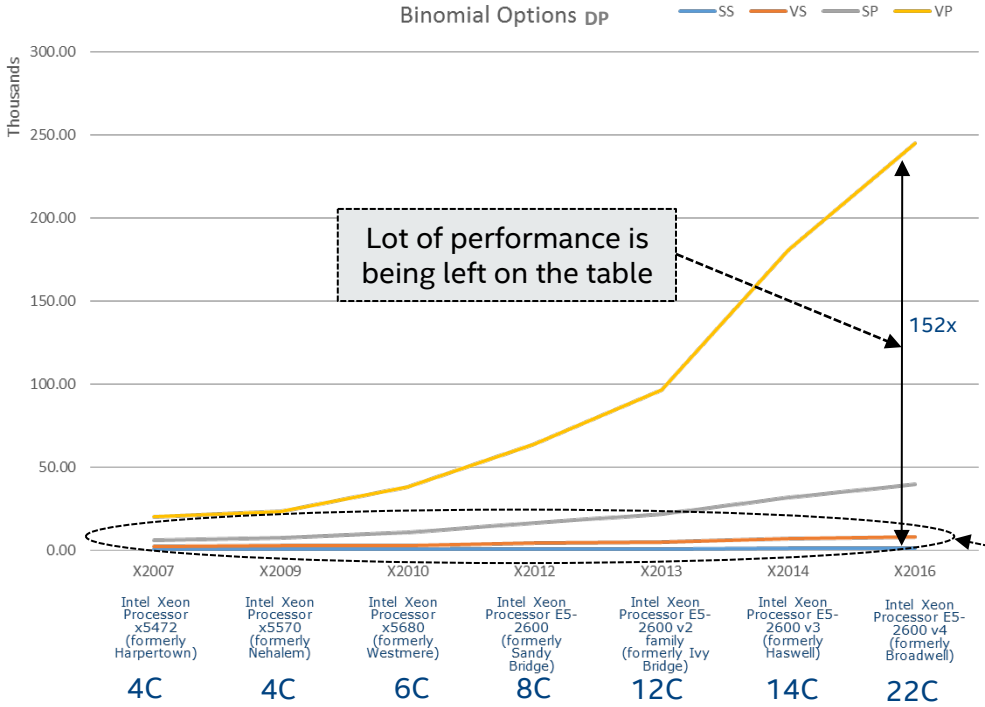

VP = Vectorized & Parallelized (MT)
SP = Scalar & Parallelized (MT)
VS = Vectorized & Single-Threaded (ST)
SS = Scalar & Single-Threaded (ST)

We believe most codes are here

Modernization (i.e. parallelization and vectorization) of your code is the solution

Tick-Tock Development Model:

