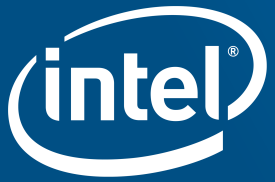

Intel® Xeon Phi™ Processor Update

Barry Davis
General Manager, Accelerated Workloads Group (AWG)
Intel Data Center Group

27 September 2017

HPC is Transforming, accelerating Deeper Insight

Advancing Science

High ROI: $515
Return Per $1 of HPC Invest[1]

Data-Driven Analytics

Diverse & New Workloads Driving Science & Industry

Modeling & Simulation

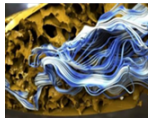

Visualization

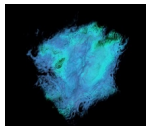

High Performance Data Analytics

Machine Learning

[1]Source: IDC HPC and ROI Study Update (September 2015)

Agenda

- HPC Market & Trends
- Yesterday and Today
- What's Next

HPC Market & trends

The Coming Flood of Data (and Compute)

BY 2020…

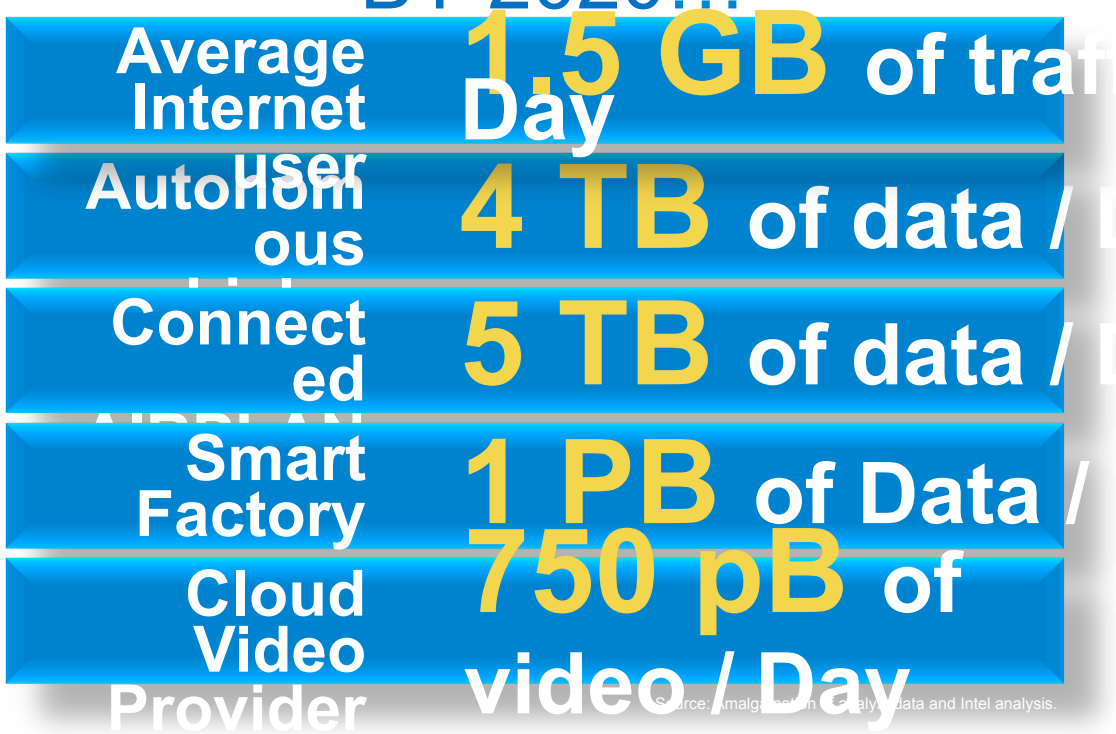

Source: Amalgamation of analyst data and Intel analysis.

HPC Market

2016

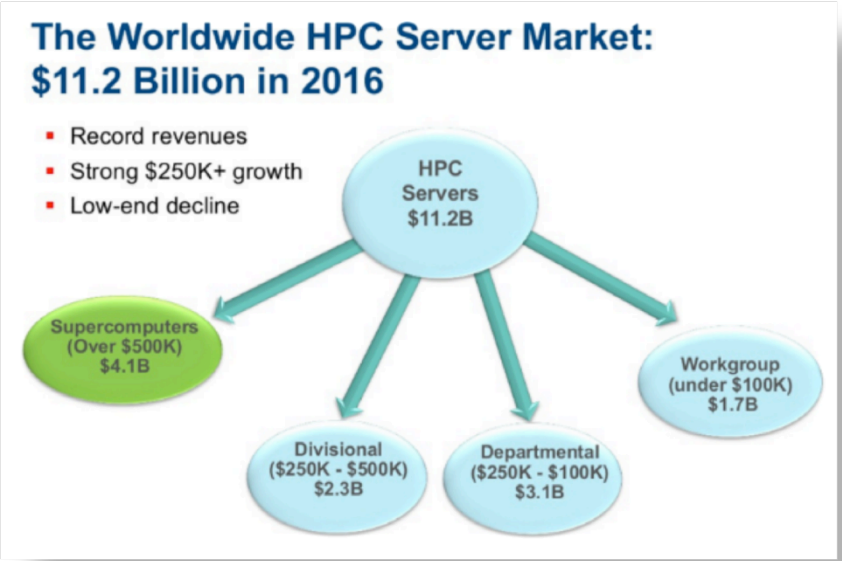

Source: https://www.top500.org/news/hpc-market-hits-record-revenues-in-2016-looks-ahead-to-double-digit-growth-in-ai

2021

HPC Growth to $14.8B, $5.4B in Supercomputing*

Exascale Systems

Optimized and Efficient Processing

Distributed computing and expanding parallelism

More real-time computing and Network Bandwidth

HPC Trends

Exascale Computing

Artificial Intelligence

Workflow Convergence

Thoughts on Exascale

ECP Goals

- Foster a rich exascale ecosystem
- Help ensure continued U.S. HPC Leadership
- Enable delivery of an exascale system in 2021 and capable exascale computers in 2022

