INTEL®: ARQUITECTURA PARA LA CONVERGENCIA DE HPC Y BIG DATA

Simón Viñals Larruga

Director de Tecnología, Intel Corp. Iberia

HPCadmintech, Madrid, Marzo 2016

Legal Disclaimers

Intel technologies features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com].

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.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual 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. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at https://www-ssl.intel.com/content/www/us/en/high-performance-computing/path-to-aurora.html.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/performance.

Intel, the Intel logo, Xeon, Intel Xeon Phi, Intel Optane and 3D XPoint are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States or other countries.

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

HPC IS FOUNDATIONAL TO INSIGHT

ONE FIFTH OF DATA CENTER SPENDING ON HW[2]

Advancing Science

And Our Understanding of the Universe

High ROI: $515

Average Return Per $1 of HPC Investment[1]

Data-Driven Analytics

Joins Theory, Experimentation, and Computational Science

[1]Source: IDC HPC and ROI Study Update (September 2015)
[2]Source: IDC 2015 Q1 World Wide x86 Sever Tracker vs IDC 2015 Q1 World Wide HPC Sever Tracker

Growing Challenges in HPC

"The Walls" System Bottlenecks

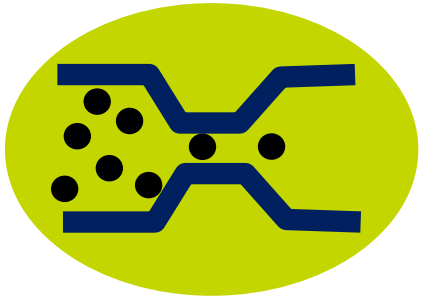

Memory | I/O | Storage
Energy Efficient Performance
Space | Resiliency |
Unoptimized Software

Divergent Infrastructure

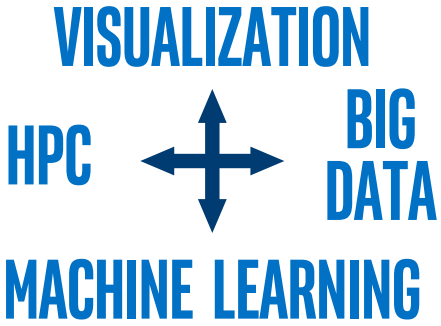

Resources Split Among
Modeling and Simulation | Big
Data Analytics | Machine
Learning | Visualization

Barriers to Extending Usage

Democratization at Every
Scale | Cloud Access |
Exploration of New Parallel
Programming Models

A Holistic Architectural Approach is Required

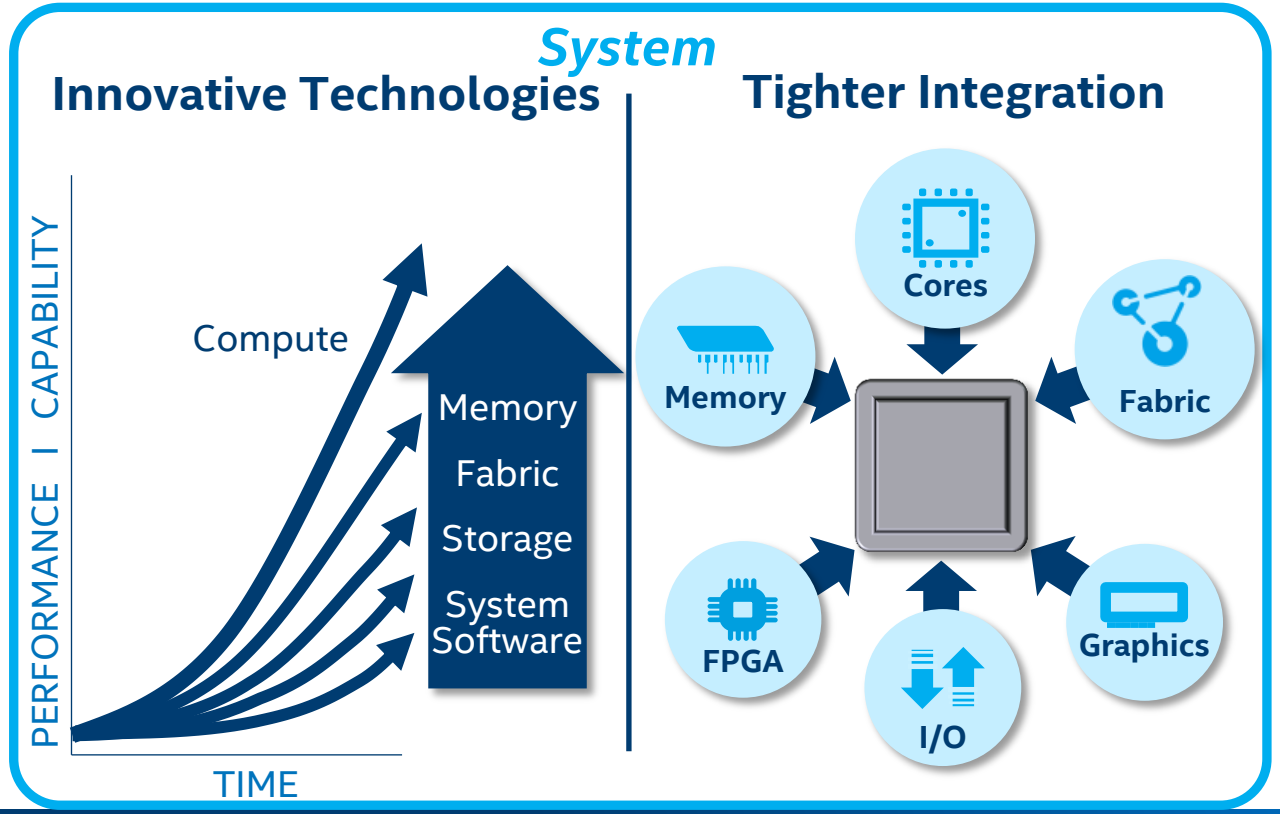

Intel® Scalable System Framework

A Holistic Design Solution for All HPC Needs

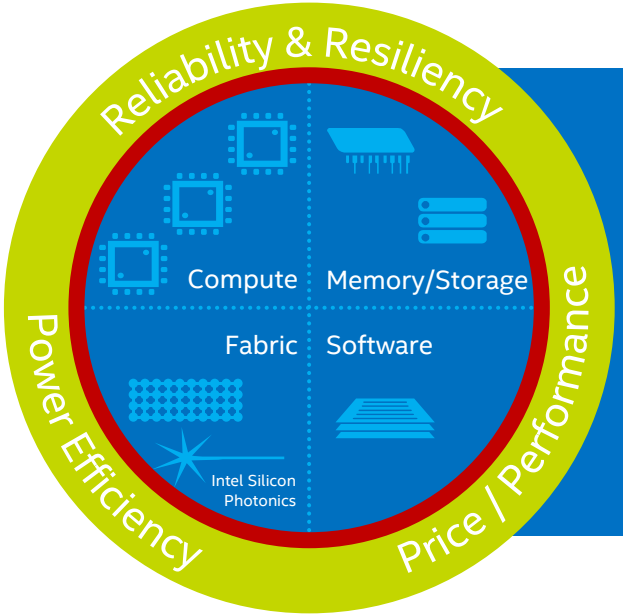

Small Clusters Through Supercomputers

Compute and Data-Centric Computing

Standards-Based Programmability

On-Premise and Cloud-Based

Intel® Xeon® Processors
Intel® Xeon Phi™ Processors
Intel® Xeon Phi™ Coprocessors
Intel® Server Boards and Platforms

Intel® Solutions for Lustre*
Intel® Optane™ Technology
3D XPoint™ Technology
Intel® SSDs

Intel® Omni-Path Architecture
Intel® True Scale Fabric
Intel® Ethernet
Intel® Silicon Photonics

HPC System Software Stack
Intel® Software Tools
Intel® Cluster Ready Program
Intel Supported SDVis

Accelerating Discovery and Innovation: BENEFITS

Intel® Scalable System Framework

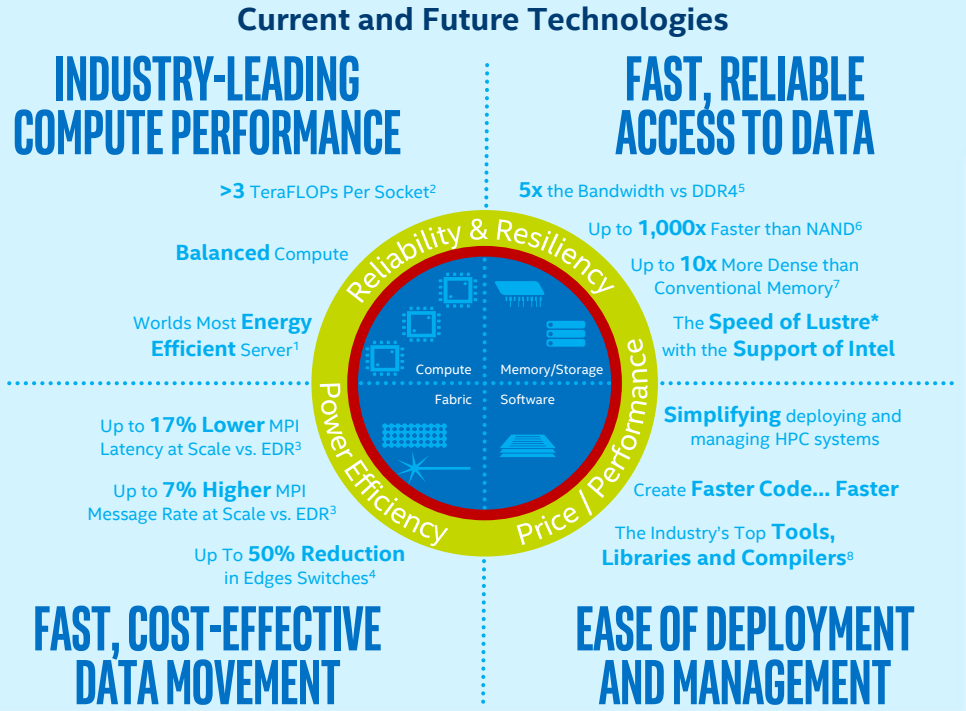

Breakthrough Performance

Standards-Based Programmability

Common Infrastructure Across Emerging Workloads

Broad Vendor Availability

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 Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate. 1Full details available at: http://www.intel.com/content/www/us/en/benchmarks/server/xeon-e5-2600-v3/xeon-e5-2600-v3-summary.html 2Over 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. 3Tests performed by Intel on Intel® Xeon® Processor E5-2697v3 dual-socket servers with 2133 MHz DDR4 memory. Intel® Turbo Boost Technology enabled and Intel® Hyper-Threading Technology disabled. Intel OPA: Open MPI 1.10.0 with PSM2. Pre-production Intel Corporation Device 24f0 – Series 100 HFI ASIC, Series 100 Edge Switch – 48 port. IOU Non-posted Prefetch disabled in BIOS. EDR: Open MPI 1.8-mellanox released with hpcx-v1.3.336-icc-MLNX_OFED_LINUX-3.0-1.0.1-redhat6.6-x86_64.tbz. Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch. 17% claim: HPCC 1.4.3 Random order ring latency. 16 nodes, 28 MPI ranks per node. 7% message rate claim: Ohio State Micro Benchmarks v. 4.4.1. osu_mbw_mr 28 MPI ranks per node, 8 byte message. osu_mbw_mr modified to use maximum rank pair communication time instead of average rank pair communication time. 4Reduction in up to 50% 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. 5Projected result based on internal Intel analysis of STREAM benchmark using a Knights Landing processor with 16GB of ultra high-bandwidth versus DDR4 memory with all channels populated. 6Performance difference based on comparison between 3D XPoint™ Technology and other industry NAND. 7Density difference based on comparison between 3D XPoint™ Technology and other industry DRAM 8Source: Evans North American Development Survey 2015 http://www.evansdata.com/

