

Innovating Toward Exascale... ...and Beyond

CHEP2015, Okinawa, Japan
April 13, 2015

Dr. William Magro

Intel Fellow, Software & Services Group
Chief Technologist, Technical Computing Software



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.
- Copyright © 2014, Intel Corporation. All rights reserved. Intel, Xeon, Xeon Phi, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

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

HPC: Unwavering Progress... Amazing Impact



Top 500 FLOPS >50% CAGR For Past Decade



One of the Most Advanced Supercomputers Ever Built

An Intel-led collaboration with ANL and Cray to accelerate discovery & innovation



>180 PFLOPS

(option to increase up to 450 PF)


>50,000 nodes

13 MW

2018 *delivery*

18X higher
performance*

>6X more energy
efficient*



Prime Contractor



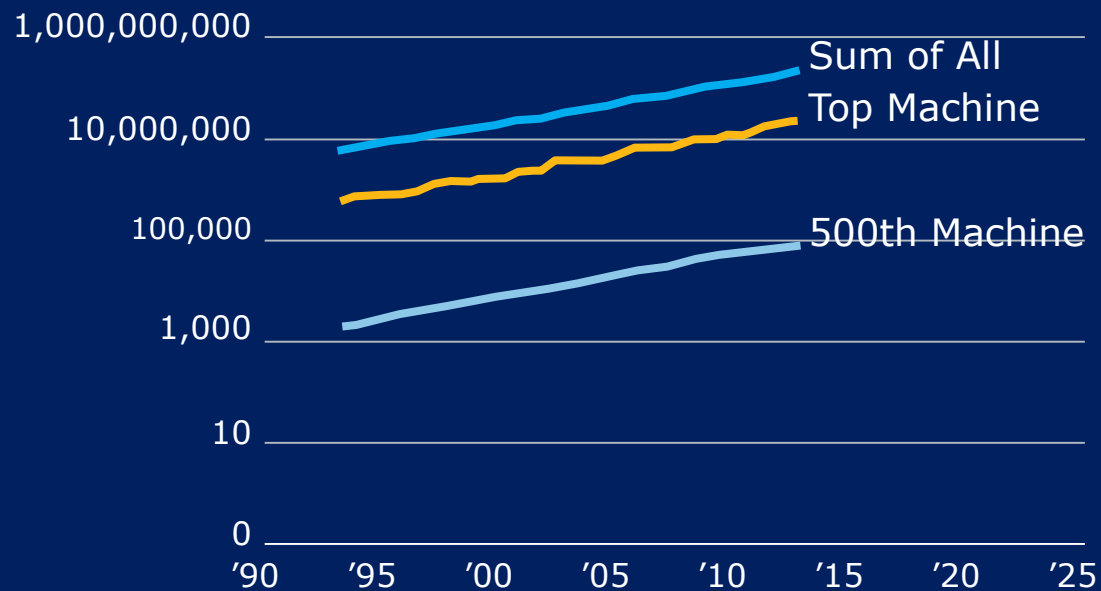
Subcontractor

Source: Argonne National Laboratory and Intel. *Versus ANL's current biggest system name MIRA (10PFs and 3.9MW).

Detailed comparison at www.intel.com/newsroom/assets/Intel_Aurora_factsheet.pdf

See "Legal Disclaimer & Optimization Notice" for important performance information. Other names and brands may be claimed as the property of others.

Supercomputing Pulls All of HPC Along



Top 500 FLOPS >50%¹ CAGR For Past Decade

**#500 system on Top 500 in
2004:**

~600 GFLOPs



10 years later...
**...one Intel® Xeon Phi™
co-processor:**

~1 TFLOPs

HPC is Evolving...Powering Discovery in New Ways

Current and future Intel innovations aimed at overcoming architectural challenges



Breaking Down “The Walls”

Memory | Fabric | Storage
Software Efficiency
Energy-efficient Performance
Space | Resiliency