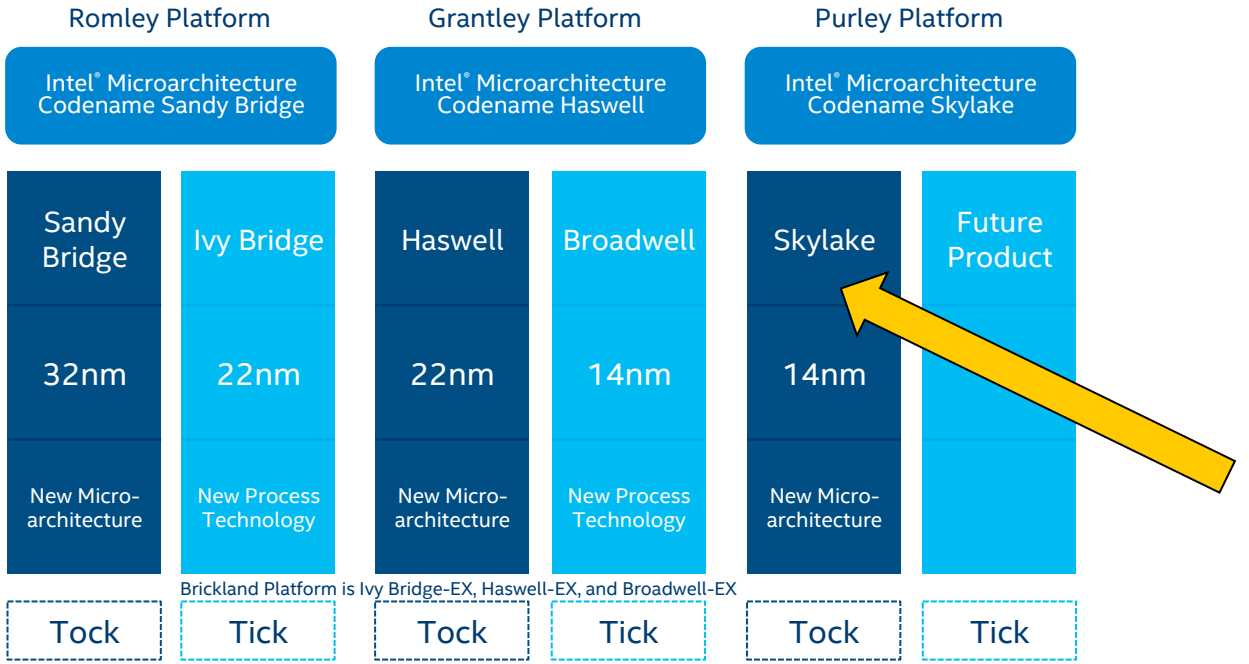

Sustained Microprocessor Innovation Leadership

Skylake Server Processor

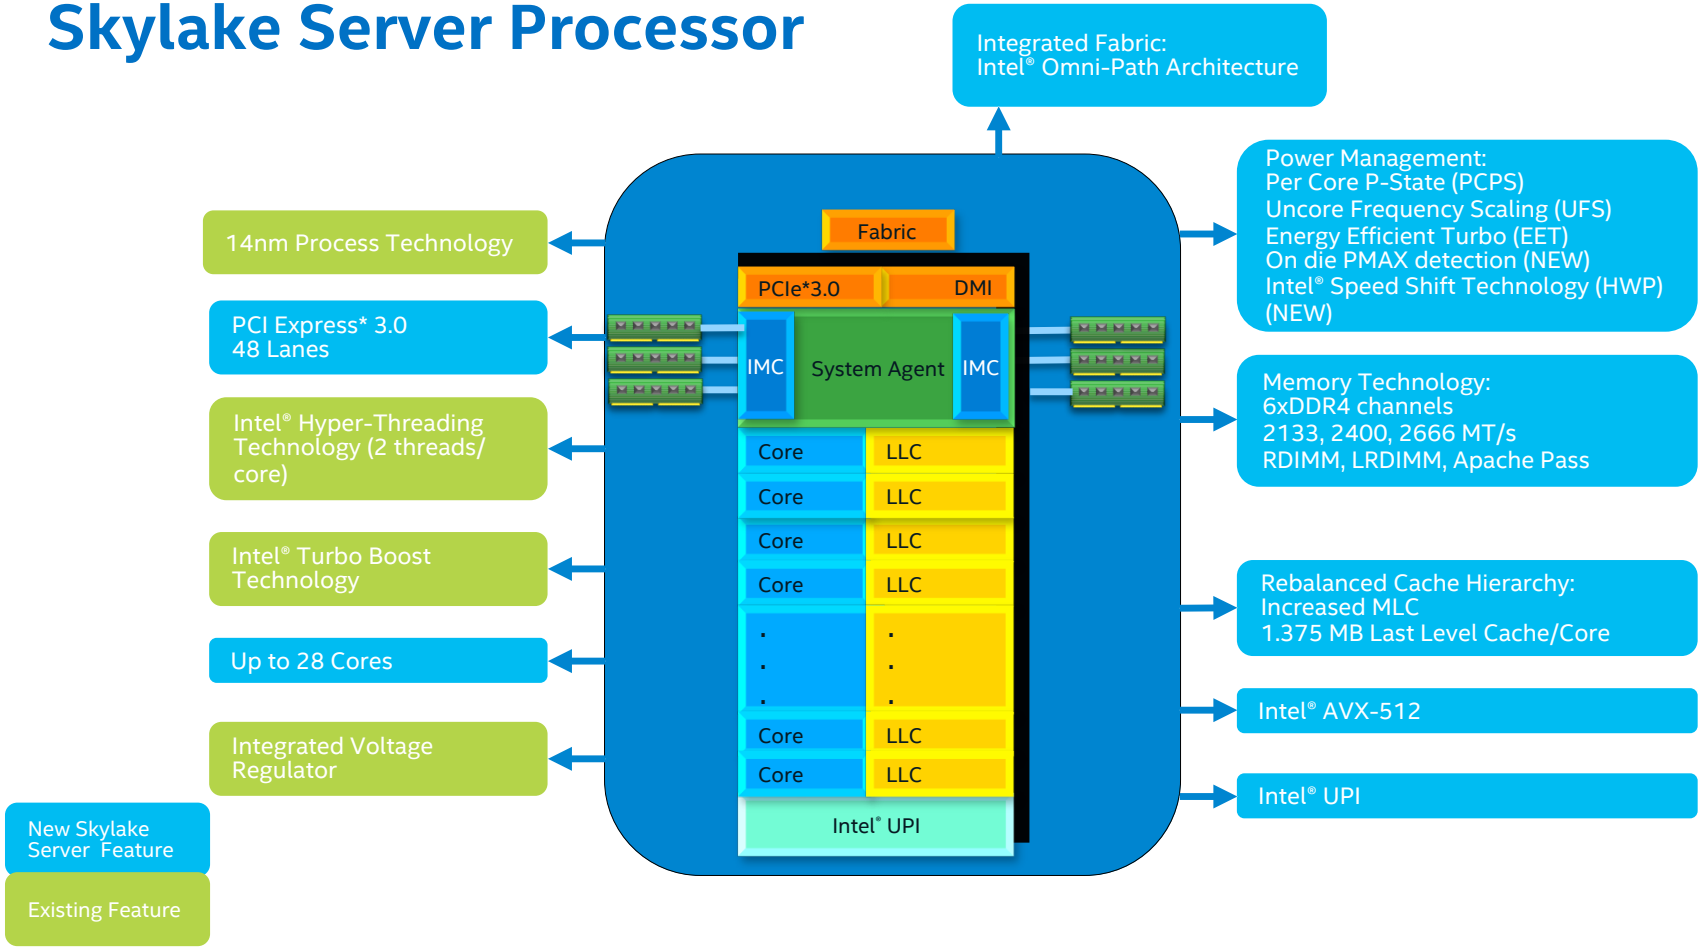

New Skylake-SP SKUing Concept*

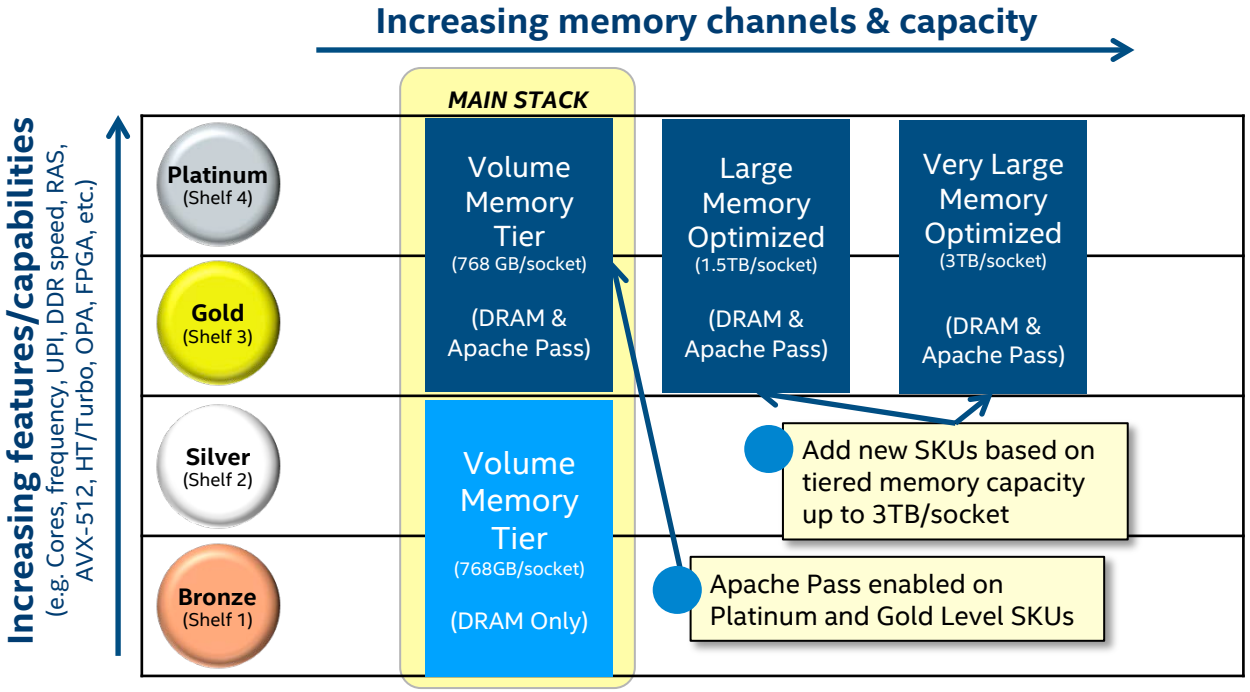

Intel will integrate memory capability tiers as another scalable feature on the Skylake-SP SKU stack

Key changes:

- Apache Pass will be enabled on Platinum and Gold level SKUs
- New tiered capacity capable SKUs in Platinum and Gold level SKU

* Example implementation shown is purely conceptual to illustrate the idea; SKUing details to be finalized via standard SKU'ing Process at future time

Intel® Xeon® Processor Scalable Family Preliminary SKU Stack

(81xx & 61xx Processors)

Platinum (Shelf 4)

81xx (Platinum)

- 2S-2UPI, 2S-3UPI, 4S-2UPI, 4S-3UPI, and 8S-3UPI capability
- 6-ch DDR4 @ 2666
- 3 UPI links @ 10.4GT/s
- Intel® Turbo Boost Technology
- Intel® HT Technology
- Intel® AVX-512 (2 512-bit FMAs)
- 48 lanes PCIe* Gen3
- Node Controller Support
- Advanced RAS

Gold (Shelf 3)

61xx (Gold)

- 2S-2UPI, 2S-3UPI, 4S-2UPI, and 4S-3UPI capability
- 6-ch DDR4 @ 2666
- 3 UPI links @ 10.4GT/s
- Intel® Turbo Boost Technology
- Intel® HT Technology
- Intel® AVX-512 (2 512-bit FMAs)
- 48 lanes PCIe* Gen3
- Node Controller Support
- Advanced RAS

105W	120W	135W	145W	165W	205W
81xx 4C-16.50M 105W x.xG	6138 20C 2.0G 120W	6152 22C 2.1G 135W	8164 26C x.xG 145W	8176 28C 2.1G 165W	8180 28C 2.5G 205W
	6130 16C 2.1G 120W	6140 18C 2.3G 135W	8160 24C 2.1G 145W	8170 26C 2.1G 165W	8168 24C 2.7G 205W
	6126 12C-19.25M 2.6G 120W	6132 14C x.xG 135W	6148 20C 2.4G 145W	6150 18C 2.7G 165W	6154 18C 3.0G 200W
	6128 6C-19.25M 3.4G 115W	6134 8C-24.75M 3.3G 130W	6142 16C 2.6G 145W	6146 12C-24.75M x.xG 165W	
			6136 12C-24.75M 3.0G 145W		
			6144 8C-24.75M x.xG 145W		

XCC Die

Core Count / Base Frequency — TDP

Changes in feature set from shelf to shelf highlighted in **red**

All SKUs, frequencies, and performance estimates are PRELIMINARY and can change without notice.