How It Works

Innovative Technologies

Compute

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

Memory/Storage

- High Bandwidth On-Package Memory
- Intel® Optane™ Technology
- Intel® Solutions for Lustre* software

Fabric

- Intel® Omni-Path Architecture
- Intel® Silicon Photonics
- Intel® True Scale Fabric

Software

- HPC System Software Stack
- Intel® Parallel Studios Software Suite
- Intel® Math Kernel Library
- Intel® Compilers

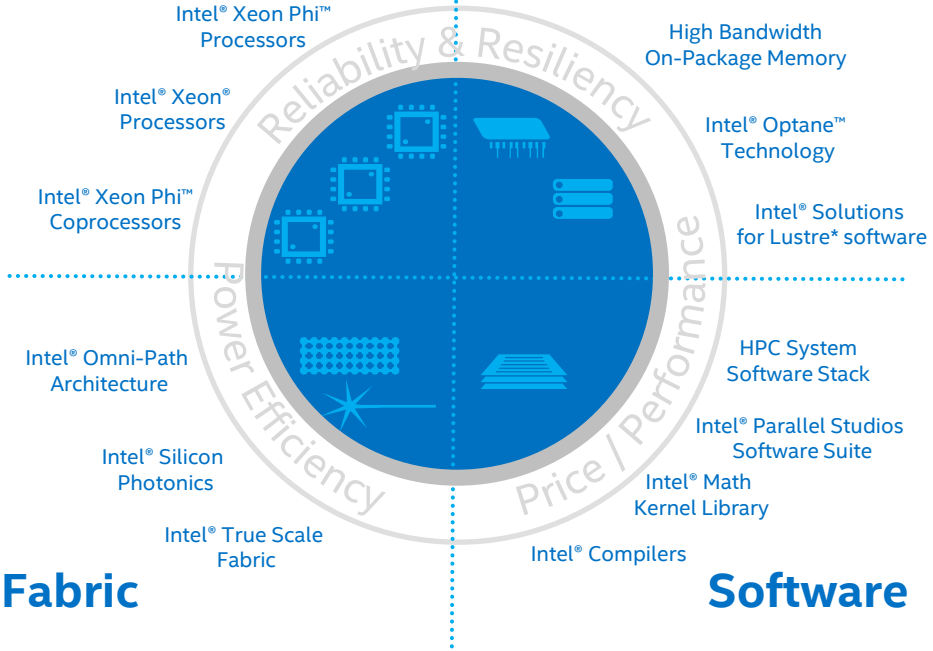

Tighter Integration
and Co-Design

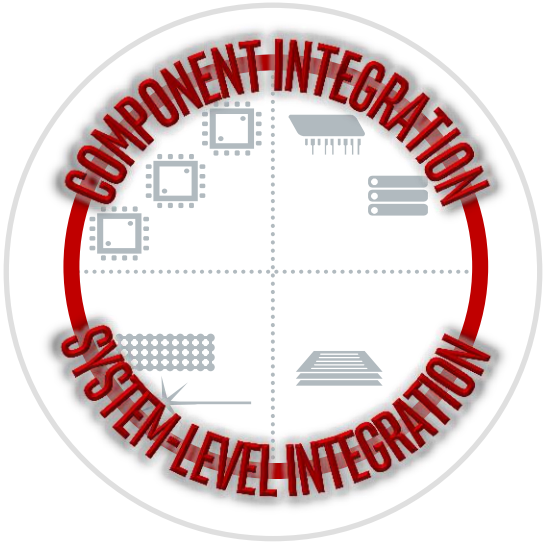

Increased System Density
Reduced System Power Consumption

Tighter Component Integration

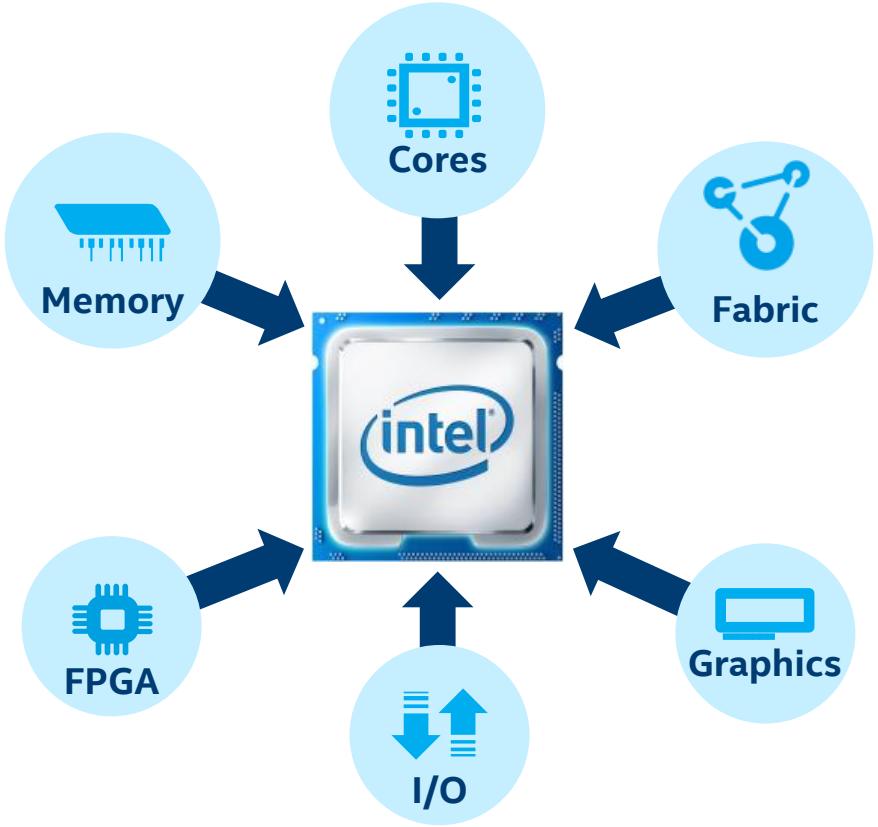

Intel® Scalable System Framework

Benefits

Bandwidth

Density

Latency

Power

Cost

Tighter System-Level Integration

Innovative Memory-Storage Hierarchy

Intel® Scalable System Framework

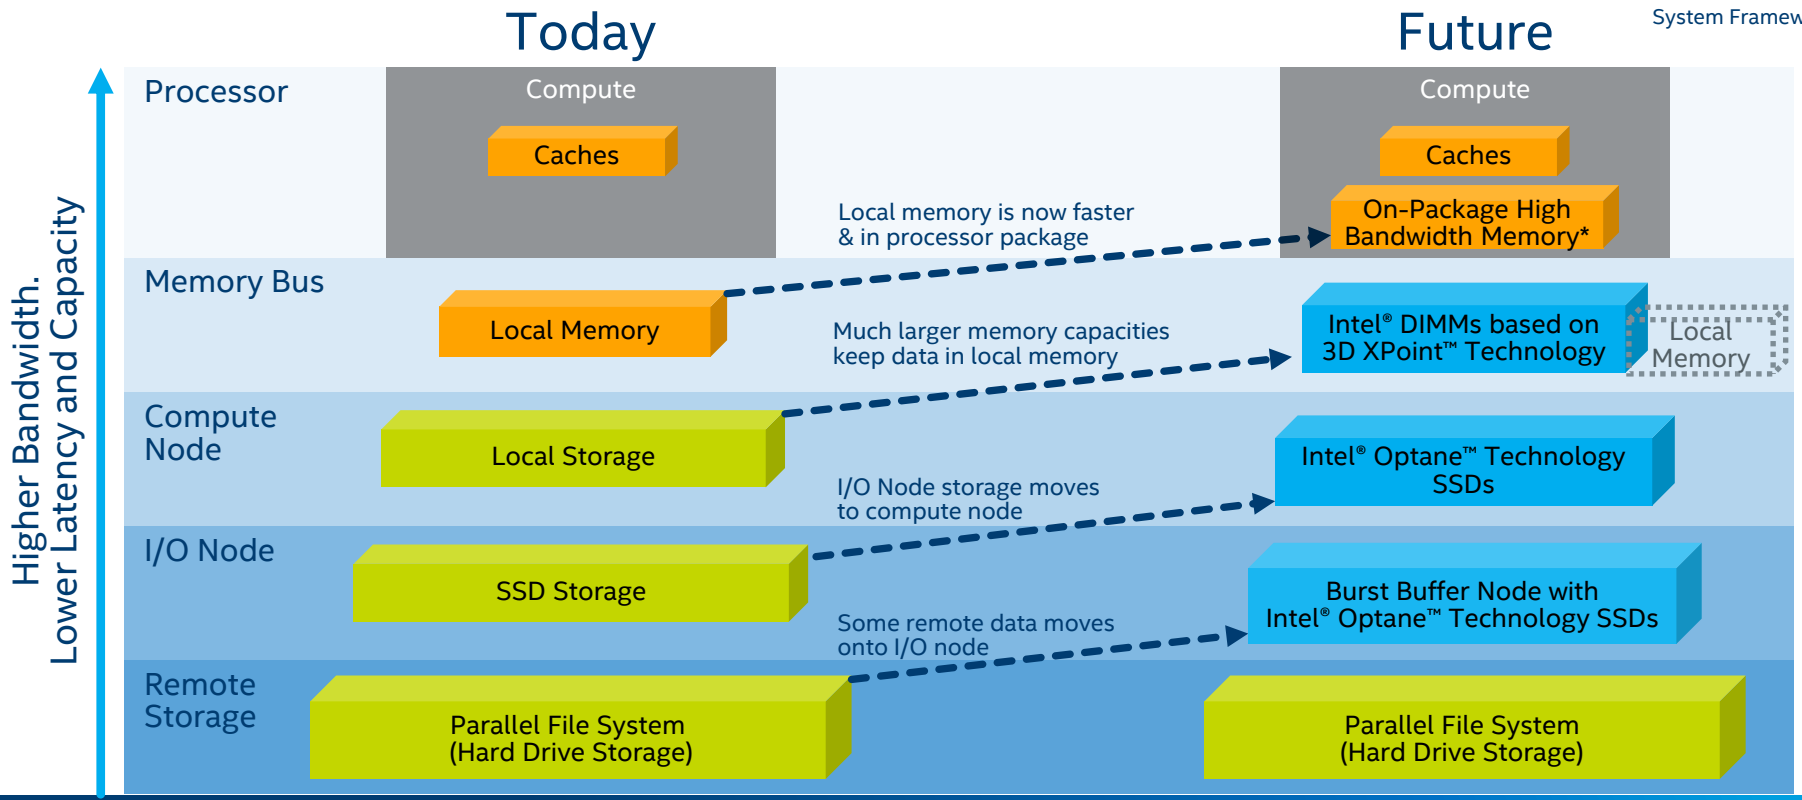

*cache, memory or hybrid mode

What Makes a Great HPC Solution?