Fast and Efficient Data Mobility

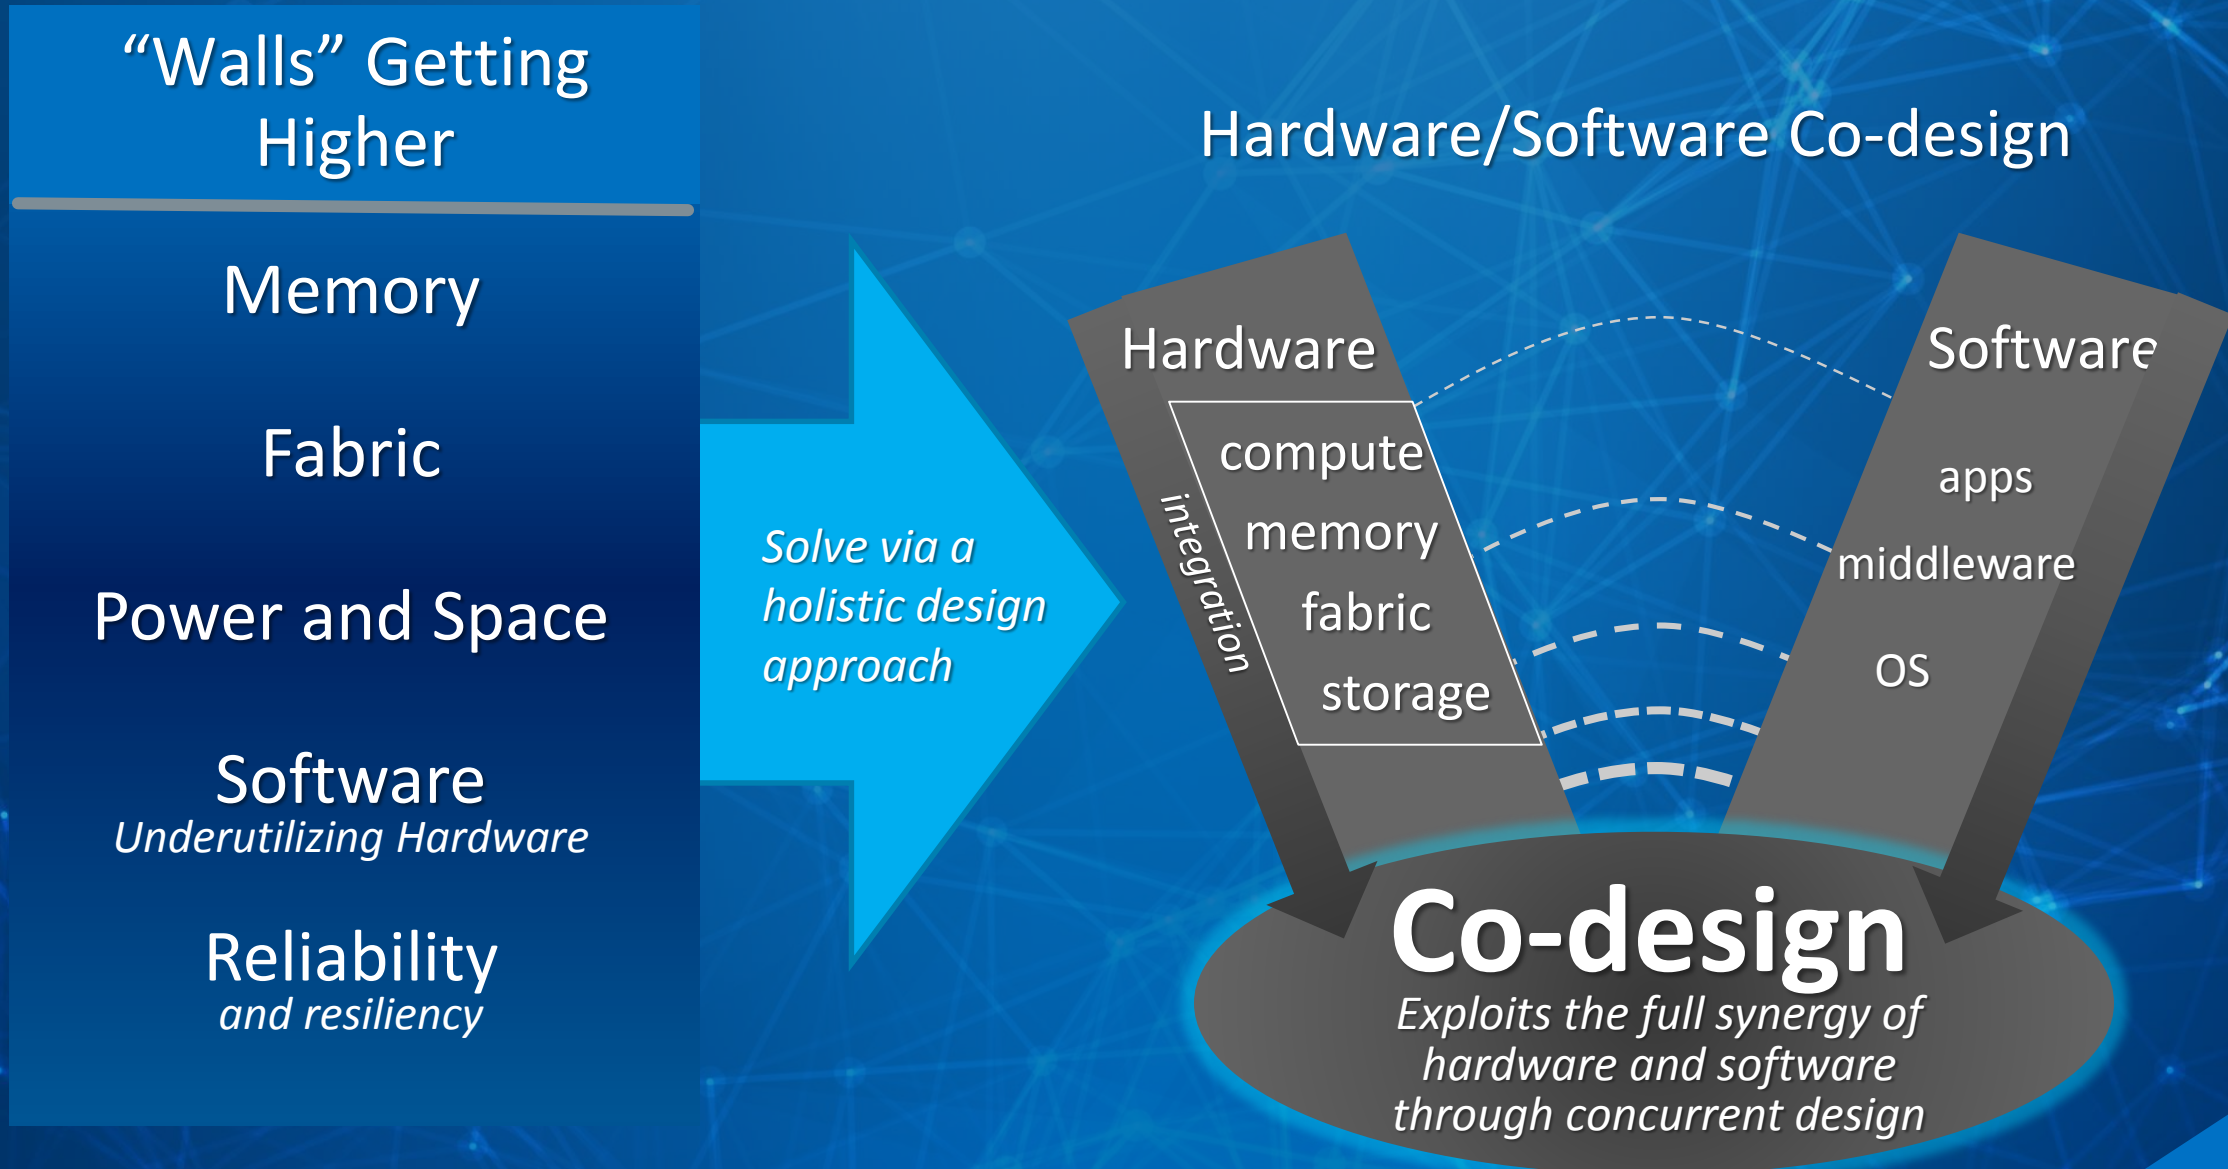
Persistent Memory Innovations
Big Data Analytics
Advanced Lustre* File System



