

ПОСТРОЕНИЕ ОПТИМАЛЬНЫХ РЕШЕНИЙ ДЛЯ ЗАДАЧ НРС И MACHINE LEARNING

Николай Местер, директор по развитию корпоративных проектов

Иван Кузьмин, руководитель отдела разработки высокопроизводительных библиотек

The HPC Opportunity

MODELING & SIMULATION




**\$515
average
return per
\$1
of HPC
investment¹**

HPC DATA ANALYTICS



**18% revenue
CAGR; >\$3
billion in
2020²**

ARTIFICIAL INTELLIGENCE



**55% revenue
CAGR; >\$47
billion in
2020³**

VISUALIZATION



**30% revenue
CAGR; >\$1.6
billion in
2020⁴**

¹ Source: Source: IDC HPC and ROI Study Update, September 2015

² Source IDC Worldside High-Performance Data Analytics Forecast 2016-2020, June 2016

³ Source: IDC Worldwide Semiannual Cognitive/Artificial Intelligence Systems Spending Guide, Oct 2016

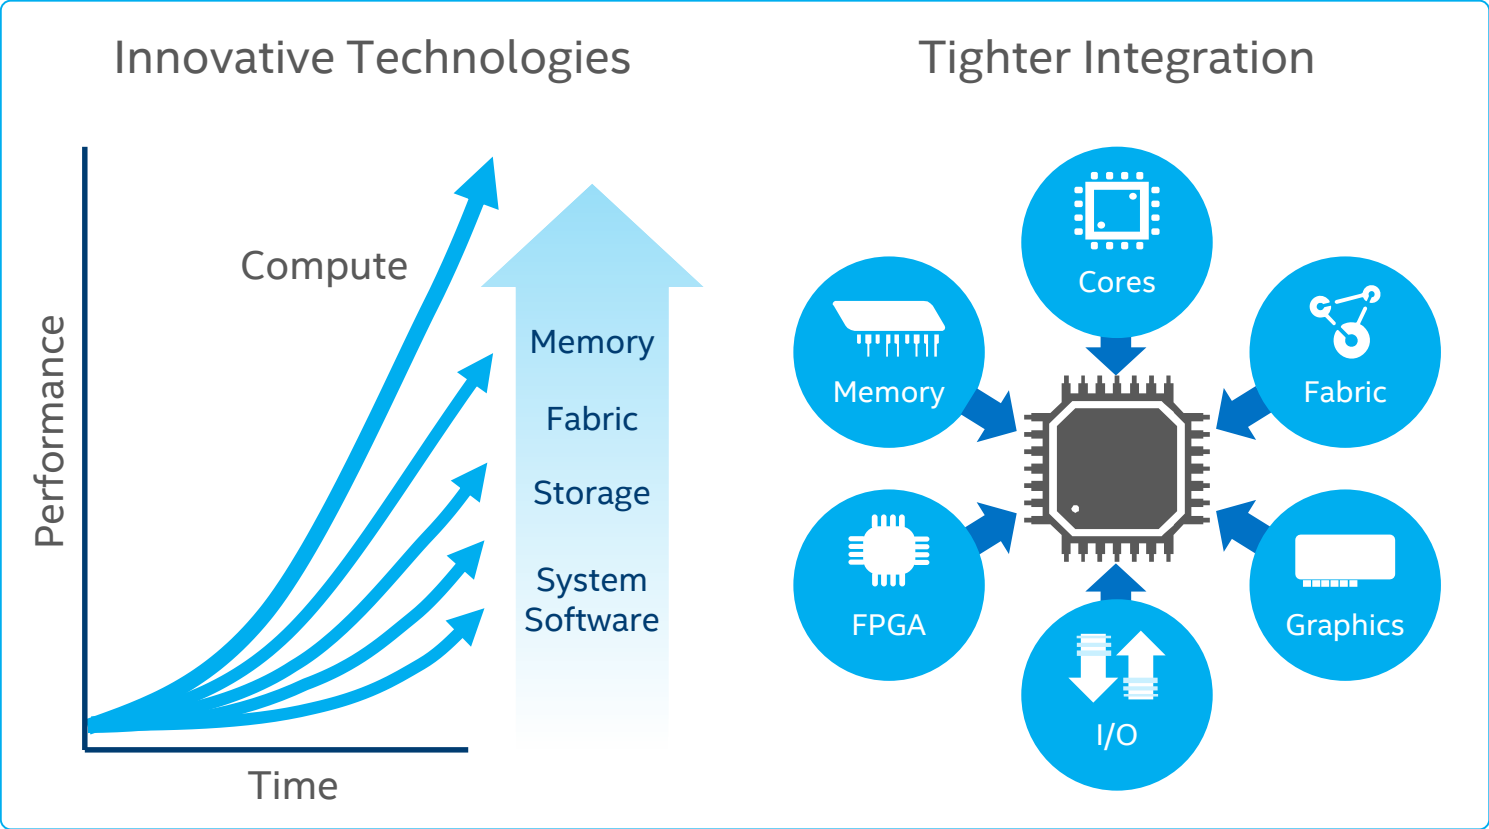
⁴ Source: MarketsandMarkets Visualization and 3D Rendering Software Market by Application, March 2016

Intel and the Intel 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. All products, dates, and figures are preliminary and are subject to change without any notice. Copyright © 2013, Intel Corporation.