Intel® Scalable System Framework

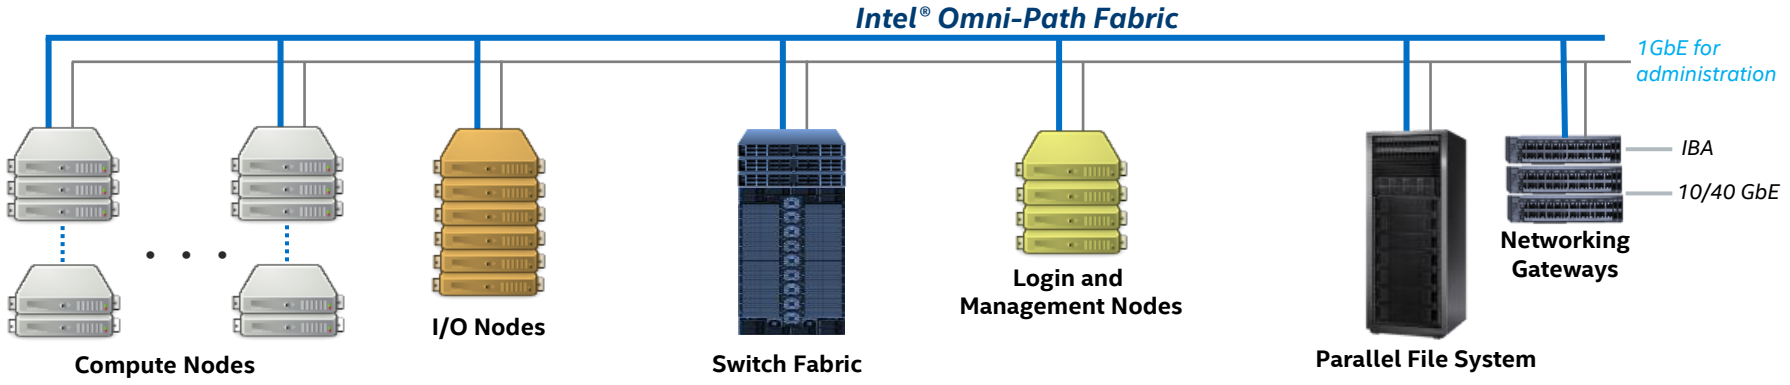

Intel® Compute
Intel® Xeon Phi™ Processors
Intel® Xeon® Processors
Intel® Optane™ Technology
Intel® Omni-Path Fabric

Burst Buffer
Intel® Xeon® Processors
Intel® Omni-Path Fabric
Intel® Optane™ Technology

Intel® Networking
Intel® Omni-Path Fabric
Intel® Silicon Photonics

Intel® Software Tools
Intel® Parallel Studio
Intel® Node Manager
Intel® Trace Analyzer

Intel® Solutions for Lustre*
Intel® Enterprise Edition for Lustre*
Intel® Foundation Edition for Lustre*
Intel® Cloud Edition for Lustre*

Reference Architecture
Intel® Cluster Ready

Actual configurations depend on specific OEM offerings and implementation.

Different Systems (Today)

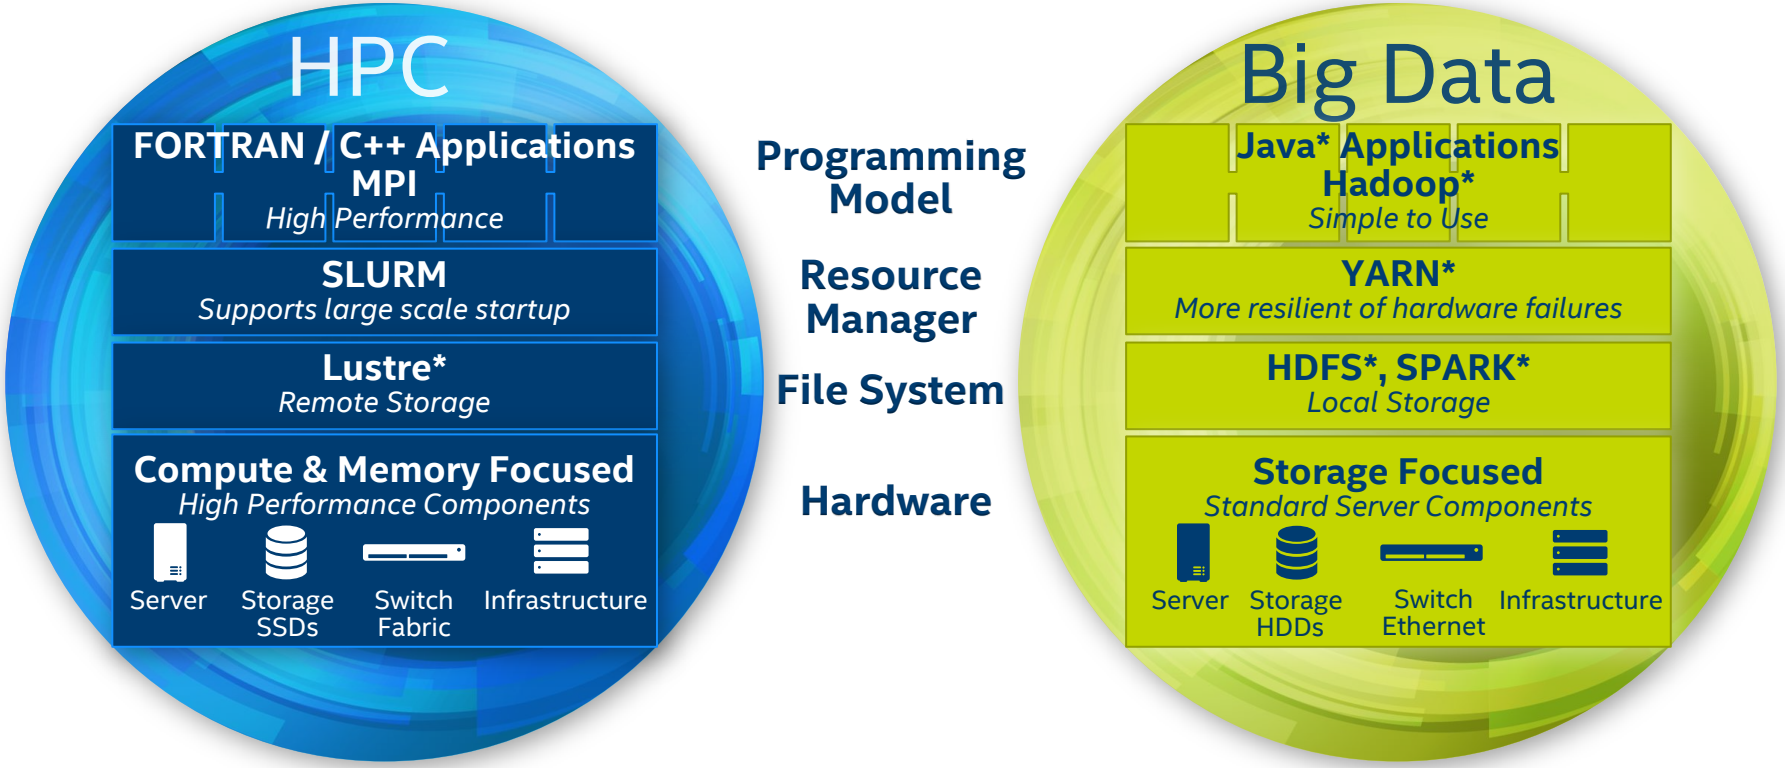

Emerging Real-Time Workflows

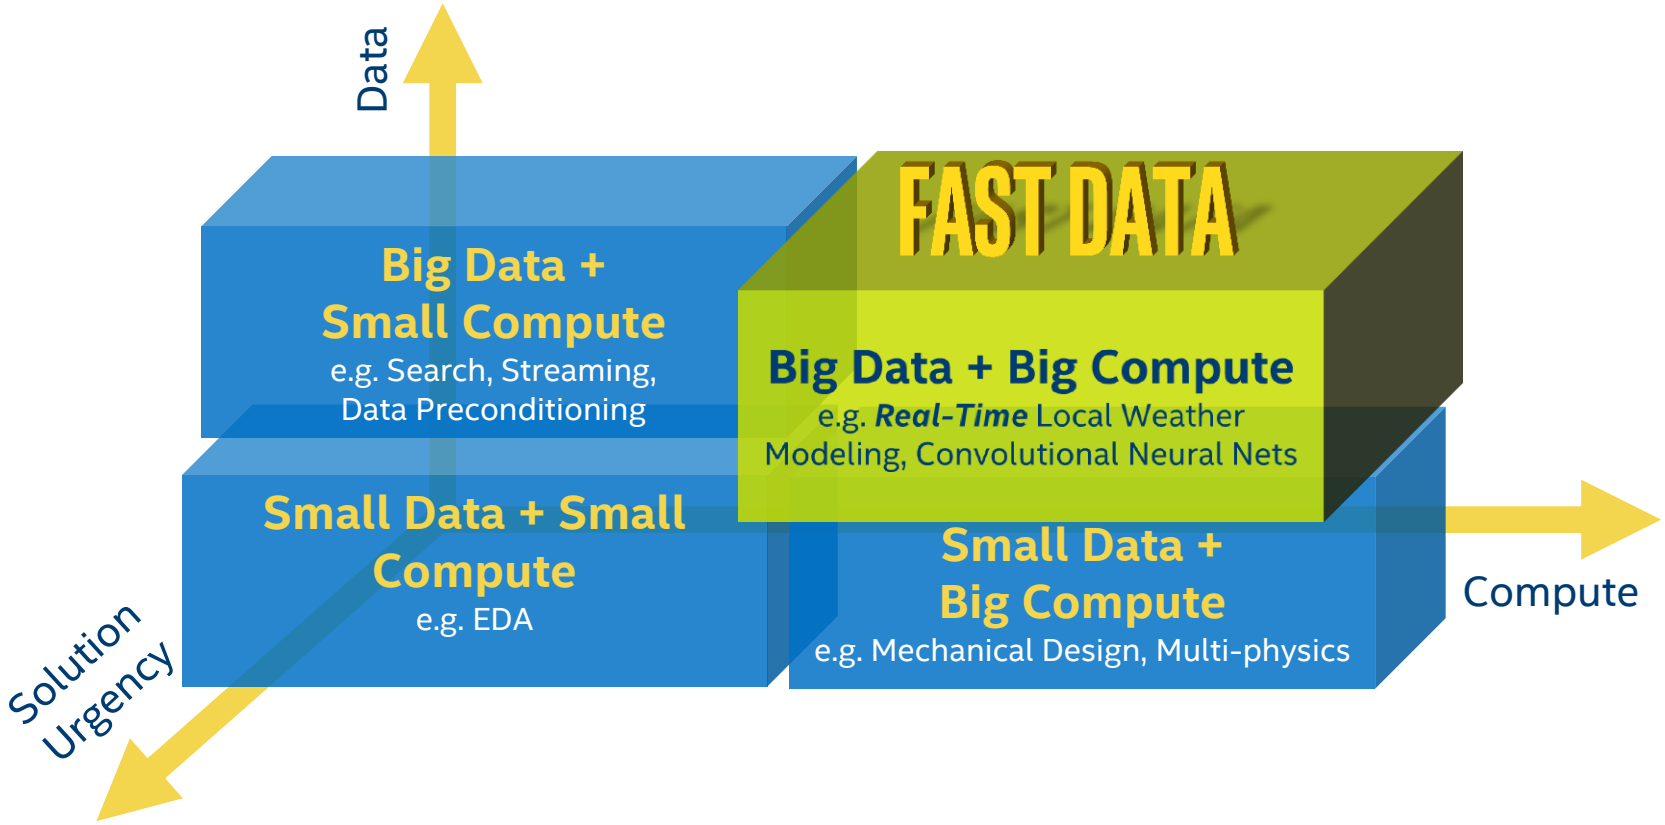

One System Architecture?

From Here …

Separate Systems

…TO?