Extending HPC's Reach


Enabling HPC at Every Scale

Breaking Down the Walls...Requires Co-design & Integration





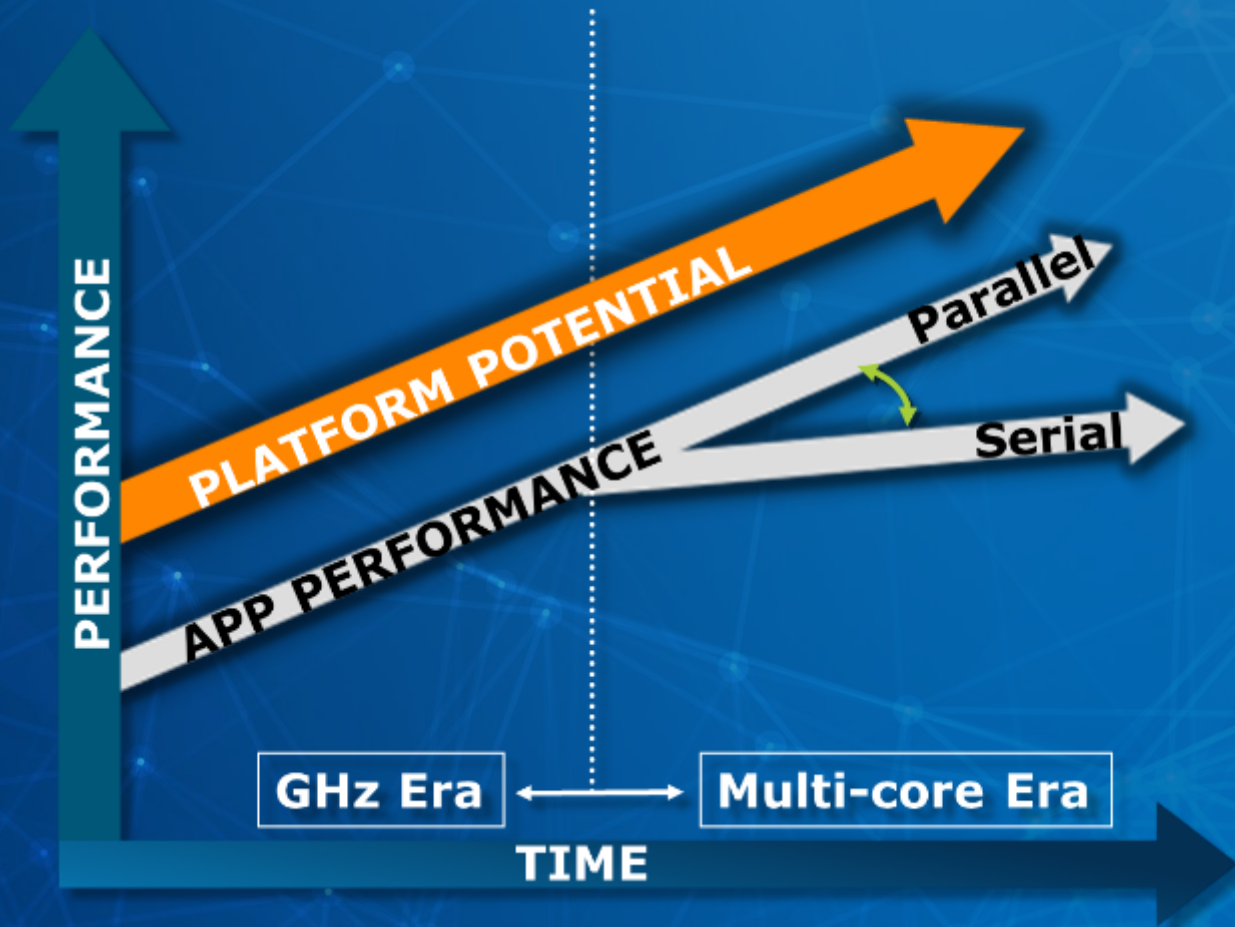
Moore's Law and Parallelism

35 YEARS OF MICROPROCESSOR TREND DATA




Original data collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond and C. Batten
Dotted line extrapolations by C. Moore

Code Modernization...Critical in the Multi-Core Era





Increasingly Parallel Code Enables Increased Processor Parallelism...and Performance



	Intel® Xeon® processor 64-bit series	Intel® Xeon® processor 5100 series	Intel® Xeon® processor 5500 series	Intel® Xeon® processor 5600 series	Intel® Xeon® processor code-named Sandy Bridge EP	Intel® Xeon® processor code-named Ivy Bridge EP	Intel® Xeon® processor code-named Haswell EX	Intel® Xeon Phi™ coprocessor Knights Corner	Intel® Xeon Phi™ processor & coprocessor Knights Landing¹
Core(s)	1	2	4	6	8	12	18	61	60+
Threads	2	2	8	12	16	24	36	244	244+
SIMD Width	128	128	128	128	256	256	256	512	512

**Product specification for launched and shipped products available on ark.intel.com.*

1. Not yet launched.

Intel® Xeon Phi™ Product Family

1 TFLOPS¹

Knights Corner



Intel® Xeon Phi™ Coprocessor
— Applications and Solutions
Catalog

3+ TFLOPS²


- Bootable Processor
- On-Pkg, High BW Memory
- Integrated Fabric

Knights Landing

2H'15
First
Commercial
Systems



+



>50 systems
providers
expected³

many more
card-based
systems

>100 PFLOPS customer system compute commits to-date³

Knights Hill

3rd Generation
Intel® Xeon Phi™
Product Family

2nd Generation
Intel Omni-Path
Architecture

10nm process
technology

¹ Claim based on calculated theoretical peak double precision performance capability for a single coprocessor. 16 DP FLOPS/clock/core * 61 cores * 1.23GHz = 1.208 TeraFLOPS
² Over 3 Teraflops of peak theoretical double-precision performance is preliminary and based on current expectations of cores, clock frequency and floating point operations per cycle. FLOPS = cores x clock frequency x floating-point operations per second per cycle. ³ Intel internal estimate

Knights Landing

Next Generation Intel® Xeon Phi™ Products

Platform Memory

Up to **384 GB** DDR4 (6 ch)

Compute

- Intel® Xeon® Processor Binary-Compatible
- **3+ TFLOPS¹, 3x ST²** (single-thread) perf. vs KNC
- **2D Mesh** Architecture
- **Out-of-Order** Cores

On-Package Memory

- Over **5x** STREAM vs. DDR4³
- Up to **16 GB** at launch

Over **60 Cores**

Integrated Intel® Omni-Path

Processor Package

Omni-Path

(optional)