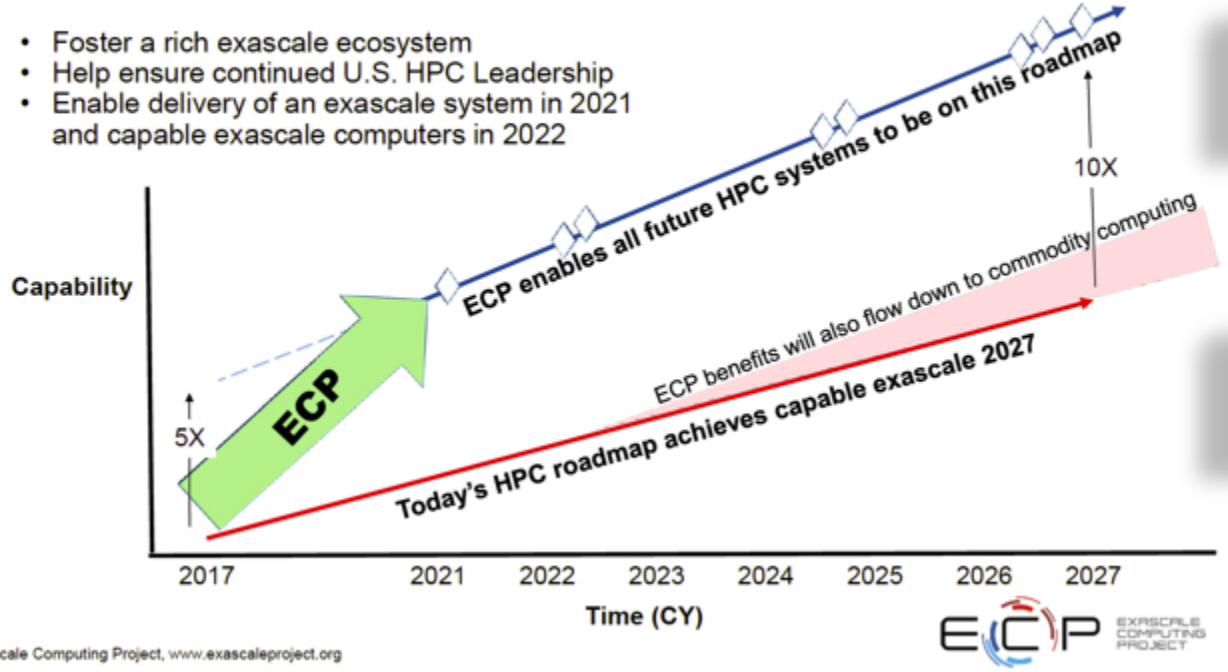

50X Application Performance Req

Needs to be a Production System

6 Exascale Computing Project, www.exascaleproject.org

Dawn of Artificial Intelligence

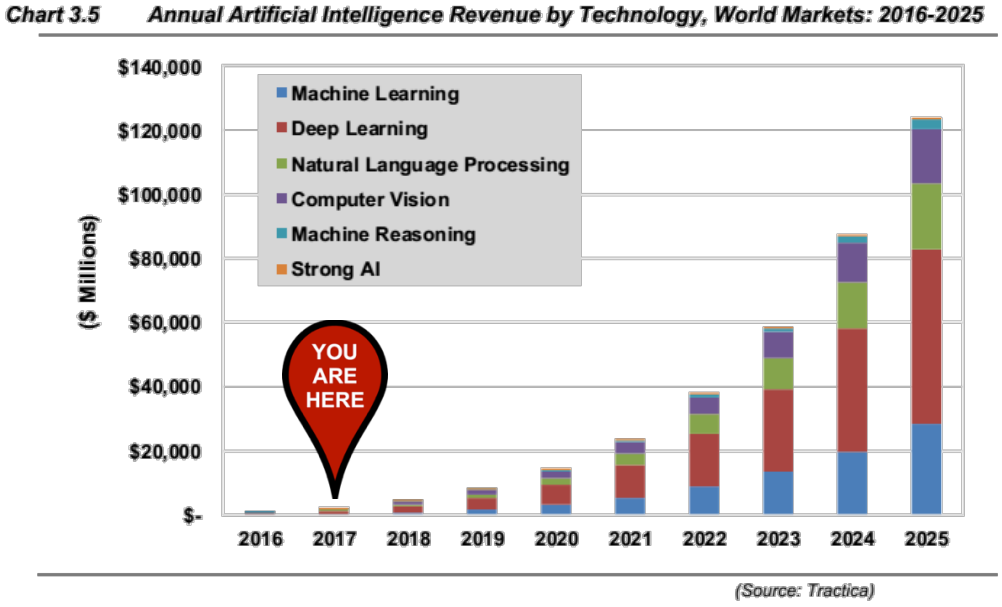

"In the past a lot of S&P 500 CEOs wished they had started thinking sooner than they did about their Internet strategy. I think five years from now there will be a number of S&P 500 CEOs that will wish they'd started thinking earlier about their AI strategy."

– Andrew Ng
AI luminary

Quote source: http://fortune.com/ai-artificial-intelligence-deep-machine-learning/

Workflow Convergence → Simulation, Data Analytics and AI

Our customers are telling us HPC is changing……

- Clear evidence → ECP must treat Mod/Sim, AI, and Analytics as first class citizens (hardware & software)

AI has become a major consumer of computing cycles and it is expected to grow

- Compute deployment both at edge and in large cloud
- Will drive economies in fabric, compute with a large focus on power and perf/W

Convergence is happening in many areas

- Cloud management and development tools/environments
- Fabric architectures
- Frameworks/topologies

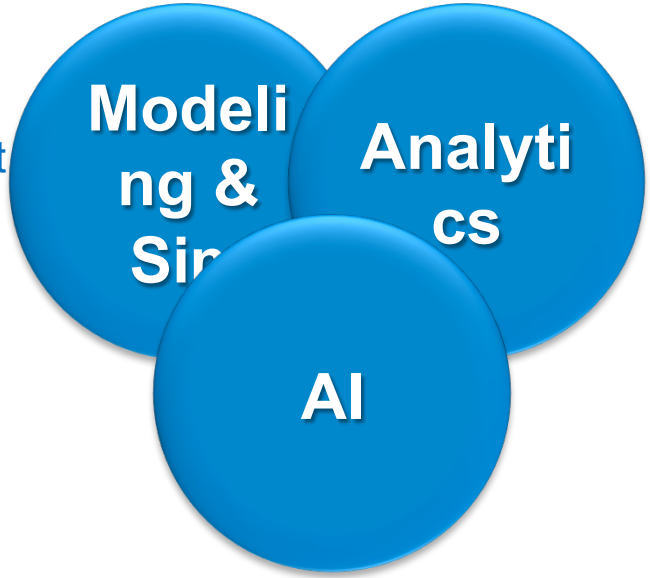

Yesterday and today

Path to Exascale Computing

General Purpose Hardware

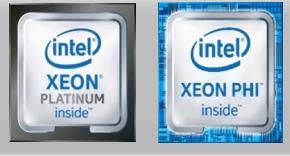

Modern Applications

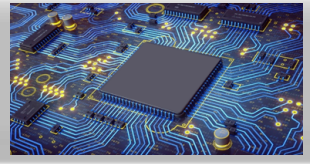

Stds Based Progr. Model → cores, caches, vectors

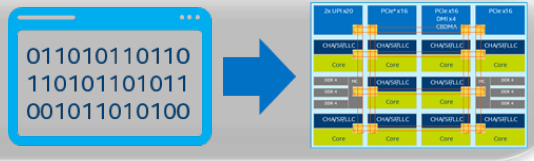

Commercial Roadmap

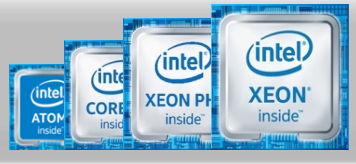

So What Are Our Results to Date?