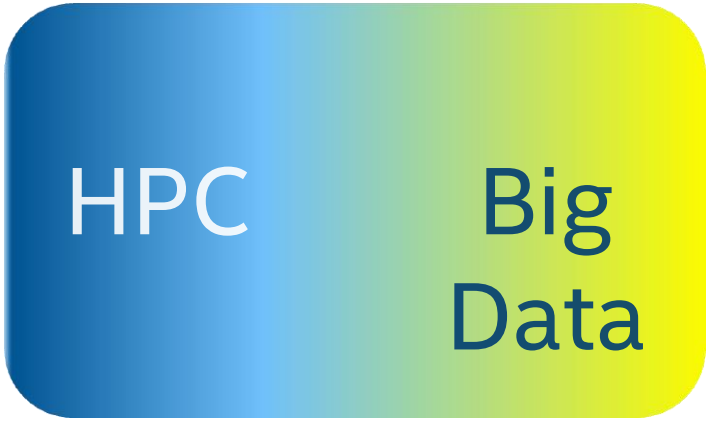

Single System Architecture

The Challenges of Moving to Single System Arch.

Many Codes on a System

Over 600

Different Applications on NERSC's Hopper System

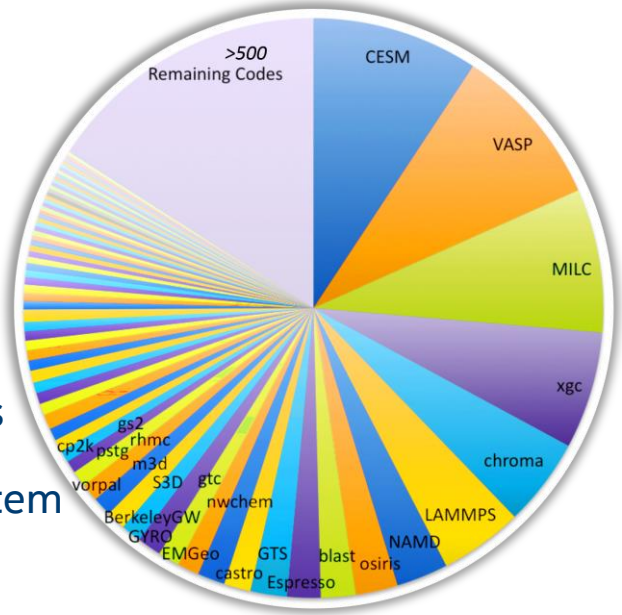

Source: NERSC Workload Analysis on Hopper. K. Antypas et al. 2014

The Challenges of Moving to Single System Arch.

Many Codes on a System

Over 600

Different Applications on NERSC's Hopper System

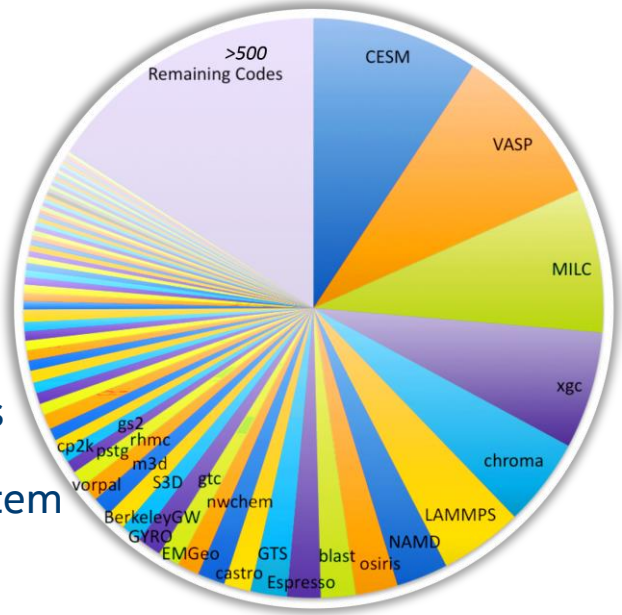

Varied Resource Needs

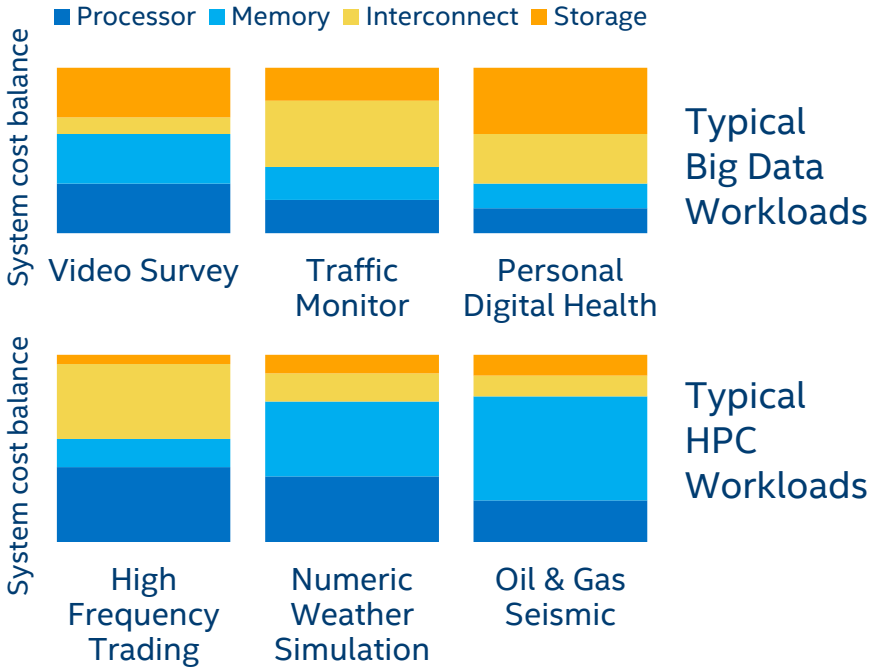

Source: NERSC Workload Analysis on Hopper. K. Antypas et al. 2014

Source: Intel for illustration purposes

Converged Architecture for HPC and Big Data

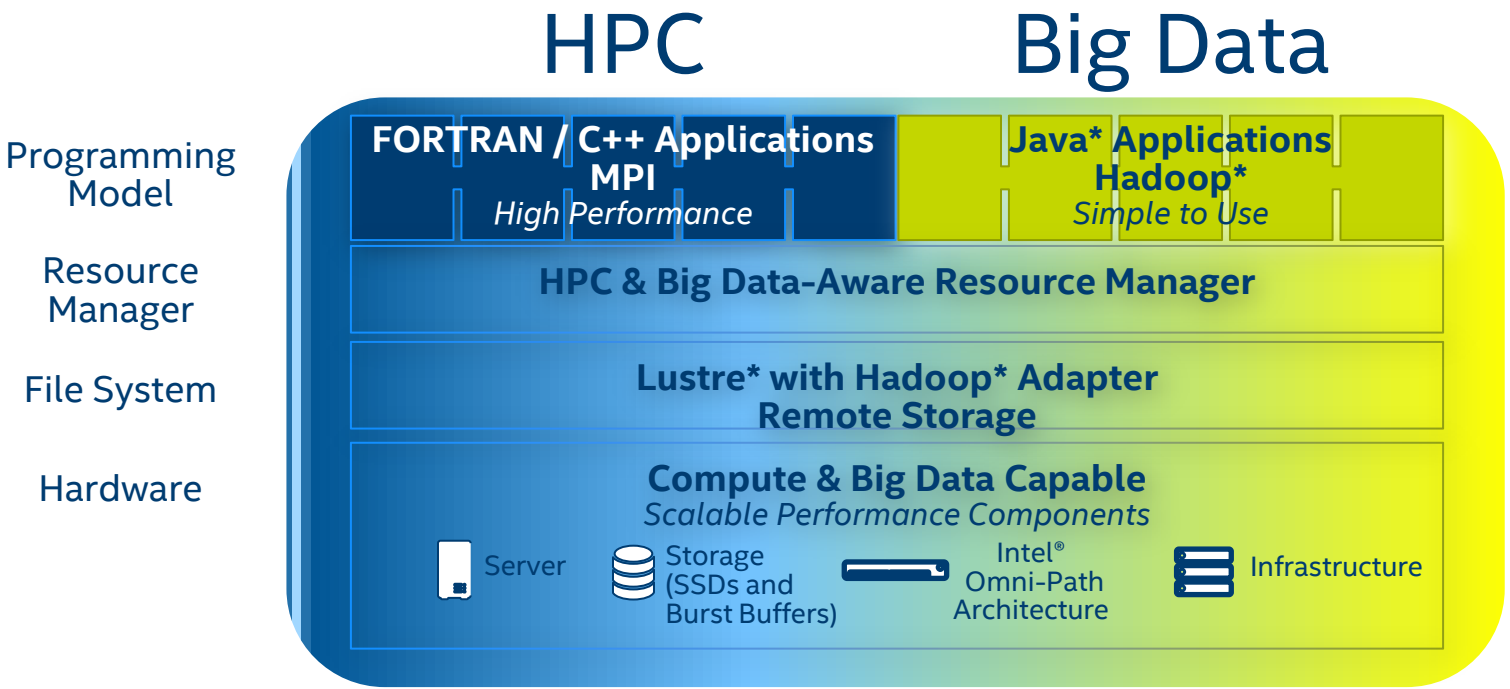

INNOVATIVE TECHNOLOGIES FOR HPC: COMPUTE

High Performance Compute

Intel® Xeon Phi™ Processors

Many-Core

Optimized for **Highly Parallel** and **Highly Vectorized** Apps

Intel® Xeon® Processors

Multi-Core

Optimized for **Serial and Parallel** Applications

Common Programming Model

Intel® Xeon® Processors

At the Heart of Intel® Scalable System Framework

Intel® Scalable System Framework

THE HEART OF THE DATA CENTER

Serial and Parallel Performance

- Up to **18 cores** – Frequency and Parallel Optimized SKUs
- Excellent Single and Multi-Thread Performance
- Up to **1,536GB** of DDR4 Memory

Technical Computing Buyers

- Choose Intel® Xeon® Processors **~95%** of the Time[1]
- Buy High Performance SKUs
- Refresh / Upgrade Regularly

1. Source: IDC HPC WW QView Q1 2015

Intel® Xeon® Processors

At the Heart of Intel® Scalable System Framework

Intel® Scalable System Framework

18
CORES
Serial and Parallel Apps

1,536
GB DDR4
For Memory Dependent Apps

95%
CHOOSE INTEL
For Technical Computing Apps[1]

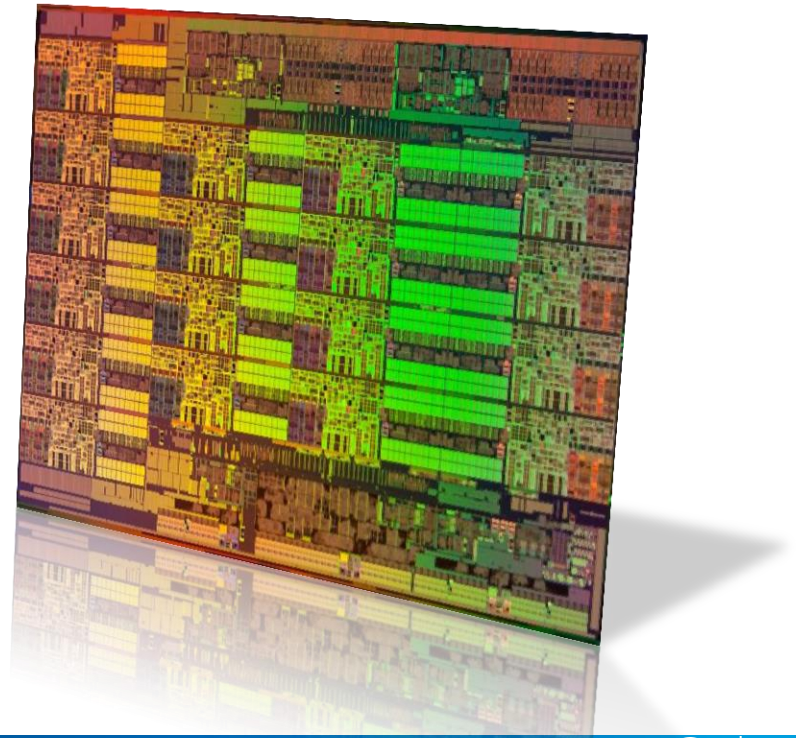

1. Source: IDC HPC WW QView Q1 2015

Introducing Intel® Xeon® Processor E5-2600 v4 Product Family (codenamed “Broadwell-EP”)

31-March

NEW microarchitecture using 14nm process technology + 20% MORE cores + 20% MORE last level cache expected to deliver 18% average performance INCREASE[1].