FREQUENCY-OPTIMIZED SKUs

POTENTIAL EXTENDED SUPPLY LIFE SKUs (5-7 year production)

Intel® Xeon® Processor Scalable Family Preliminary SKU Stack

(81xxF & 61xxF processors with integrated Intel® Omni-Path Fabric)

- Single on-package OPA Host Fabric Interface (HFI) port
- Fabric component interfaces to SKL using 16 PCIe* lanes
- Fabric PCIe* lanes are incremental to the existing 48 PCIe* lanes
- Same socket for SKL and SKL-F processors

81xxF (Platinum)

- ≥ 24 cores
- 2S-2UPI capability
- 6-ch DDR4 @ 2666
- 2 UPI links @ 10.4GT/s
- Intel® Turbo Boost Technology
- Intel® Hyper-Threading Technology
- Intel® AVX-512 (2 512-bit FMAs)
- 48 lanes PCIe* Gen3
- Core RAS

61xxF (Gold)

- ≤ 22 cores
- 2S-2UPI capability
- 6-ch DDR4 @ 2666
- 2 UPI links @ 10.4GT/s
- Intel® Turbo Boost
- Intel® Hyper-Threading
- Intel® AVX-512 (2 512-bit FMAs)
- 48 lanes PCIe* Gen3
- Core RAS

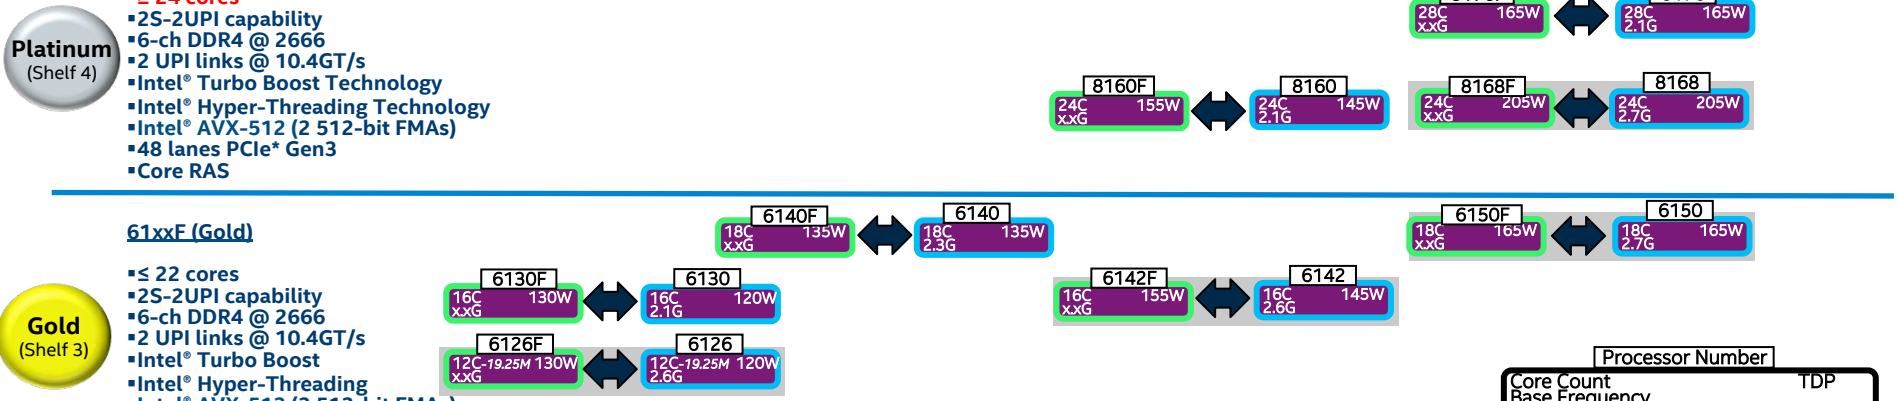

All SKUs, frequencies, and performance estimates are PRELIMINARY and can change without notice.

Skylake + FPGA on Purley

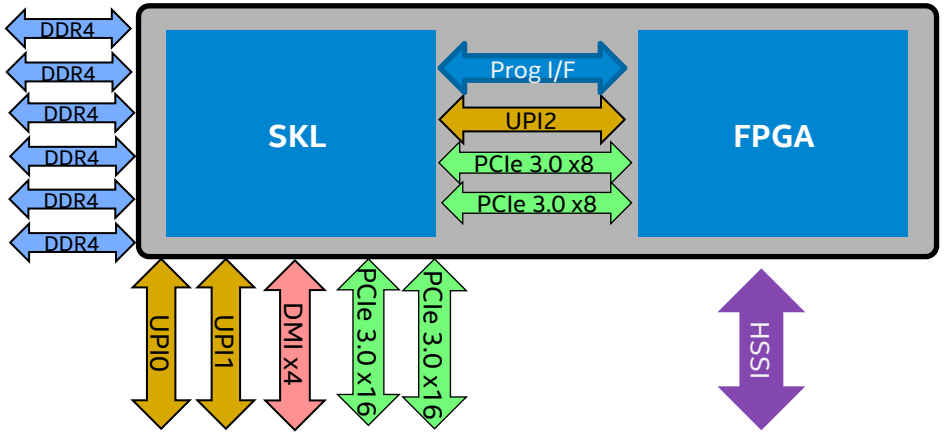

- Power for FPGA is drawn from socket & requires modified Purley platform specs
- Platform Modifications include Stackup, Clock, Power Delivery, Debug, Power up/down sequence, Misc IO pins (see BOM cost section)

Cores	Up to 28C with Intel® HT Technology	
FPGA	Altera® Arria 10 GX 1150	
Socket TDP	Shared socket TDP of 165W combined, or Up to 165W SKL & Up to 90W FPGA	
Socket	Socket P	
Scalability	Up to 2S – with SKL-SP or SKL + FPGA SKUs	
PCH	Lewisburg: DMI3 – 4 lanes; 14xUSB2 ports Up to: 10xUSB3; 14xSATA3, 20xPCIe*3 New: Innovation Engine, 4x10GbE ports, Intel® QuickAssist Technology	
	For CPU	**For FPGA**
Memory	6 channels DDR4 RDIMM, LRDIMM, Apache Pass DIMMs	Low latency access to system memory via UPI & PCIe interconnect
	2666 1DPC, 2133, 2400 2DPC	
Intel® UPI	2 channels (10.4, 9.6 GT/s)	1 channel (9.6 GT/s)
PCIe*	PCIe* 3.0 (8.0, 5.0, 2.5 GT/s)	PCIe* 3.0 (8.0, 5.0, 2.5 GT/s)
	32 lanes per CPU Bifurcation support: x16, x8, x4	16 lanes per FPGA Bifurcation support: x8
High Speed Serial Interface (Different board design based on HSSI config)	N/A	2xPCIe 3.0 x8
		Direct Ethernet (4x10 GbE, 2x40 GbE, 10x10 GbE, 2x25 GbE)

SKL+FPGA Usage Cases & Target Workloads

FPGA for Look Aside Acceleration

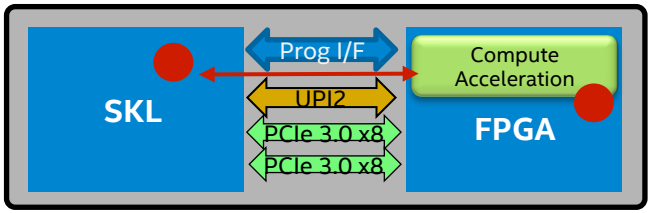

- Deep Learning classification/scoring*
- Video & Media Transcoding
- Data Analytics* (ex: recommendation engine)
- Genomics*
- Compression*
- Memory copy routines

FPGA for Low Latency Inline Processing

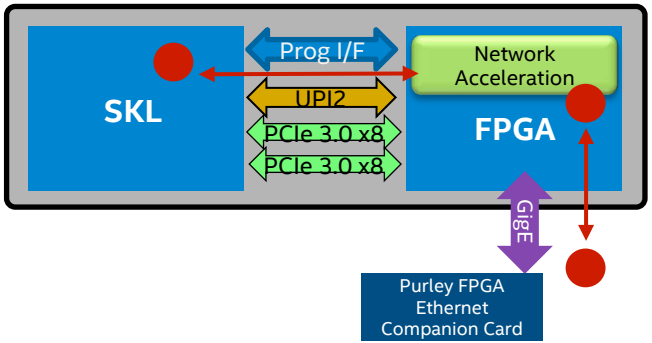

- Virtual Network Functions* (Switch, Router, Load Balancer)
- FSI market data pre-filtering, Risk Analytics, Option Pricing
- Image Processing
- Automotive video input
- Accelerator Chaining: Custom Compression, Security

*Support for Application Development available using Library Approach

SKL+FPGA Customer Profile

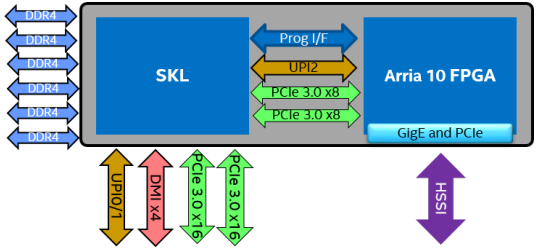

Application Development Method	Customer Profile
New - Library Approach	**Target Customer:** No previous FPGA or hardware design experience; focus on end user application tuning **Target Segment:** Cloud, Enterprise (Health & Science, Analytics)
New - Library Approach **New - Extended IA Flow**	**Target Customer:** Customer with RTL expertise (discrete FPGA or ASIC), focus on networking acceleration **Target Segment:** Cloud, Networking, Enterprise (Government, FSI)
Traditional Flow (RTL or OpenCL)	**Target Customer:** Customer with RTL expertise (discrete FPGA or ASIC) or OpenCL expertise **Target Segment:** Networking, Enterprise (FSI, Health & Science, Gov't)