Supercomputing: No Longer Restricted To A Few ‘Superpowers’

$46 million [1]

In the late 1990s, the United States built ASCI Red, the first computer to break one trillion floating-point calculations per second barrier

$2438-$4876 [2]

Twenty years later, 3 Teraflops of double precision compute could be delivered in one compute node one Intel® Xeon Phi™ 7200 family processor

1. Source: http://www.extremetech.com/extreme/125271-the-history-of-supercomputers/6
2. Based on Recommended Customer Price http://ark.intel.com/compare/94033,94034,94035,95828,95829,95830,95831,94709

Intel-SURFsara Research Collaboration*

INTEL-SURFSARA RESEARCH COLLABORATION *

Research Goals:

- Fastest time-to-train and time per iteration for deep neural network models
- Generalization of methodology to various Intel CPU architectures

Recent Breakthroughs <blog link> **

- **Achieving Deep Learning Training in less than 40 Minutes on ImageNet-1K**
- **Highest-to-date Accuracy and Training Time on ImageNet-22K and Places-365**

Scaling up Deep Learning on Stampede 2*
Resnet-50 on Imagenet-1K

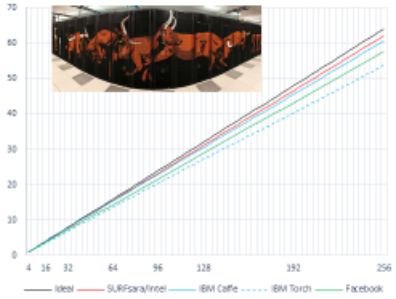

System specs
- Intel® Xeon Phi™ Processor 7250 (KNL)
- 96GB RAM
- Cache-quadrant mode

All models converge at Top1/5 > **74/92%**

1-256 node runs
- batch size of 32 per node
- scaling efficiency of **97%**
- Total time to train: **70 minutes**

512-node runs
- batch of 24 per node: **46 minutes**

768-node runs
- batch of 24 per node: **39 minutes**

SURFsara/Intel: Scaling beyond Imagenet-1K
Industry leading accuracy on large datasets *

Imagenet-22K
14.2 million images; 21841 object categories

Architecture: Wide Residual Networks trained in a large-scale distributed fashion

Network	Global batch size	# nodes	Accuracy (top1/top5)
Resnet-101 (IBM)	5120	256	33.8
WRN-50-2	7680	240	35.97/63.87
WRN-50-2	6400	200	**36.91/65.08**

Places-365
8 million images, 365 scene categories

Architecture: Wide Residual Networks trained in a large-scale distributed fashion

Network	Global batch size	# nodes	Accuracy (top1/top5)
WRN-50-2	6720	256	**57.26/87.46**
Resnet-152 baseline	N/A	N/A	--/85.07

Scaling up Deep Learning on MareNostrum 4* **
Resnet-50 on Imagenet-1K

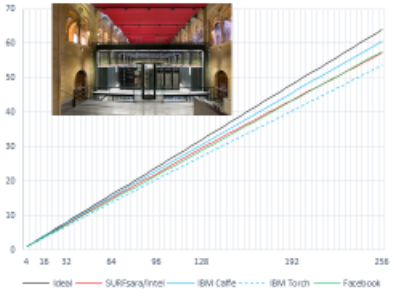

System specs
- 2S Intel® Xeon® Processor 8160 (SKX)
- 96GB RAM

All models converge at Top1/5 > **74/92%**

1-256 node runs
- batch size of 32 per node
- scaling efficiency of **90%**
- Total time to train: **70 minutes**

512-node run
- batch of 16 per node: **44 minutes**

** We acknowledge PRACE for awarding us access to resource MareNostrum 4 based in Spain at Barcelona Supercomputing Center

Lawrence Berkeley National Laboratory & Stanford University

Cori Supercomputer: First 15 Peta FLOP Deep Learning System

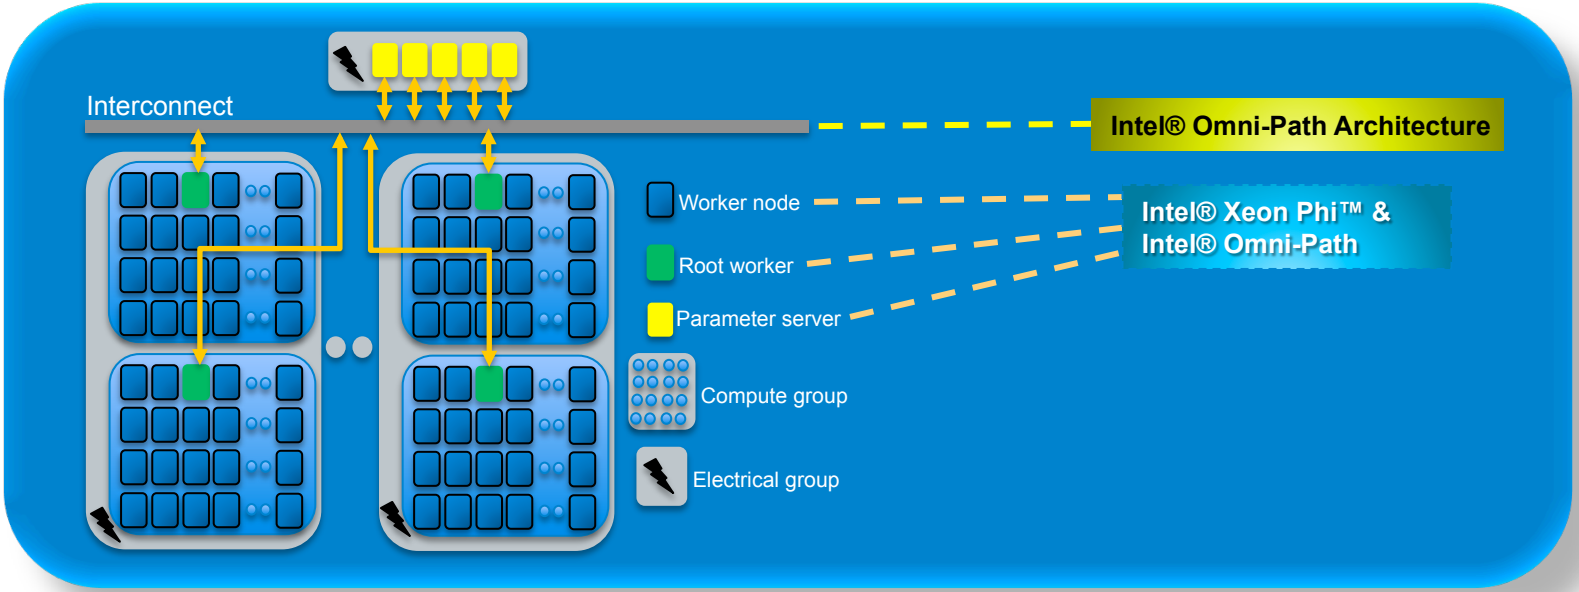

First 15-petaflop DL System For Solving Scientific Pattern Classification Problems On HPC Architectures