Support for up to 2400 MT/s with DDR4 memory for greater I/O throughput.

New and increased Resource Monitoring and Allocation capabilities, providing an optimum data center Orchestration and Virtualization experience.

Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. 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. Configurations: Intel Performance Projections as of August 2015. For more information go to http://www.intel.com/performance/datacenter.

A High Performance Compute Foundation

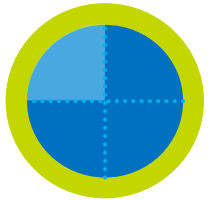

Intel's HPC Scalable System Framework

EXTREME PARALLEL PERFORMANCE

Parallel performance

- **72** cores; 2 VPU/core; 6 DDR4 channels with 384GB capacity
- **>3** Teraflop/s per socket[1]

Integrated memory

- **16GB; 5X** bandwidth vs DDR4[2]
- **3 configurable** modes (memory, cache, hybrid)

Integrated fabric

- **2** Intel Omni-Path Fabric Ports (more configuration options)

Market adoption

- **>50** systems providers expected[3]
- **>100** PFLOPS customer system compute commits to-date[3]
- Software development kits shipping Q4'15,
  Hardware shipping Q3'16

Learn More:

https://software.intel.com/en-us/articles/what-disclosures-has-intel-made-about-knights-landing?wapkw=knights+landing

[1] Source: Intel internal information. [2]Projected result based on internal Intel analysis of STREAM benchmark using a Knights Landing processor with 16GB of ultra high-bandwidth versus DDR4 memory only with all channels populated. 2Over 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.

Intel® Xeon Phi™ Processors

The Faster Path to Discovery

Intel® Scalable System Framework

HIGHLY PARALLEL PERFORMANCE

Programmability

- Binary-Compatible with Intel® Xeon® Processors

Parallel Performance

- Up to 72 Cores; 2 VPU/Core
- >3 Teraflop/s Per Socket[1]

Integrated Memory and Fabric

- Up to 16GB On-Package; 5X Bandwidth vs DDR (over 400GB/s)[2]
- 2 Intel Omni-Path Fabric Ports (More Configuration Options)
- 6 DDR4 Channels with up to 384GB Memory Capacity

Market Adoption

- Over 50 Systems Providers Expected[3]
- >100 PFLOPS Customer System Compute Commits To-Date[3]

1. 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. 2. Projected result based on internal Intel analysis of STREAM benchmark using a Knights Landing processor with 16GB of ultra high-bandwidth versus DDR4 memory only with all channels populated. 3. Source: Intel internal information

Intel® Xeon Phi™ Processors

The Faster Path to Discovery

Intel® Scalable System Framework

>3TF
PEAK DP
Performance[1]

3X
FASTER
Single Thread Perf. VS. KNC[2]

5X
BANDWIDTH
MCDRAM VS. DDR4[3]

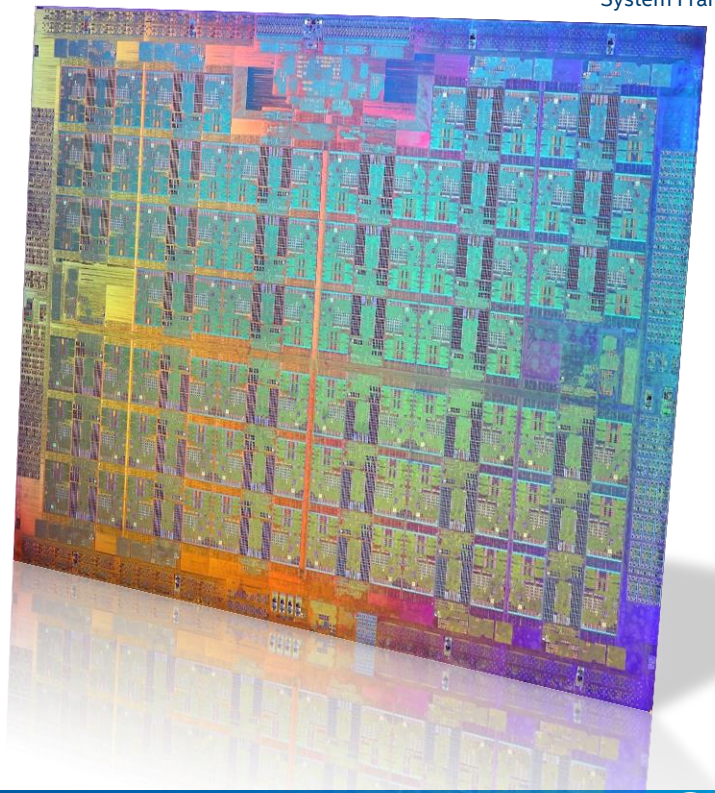

1. Over 3 Teraflops of peak theoretical double-precision performance is preliminary and based on current expectations of cores, clock frequency and floating point operations per cycle. 2. Projected peak theoretical single-thread performance relative to 1st Generation Intel® Xeon Phi™ Coprocessor 7120P 3. Projected result based on internal Intel analysis of STREAM benchmark using a Knights Landing processor with 16GB of ultra high-bandwidth versus DDR4 memory with all channels populated.

INNOVATIVE TECHNOLOGIES FOR HPC: NETWORKING

Intel® High Performance Interconnects

Intel® Omni-Path Fabric

HPC's Next Generation Fabric

Better System Scaling

***48 port** Switch chip arch.*

Excellent ***Application*** Scaling

***100 Gbps** Line speed*

Intel® True Scale Fabric

All the Performance
at a
Fraction of the Investment

Host fabric adapters
Edge and Director Switches
Management Software

Intel® Ethernet

1 billion Intel® Ethernet Ports Shipped and More than **30 Years** of Innovation

1/10/40/100GbE Solutions
I/O Virtualization
Multi-Host Controllers

100+ Switches & Server Platforms Available at Launch

A High Performance Fabric

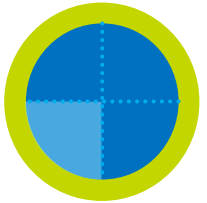

Intel's HPC Scalable System Framework

Better scaling vs EDR

- **48** radix chip ports
- **73%** higher switch MPI message rate[2]
- **60%** lower switch fabric latency[3]

Configurable / Resilient

- **Job prioritization** (Traffic Flow Optimization)
- **No-compromise resiliency** (Packet Integrity Protection and Dynamic Lane Scaling)

Market adoption

- **>100** OEM designs[1]
- >**100ku** nodes in opportunity pipeline[1]

Learn More:

Intel OPA WEBINAR

[1] Source: Intel internal information. Design win count based on OEM and HPC storage vendors who are planning to offer either Intel-branded or custom switch products, along with the total number of OEM platforms that are currently planned to support custom and/or standard Intel® OPA adapters. Design win count as of July 1, 2015 and subject to change without notice based on vendor product plans. [2] Based on Prairie River switch silicon maximum MPI messaging rate (48-port chip), compared to Mellanox CS7500 Director Switch and Mellanox SB7700/SB7790 Edge switch product briefs (36-port chip) posted on www.mellanox.com as of July 1, 2015. [3] Latency reductions based on Mellanox CS7500 Director Switch and Mellanox SB7700/SB7790 Edge switch product briefs posted on www.Mellanox.com as of July 1, 2015, compared to Intel® OPA switch port-to-port latency of 100-110ns that was measured data that was calculated from difference between back to back osu_latency test and osu_latency test through one switch hop. 10ns variation due to "near" and "far" ports on an Eldorado Forest switch. All tests performed using Intel ® Xeon® E5-2697v3, Turbo Mode enabled. Up to 60% latency reduction is based on a 1024-node cluster in a full bisectional bandwidth (FBB) Fat-Tree configuration (3-tier, 5 total switch hops), using a 48-port switch for Intel Omni-Path cluster and 36-port switch ASIC for either Mellanox or Intel® True Scale clusters. Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance

Intel® Omni-Path Architecture

HPC's Next-Generation Fabric

Intel® Scalable System Framework

26% MORE Servers than EDR[1]

60% LOWER Cooling Costs[2]

2.3X GREATER Fabric Scalability[3]

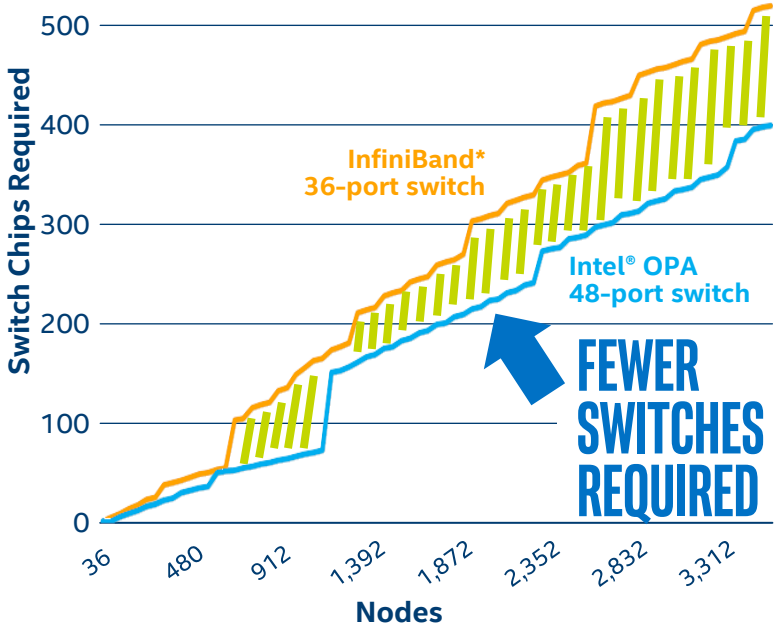

1. Assumes a 750-node cluster, and number of switch chips required is based on a full bisectional bandwidth (FBB) Fat-Tree configuration. Intel® OPA uses one fully-populated 768-port director switch, and Mellanox EDR solution uses a combination of 648-port director switches and 36-port edge switches. Mellanox component pricing from www.kernelsoftware.com, with prices as of November 3, 2015. Compute node pricing based on Dell PowerEdge R730 server from www.dell.com, with prices as of May 26, 2015. Intel® OPA pricing based on estimated reseller pricing based on Intel MSRP pricing on ark.intel.com. 2. Assumes a 750-node cluster, and number of switch chips required is based on a full bisectional bandwidth (FBB) Fat-Tree configuration. Intel® OPA uses one fully-populated 768-port director switch, and Mellanox EDR solution uses a combination of director switches and edge switches. Mellanox power data based on Mellanox CS7500 Director Switch, Mellanox SB7700/SB7790 Edge switch, and Mellanox ConnectX-4 VPI adapter card installation documentation posted on www.mellanox.com as of November 1, 2015. Intel OPA power data based on product briefs posted on www.intel.com as of November 16, 2015. Intel® OPA pricing based on estimated reseller pricing based on Intel MSRP pricing on ark.intel.com. 3 Number of switch chips required, switch density, and fabric scalability are based on a full bisectional bandwidth (FBB) Fat-Tree configuration, using a 48-port switch for Intel® Omni-Path Architecture and 36-port switch ASIC for either Mellanox or Intel® True Scale Fabric. *Other names and brands may be claimed as the property of others. 2.3X fabric scalability based on a 27,648-node cluster configured with the Intel® Omni-Path Architecture using 48-port switch ASICs, as compared with a 36-port switch chip that can support up to 11,664 nodes.