Knights Landing and 2s Skylake-SP

KNL: first Xeon Phi bootable host processor with on-package high BW memory for highly-parallel applications that achieves maximum FLOPs

- Optimized for highly parallel workloads with higher cores/threads per socket, achieving greater compute density
- Achieves greater energy efficiency per flop for highly parallel workloads (performance/watt)
- Best suited for memory BW sensitive applications (plus DDR for large memory footprint)
- Life Sciences, Energy, Simulation, Risk, HPDA, Machine Learning

2s Skylake-SP: significant performance combined with a rich feature set resulting in compelling benefits across a broad variety of usage models

- Best general purpose CPU with great parallel and serial performance; legacy OOB perf.
- Best suited for workloads where response time and serial performance is important (per core performance/lower latency)
- Better return on workloads where software licensing cost is on a per core or per thread basis and is a major factor in TCO
- Broad range of E&G, Cloud, Networking and Storage workloads

Common source code and programming model across Xeon Phi and Xeon, enabling a single code investment on IA. Common integration of OPA fabric and AVX-512 vector engine.

Intel® Scalable System Framework

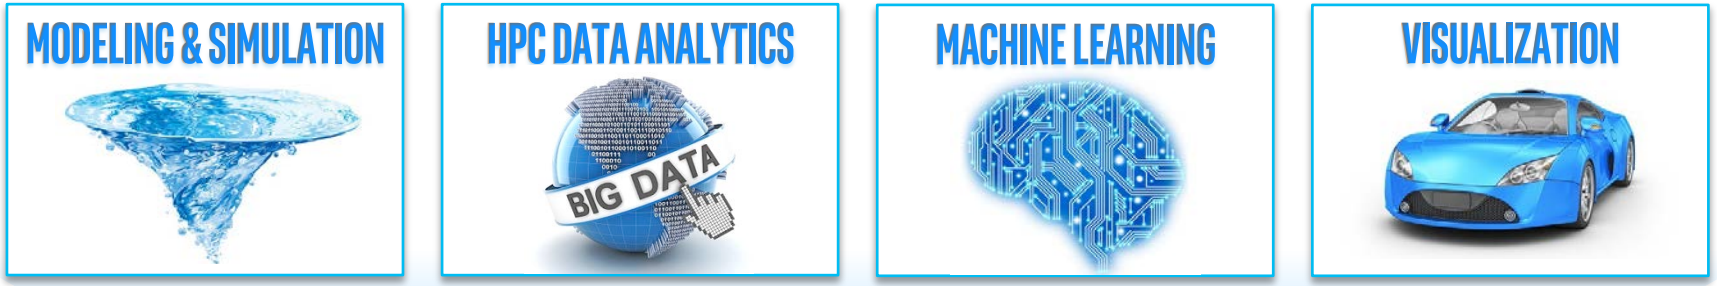

MANY WORKLOADS – ONE FRAMEWORK

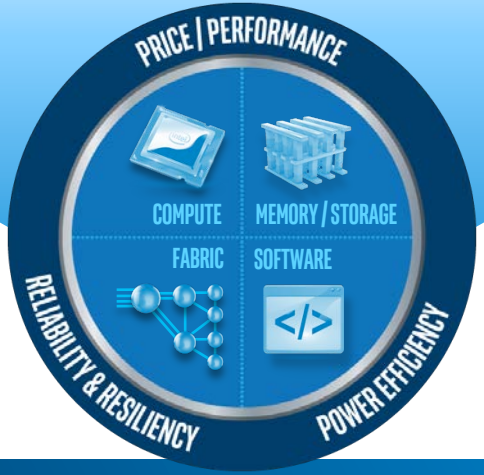

A Flexible Framework for Today & Tomorrow

Enabling Breakthrough System Performance

Backup

CONFIGURATION DETAILS

Configuration details: LAMMPS Coarse-Grain Water Simulation*

Intel® Xeon® processor E5-2697 v4 : Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz , 18 Cores/Socket, 36 Cores, 72 Threads (HT and turbo on), DDR4 128GB, 2400 MHz, Red Hat 6.7, Wildcat Pass Motherboard, BMC 1.33.9832, FRU/SDR Package 1.09, 1.0 TB SATA Western Digital* 1003FZEX-00MK2A0 System Disk, 448W mean power consumption for LAMMPS water simulation, Scalability tests performed on nodes with Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16.

Intel® Xeon Phi™ processor 7250 : Intel® Xeon Phi™ processor 7250, 68 core, 272 threads, 1400 MHz core freq. (turbo on), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode, Adams Pass Motherboard, BMC 12.951, FRU/SDR Package 1.1, 1.0 TB SATA drive Western Digital* 1003FZEX-00MK2A0 System Disk, 378W mean power consumption for LAMMPS water simulation, Scalability tests performed on nodes with Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16

NVIDIA Tesla K80*: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz , 18 Cores/Socket, 36 Cores, 72 Threads (HT and turbo on), DDR4 128GB, 2400 MHz, Red Hat 7.2, Super Micro* SuperServer 1028GR-TR, Bios Version 2.0a, Super Micro* X10DRG-H Motherboard, CSE-118GHTS-R1K66BP FRU, 500GB SATA Seagate* ST9500423AS System Disk, NVIDIA Tesla* K80 GPU, NVIDIA CUDA* 7.5.17 (Driver: 352.39), ECC enabled, persistence mode enabled. Number of MPI tasks on host varied to give best performance. CUDA MPS* used where possible. 608W mean power consumption for LAMMPS water simulation.

LAMMPS CONFIGURATION: 22 Mar 2016 (Git Hash: 154eb1f886fde), Intel® Compiler 16.0.2, Intel® MPI 5.1.2.150, Optimization Flags: "-O2 -fp-model fast=2 -no-prec-div -qoverride-limits"

Configuration details: Embree*

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7520, MCDRAM 16 GB RAM 96 GB 6*16GB 2400MHz, Reg ECC DDR4 BIOS Configuration: L2 HWP patches (same as targeted for production) Software Details: Intel Compiler Version 16.0.1, ISPC Compiler Version 1.9, Embree Version 2.9.0, OS / Kernel CentOS 7.1

NVIDIA* GPU: NVIDIA Optix* System Xeon DP Intel® Xeon® processor E5-2699 v3 LGA2011 2.3GHz 45MB 145W Dual socket 18 core RAM 128 GB total 8*16GB 2133MHz Reg ECC DDR4 BIOS SE5C610.86B. 01.01.0005.101720141054 Intel SSDSA2M160G2GC, 1x160 GB SATA SSD, NVIDIA* GeForce* GTX* Titan X 3072 CUDA Cores 12GB memory Software Details: CUDA Version 7.5. OptiX Version 3.9.0 NVIDIA Driver Version 346.46 OS / Kernel CentOS release 6.6 / 2.6.32-504.23.4.el6.x86_64

Configuration details: Monte Carlo / Black-Scholes / Binomial Tree

NVIDIA Tesla K80*: 2XGK210B PCI Express GEN3 Dual GPU 2496 Processor cores Base Clock 560MGH Boost Range 562-875MHz 12GB GDDR5 Memory Clock 2.5GHz. Red Hat 6.7 (Santiago) **HOST: Intel® Xeon® processor E5-2697 v4:** Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo OFF), 1600 MHz uncore freq. MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode.

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Configuration details: STAC-A2

STAC SUT ID INTC160428 - Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), MCDRAM 16 GB 7.2 GT/s, DDR4 96GB 2400 MHz, CentOS 7.2, quadrant cluster mode, flat memory mode. See www.STACresearch.com/INTC160428.

STAC SUT ID IBM150305 - IBM POWER8™ : IBM Power System* sever, 2x 12-core POWER8* @ 3.52 GHz, 24 cores / 192 Threads (only 96 used), 1 TB DDR3, RH 7.0, IBM XL C/C++ for Linux v13.1. See www.STACresearch.com/IBM150305.