Deep Learning at 15PF†
(with NERSC, Stanford, and Univ of Montreal)

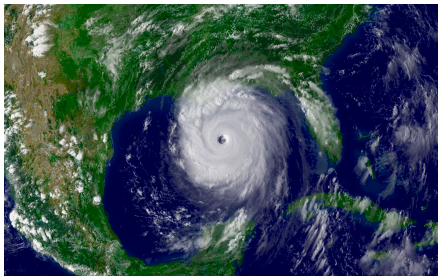

Scientific Achievement

- Signal vs. Background classification for LHC datasets exceeds physics cuts
- Pattern discovery for Climate data

Methods Achievement

- Hybrid parameter update strategy
- Supervised and semi-supervised architectures

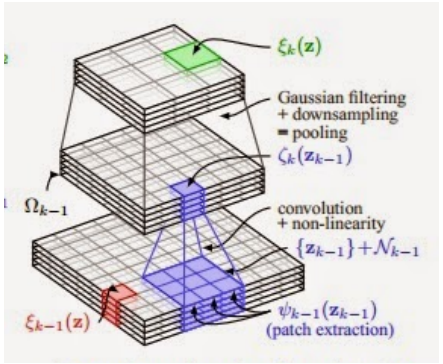

CS Achievement

- IntelCaffe + MLSL optimized on KNL
- ~2TF peak on single KNL node
- ~15 PF peak on ~9300 nodes

† "Petascale Deep Learning" Thorsten Kurth, Jian Zhang, Nadathur Satish, Ioannis Mitliagkas, Evan Racah, Mostofa Patwary, Tareq Malas, Narayanan Sundaram, Wahid Bhimji, Mikhail Smorkalov, Jack Deslippe, Mikhail Shiryaev, Srinivas Sridharan, Prabhat, and Pradeep Dubey, accepted at Supercomputing 2017

Celeste: 1st Julia application to hit 1PF†

(In collaboration with NERSC, UCB, MIT and Julia Computing)

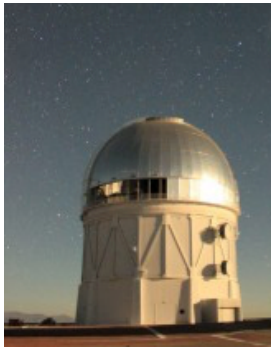

Scientific Achievement

- First catalog with parameter & *uncertainty* estimates for over 300M obje
- 55 TB SDSS dataset processed in 15 minutes
- DESI instrument will use catalog for target selection

Methods Achievement

- Bayesian Inference on world's largest generative model (in science)
- Joint estimation of Billions of parameters

CS Achievement

- Code written in Julia, optimized for execution on KNL
- Code scaled on 9300 Cori KNL nodes

† Cataloging the Visible Universe through Bayesian Inference at Petascale in Julia; https://www.youtube.com/watch?v=uecdcADM3hY&feature=youtu.be

Intel® Xeon Phi™ Processor Top500 Listings

June 2017 Top500 List[1] has 13 Intel Xeon Phi deployments
– over 57.5 PetaFlops

Seven listings in Top 50:

#6: Cori (NERSC, USA); Cray XC – 14 PFs

#7: Oakforest PACS (JCAHPC, Japan); Fujitsu CX1640 M1 – 13.5 PFs

#12: Stampede2 (TACC, USA); Dell PowerEdge – 6.8 PFs

#14: Marconi (CINECA, Italy); Lenovo – 6.2 PFs

#16: Theta (Argonne National Lab, USA); Cray XC40 – 5.8 PFs

#29: Onyx (ERDC DSRC, USA); Cray XC40 – 3.4 PFs

#37: Camphor 2 (ACCMS, Kyoto Univ, Japan); Cray XC40 – 3.0 PFs

[1]Top 500 Results: https://www.top500.org/lists/2017/06/

Where Does This Leave Us?

Based on our progress ….

- Continuing to deliver better, modern code and scale application performance
- Incorporating new capabilities and technologies into the HPC Workflow
- On track for Exascale in 2021

"Now this is not the end. It is not even the beginning of the end. But it is, perhaps, the end of the beginning."

– Winston Chrchill

We've Made Great Progress…But It's Time To Double Down!

What’s next

Today…Highly parallel

Intel® Xeon® Processors are increasingly parallel with larger vectors…and benefits from modern code

Intel® Xeon Phi™ Processors are extremely parallel and use general purpose programming

Up to 72 cores (288 threads)

V[512]

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

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to http://www.intel.com/performance

When to Use Intel® Xeon Phi™

Why Xeon Phi™?

Improve Performance

-AND/OR-

Improve ROI

Unlock Potential

Which Apps?[1]

- Scalable to >60 cores

AND

- Heavily Vectorized

-OR-

- Local memory BW bound

If yes…

Optimized for Highly-Parallel Applications

If no…

Commonly-Used Parallel Processor*

Intel® Xeon Phi™ is optimal for applications that scale to >60 cores and are highly threaded or memory bandwidth bound

[1]Performance results on Intel® Xeon Phi™ will vary depending on app characteristics. For more information, see: https://software.intel.com/sites/default/files/article/383067/is-xeon-phi-right-for-me.pdf

Intel® Xeon® Processor Scalable Family (Skylake)

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

Most agile
AI platform

Built-in ROI

Begin your AI journey today using existing, familiar infrastructure

- Increase datacenter utilization without additional, unique investment
- Use common AI frameworks and/or BigDL for scale deep learning training on Spark Hadoop

Potent Performance

Up to 2.2X deep learning training & inference perf vs. prior gen[1]

- ✓ More cores, threads (up to 28 & 56)
- ✓ Intel® AVX-512
- ✓ Optimized SW: cut training time from days to hours with up to 113x perf vs. prior gen[2]
- ✓ Faster I/O
- ✓ 6 mem channels
- ✓ INT8 support

Production-ready

Robust support for full range of AI deployments

- Runs all combinations of AI workloads:
  - ❑ Classic ML
  - ❑ Deep Learning
  - ❑ Reasoning
  - ❑ Emerging AI
  - ❑ Analytics
  - ❑ More
- Server-class reliability, hardware enhanced security and manageability

[1,2]Configuration details on slide: 40
Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit: http://www.intel.com/performance Source: Intel measured as of November 2016
Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.
Notice Revision #20110804