Intel® Omni-Path Architecture

HPC's Next-Generation Fabric

Intel® Scalable System Framework

26% MORE Servers than EDR[1]

60% LOWER Cooling Costs[2]

2.3X GREATER Fabric Scalability[3]

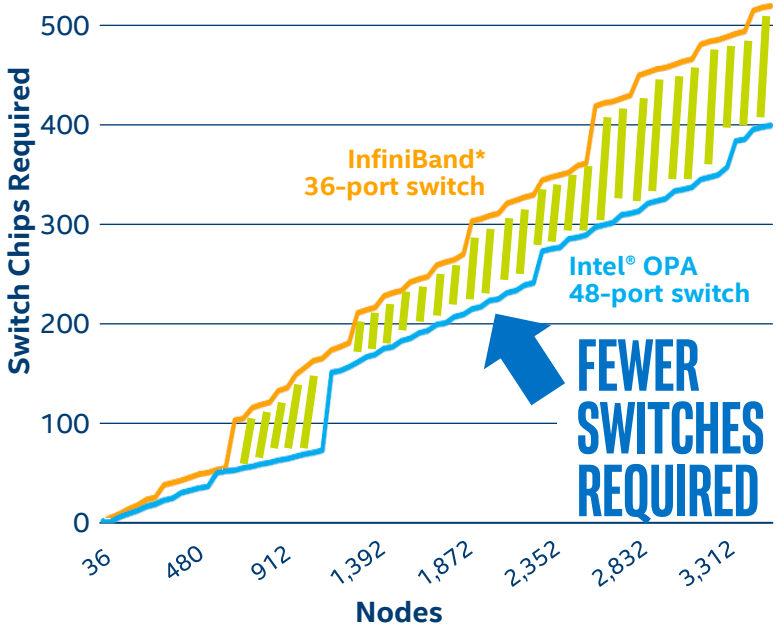

1. Assumes a 750-node cluster, and number of switch chips required is based on a full bisectional bandwidth (FBB) Fat-Tree configuration. Intel® OPA uses one fully-populated 768-port director switch, and Mellanox EDR solution uses a combination of 648-port director switches and 36-port edge switches. Mellanox component pricing from www.kernelsoftware.com, with prices as of November 3, 2015. Compute node pricing based on Dell PowerEdge R730 server from www.dell.com, with prices as of May 26, 2015. Intel® OPA pricing based on estimated reseller pricing based on Intel MSRP pricing on ark.intel.com. 2. Assumes a 750-node cluster, and number of switch chips required is based on a full bisectional bandwidth (FBB) Fat-Tree configuration. Intel® OPA uses one fully-populated 768-port director switch, and Mellanox EDR solution uses a combination of director switches and edge switches. Mellanox power data based on Mellanox CS7500 Director Switch, Mellanox SB7700/SB7790 Edge switch, and Mellanox ConnectX-4 VPI adapter card installation documentation posted on www.mellanox.com as of November 1, 2015. Intel OPA power data based on product briefs posted on www.intel.com as of November 16, 2015. Intel® OPA pricing based on estimated reseller pricing based on Intel MSRP pricing on ark.intel.com. 3 Number of switch chips required, switch density, and fabric scalability are based on a full bisectional bandwidth (FBB) Fat-Tree configuration, using a 48-port switch for Intel® Omni-Path Architecture and 36-port switch ASIC for either Mellanox or Intel® True Scale Fabric. *Other names and brands may be claimed as the property of others. 2.3X fabric scalability based on a 27,648-node cluster configured with the Intel® Omni-Path Architecture using 48-port switch ASICs, as compared with a 36-port switch chip that can support up to 11,664 nodes.

INNOVATIVE TECHNOLOGIES FOR HPC: MEMORY AND STORAGE

High Performance Memory and Storage

High-Bandwidth Memory

Configurable Modes

Integrated into the Processor

Intel® Optane™ Technology

(built with 3D XPoint™ Technology)

SSDs

DIMMs

Intel Solutions for Lustre* Software

The **Most Widely Used** File System for **HPC**

New Technologies Are Bringing Memory Closer to Compute

Bringing Memory Back Into Balance

High Bandwidth, On-Package Memory

Intel® Scalable System Framework

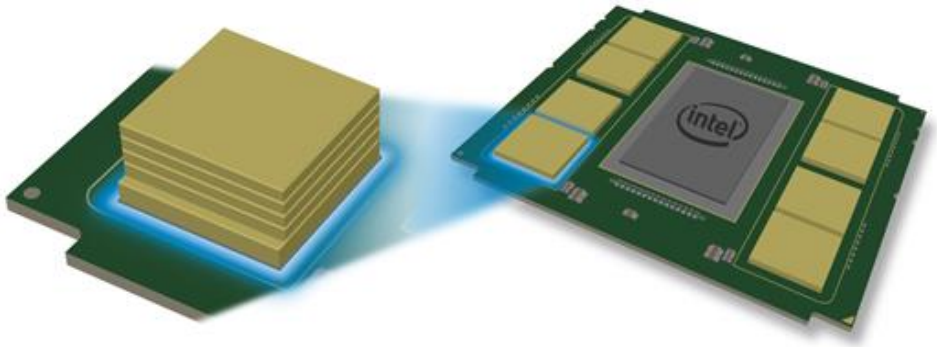

Up to 16GB with Knights Landing

5x the Bandwidth vs DDR4[1], >400 GB/s[1]

>5x More Energy Efficient vs GDDR5[2]

>3x More Dense vs GDDR5[2]

3 Modes of Operation

Flat Mode: Acts as Memory

Cache Mode: Acts as Cache

Hybrid Mode: Mix of Cache and Flat

[1] Projected result based on internal Intel analysis of STREAM benchmark using a Knights Landing processor with 16GB of ultra high-bandwidth versus DDR4 memory with all channels populated.

[2] Projected result based on internal Intel analysis comparison of 16GB of ultra high-bandwidth memory to 16GB of GDDR5 memory used in the Intel® Xeon Phi™ coprocessor 7120P.

Bringing Memory Back Into Balance

High Bandwidth, On-Package Memory

Intel® Scalable System Framework

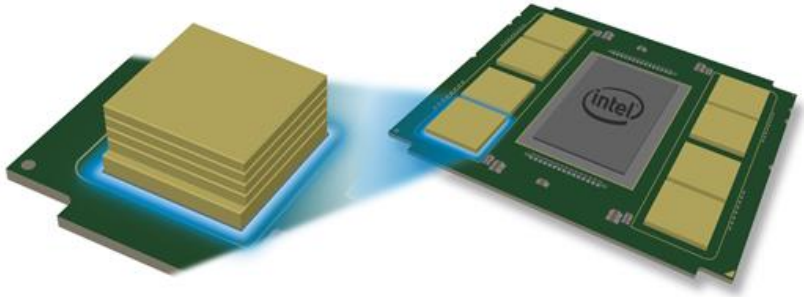

5X BANDWIDTH VS. DDR4[1]

>5X ENERGY EFFICIENT VS. GDDR5

>3X DENSITY VS. GDDR5[2]

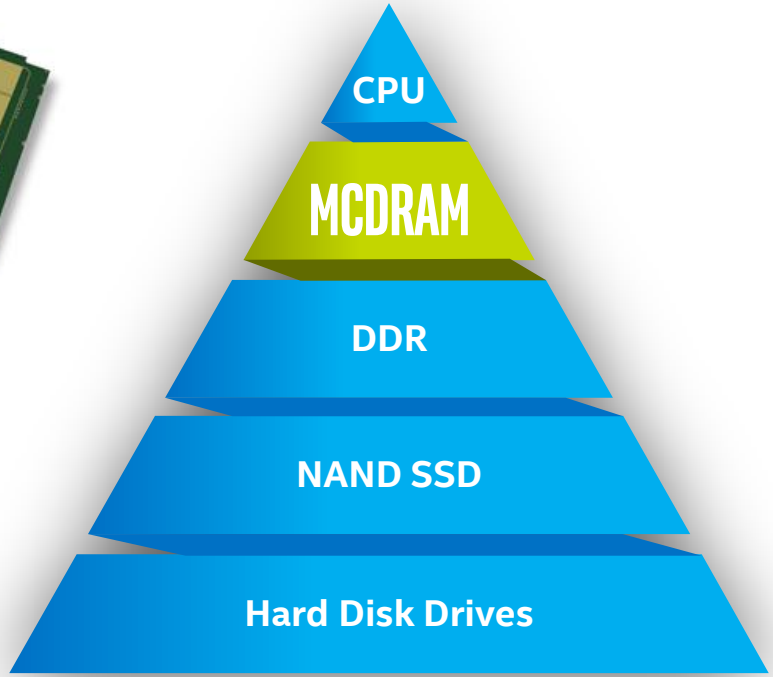

[1] Projected result based on internal Intel analysis of STREAM benchmark using a Knights Landing processor with 16GB of ultra high-bandwidth versus DDR4 memory with all channels populated.

[2] Projected result based on internal Intel analysis comparison of 16GB of ultra high-bandwidth memory to 16GB of GDDR5 memory used in the Intel® Xeon Phi™ coprocessor 7120P.