A Holistic Architectural Approach is Required

System




Application

Modernized Code



Key Elements of Intel® SSF



INTEL® SCALABLE SYSTEM FRAMEWORK



**MARKET
LEADING¹**



**HIGHLY
PARALLEL**




**COST
ADVANTAGE**



**INTEL
SUPPORTED**



**FLEXIBILITY
& STABILITY**



**EXTREME
SCALABILITY**

Intel® Xeon® Scalable Processor Enables Amazing Discoveries through HPC




**Origins of the
Universe**




**Weather
Forecasting**



**Energy
Research**



**Material
Science**



Physics



**Personalized
Healthcare**

INTEL® XEON® SCALABLE PROCESSOR

Наибольшее количество улучшений платформы за декаду



- ✓ Существенное улучшение производительности на ядро
- ✓ **Intel® Advanced Vector Extension 512 (Intel® AVX-512):**
 - ✓ до 2X FLOPs/second peak performance по сравнению с предыдущей архитектурой
- ✓ Ускорение IO за счет **Intel® Omni-Path Architecture (Fabric)**
- ✓ **Интегрированная технология Integrated Intel® QuickAssist Technology** (crypto & compression offload)
- ✓ **Улучшенные RAS свойства** (Reliability, Availability, Serviceability)
- ✓ **Intel® Resource Director Technology (Intel® RDT)** для эффективности Efficiency и TCO

A GLIMPSE INSIDE THE INTEL® XEON® SCALABLE PROCESSOR PLATFORM



Fabric
Intel® Omni-Path
Architecture




Networking
Intel® Ethernet



Accelerators
Intel® QuickAssist
Intel® AVX-512



SSDs
Intel® Optane™ SSD
DC P4800X
Intel® SSD DC P4600
Intel® VMD



Complementary
Intel® FPGA
Intel® Xeon Phi™
Intel® Silicon Photonics

INTEGRATED OPTIONS

Workload optimized frameworks & telemetry

(e.g. Caffe*, Intel® DAAL, Intel® MKL, DPDK, SNAP*, SPDK)