STAC SUT ID NVDA141116 - NVidia® Tesla® K80 : Supermicro* SYS-2027GR-TRHF, Intel Xeon E5-2690 v2, 3.00GHz, 128GB DDR3, 2XGK210B PCI Express GEN3 Dual GPU 2496 Processor cores Base Clock 560MHz Boost Range 562-875MHz 12GB GDDR5 Memory Clock 2.5GHz. NVIDIA CUDA* 6.5 (Driver 340.58), CentOS 6.6 + Intel® Xeon® processor E5-2690 v2: 10 Cores/Socket, 20 Cores (HT off), DDR3 128GB. See www.STACresearch.com/NVDA141116.

Configuration details: CP2K* Linear Scaling (LS) Density Function Theory (DFT)

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 cores/socket, 36 cores, 72 threads (HT and Turbo ON), DDR4 64 GB, 2400 MHz, RHEL 6.6

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7260, 68 core (272 threads), 1.4 GHz base core freq. (Turbo ON), 1.7 GHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 86B.01.01.0124, DDR4 96 GB 2400 MHz, quadrant cluster mode, MCDRAM cache memory mode, RHEL 6.6, MPSP 1.3.0, Intel Compiler 2017

NVIDIA Tesla K80*: 2XGK210B PCI Express GEN3 Dual GPU 2496 Processor cores Base Clock 560MGH Boost Range 562-875MHz 12GB GDDR5 Memory Clock 2.5GHz. Red Hat 6.6 **HOST:** Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 cores/socket, 36 cores, 72 threads (HT and Turbo ON), DDR4 64 GB, 2400 MHz, RHEL 6.6

Building CP2K

LIBXSMM: git clone https://github.com/hfp/libxsmm.git; git checkout tags/1.4.3
- Automatically selected when using ARCH files as mentioned below

CP2K: make ARCH=Linux-x86-64-intel VERSION=psmp AVX=3 MIC=1 LIBXSMM=2
Either CP2K intel branch (git clone --branch intel https://github.com/cp2k/cp2k.git cp2k.git)
Or master CP2K 4.0-development (https://github.com/cp2k/cp2k.git cp2k.git)
- https://github.com/cp2k/cp2k/raw/intel/cp2k/arch/Linux-x86-64-intel.x
- https://github.com/cp2k/cp2k/raw/intel/cp2k/arch/Linux-x86-64-intel.psmp

Runing CP2K: CP2K/intel@KNL: CP2K_RECONFIGURE=1 (huge pages), CP2K_STACKSIZE=10000 (SMMs/batch),
I_MPI_PIN_DOMAIN=auto, I_MPI_PIN_ORDER=scatter, 64 ranks, 1 thread/core

Configuration details : Qphix*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo OFF) , 18 Cores/Socket, 36 Cores, 72 Threads (HT ON), DDR4 128GB, 2400 MHz, Red Hat 6.5, Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode, Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16

NVIDIA* GPU: 1x NVIDIA Titan X* GPU, ECC enabled, persistence mode enabled, full GPU Boost (auto), NVIDIA CUDA* 7.5.17, QUDA* v0.8. Host CPU was 2S Intel® Xeon® processor E5-2697 v3.

Configuration details: Trinity MILC*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

NVIDIA Titan X*: Persistence mode enabled, full GPU Boost (auto), NVIDIA CUDA* 7.5. Host CPU was 2S Intel® Xeon® processor E5-2697 v3.

Configuration details: AMBER 16 IMPLICIT*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on, Turbo on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. Turbo on, 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Configuration details: AMBER 16 EXPLICIT*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on, Turbo on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. Turbo on, 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Configuration details: ROME*/SML

Intel® Xeon® processor E5-2697 v4: Dual socket Intel® Xeon® processor E5-2697 v4, @2.3GHz 145W, 18 cores/socket HT enabled, 128GB RAM, Red Hat* Enterprise Linux Server release 6.7 (Santiago)

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. MCDRAM 16 GB 7.2 GT/s, DDR4 96GB 2400 MHz, Red Hat Enterprise Linux Server release 6.7 (Santiago), quad cluster mode, MCDRAM cache mode

Workload: provided by Intel® PCCSB, Contact Youdong Mao youdong_mao@dfci.harvard.edu. (Performance data is based on 30 iterations). Workload Descriptions:
Inflammasome data: 16306 images of NLRC4/NAIP2 inflammasome with a size of 2502 pixles
RP-a: 57001 images of proteasome regulatory particles (RP) with a size of 1602 pixels
RP-b: 35407 images of proteasome regulatory particles (RP) with a size of 1602 pixels

Configuration details: RELION*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT off), DDR4 128GB, 2400 MHz, CentOS release 6.7 (Final)

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo ON), MCDRAM 16 GB 6.4 GT/s, BIOS GVPRCRB1.86B.0010.R00.1603251732, DDR4 96GB 2133 MHz, Red Hat 7.2 (Maipo), quad cluster mode, MCDRAM cache mode

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), MCDRAM 16 GB 7.2 GT/s, BIOS GVPRCRB1.86B.0010.D42.1604182214, DDR4 96GB 2400 MHz, Red Hat 7.2 (Maipo), quad cluster mode, MCDRAM cache mode. **Workload:** provided by Peking University. Contact Yanan Zhu <ynzhu@pku.edu.cn>

Configuration details: GROMACS

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON), 18 Cores/Socket, 36 Cores, 72 Threads (HT on), Wildcat Pass, DDR4 128GB, 2400 MHz, BMC ver. 1.33.9832, Red Hat 7.2, BIOS 86B0271.R00, FRU/SDR Package 1.09, kernel 3.10.0-327.el7.x86_64, 1 1.0 TB SATA drive WD1003FZEX-00MK2A0, Idle Power measurement 89W

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. 6 x 16 GB 2400 MHz DDR4, BMC ver. 12.951, Red Hat 7.2, BIOS 10R00, FRU/SDR Package 1.1, kernel 3.10.0-327.el7.x86_64, 1 1.0 TB SATA drive WD1003FZEX-00MK2A0, Idle Power measurement 125W

Configuration details: Nanoscale Molecular Dynamics program*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo OFF) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo ON), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 09D10, DDR4 96GB 2133 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Configuration details: Nanoscale Molecular Dynamics program* (perf./watt)

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo OFF), 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 8x16GB 2400 MHz, BIOS 86B0271.R00, Motherboard Wildcat Pass, BMC 1.33.9832, FRU/SDR package 1.09, Red Hat 7.2 kernel 3.10.0-327.el7.x86_64, System Disk 1 1.0 TB SATA drive WD1003FZEX-00MK2A0, coprocessor N/A, Idle Power measurement 129W, energy usage to complete benchmark calculation in Joules: APOA1 – 4,565; STMV – 61,138.

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), 1700 MHz uncore freq., DDR4 6x16GB 2400 MHz quadrant cluster mode, MCDRAM 16 GB 6.4 GT/s flat memory mode, BIOS 10R00, Motherboard Adams Pass, Sleds per Chassis 1, BMC 12.951, FRU/SDR package 1.1, Red Hat 7.2 kernel 3.10.0-327.el7.x86_64, System Disk 1 1.0 TB SATA drive WD1003FZEX-00MK2A0, coprocessor N/A, Idle Power measurement 89W, energy usage to complete benchmark calculation in Joules: APOA1 – 3,899; STMV – 43,218.

Configuration details: BAW

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7210 (64 cores): Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo OFF), 1600 MHz uncore freq. MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode.

Intel® Xeon Phi™ processor 7250 (68 cores): Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Configuration details: TACC LBS3D

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON), 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat* 7.2

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq., (Turbo On), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10.R00, DDR4 16GB 2400 MHz, Red Hat 7.2, Quadrant cluster mode, MCDRAM+ DDR4 Flat mode

Additional details for LB3D configuration: 34x2/4T MPIxOMP decomposition on Intel Xeon Phi Processor based system. 36x1 MPIxOMP for Intel Xeon processor based platforms.

Configuration details: NASA OVERFLOW

Intel® Xeon® processor E5-2680 v3: Dual Socket Intel® Xeon® processor E5-2680 v3 2.5 GHz (Turbo ON) , 12 Cores/Socket, 24 Cores, 48 Threads (HT on), DDR4 128GB, 2133 MHz, SUSE Linux

Intel® Xeon® processor E5-2680 v4: Dual Socket Intel® Xeon® processor E5-2680 v4 2.4 GHz (Turbo ON) , 14 Cores/Socket, 28 Cores, 56 Threads (HT on), DDR4 128GB, 2400 MHz, SUSE Linux

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo On), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10.R00, DDR4 16GB 2400 MHz, SLES 12 SP1, Quadrant cluster mode, MCDRAM cache mode

Additional details for OVERFLOW configuration: 34x2 (DLR/F6) and 8x8 (NASrotor) MPIxOMP decomposition on Intel Xeon Phi Processor based system. 24x1 and 28x1 MPIxOMP for Intel Xeon processor based platforms.