Bridging the Memory-Storage Gap

Intel® Optane™ Technology Based on 3D XPoint™

Intel® Scalable System Framework

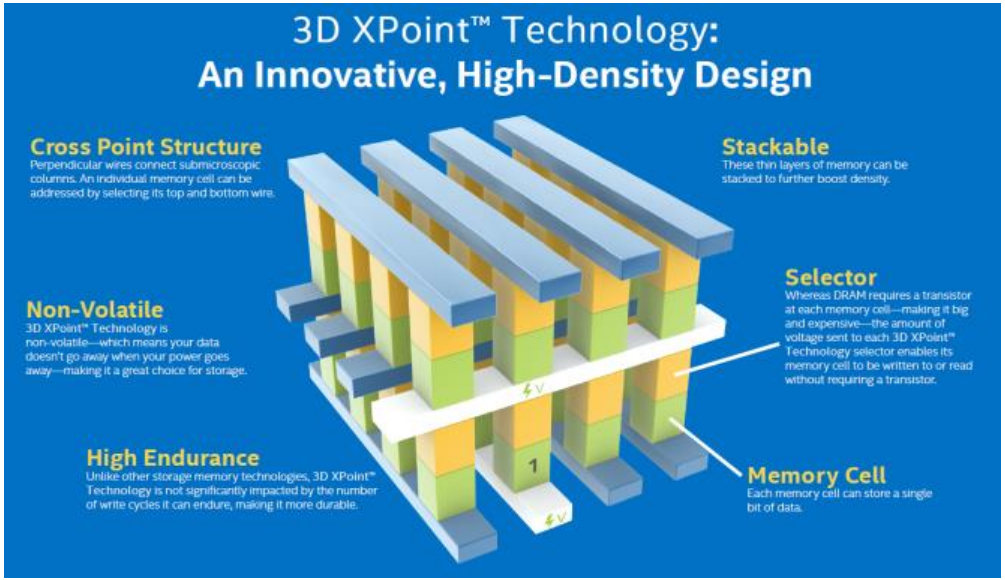

SSD

Intel® Optane™ SSDs 5-7x Current Flagship NAND-Based SSDs (IOPS)[1]

DRAM-like performance

Intel® DIMMs Based on 3D-XPoint™

1,000x Faster than NAND[1]

1,000x the Endurance of NAND[2]

Hard drive capacities

10x More Dense than Conventional Memory[3]

[1] Performance difference based on comparison between 3D XPoint™ Technology and other industry NAND
[2] Density difference based on comparison between 3D XPoint™ Technology and other industry DRAM
[2] Endurance difference based on comparison between 3D XPoint™ Technology and other industry NAND

A TIMELINE OF MEMORY CLASS INTRODUCTIONS

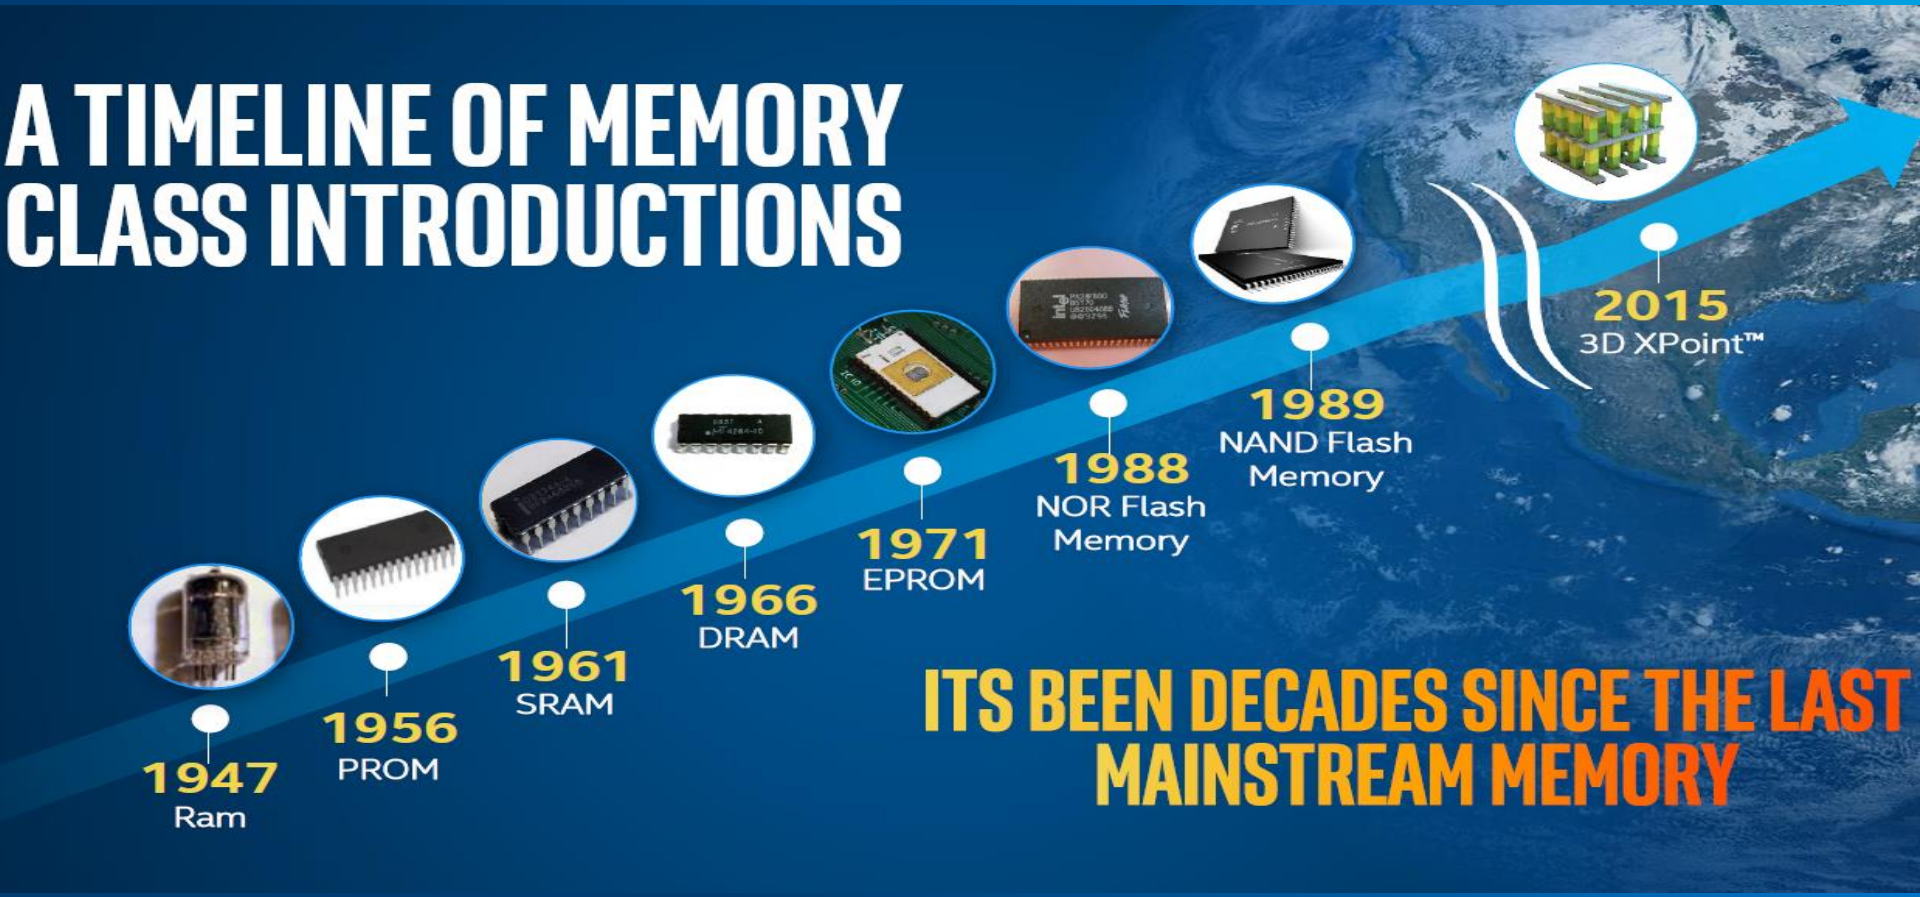

ITS BEEN DECADES SINCE THE LAST MAINSTREAM MEMORY

EL GAP DE LAS TECNOLOGÍAS ACTUALES DE MEMORIA

NAND

DENSIDAD
NO-VOLÁTIL
BARATA
LENTA

PRECIO/ CAPACIDAD

GAP

VELOCIDAD/ VOLATILIDAD

DRAM

MENOS DENSIDAD
VOLÁTIL
CARA
RÁPIDA

EL GAP DE LAS TECNOLOGÍAS ACTUALES DE MEMORIA: SOLUCIONADO

NAND	3D XPoint™ TECHNOLOGY	DRAM
		
DENSIDAD	DENSIDAD	MENOS DENSIDAD
NO-VOLÁTIL	NO-VOLÁTIL	VOLÁTIL
BARATA	BARATA	CARA
LENTA	RÁPIDA	RÁPIDA

REVOLUCIONARIA TECNOLOGÍA 3D XPoint™

1000 VECES MÁS RÁPIDA
QUE NAND

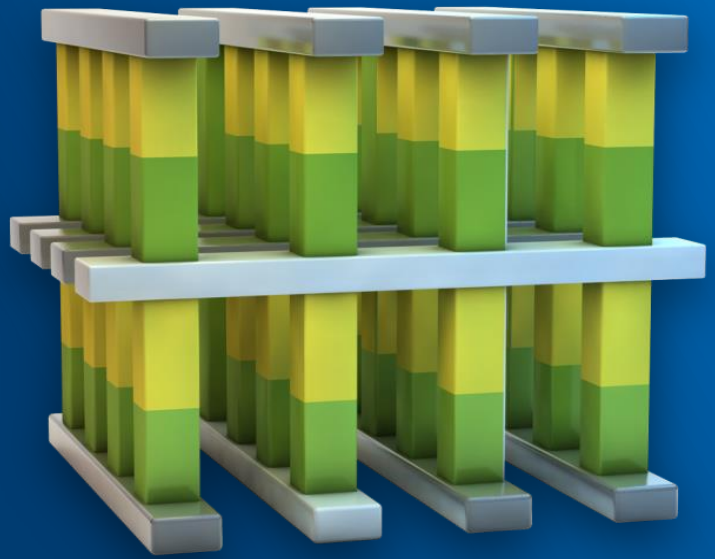

1000 VECES MÁS DURADERA
QUE NAND

10 VECES MÁS DENSIDAD
QUE DRAM

Bridging the Memory-Storage Gap

Intel® Optane™ Technology

Intel® Scalable System Framework

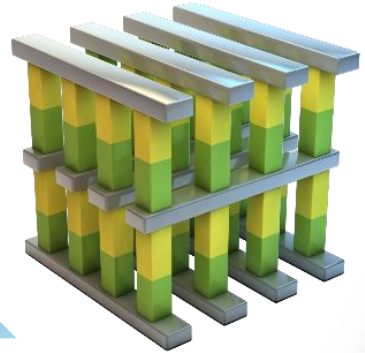

1000X
FASTER
Than NAND[1]

1000X
ENDURANCE
Of NAND[2]

10X
DENSER
Than DRAM[3]

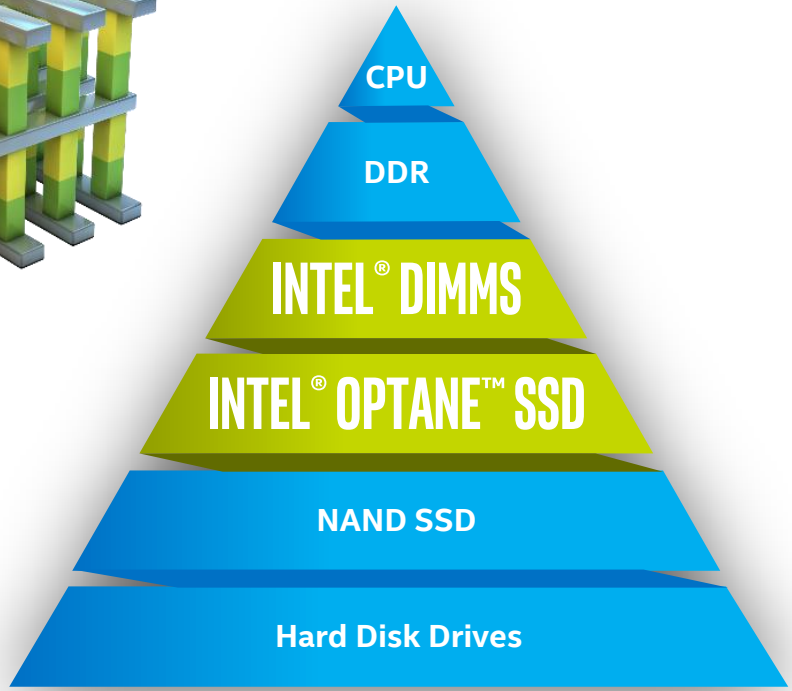

[1] Performance difference based on comparison between 3D XPoint™ Technology and other industry NAND
[2] Density difference based on comparison between 3D XPoint™ Technology and other industry DRAM
[2] Endurance difference based on comparison between 3D XPoint™ Technology and other industry NAND