PERFORMANCE



SECURITY



AGILITY



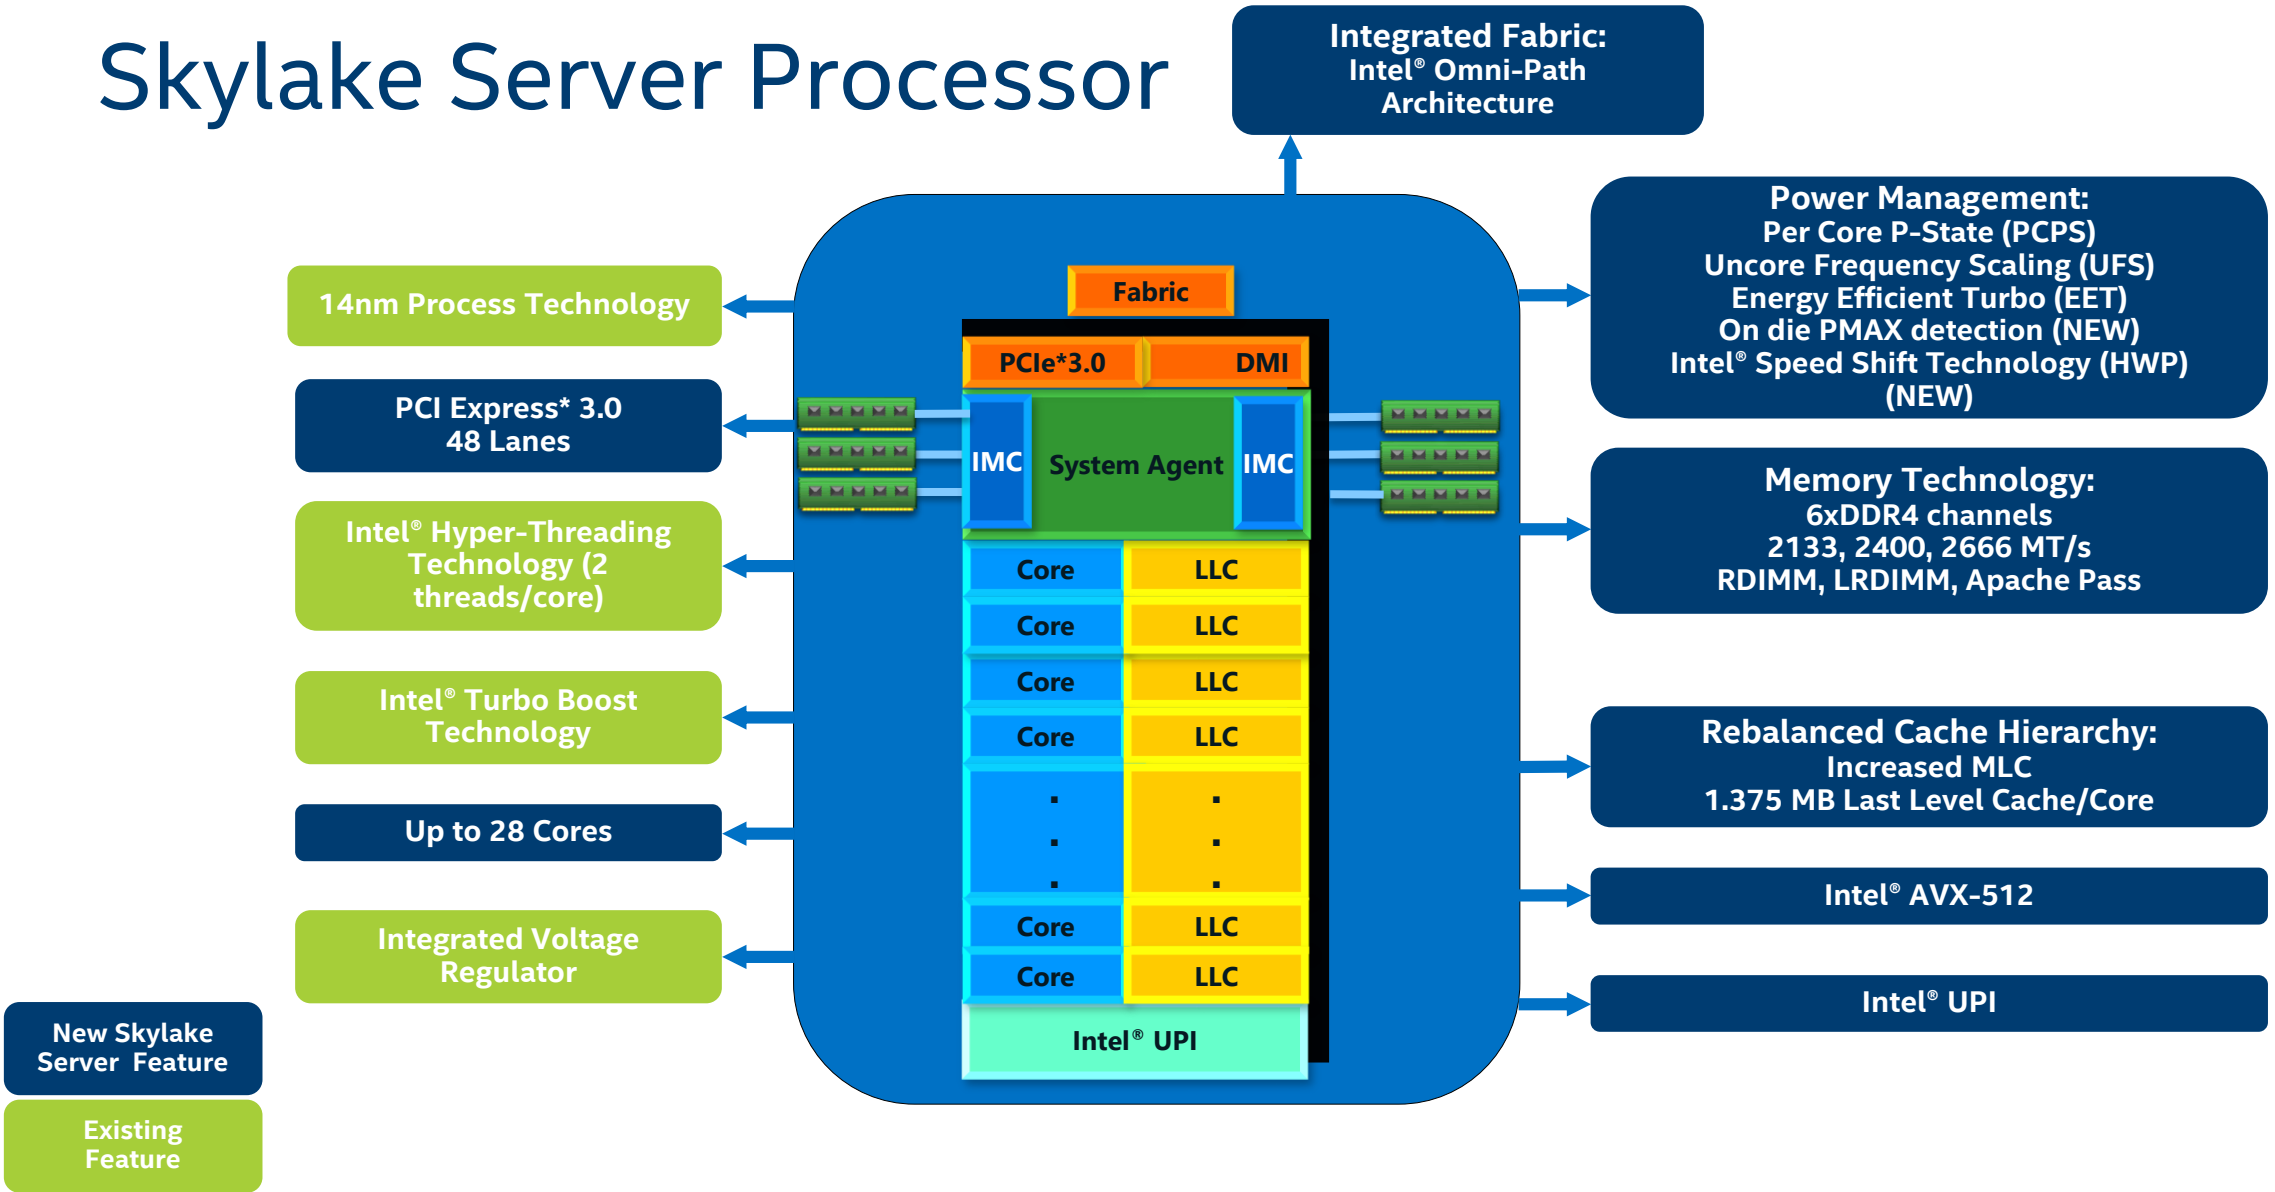
**ADVANCING VIRTUALLY EVERY ASPECT: BRAND NEW CORE,
CACHE, ON-DIE INTERCONNECTS, MEMORY CONTROLLER & MORE**