Configuration details: OpenFOAM

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode; MCDRAM cache memory mode used for the 20M Cell Motorbike benchmark

OpenFOAM software: Development version from GitHub (https://github.com/OpenFOAM/OpenFOAM-dev, version-3.0.0-480-gf8f835c*), Intel® MPI Library Version 5.1.3 Build 20160120, Intel® Compiler Version16.0.1 Build 20151021.

NVIDIA vs Intel® Xeon Phi™ System Configuration and Pricing

NVIDIA System: (price quoted directly from Colfax website)
$13,737.59 with NVIDIA Tesla K80
$10,737.60 with NVIDIA TitanX

Node count:	1
Vendor:	Colfax
CPU:	2S Intel® Xeon® E5-2697v4
Memory:	128GB DDR4 2400
Storage:	1 TB SATA SSD
Accelerator:	1 NVIDIA Tesla K80 or 1 NVIDIA TitanX for OptiX

KNL System: (retail and recommended customer prices used)
$7,294.43 for Intel® Xeon Phi™ processor 7250
$4,994.43 for Intel® Xeon Phi™ processor 7210

Node count:	1
Platform:	Intel® Adams Pass with 1 populated sled
CPU:	Intel® Xeon Phi™ processor 7250 or Intel® Xeon Phi™ processor 7210
Memory:	96GB DDR4 2400
Storage:	1 TB SATA SSD

Configuration details: OpenLB

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 Cores/Socket, 36 Cores, 72 Threads (HT and Turbo ON), DDR4 128GB, 2400 MHz, CentOS release 6.7, Composer 2016.2.181

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (HT and Turbo ON). MCDRAM 16 GB 6.4 GT/s, DDR4 96GB 2133 MHz, Red Hat 7.2, Quadrant cluster mode, MCDRAM flat memory mode, Composer 2016.2.181

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (HT and Turbo ON). MCDRAM 16 GB 7.2 GT/s, DDR4 96GB 2400 MHz, Red Hat 7.2,Quadrant cluster mode, MCDRAM flat memory mode, Composer 2016.2.181

Configuration details: HiFUN

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON), 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, RHEL 6.5

Intel® Xeon Phi™ processor 7250: IIntel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode.

Configuration details: GE Tacoma

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON), 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, RHEL 7.2

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210, 64 cores, 256 threads, 1400 MHz core freq., 1600 MHz uncore freq., MCDRAM 16 GB 6400 MHz, BIOS 10.R01, DDR4 98 GB 2134 MHz, Quadrant cluster mode, MCDRAM flat mode

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1466 MHz core freq., 1700 MHz uncore freq., MCDRAM 16 GB 7200 MHz, BIOS 10.R00, DDR4 98 GB 2400 MHz, Quadrant cluster mode, MCDRAM flat mode.

Additional details for TACOMA configuration: 64x4 MPIxOMP decomposition on Intel Xeon Phi Processor based systems. 32x1 MPIxOMP for Intel Xeon processor E5-2697v4.

Configuration details: NEMO

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, Turbo OFF , 18 Cores/Socket, 36 Cores, 36 Threads (HT off), DDR4 128GB, 2400 MHz, Red Hat 7.2

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq., Turbo OFF, 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10R00, DDR4 96GB 2400 MHz, quad cluster mode, MCDRAM cache memory mode, Memory mode = quadrant, Red Hat 7.2

NEMO Recipe

1. You should be registered at NEMO web-site http://www.nemo-ocean.eu/
2. Obtain NEMO code:

svn co –r 6609 https://forge.ipsl.jussieu.fr/nemo/svn/branches/2016/dev_v3_6_STABLE_OMP/NEMOGCM

3. Obtain XIOS code and built it using following instruction:

http://www.nemo-ocean.eu/Using-NEMO/User-Guides/Basics/XIOS-IO-server-installation-and-use

4. Create custom .fcm file in NEMOGCM/ARCH directory based on avail configurations.
5. Add paths to NetCDF and XIOS in configuration.
6. Replace with “-r8 -O3 -openmp -xMIC-AVX512” (to build binary for the Intel® Xeon Phi™ processor) and with “-r8 -O3 -openmp -xCORE-AVX2” (to build binary for BDW) %FCFLAGS and change “%CPP” to “icc –E”, “%FC” to “mpiifort”, “%LD” to “mpiifort”, and “%LDFLAGS” to “-lstdc++ -lifcore”.
7. Use this instruction to build and run NEMO:

http://www.nemo-ocean.eu/Using-NEMO/User-Guides/Basics/NEMO-Quick-Start-Guide

8. GYRE is a default NEMO workload, avail in NEMO package. default grid resolution is 25, you can change it to make bigger or smaller grid. All instructions placed here - http://www.nemo-ocean.eu/Using-NEMO/Configurations/GYRE (you should register and login to view most of technical information about NEMO).

(continued next slide)

9. Create 3 workloads by modifying namelist_ref and namelist_cfg files:
    1. Switch creating mesh files to off by changing “nn_msh” to 0 in namelist_ref file.
    2. Enable benchmark mode by changing “nn_bench” to 1 in namelist_ref file.
    3. To create GYRE 30, GYRE 50 and GYRE 70 workloads set following params in namelist_cfg file:

Parameter	GYRE 30	GYRE 50	GYRE 70
jp_cfg	30	50	70
jpidta	902	1500	2102
jpjdta	602	1000	1402
jpkdta	22	22	22
jpiglo	902	1500	2102
jpjglo	602	1000	1402

Configuration details: Danish Meteorological Institute HIROMB-BOOS-Model*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.7

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Configuration details: MPAS Ocean 4.0*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, Turbo mode ON , 18 Cores/Socket, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 7.2. BIOS 86B0271.R00. Wildcat Pass Platform BMC version 1.33.9832 FRU/SDR Package 1.09. 1 1-TB SATA disk (Western Digital WD1003FZEX-00MK2A0) installed.

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. Turbo mode ON, 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10R00, DDR4 96GB 2400 MHz, Red Hat 7.2, quad cluster mode, MCDRAM Cache memory mode. Adams Pass Platform BMC version 12.9511 FRU/SDR Package 1.10. 1 1-TB SATA disk (Western Digital WD1003FZEX-00MK2A0) installed

Recipe for The Intel® Xeon Phi™ processor:

1. Building all the 3rd party libraries (NetCDF, ParallelNetCDF, PIO) are same as Intel Xeon.
2. Compilation of MPAS: In the Makefile, for ifort target, -xMIC-AVX512 is used as additional flag.
    "FFLAGS_OPT = -O3 –xCORE-AVX2 -convert big_endian -FR"
    "CFLAGS_OPT = -O3 –xCORE-AVX2 "
    "CXXFLAGS_OPT = -O3 –xCORE-AVX2 "
    Command: make ifort CORE=ocean MODE=forward
3. Running instructions:

Intel® Xeon Phi™ processor memory mode: cache, cluster mode: quadrant.
Following environment variables should be set:
export I_MPI_PIN_DOMAIN=core
export I_MPI_FABRICS=shm
ulimit -s unlimited
1 MPI rank/core is used for 68 core Intel® Xeon Phi™ processor 7250. “mpirun” command is same as the Intel Xeon processor.