Intel® Solutions for Lustre* Software

The Speed of Lustre* with the Support of Intel

Intel® Scalable System Framework

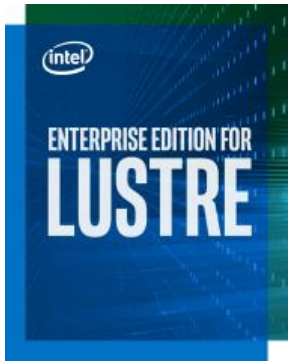

EXTREME SCALE STORAGE FOR HPC

Intel® Enterprise Edition for Lustre* Software v2.4

- Support for “Distributed Namespace” (DNE) Feature to Scale Out the Metadata Performance of Lustre*
- Support for the Latest OS: Red Hat* 6.5-7 and SUSE* 11sp3-12
- Parallel Read IO Performance & HSM Scalability Improvements

Intel® Cloud Edition for Lustre* Software v1.2

- Support for Over-the-Wire and Storage Encryption
- Disaster Recovery from File System Snapshots
- Simplified File System Mounting on Clients
- Support for Intel® Xeon® Processor E5-2600 v3 Product Family-Based Instances

Intel® Foundation Edition for Lustre* Software v2.7.1

- Delivers the Latest Functions and Features
- Fully Supported by Intel

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

Intel® Solutions for Lustre* Software

The Speed of Lustre* with the Support of Intel

Intel® Scalable System Framework

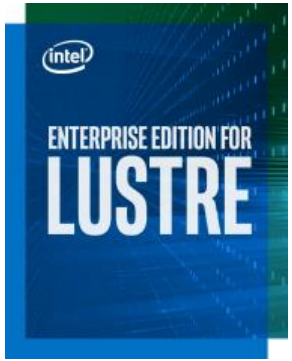

INTEL® MANAGER FOR LUSTRE
EASE OF DEPLOYMENT AND MANAGEMENT

24/7 INTEL SUPPORT
HIGH AVAILABILITY/RELIABILITY

HADOOP* ADAPTER
BIG DATA AND HPDA

EXTREME SCALE STORAGE FOR HPC

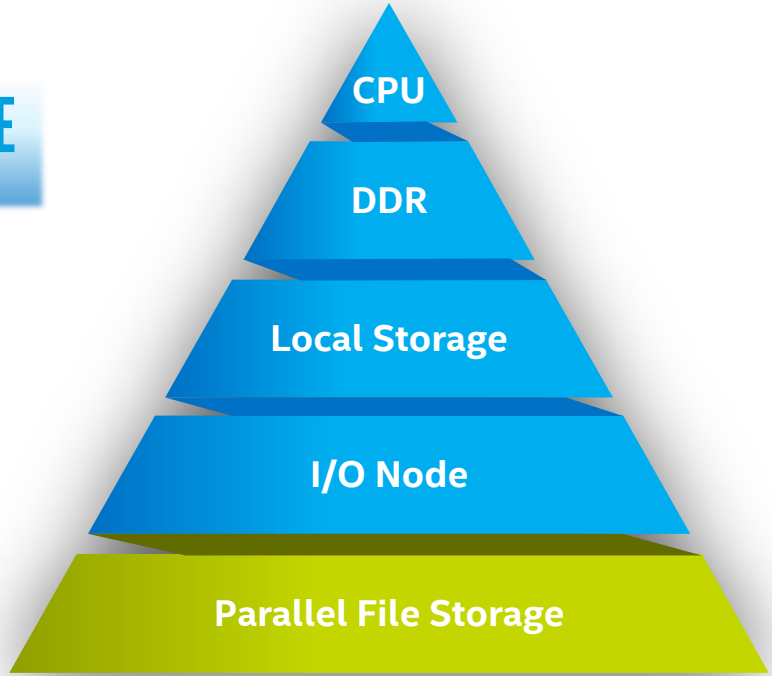

*Other names and brands may be claimed as the property of others

INNOVATIVE TECHNOLOGIES FOR HPC: CODE MODERNIZATION

Efficient, Scalable & Portable App Performance

Take Advantage of

1. Thread, Vector and Data Level Parallelism 
2. High Bandwidth Memory and NVM 
3. Increased Bandwidth and Reduced Latency Fabric 

Intel Is Helping You Develop Modern Code

Intel® Parallel Computing Centers

Collaborating to Accelerate the Pace of Discovery

>50 Centers

5 Focusing on Lustre*

17 Countries

>90 Codes

2 User Groups

https://software.intel.com/en-us/ipcc

Intel® Modern Code Developer Community

Developer Zone

Software Tools, Training Webinars, How-To Guides, Parallel Programming BKMs, Technical Content, Support Forums, Remote Access to Hardware

Experts

Black Belts and Intel Engineer experts
F2F, Conferences and Tradeshows

software.intel.com/moderncode

HPC Developer Conferences
https://hpcdevcon.intel.com/

Community, ISV and Proprietary Codes

Intel® Software Solutions

Intel® Software Defined Visualization

Low Cost
No Dedicated Viz Cluster

Excellent Performance
Less Data Movement, I/O
Invest Power, Space, Budget in Greater Compute Capability

High Fidelity
Work with Larger Data Sets – Not Constrained by GPU Memory

Intel® Parallel Studio

Faster Code
Boost Application Performance on Current and Next-Gen CPUs

Create Code Faster
Utilizing a Toolset that Simplifies Creating Fast and Reliable Parallel Code

HPC System Software Stack

An Open Community Effort
Broad Range of Ecosystem Partners
Open Source Availability

Benefits the Entire HPC Ecosystem
Accelerate Application Development
Turnkey to Customizable

Open Software Available Today!

Bringing Your Data into Focus

Intel-Supported Software Defined Visualization (SDVis)

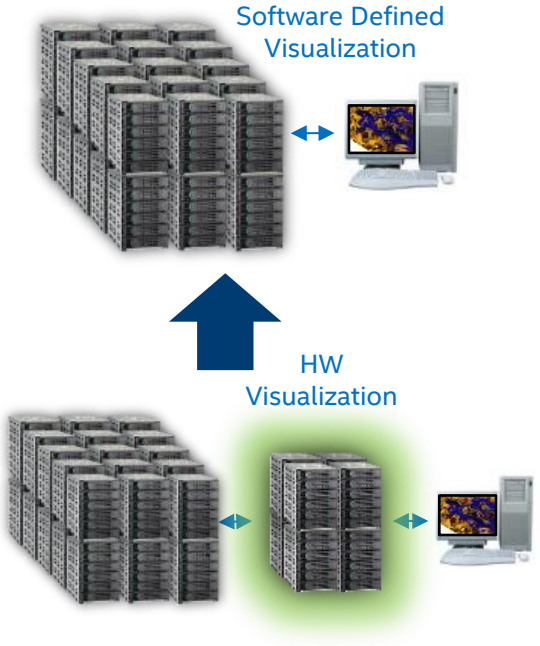

Image Rendered by OSPRay

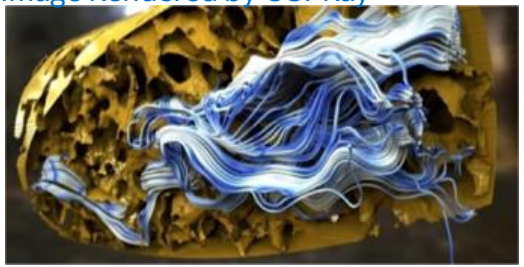

Standard OpenGL Image

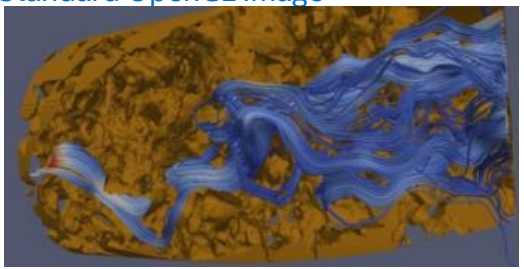

Embree

- CPU Optimized Ray Tracing Algorithms
- 'Tool kit' for Building Ray Tracings Apps
- Broadly Adopted by 3rd Party ISVs
- Web Site: http://embree.github.io

OSPRay[1]

- Rendering Engine Based on Embree
- Library and API Designed to Ease the Creation of Visualization Software
- Web Site: http://ospray.org

OpenSWR[1]

- Rasterization Visualization on CPUs
- Good Enough to Replace HW GPU
- Supports ParaView, Visit, VTK
- Web Site: http://openswr.org

[1] Currently available in alpha

Intel's Commitment to the HPC Community

Intel® Modern Code Developer Community

An Online Community
to Reach **400,000** Developers and Partners with Tools, Trainings and Support

Hands on Training
for 10,000 Developers and Partners

Remote Access
to Intel® Xeon® Processor and Intel® Xeon Phi™ Coprocessor-based Clusters

Intel® Supporting the Software Community

Leading Contributor to Multiple **Open Source Projects**, Including Linux*, Luster*, OpenHPC, Embree

Working with ISVs and the Community to Help Modernize Codes Across the Ecosystem

Intel® Parallel Compute Centers

Focused on Modernizing Community Code
50+ Intel® PCC Modernizing More Than 90 HPC Computing Codes Across 16 Domains

A Global Effort
Located in 15 Countries and Four Continents

Join the Online Community Today!

A Global Online Community

Intel® Modern Code Developer Community

Developer zone

- Modern Code Zone
- Software Tools, Training Webinars
- How-to guides, Parallel Programming BKMs
- Remote Access to Hardware
- Support Forums

Topics

- Vectorization/Single Instruction, Multiple Data (SIMD)
- Multi-Threading
- Multi Node/Clustering
- Take Advantage of On-Package High-Bandwidth Memory
- Increase Memory and Power Efficiency

Experts

- Black Belts, & Intel Engineer Experts
- Technical Content, Training -Webinars, F2F, Forum Support
- Conference and Tradeshows: Keynotes, Presentations, BOFs, Demos, Tutorials

software.intel.com/moderncode

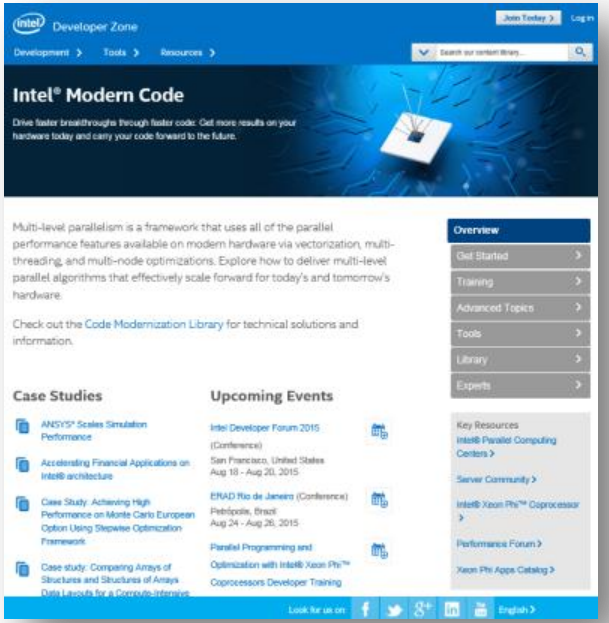

Accelerate Your Discovery and Innovation

Breakthrough Performance

Standard-Based Programmability

Common Infrastructure Across Workloads

Broad Vendor Availability


intel®
experience
what's inside™