Intel® Advanced Vector Extensions 512 (Intel® AVX-512)
Intel® Volume Management Device (Intel® VMD)
Intel® Data Analytics Acceleration Library (Intel® DAAL)

Intel® Math Kernel Library (Intel® MKL)
Storage Performance Development Kit (SPDK)

Data Plane Development Kit (DPDK)
Intel® Resource Director Technology (Intel® RDT)

Skylake Server Processor



INTEL® XEON® SCALABLE PROCESSORS

THE FOUNDATION FOR AGILE, SECURE WORKLOAD-OPTIMIZED HYBRID CLOUD



BEST

SCALABLE PERFORMANCE
HARDWARE-ENHANCED SECURITY
ADVANCED RAS

OPTIMIZED FOR WIDEST RANGE
OF EVOLVING/MULTI WORKLOADS

MISSION-CRITICAL, VIRTUALIZATION/CONSOLIDATION,
REAL-TIME ANALYTICS AND ARTIFICIAL INTELLIGENCE

MAINSTREAM



GREAT

SCALABLE PERFORMANCE
MEMORY PERFORMANCE
ADVANCED RAS

WORKLOAD-OPTIMIZED +
EFFICIENCY & AGILITY

PERFORMANCE FOR GENERAL-PURPOSE
COMPUTE, STORAGE AND NETWORKING



GOOD

SCALABLE PERFORMANCE
AT LOW POWER
STANDARD RAS

MODERATE TASKS

INTEL® TURBO BOOST TECHNOLOGY AND
INTEL® HYPER-THREADING TECHNOLOGY
FOR MODERATE WORKLOADS

EFFICIENT



ENTRY

SCALABLE PERFORMANCE
HARDWARE-ENHANCED SECURITY
STANDARD RAS

LIGHT TASKS

ENTRY PERFORMANCE, PRICE SENSITIVE
FOR LIGHT WORKLOADS

ENTRY

INTEL® XEON® SCALABLE PROCESSOR

THE FOUNDATION FOR AGILE, SECURE WORKLOAD-OPTIMIZED HYBRID CLOUD



UP TO 28 CORES

UP TO 2, 4 & 8 SOCKET SUPPORT WITH UP TO 3 UPI LINKS

DDR4 2666 MHz WITH UP TO 1.5 TB TOPLINE MEMORY CHANNEL BANDWIDTH

HIGHEST ACCELERATOR THROUGHPUT

MAINSTREAM

INTEL® XEON® PLATINUM 81XX PROCESSORS

- Most cores (28C, 56T)
- Highest frequency (up to 3.6 GHz, 4C)
- Most Socket flexibility (2, 4, 8+)
- Elite IO/memory (48 PCIe 3.0 lanes, 6 memory channels)
- Most/fastest 3 UPI (Ultra Path Interconnects, 10.4 GT/s)
- Fastest Memory (DDR4-2666 MHz)
- Highest Intel® AVX-512 with 2 FMA per core
- Intel® Turbo Boost and Intel® Hyper-Threading Technology
 - New Advanced RAS
- Node Controller Support to assist in scaled node management

RESULT:

- Best workload-optimized performance for general purpose compute across virtualized environments and hybrid cloud deployments
- Best choice for mission-critical, traditional DC apps, real-time analytics, and AI/deep learning inference workloads; smart path for DL training
- Best performance for the most demanding storage and networking workloads

RAS: Reliability, Availability, and Serviceability

Intel® Xeon Phi™ Processor – TCO Solution for HPC & AI

A Key Element of HPC, AI, and Mixed Workload Clusters



Total Cost of
Ownership

Price Performance
Power Efficiency
Performance



Optimized for HPC
& AI

Highly-Parallel
No PCIe Bottlenecks
Scalability



Complements
Intel® Xeon®

Common Programming
Mixed Clusters
Runs x86 code

Reduces total cost of ownership, designed for HPC & AI, protects investment

Intel® Xeon Phi™ Processor Architecture



Self-Boot Processor

Binary-compatibility with Xeon, 3+ TFLOPS¹ (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:
(up to 36)



Enhanced Intel® Atom™ cores based on Silvermont Microarchitecture



Tile



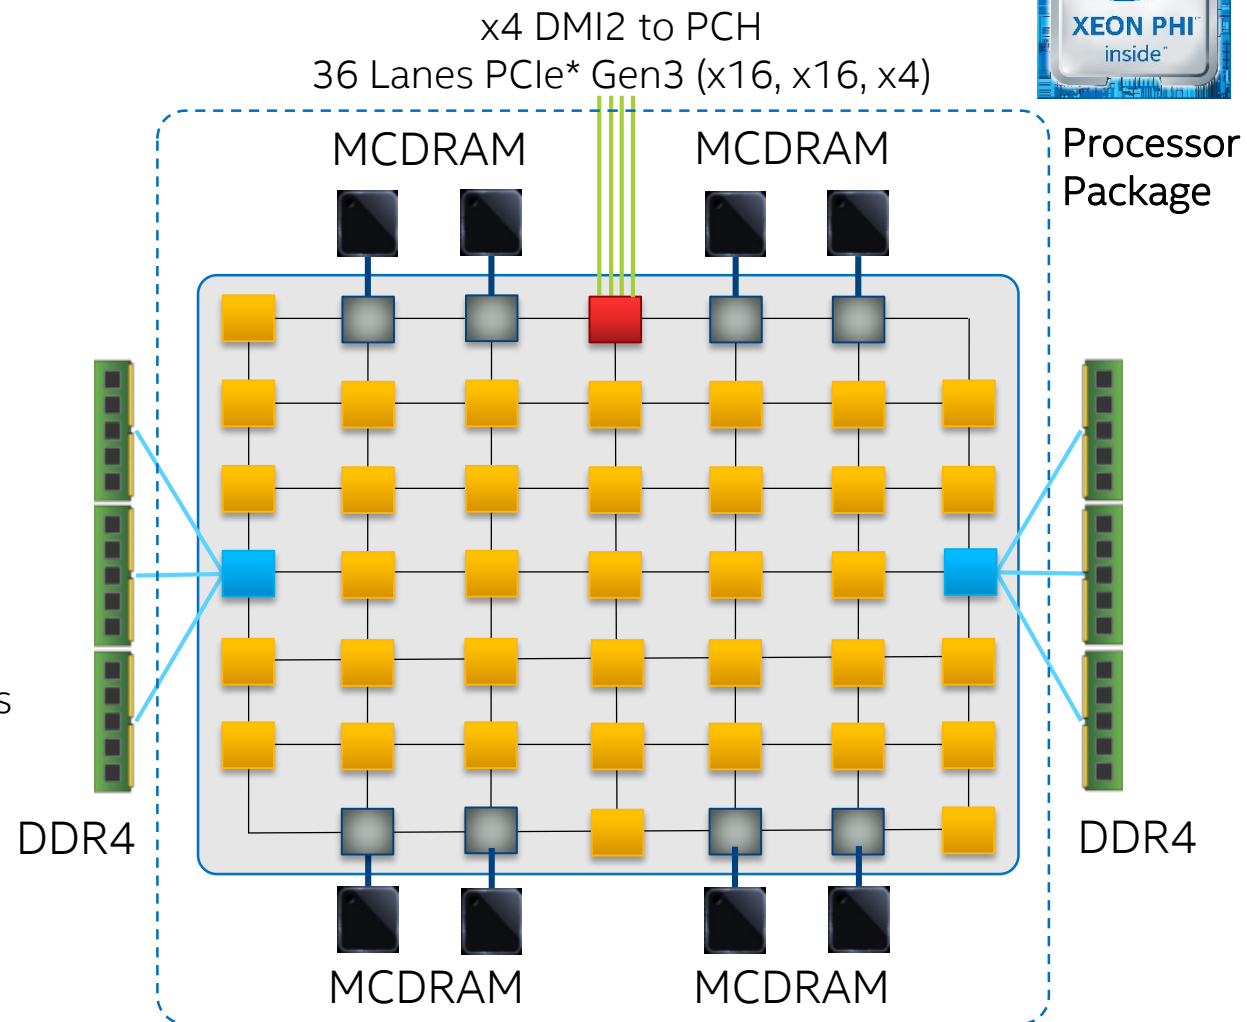
EDC (Embedded DRAM Controller)