EC 30 to 60 benchmark		
	Time (sec)	Sys. Power (Watts avg.)
Intel® Xeon® processor E5-2697 v4 (36 core)	4447	408.78
Intel® Xeon Phi™ processor 7250 (68 core)	3621	256.50

Performance Increase	**Watt Performance**
1.228113781	1.59

1.228 x 1.59 =1.95X

Configuration details: Non-Hydrostatic Icosahedral Model*
Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz Turbo OFF , 18 Cores/Socket, 36 Cores, 72 Threads HT on, DDR4 128GB, 2400 MHz, Red Hat 6.5
Intel® Xeon Phi™ processor 7210 : Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo OFF), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 09D10, DDR4 96GB 2133 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode
Configuration details: Non-Hydrostatic Icosahedral Model*
Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo OFF), 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5
Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo OFF), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 09D10, DDR4 96GB 2133 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode
Intel® OPA: Series 100 HFI ASIC (B0 silicon), Series 100 Edge Switch – 48 port (B0 silicon). Intel® OPA fabric software revision 10.0.1.0.50 (applies to cluster results only)
Configuration details: GNAQPMS
Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, Turbo ON , 18 Cores/Socket, 36 Cores, 72 Threads (HT ON), DDR4 128GB, 2400 MHz, CentOS release 6.7 (Final)
Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq., Turbo ON, 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 10D42, DDR4 96GB 2133 MHz, Red Hat 7.2 (Maipo), SNC4 cluster mode, MCDRAM cache memory mode, MPSP 1.2.2
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq., Turbo ON, 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10R00, DDR4 96GB 2400 MHz, Red Hat 7.2 (Maipo), SNC4 cluster mode, MCDRAM cache memory mode, MPSP 1.2.2
Code: It was an MPI application and optimized using OpenMP
Configuration details: Weather & Research Forecast Model*
Intel® Xeon® processor E5-2697 v4 : Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz , 18 Cores/Socket, 36 Cores, 72 Threads (HT and turbo on), DDR4 128GB, 2400 MHz, Red Hat 6.7
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (turbo on), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode
Configuration details: HOMME Atmospheric Dynamical Core*
Intel® Xeon® processor E5-2697 v4 : Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, Turbo mode ON , 18 Cores/Socket, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 7.2. BIOS 86B0271.R00. Wildcat Pass Platform BMC version 1.33.9832 FRU/SDR Package 1.09 1 1-TB STAT disk and 1 800GB SSD disk installed. Run details: Used 36 MPI ranks with 2 horizontal OpenMP threads per rank.
Intel® Xeon Phi™ processor 7250 : Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. Turbo mode ON, 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10R00, DDR4 96GB 2400 MHz, Red Hat 7.2, quad cluster mode, MCDRAM flat memory mode. Adams Pass Platform BMC version 12.9511 FRU/SDR Package 1.10. 1 1-TB SATA disk installed. Run details: Used 64 MPI ranks with 2 horizontal OpenMP threads per rank.
Configuration details: POP
Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 Cores/Socket, 36 Cores, 72 Threads (HT and Turbo ON), DDR4 128GB, 2400 MHz, Oracle Linux Server release 6.7
Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core , 256 threads, 1300 MHz core freq. (HT and Turbo ON), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 10D28, DDR4 96GB 2133 MHz, Red Hat 7.2, quad cluster mode, MCDRAM cache memory mode, MPSP 1.2.2; MKL:
11.3.2
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (HT and Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10D28, DDR4 96GB 2400 MHz, Red Hat 7.2, quad cluster mode, MCDRAM cache memory mode, MPSP 1.2.2; MKL:
11.3.2

Configuration details: MASNUM WAVE

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT ON), DDR4 128GB, 2400 MHz, CentOS release 6.7(Final)
Intel® Xeon Phi™ processor 7210 (64 cores): Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo ON), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 10D42, DDR4 96GB 2133 MHz, Red Hat 7.2 (Maipo), SNC4 cluster mode, MCDRAM cache memory mode
Intel® Xeon Phi™ processor 7250 (68 cores): Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10R00, DDR4 96GB 2400 MHz, Red Hat 7.2 (Maipo), SNC4 cluster mode, MCDRAM cache memory mode

Trinity Baseline Configurations

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 36 Threads (HT off), DDR4 128GB, 2400 MHz, Red Hat 6.5
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo on: 1500 MHz), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, DDR4 96GB 2400 MHz, SUSE Linux 12, various cluster and memory modes

Workload	Xeon Phi 7250 Config	Xeon E5-2697 v4 Config
AMG	SNC4 Cache; 60x60x60; 272 ranks	109x109x109; 36 ranks
MiniFE	Quad Flat; 307x307x307; 136 ranks	244x244x244; 36 ranks
UMT	SNC4 Cache; 7x7x7; 272 ranks	7x7x7; 36 ranks
SNAP	SNC4 Flat; 32x64x68; 136 ranks	32x24x48; 36 ranks
GTC	Quad Flat; npartdom=2, micell=200; 128 ranks	Npartdom=1; micell=100; 36 ranks
MILC	SNC4 Flat; 16x32x32x34; 136 ranks	16x16x16x36; 36 ranks
MiniGhost	SNC4 Flat; 268x268x272; 136 ranks	452x453x456; 36 ranks
MiniDFT	Quad Flat; Single Node Workload; 68 ranks	Single Node Workload; 36 ranks

Configuration details: CP2K* Linear Scaling (LS) Density Function Theory (DFT)

Intel® Xeon® processor E5-2695 v4: Dual Socket Xeon® processor E5-2695 v4 2.1 GHz, 18 cores/socket, 36 cores, 72 threads (HT and Turbo ON), DDR4 64 GB, 2400 MHz, SUSE 11.3
Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210, 64 core (256 threads), 1.3 GHz base core freq. (Turbo ON), 1.6 GHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 86B.01.01.0124, DDR4 96 GB 2133 MHz, SNC4 cluster mode, MCDRAM cache memory mode, RHEL 7.2, MPSP 1.3.0, Intel Compiler 2017

Configuration details: BerkeleyGW (Sigma Phase) Benzene

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON), 18 Cores/Socket, 36 Cores, 72 Threads (HT on), Wildcat Pass, DDR4 128GB, 2400 MHz, BMC ver. 1.33.9832, Red Hat 7.2, BIOS 86B0271.R00, FRU/SDR Package 1.09, kernel 3.10.0-327.el7.x86_64, 1 1.0 TB SATA drive WD1003FZEX-00MK2A0, Idle Power measurement 89W
Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7250 64 core, 256 threads, 1300 MHz core freq. 6 x 16 GB 2400 MHz DDR4, Memory mode = cache, Cluster mode=quadrant, BMC ver. 12.951, Red Hat 7.2, BIOS 10R00, FRU/SDR Package 1.1, kernel 3.10.0-327.el7.x86_64, 1 1.0 TB SATA drive WD1003FZEX-00MK2A0, Idle Power measurement 122W
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. 6 x 16 GB 2400 MHz DDR4, Memory mode = cache, Cluster mode=quadrant, BMC ver. 12.951, Red Hat 7.2, BIOS 10R00, FRU/SDR Package 1.1, kernel 3.10.0-327.el7.x86_64, 1 1.0 TB SATA drive WD1003FZEX-00MK2A0, Idle Power measurement 125W

Configuration details: PWmat

Intel® Xeon® processor E5-2697 v4: 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, CentOS release 6.7
Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq., HT ON, Turbo ON, 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 10R00, DDR4 96GB 2133 MHz, Red Hat 7.2, SNC4 cluster mode, MCDRAM cache memory mode, MKL: 11.3.2
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq., HT ON, Turbo ON, 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10R00, DDR4 96GB 2400 MHz, Red Hat 7.2, SNC4 cluster mode, MCDRAM cache memory mode, MKL: 11.3.2

Building CP2K
LIBXSMM: git clone https://github.com/hfp/libxsmm.git; git checkout tags/1.4.3
Automatically selected when using ARCH files as mentioned below
CP2K: make ARCH=Linux-x86-64-intel VERSION=psmp AVX=3 MIC=1 LIBXSMM=2
Either CP2K intel branch (git clone --branch intel https://github.com/cp2k/cp2k.git cp2k.git)
Or master CP2K 4.0-development (https://github.com/cp2k/cp2k.git cp2k.git)
https://github.com/cp2k/cp2k/raw/intel/cp2k/arch/Linux-x86-64-intel.x
https://github.com/cp2k/cp2k/raw/intel/cp2k/arch/Linux-x86-64-intel.psmp
Runing CP2K
CP2K/intel@KNL: CP2K_RECONFIGURE=1 (huge pages), CP2K_STACKSIZE=10000 (SMMs/batch),
I_MPI_PIN_DOMAIN=auto, I_MPI_PIN_ORDER=scatter, 64 ranks, 2 thread/core

Configuration details: Quantum ESPRESSO*
Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 cores/socket, 36 cores, 72 threads (HT and Turbo ON), DDR4 64 GB, 2400 MHz, RHEL 6.7
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7260, 68 core (272 threads), 1.4 GHz base core freq. (Turbo ON), 1.7 GHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS GVPRCRB1.86B.0010.R00, DDR4 96 GB 2400 MHz, quadrant cluster mode, MCDRAM cache memory mode, RHEL 6.7, MPSP 1.3.0, Intel Compiler 2017
Building Quantum ESPRESSO
- Configure for Intel Compiler and MKL according to http://www.quantum-espresso.org/wp-content/uploads/Doc/user_guide/node14.html#SECTION00037670000000000000
- Modify defines in make.sys to
  -D__INTEL -D__OPENMP -D__DFTI -D__MPI -D__PARA -D__ELPA -D__SCALAPACK
- Add -xMIC-AVX512 to CFLAGS and FFLAGS