AI

Datacenter

All purpose

Intel® Xeon® Processor Family

most agile AI Platform

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

Highly-parallel

Intel® Xeon Phi™ Processor (Knights Mill†)

Faster DL Training

Scalable performance optimized for even faster deep learning training and select highly-parallel datacenter workloads*

Flexible acceleration

Intel® FPGA

Enhanced DL Inference

Scalable acceleration for deep learning inference in real-time with higher efficiency, and wide range of workloads & configurations

Deep Learning

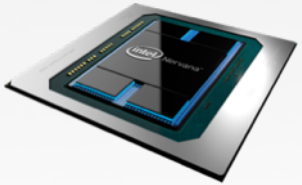

Crest Family†

Deep learning by design

Scalable acceleration with best performance for intensive deep learning training & inference

†Codename for product that is coming soon
All performance positioning claims are relative to other processor technologies in Intel's AI datacenter portfolio
*Knights Mill (KNM); select = single-precision highly-parallel workloads generally scale to >100 threads and benefit from more vectorization, and may also benefit from greater memory bandwidth e.g. energy (reverse time migration), deep learning training, etc.
All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

What is Knights Mill?

First Intel product targeted specifically at Deep Learning training workloads

- Up to 4x DL Peak performance over Intel® Xeon Phi™ Processors 7200 Series[1]

Built on top of 2nd generation Intel® Xeon Phi™ processor

- Improved efficiency
- Optimized for scale-out
- Enhanced variable precision
- Flexible, high capacity memory

Knights Mill exploits 4 levels of parallelism

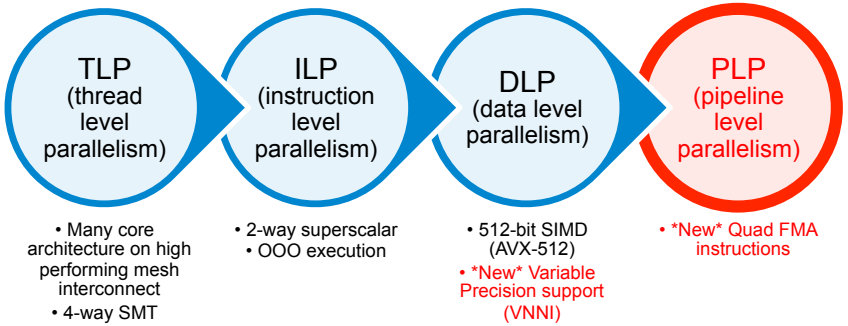

[1]Intel internal estimate: Performance estimate wrt Xeon Phi™ 7290 SKU SGEMM. Performance Calculation= AVX freq X Cores X Flops per Core X Efficiency

Goals Moving Forward

Hasten Pace Of Architectural Innovation And Increase Cadence Of New Products

Deliver Improved, Real World Application Performance

Single Platform Scalable to Multiple Workloads – Mod/Sim, AI (Machine & Deep Learning), Analytics → On-prem and in the Cloud