IMC (Integrated Memory Controller)




IIO (Integrated I/O Controller)



¹Theoretical peak performance

Intel® Xeon Phi™ Product Family x200

Intel® Xeon Phi™ Processor



Host Processor in Groveport Platform

Self-boot Intel® Xeon Phi™ processor

Intel® Xeon Phi™ Target Segments & Applications


Deep Learning Training




Material Science: **VASP***, **NWCHEM***, **GTC-P***



QCD: **QPHIX***, **MILC***, **CHROMA***, **CCS QCD***



CFD/Mfg: **OPENFOAM***, **CLOVERLEAF***, **LSTC LS-DYNA***, **CONVERGENT SCIENCE CONVERGE CFD***



Weather/Climate/Cosmology: **WRF***, **NEMO***, **WALLS***



Energy: **ISO3DFD***



FSI: **STAC A2***, **MONTE CARLO***, **BLACK SCHOLES***, **BINOMIAL OPTIONS***



MD: **LAMMPS***, **NAMD***, **GROMACS***, **AMBER***

Features Driving Perf & Perf/\$/W

16GB MCDRAM

High memory (MCDRAM) BW (≤ 490 GB/s)

Intel® AVX-512 ER

High system memory (≤ 400 GB)

High number of physical cores (≤ 72)

High number of threads (≤ 288)

Lower system price (~\$4700)¹

Lower system price (~\$4700)¹

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

¹Software 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 Slides 40-52.




Intel® Xeon Phi™ Utilization Value

Homogenous
"Large" Core



Homogenous
"Small" Core



GPU


General Purpose

Workload Optimized

Intel® Xeon Phi™ Utilization Benefits

- Runs optimized applications best
- Runs all x86 applications
- Doesn't reduce resources for some applications

Example Supercomputer Cluster
Top 25 Applications



GPU Utilization Limitations

- Requires coding to run application, requires optimization to run best
- Doesn't run x86 applications
- Dedicated resource reduces cluster performance for some applications

World's Most Responsive Data Center SSD¹

Delivering an **industry leading combination of low latency, high endurance, QoS and high throughput**, the Intel® Optane™ SSD is the first solution to **combine the attributes of memory and storage**. This innovative solution is optimized to **break through storage bottlenecks** by providing a new data tier. It accelerates applications for **fast caching and storage, increasing scale per server** and reducing transaction cost. Data centers based on the latest Intel® Xeon® processors can now also **deploy bigger and more affordable datasets** to gain new insights from larger memory pools.




1. Responsiveness defined as average read latency measured at Queue Depth 1 during 4k random write workload. Measured using FIO 2.15. Common configuration - Intel 2U Server System, OS CentOS 7.2, kernel 3.10.0-327.el7.x86_64, CPU 2 x Intel® Xeon® E5-2699 v4 @ 2.20GHz (22 cores), RAM 396GB DDR @ 2133MHz. Intel drives evaluated - Intel® Optane™ SSD DC P4800X 375GB and Intel® SSD DC P3700 1600GB. Samsung* drives evaluated - Samsung SSD PM1725a, Samsung SSD PM1725, Samsung PM963, Samsung PM953. Micron* drive evaluated - Micron 9100 PCIe* NVMe* SSD. Toshiba* drives evaluated - Toshiba ZD6300. Test - QD1 Random Read 4K latency, QD1 Random RW 4K 70% Read latency, QD1 Random Write 4K latency using FIO 2.15.

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

Breakthrough Performance

4K 70/30 RW Performance at Low Queue Depth




- ✓ **5-8x faster** at low Queue Depths¹
- ✓ Vast majority of **applications generate low QD** storage workloads

1. Common Configuration - Intel 2U PCSD Server ("Wildcat Pass"), OS CentOS 7.2, kernel 3.10.0-327.el7.x86_64, CPU 2 x Intel® Xeon® E5-2699 v4 @ 2.20GHz (22 cores), RAM 396GB DDR @ 2133MHz. Configuration - Intel® Optane™ SSD DC P4800X 375GB and Intel® SSD DC P3700 1600GB. Performance - measured under 4K 70-30 workload at QD1-16 using fio-2.15.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance.

Predictably Fast Service

Read QoS in Mixed Workload



✓ up to **60X** better at 99% QoS¹


✓ Ideal for critical applications with aggressive latency requirements

1. Common Configuration - Intel 2U PCSD Server ("Wildcat Pass"), OS CentOS 7.2, kernel 3.10.0-327.el7.x86_64, CPU 2 x Intel® Xeon® E5-2699 v4 @ 2.20GHz (22 cores), RAM 396GB DDR @ 2133MHz. Configuration - Intel® Optane™ SSD DC P4800X 375GB and Intel® SSD DC P3700 1600GB. QoS - measures 99% QoS under 4K 70-30 workload at QD1 using fio-2.15.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance.