- **1st** Intel processor to integrate

I/O

Up to **36 PCIe 3.0** lanes

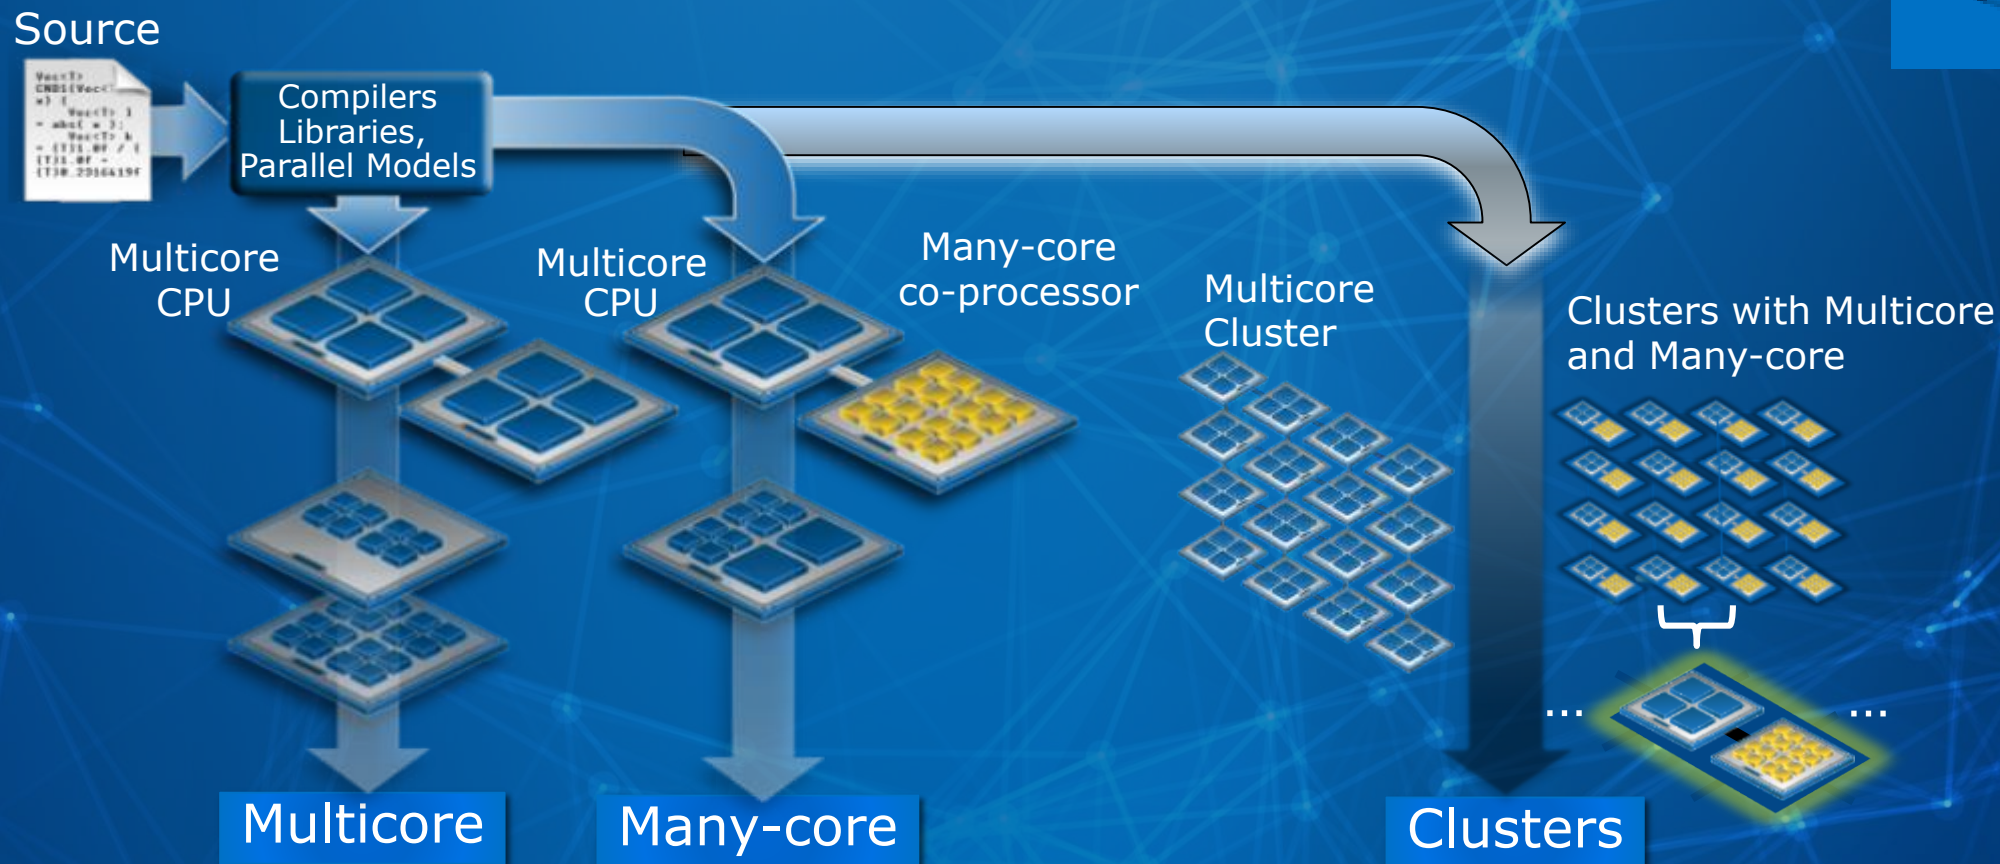
Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