Running Quantum ESPRESSO
Intel® Xeon Phi™ processor 7250
mpirun -n 68 <PATH_TO_BINARY>/pw.x -nk 2 -nt 34 -nd 25 -i ausurf.in
Intel® Xeon® processor E5-2697 v4:
mpirun -n 36 <PATH_TO_BINARY>/pw.x -nk 1 -nt 36 -nd 36 -i ausurf.in

Configuration details : Qphix*
Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo OFF) , 18 Cores/Socket, 36 Cores, 72 Threads (HT ON), DDR4 128GB, 2400 MHz, Red Hat 6.5, Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode, Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16

Configuration details: Cloverleaf*
Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5
Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo OFF), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 09D10, DDR4 96GB 2133 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode
Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Configuration details: VLPL-S

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads, HT ON, DDR4 128GB, 2400 MHz, CentOS release 6.7(Final)

Intel® Xeon Phi™ processor 7210): Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq., Turbo ON, 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 10D42, DDR4 96GB 2133 MHz, Red Hat 7.2 (Maipo), SNC4 cluster mode, MCDRAM cache memory mode

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq., Turbo ON, 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10R00, DDR4 96GB 2400 MHz, Red Hat 7.2 (Maipo), SNC4 cluster mode, MCDRAM cache memory mode

Configuration details: Soft Sphere Simulation

Intel® Xeon® processor E5-2697 v4: Intel® Xeon® Dual Socket processor E5-2697 v4 2.3 GHz, 18 Cores/Socket, 36 Cores, 72 Threads (HT and Turbo ON), DDR4 128GB, 2400 MHz, Oracle Linux Server release 6.7

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (HT and Turbo ON), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 10D28, DDR4 96GB 2133 MHz, Red Hat 7.2, quad cluster mode, MCDRAM cache memory mode.

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (HT and Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10D28, DDR4 96GB 2400 MHz, Red Hat 7.2, quad cluster mode, MCDRAM cache memory mode.

Configuration details: PETSC – Portable, Extensible Toolkit For Scientific Computation*

Intel® Xeon® processor E5-2697 v4 : Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo OFF) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7250 : Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. Turbo mode ON, 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10R00, DDR4 96GB 2400 MHz, Red Hat 7.2, quad cluster mode, MCDRAM flat memory mode. Adams Pass Platform BMC version 12.9511 FRU/SDR Package 1.10. 1 1-TB SATA disk installed.

Configuration & Recipe details: YASK HPC Stencils, AWP-ODC Kernel

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat Enterprise Linux Server release 7.2

Recipe:

- Download code from https://github.com/01org/yask and install per included directions
- make stencil=awp arch=hsw cluster=x=2,y=2,z=2 fold=y=8 omp_schedule=guided mpi=1
- ./stencil-run.sh -arch hsw -ranks 2 -bx 74 -by 192 -bz 20 -pz 2 -dx 512 -dy 384 -dz 768

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250, 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 86B.0010.R00, DDR4 96GB 2400 MHz, quad cluster mode, MCDRAM flat memory mode, Red Hat Enterprise Linux Server release 6.7

Recipe:

- Download code from https://github.com/01org/yask and install per included directions
- make stencil=awp arch=knl INNER_BLOCK_LOOP_OPTS='prefetch(L1,L2)'
- ./stencil-run.sh -arch knl -bx 128 -by 32 -bz 32 -dx 1024 -dy 384 -dz 768

Configuration & Recipe details: YASK HPC Stencils, iso3DFD Kernel

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat Enterprise Linux Server release 7.2

Recipe:

- Download code from https://github.com/01org/yask and install per included directions
- make stencil=iso3dfd arch=hsw mpi=1
- ./stencil-run.sh -arch hsw -ranks 2 -bx 256 -by 64 -bz 64 -dx 768 -dy 1024 -dz 768

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250, 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 86B.0010.R00, DDR4 96GB 2400 MHz, quad cluster mode, MCDRAM flat memory mode, Red Hat Enterprise Linux Server release 6.7

Recipe:

- Download code from https://github.com/01org/yask and install per included directions
- make stencil=iso3dfd arch=knl
- ./stencil-run.sh -arch knl -bx 192 -by 96 -bz 96 -dx 1536 -dy 1024 -dz 768

Configuration details: SeisSol Seismic Solver*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 Cores/Socket, 36 Cores, 72 Threads (HT and Turbo ON), DDR4 128GB, 2400 MHz, RHEL 6.7

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (HT and Turbo ON), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 09D10, DDR4 96GB 2133 MHz, RHEL 6.7, quad cluster mode, MCDRAM flat memory mode, MPSP 1.2.2; MKL:
11.3.2

Intel® Xeon Phi™ processor 7250 (68 cores): Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (HT and Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, SUSE 12, quad cluster mode, MCDRAM flat memory mode, MPSP 1.2.2; MKL: 11.3.2

Configuration details: SPECFEM3D_GLOBE

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo OFF) , 18 Cores/Socket, 36 Cores, 72 Threads (HT ON), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo OFF), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode, Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16

Configuration details: 3D Isotropic Finite Difference*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz, 18 Cores/Socket, 36 Cores, 72 Threads (HT ON, Turbo OFF), DDR4 128GB, 2400 MHz, Oracle Linux Server release 6.7

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (HT and Turbo ON), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 10D28, DDR4 96GB 2133 MHz, CentOS 7.2, quad cluster mode, MCDRAM flat memory mode, MPSP 1.3.1

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (HT and Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 10D28, DDR4 96GB 2400 MHz, CentOS 7.2, quad cluster mode, MCDRAM flat memory mode, MPSP 1.3.1

Configuration details: Sparse Matrix Vector Multiply using SpMV*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo off), 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo off), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 09D10, DDR4 96GB 2133 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo off), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Configuration details: High Performance Conjugate Gradients*

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON) , 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7210: Intel® Xeon Phi™ processor 7210 64 core, 256 threads, 1300 MHz core freq. (Turbo ON), 1600 MHz uncore freq., MCDRAM 16 GB 6.4 GT/s, BIOS 09D10, DDR4 96GB 2133 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode

Configuration details: High Performance Conjugate Gradients* (cluster)

Intel® Xeon® processor E5-2697 v4: Dual Socket Intel® Xeon® processor E5-2697 v4 2.3 GHz (Turbo ON), 18 Cores/Socket, 36 Cores, 72 Threads (HT on), DDR4 128GB, 2400 MHz, Red Hat 6.5

Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), 1700 MHz uncore freq., MCDRAM 16 GB 7.2 GT/s, BIOS 09D10, DDR4 96GB 2400 MHz, Red Hat 6.7 (Santiago), quad cluster mode, MCDRAM flat memory mode. Cluster results exploited Intel® OPA Fabric.

Intel® OPA: Series 100 HFI ASIC (B0 silicon), Series 100 Edge Switch – 48 port (B0 silicon). Intel® OPA fabric software revision 10.0.1.0.50 (applies to cluster results only)

Intel® MPI version was 5.1.3.181. OpenMP from Intel Compiler 16.0 update 1. MKL is not used/required for running HPCG. For optimal performance, KNL should be booted in Quadrant Cluster Mode, Flat Memory Mode, Turbo Mode enabled.

Running on 68-cores Xeon Phi 7250 node with MCDRAM on NUMA node 1: export KMP_PLACE_THREADS=17c,2t export KMP_AFFINITY=granularity=fine,compact #> mpiexec.hydra -n 4 –hosts knl7250 numactl --membind=1 xhpcg_knl --n=160

For multiple OPA linked nodes following OPA parameters used:

export PSM2_MQ_RNDV_HFI_WINDOW=4194304 export PSM2_MQ_EAGER_SDMA_SZ=65536 export PSM2_MQ_RNDV_HFI_THRESH=200000 export PSM2_IDENTIFY=1 export I_MPI_FABRICS=shm:tmi

export I_MPI_TMI_PROVIDER=psm2 export I_MPI_FALLBACK=0

Running on several nodes made as follows: export KMP_PLACE_THREADS=17c,2t export KMP_AFFINITY=granularity=fine,compact #> mpiexec.hydra -n 8 -hosts knl1,knl2 -ppn 4 numactl --membind=1 xhpcg_knl --n=160