Responsive Under Load

Average Read Latency under Random Write Workload




✓ up to **40X faster response time** under workload¹

✓ Consistently **amazing response time under load**

1. Responsiveness defined as average read latency measured at queue depth 1 during 4k random write workload. Measured using FIO 2.15. Common Configuration - Intel 2U PCSD Server ("Wildcat Pass"), OS CentOS 7.2, kernel 3.10.0-327.el7.x86_64, CPU 2 x Intel® Xeon® E5-2699 v4 @ 2.20GHz (22 cores), RAM 396GB DDR @ 2133MHz. Configuration - Intel® Optane™ SSD DC P4800X 375GB and Intel® SSD DC P3700 1600GB. Latency - Average read latency measured at QD1 during 4K Random Write operations using fio-2.15.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance.



ACCELERATING THE FUTURE FOR ANALYTICS AND AI

WITH INTEL TECHNOLOGIES

IVAN KUZMIN

OUR COMMITMENT IS TO
ENABLE THE BEST
ANALYTICS EXPERIENCE,
FROM HW TO SW

ANALYTICS APPS



ANALYTICS PIPELINE



DATA PLATFORM

cloudera®



INFRASTRUCTURE



ML/DL

PERFORMANCE AND SECURITY

ENABLING BEST PERFORMANCE AND SECURITY

PERFORMANCE

4X

MLlib* thru Intel® for Intel Math Kernel Library

5.6X

HBase off-heap read

8X

HDFS Erasure Coding codec thru Intel® Intelligent Storage Acceleration Library

SECURITY

1.28X

Spark shuffle file encryption

1.35X

Spark* shuffle RPC encryption



ENABLING BEST SCALABILITY FOR MACHINE LEARNING

>10X

SCALABILITY IMPROVEMENT FOR
**CUSTOMER ANALYSIS USING
K-MEANS CLUSTERING**

<https://github.com/intel-analytics/SparseSpark>

>10X

SCALABILITY IMPROVEMENT FOR
**TOPIC MODELING USING LATENT
DIRICHLET ALLOCATION**


<https://github.com/intel-analytics/TopicModeling>

ADVANCE SOLUTIONS FOR NEW NEEDS (DEEP LEARNING)

**DISTRIBUTED
DATA**




**EFFICIENT
SCALABILITY**



**LOWER
TCO**



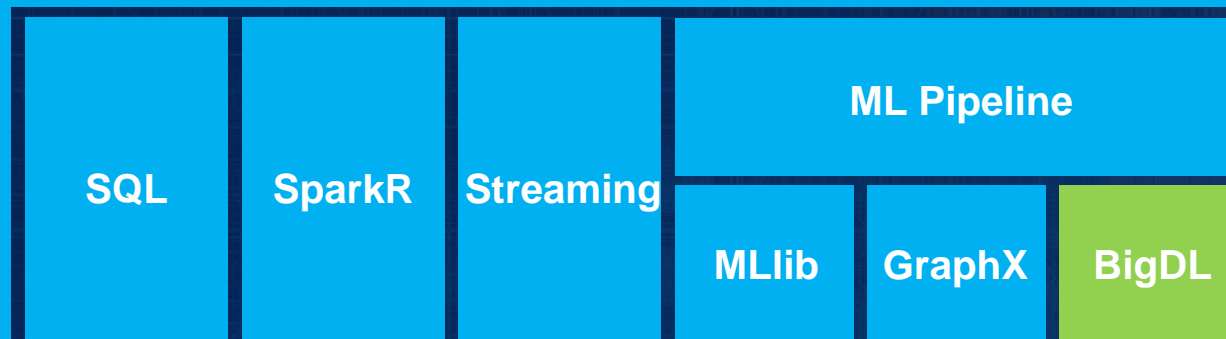
**HIGH
PERFORMANCE**



BIGDL

ANSWERING THE NEW NEEDS

DATA FRAME



Spark Core



Working with Xeon clusters

Feature Parity with Caffe* and Torch*

Drives Spark/Hadoop as Unified Data Analytics Platform, Giving Easy Experience and Lower TCO

Deep Learning on Big Data Platform, Enabling Efficient Scale-Out

High Performance at Single-Node

BIGDL READY FOR WIDE ADOPTION

cloudera®



databricks®



Alibaba Cloud
aliyun.com



amazon
webservices™

CRAY
THE SUPERCOMPUTER COMPANY



Microsoft
Azure



Lightbend

INTEL® NERVANA™ PORTFOLIO

EXPERIENCES



TOOLKITS

Intel® DL
Training &
Deployment

Intel® Nervana™
DL Software &
Cloud

Intel®
Computer
Vision SDK

Intel® GO™
Automotive
SDK

Movidius
Fathom

FRAMEWORKS



LIBRARIES



Intel Distribution

Intel® DAAL



Intel® Nervana™ Graph*

Intel® MKL

MKL-DNN

Intel® MLSL

HARDWARE



Compute

Memory/Storage

Networking

Computer Vision

END
TO
END
AI


*Future

AI ON INTEL: UNLEASHING THE NEXT WAVE

A man and a woman are sitting at a desk, looking at a laptop. The man is pointing at the screen, and the woman is looking at a tablet. The entire image has a blue overlay.