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 complete information visit <http://www.intel.com/performance>.

Intel® Parallel Studio XE

Portable, Standards-Based Programming



Unlike accelerators, optimizations for Intel® Xeon Phi™ and Intel® Xeon® products share the same languages, directives, libraries, and tools.

Simplifying Vectorization

OpenMP 4: vectorization directive

```
#pragma omp simd
for(int ray=0; ray < N; ray++) {
    float Color = 0.0f, Opacity = 0.0f;
    int len = 0;
    int upper = raylen[ray];
    while (len < upper) {
        int voxel = ray + len;
        len++;
        if(visible[voxel] == 0) continue;
        float O = opacity[voxel];
        if(O == 0.0) continue;
        float Shading = O + 1.0;
        Color += Shading * (1.0f - Opacity);
        Opacity += O * (1.0f - Opacity);
        if(Opacity > THRESH) break;
    }
    color_out[ray] = Color;
}
```

Intel® Advisor XE: Vectorization Advisor

Intel Advisor XE 2016

Summary Survey Report Suitability Report Correctness Report Memory Access Patterns

Filter by Loop Type Vectorized Not Vectorized Filter by Source All Filter by Module All

Function Call Sites and Loops	Self Time	Total Time	Memory analysis	Compiler Vectorization	Gain Estimate	Vectorized Loops
▸ [loop at mmult_se...	10.040s	10.040s	<input type="checkbox"/>	Vectorized ...	2.19727	SSE2
▸ [loop at mmult_serial.cp...	0.000s	10.100s		Scalar		SSE2
▸ [loop at mmult_serial.cp...	0.000s	10.100s		Scalar		
▸ [loop in __libc_start_mai...	0.000s	10.100s		Scalar		

Top Down Source Assembly Assistance Recommendations

Function Call Sites and Loops	Total Time %	Total Time	Self Time	Hot Loops	Vector... Loops	Location
						Source Loc... Module
▼ Total	100.0%	10.100s	0s			
▼ __libc_start_main	100.0%	10.100s	0s			libc-2.12.so
▼ [loop in __libc_st...	100.0%	10.100s	0s			libc-2.12.so
▼ main	100.0%	10.100s	0s			mmult_seri... 1_mmult_serial
▼ [loop at mm...	100.0%	10.100s	0s	<input checked="" type="checkbox"/>		mmult_seri... 1_mmult_serial
▼ [loop at m...	100.0%	10.100s	0s	<input checked="" type="checkbox"/>	SSE2	mmult_seri... 1_mmult_serial
▼ multiply_d	100.0%	10.100s	0.0600s	<input checked="" type="checkbox"/>		mmult_seri... 1_mmult_serial
▸ [loop ...	99.4%	10.040s	10.0400s	<input checked="" type="checkbox"/>	SSE2	mmult_seri... 1_mmult_serial


Intel® Parallel Computing Centers

Over 50 centers worldwide

Eight working in high-energy physics



筑波大学
University of Tsukuba



YEARS/ANS CERN



Learn more @ <https://software.intel.com/en-us/ipcc>



Intel® Omni-Path: the Next-Generation Fabric

- Designed for Maximum Scalability
- Host and Fabric Optimized for HPC
- Flexible Configurations
- End-to-End Solution

INTEGRATION



Starting with
Knights Landing



Future 14nm
generation

★ Coming in '15

✓ PCIe
Adapters

✓ Edge
Switches

✓ Director
Systems


✓ Intel Silicon
Photonics

✓ Open
Software
Tools*

Intel® Omni-Path Architecture

Benefits for “Every” Scale


Better
System
Scaling



Better
Application
Scaling



56% Lower Latency²



Small Clusters

Higher Port Density

48 ports supports up to 12 add'l nodes by only adding CABLES¹

1.3x

Maximize **SINGLE SWITCH** investment

“Mainstream” Clusters

Fewer Switches²

up to **1/2**

Supercomputers

Higher Scaling

Over 27k **NODES** in a 2-tier 5-hop FABRIC³

2.3x

¹ A 48-port Fat-Tree full bisectional bandwidth (FBB) topology based on a 36-port switch chip requires four (4) additional Edge Switches and 54 additional cables.