Some Food For Thought

Application Performance: Memory BW vs FLOPS

Application Performance: Memory BW

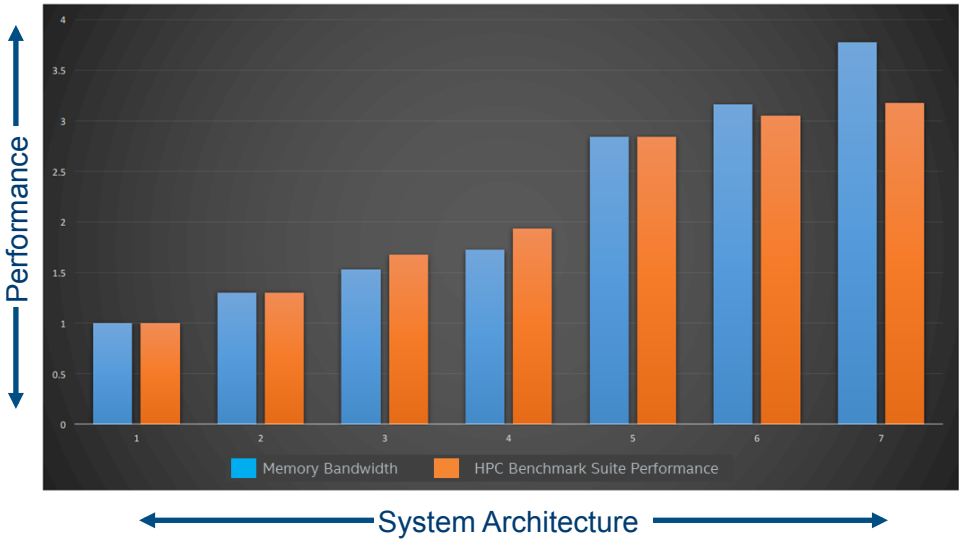

Application Performance: FLOPS

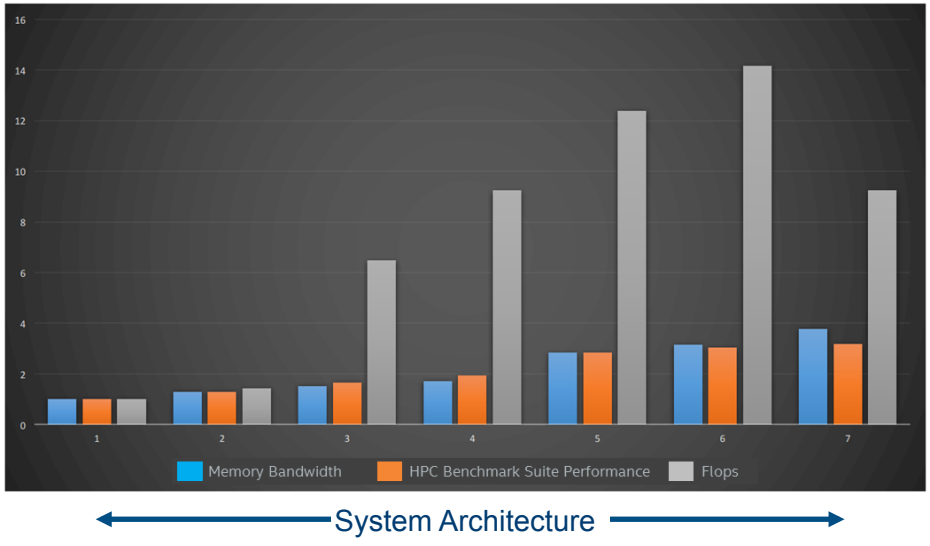

Application Performance comes from Memory BW not FLOPS

Source: Supercomputing OEM

Prior Direction

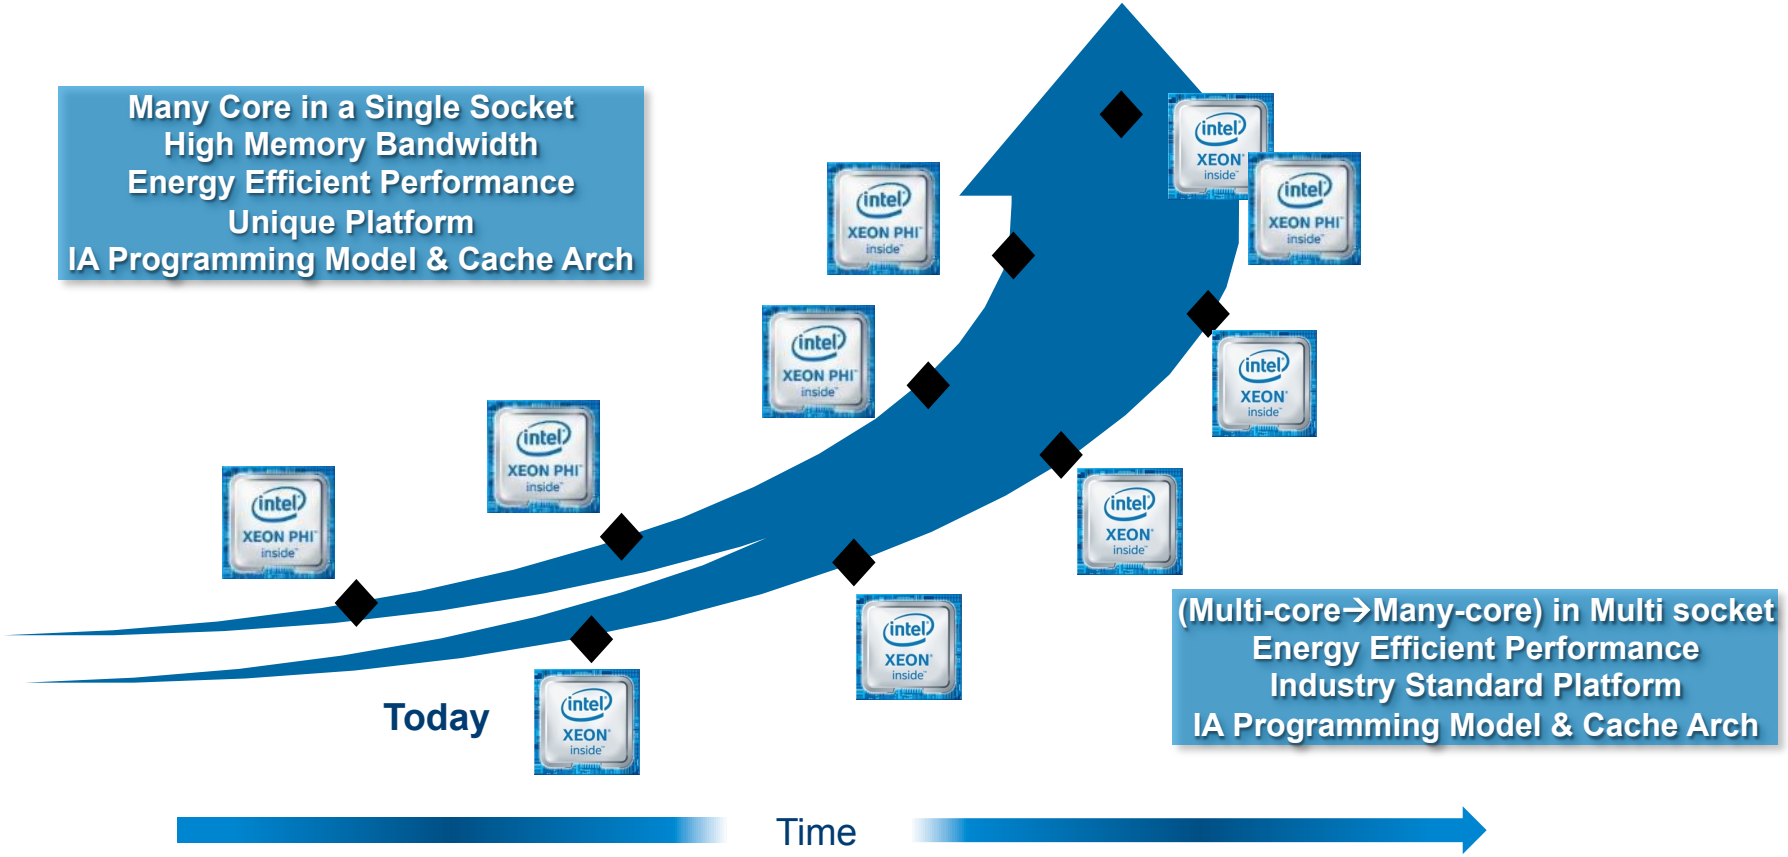

Today's Direction

(Multi-core→Many-core) in Multi socket
Energy Efficient Performance
Industry Standard Platform
IA Programming Model & Cache Arch
High Memory Bandwidth
Parallel Workload Optimizations