² Latency reductions based on Mellanox CS7500 Director Switch and Mellanox SB7700/SB7790 Edge switches compared to preliminary Intel simulations for Intel® Omni-Path switches. Fewer switches claim based on a 1024-node full bisectional bandwidth (FBB) Fat-Tree configuration, using a 48-port switch for Intel® Omni-Path cluster and 36-port switch ASIC for either Mellanox or Intel® True Scale clusters. ³ Actual number is 27,628 nodes based on a cluster configured with the Intel® Omni-Path Architecture using 48-port switch ASICs, as compared with a 36-port switch chip that can support 11,664 nodes. See “Legal Disclaimer & Optimization Notice” for important performance information.

Intel® Omni-Path Architecture: Enhanced Switching Fabric



High messaging rates: Designed to support high MPI traffic from each node



Low latency: Extremely low port-to-port switch latency



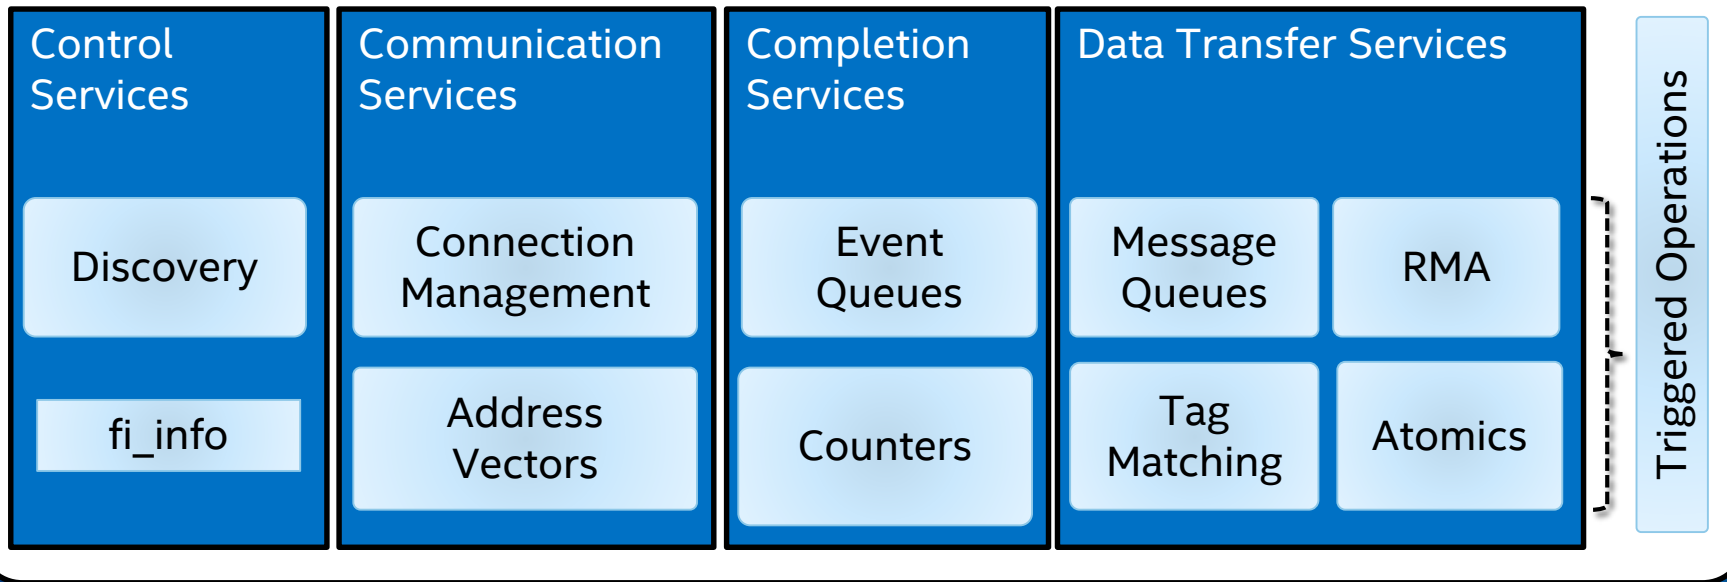
End-to-End Reliability: Built-in error detection & correction



Consistent end-to-end latency: Enables higher MPI application performance


OpenFabrics* Scalable Fabric Interface: A New Software Framework to Unleash Fast Fabrics

Open Fabrics Interfaces (OFI)



- Better semantic match for HPC: high performance
- Lower software complexity: high productivity