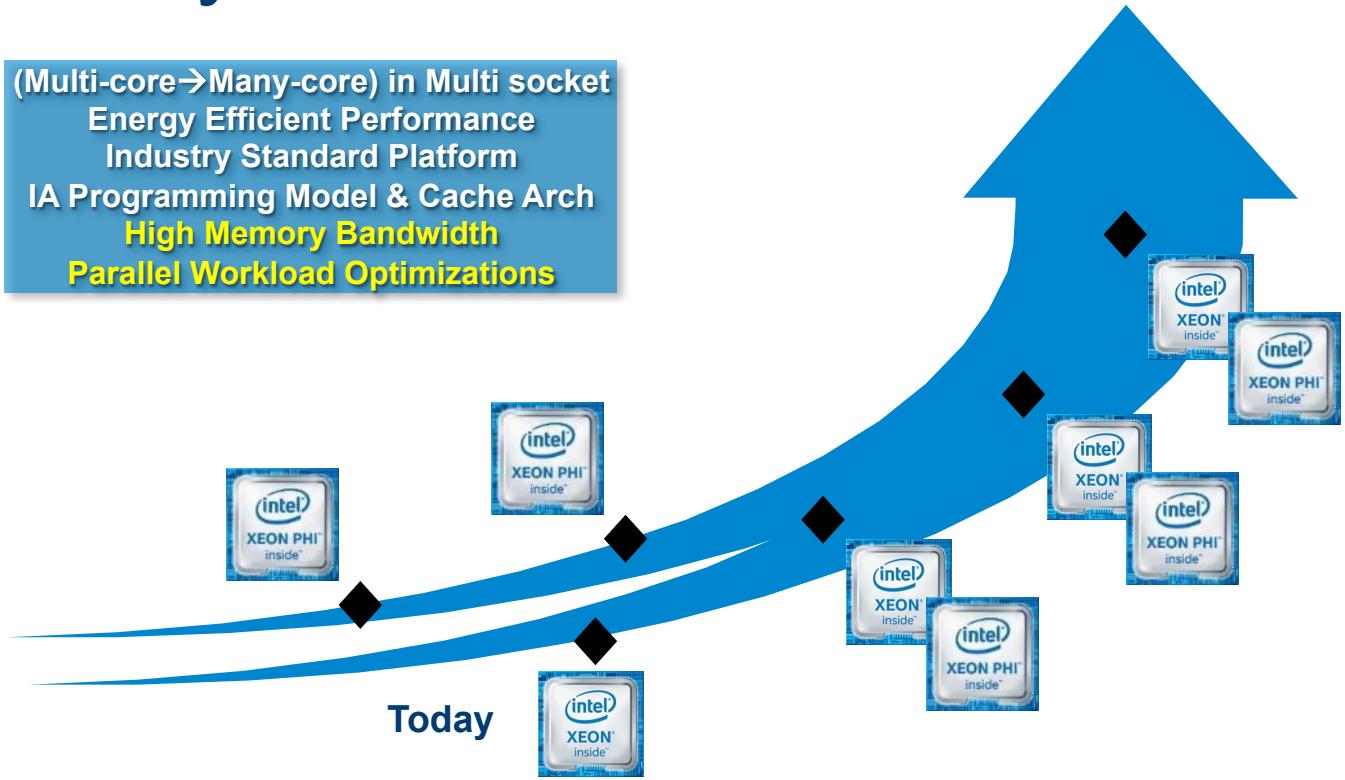

Shameless advertisements

KNL Developer Program: XeonPhiDeveloper.com

Highly-Parallel Performance

Software Tools & Libraries

Support & Training

Worldwide Partners*

COLFAX Customized Solutions

HPC SOLUTIONS

DAEHAN 대한컨설팅그룹 CONSULTING GROUP

The Right Answer in Data Center PARATERA 并行

MEGWARE

intel XEON PHI inside

Order your turnkey Intel® Xeon Phi™ Platform today!
Cutting edge platform capabilities, performance to deliver multi-threaded, vectorized software for today's HPC workloads!

PARALLEL STUDIO XE

*Names and brands may be claimed as the property of others

Intel® Xeon Phi: Remote Access Worldwide

- Try open, broad remote access to Intel® Xeon Phi™ processor based platforms for free (limited time) from our worldwide partners
- Cluster includes Intel Parallel Studio XE Cluster Edition, Intel Omni-Path Architecture and Intel SSDs
- Utilize key workloads/ benchmarks or bring your code
- Test scalability beyond single node
- Try before you buy!

Go to www.remotexeonphi.com – Portal with multiple partners. Free (limited time), broad and open access for developers and customers

Dell*: Click here to access

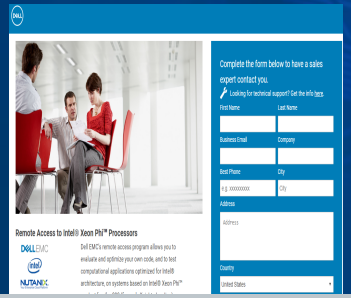

Penguin*: Click here to access

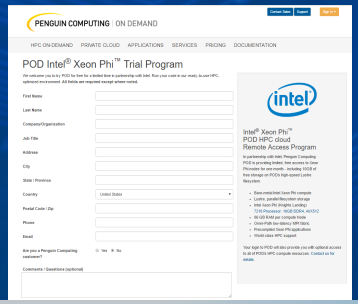

Fujitsu*: Click here to access

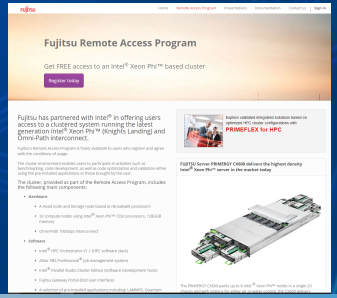

Colfax*: Click here to access

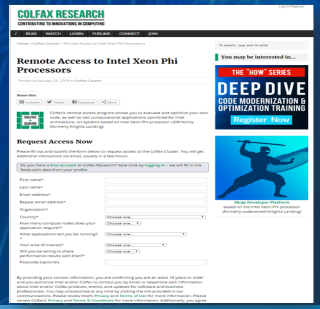

Rescale*: Click here to access

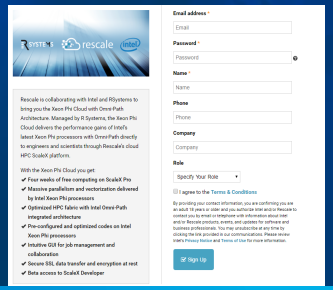

*Names and brands may be claimed as the property of others

INTEL® HPC DEVELOPER CONFERENCE

November 11-12, 2017 • Denver, CO

Key Topics:

- Parallel Programming
- High Productivity Languages
- Artificial Intelligence
- Systems
- Enterprise
- Visualization Development

All Technical Content:

- **77** Lectures (30 minutes)
- **15** Tutorials (90 minutes)
- **30** Poster Chats (60 minutes)

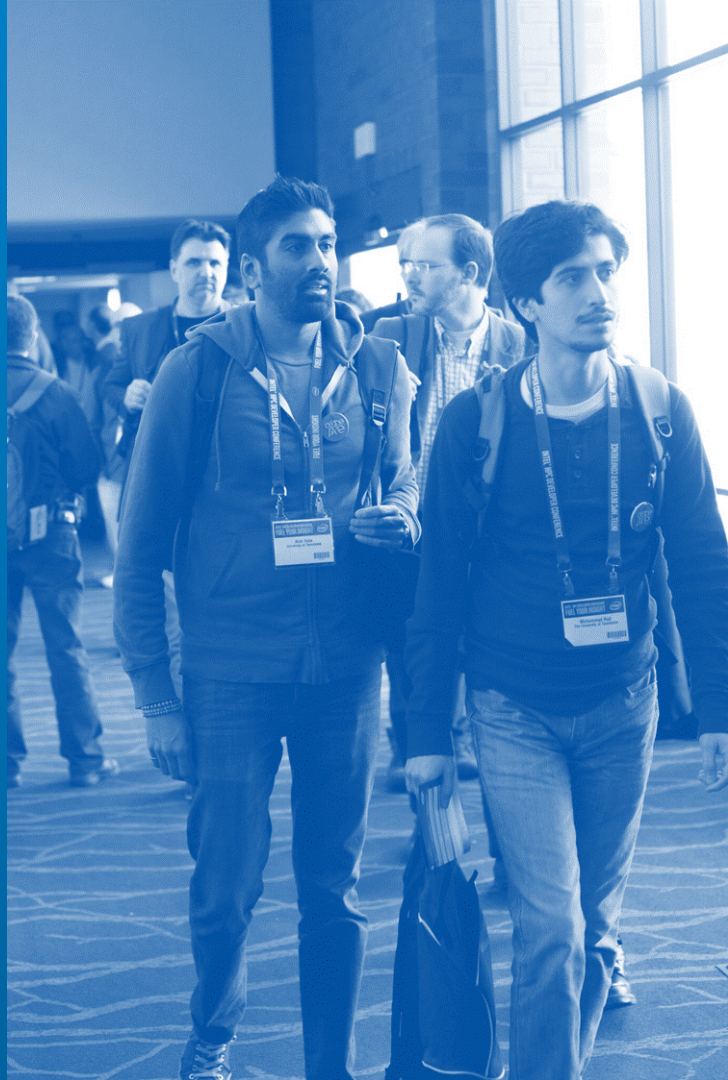

Register Now!

intel.com/hpcdevcon

In Summary ……

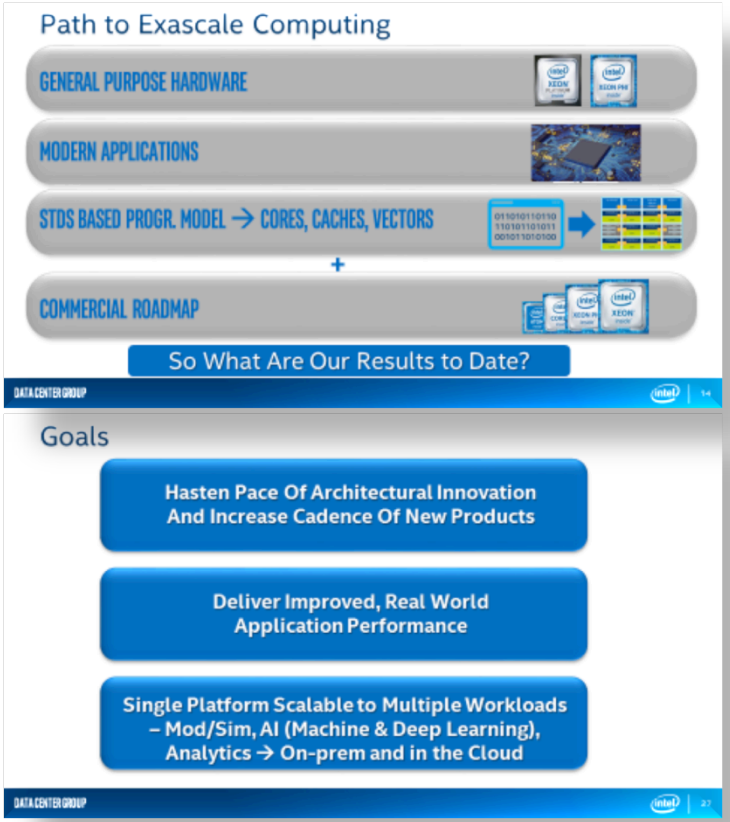

Delivering Exascale and Mainstream Performance in a Common Investment

Industry Leading Platform Features and Architectural Capabilities

Common SW Architecture and Programming Model → Always

Consistent Workload Tuning and Modernization Techniques → Always

Common platform for All Workloads → Mod/Sim, AI, Analytics

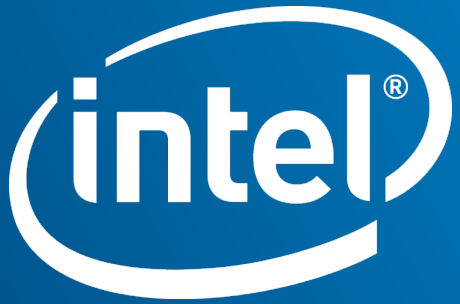
intel®

Improving Data Visualization with Intel® Solutions

Optimizing for Visualization

- Optimized for parallel processing and latest instruction sets
- OpenSWR*, Embree*, & OSPRay* all available
- Used by ParaView*, VisIt*, VMD*, CEI* EnSight*, and more…

Reducing Cost for Visualization

- Lower cost of host vs cost of host + card for GPUs
- Single host can address up to 384GB memory
- Additional use as general purpose compute platform

Pre-Configured Appliance Solution Available Now

- 1.58x to 1.91x better performance than GPUs[1]
- Supports data sets up to 1.5TB
- Standard configuration price $79,000[2]
- More Info: http://sdvis.xeonphi.com/

For in-situ, post-processing, and professional rendering visualization needs

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

[2]Pricing as of June 15, 2017. Pricing is subject to change without notice

[1]Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to www.intel.com/benchmarks. Configurations: See Slide 54.

Configuration details
64-node CPU system Intel® Xeon® 6148 Gold processor with 10GB Ethernet

Benchmark Segment	AI/ML
Benchmark type	Training
Benchmark Metric	Images/Sec or Time to train in seconds
Framework	Caffe
Topology	Resnet
# of Nodes	64
Platform	Wolfpass (Skylake)
Sockets	2S
Processor	Xeon Processor code named Skylake, B0, ES2*, 24c, 2.4GHz, 145W, 2666MT/s, QL1K CPUID=0x50652
BIOS	SE5C620.86B.01.00.0412.020920172159
Enabled Cores	24 cores / socket
Platform	Wolfpass (Skylake)
Slots	12
Total Memory	192GB
Memory Configuration	12x16GB DDR4 2R, 1.2V, RDIMM, 2666MT/s
Memory Comments	Micron MTA 18ASF2G72PDZ-2G6B1
SSD	800GB Model: ATA INTEL SSDSC2BA80 (scsi)
OS	Oracle Linux Server 7.3, Linux kernel 3.10.0-514.6.2.0.1.el7.x86_64.knl1

Other Configurations	Intel Corporation Ethernet Connection X722 for 10GBASE-T (rev 03)
HT	ON
Turbo	ON
Computer Type	Server
Framework Version	https://github.com/intel/caffe Release 1.0.3
Topology Version	ResNet-50 Facebook variant, refer to models/intel_optimized_models/multinode/ resnet_50_256_nodes_8k_batch
Dataset, version	Imagenet, ILSVRC 2012 (/lfs/lfs09/ lmdb_raw_compressed)
Performance command	caffe train -solver solver.prototxt -engine MKLDNN
Data setup	Convert to LMDB from raw ILSVRC 2012 images
Compiler	gcc 4.8.5 20150623
MPI Library version	Version 2017 Update 3 Build 20170405
MLSL Library version	f068d07deccf175a0bc8fe51f0eedce37a2a4470
MKL Library version	mklml_lnx_2018.0.20170720
MKL DNN Library Version	v0.9.0

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. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice. Roadmap not reflective of exact launch granularity and timing – please refer to ILU guidance. The products described in this document 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 processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families: Go to: Learn About Intel® Processor Numbers

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. Material in this presentation is intended as product positioning and not approved end user messaging.

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**.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at www.intel.com.

Intel, Intel Inside, the Intel logo, Xeon and Xeon Phi are 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.

Legal Disclaimer & Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS”. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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.

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