Big Data Meets High Performance: New Intel® Data Analytics Acceleration Library




- Highly-optimized building blocks
- Supports all data analysis stages
- Supports batch, streaming, and distributed processing
- Works with popular platforms (Hadoop, Spark) and tools (R, Python, Matlab)
- Flexible data interfaces (CSV, MySQL, HDFS, RDD (Spark))
- Handles sparse and noisy data
- C++ and Java APIs

Get the beta @ bit.ly/psxe2016beta


Innovation Across Key System Ingredients

CPU



Intel® Xeon® Processors
Intel® Xeon Phi™ Product
Family

Software & Tools



Intel® Parallel Studio
Intel® Enterprise Edition
for Lustre® software

Fabric



Intel® Omni-Path
Architecture

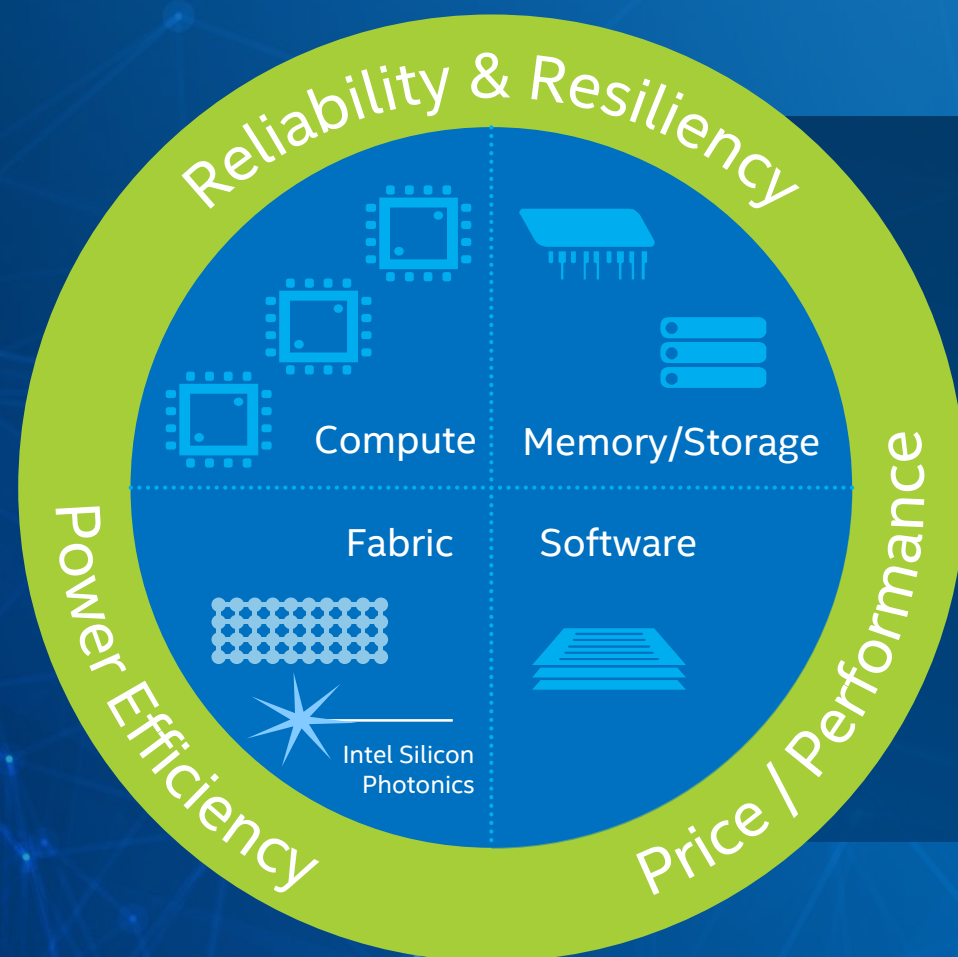
Storage



Intel® Solid-State
Drives (NVMe)

Intel's HPC Scalable System Framework

Hardware and software building blocks that integrate into powerful, compatible systems



Powering HPC at every scale

Compute- and data-centric computing

Standards-based programmability

Application compatibility

Intel® Xeon® Processors
Intel® Xeon Phi™ Coprocessors
Intel® Xeon Phi™ Processors

Intel® True Scale Fabric
Intel® Omni-Path Fabric
Intel® Ethernet

Intel® SSDs
Intel® Lustre*-based Solutions
Intel® Silicon Photonics Technology

Intel® Parallel Studio Developer Tools
Intel® Cluster Ready Program

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

Innovating Toward Exascale...and Beyond

Innovative technologies in a scalable system framework

A co-design approach that optimizes workload performance

Powerful software tools to unlock performance & productivity

A thriving, open, and enabled ecosystem

We're on this journey together

Aurora