**WE KNOW
THE FUTURE
BECAUSE
WE'RE
BUILDING IT
LET'S COLLABORATE**

Software.intel.com/BigDL
Software.intel.com/AI
MEET US TODAY AND TOMORROW



LEGAL NOTICES & DISCLAIMERS

- Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.
- No computer system can be absolutely secure.
- 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>.

Configurations:

- 4.3X for Spark MLlib thru Intel Math Kernel Library (MKL)
 - Spark-Perf (same for before and after): 9 nodes each with Intel® Xeon® processor E5-2697A v4 @ 2.60GHz * 2 (16 cores, 32 threads); 256 GB ; 10x SSDs; 10Gbps NIC
- 5.6x for HBase off heaping read in micro workload (PE) and 1.3x in real Alibaba production workload
 - PE (same for before and after): Intel® Xeon® Processor X5670 @ 2.93Hz *2 (6 cores, 12 threads); RAM: 150 GB; 1Gbps NIC
 - Alibaba (same for before and after): 400 nodes cluster with Intel® Xeon® processors

Intel, the Intel logo, Xeon, Xeon phi, Atom, Core, Nervana, Iris, RealSense, GO, etc. are trademarks of Intel Corporation in the U.S. and/or other countries.

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

© 2017 Intel Corporation

LEGAL NOTICES & DISCLAIMERS

Configurations:

- 2.7X performance for big data decision support workload (MKL)

SUT		1 -- BDX (E5-2699v4) + old S/W Stack A	2 -- SKX (Platinum 8160) + old S/W Stack A	3 -- BDX (E5-2699v4) + new, optimized S/W Stack B	4 -- SKX (Platinum 8160) + new, optimized S/W Stack B
H/W	# of Nodes	4	4	4	4
	# of Master Nodes	1	1	1	1
	# of Worker Nodes	3	3	3	3
	CPU	Xeon E5-2699v4	Xeon Platinum 8160	Xeon E5-2699v4	Xeon Platinum 8160
	# of CPU vCores	88	96	88	96
	# of Cores per CPU	22	24	22	24
	# of Sockets	2	2	2	2
	Hyperthread	2	2	2	2
	Clock	2.2 GHz (3.60 GHz Max)	2.1 GHz (3.70 GHz Max)	2.2 GHz (3.60 GHz Max)	2.1 GHz (3.70 GHz Max)
	Cache	55MB Smart Cache	33 MB L3 Cache	55MB Smart Cache	33 MB L3 Cache
	Memory	384GB DDR4 (12 x 32GB, 2666 MT/s)	384GB DDR4 (24 x 16GB, 2133 MT/s)	384GB DDR4 (12 x 32GB, 2666 MT/s)	384GB DDR4 (24 x 16GB, 2133 MT/s)
	Data Storage	800GB * 8 SATA SSD	800GB * 8 SATA SSD	800GB * 8 SATA SSD	800GB * 8 SATA SSD
	Network	10Gbe	10Gbe	10Gbe	10Gbe
BIOS	BIOS version	SE5C610.86B.01.01.0018.0720 20161249	SE5C620.86B.01.00.0470.0407 20170855	SE5C610.86B.01.01.0018.0720 20161249	SE5C620.86B.01.00.0470.0407 20170855
S/W	Hive	Apache Hive 2.0.0	Apache Hive 2.0.0	Upstream, commit id: 3330403a7dea206c85828dfd5319a7ae0d9bcef5	Upstream, commit id: 3330403a7dea206c85828dfd5319a7ae0d9bcef5
	Spark	Apache Spark 1.6.3	Apache Spark 1.6.3	Apache Spark 2.0.2	Apache Spark 2.0.2
	Hadoop	Apache Hadoop 2.7.3	Apache Hadoop 2.7.3	Apache Hadoop 2.7.3	Apache Hadoop 2.7.3
	OS Version	CentOS 7.3	CentOS 7.3	CentOS 7.3	CentOS 7.3
	JDK	Oracle JDK 1.8.0_121	Oracle JDK 1.8.0_121	Oracle JDK 1.8.0_121	Oracle JDK 1.8.0_121

LEGAL DISCLAIMERS

- Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.
- No computer system can be absolutely secure.
- 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>.

Configurations:

- 4.3X for Spark MLlib thru Intel Math Kernel Library (MKL)
 - Spark-Perf (same for before and after): 9 nodes each with Intel® Xeon® processor E5-2697A v4 @ 2.60GHz * 2 (16 cores, 32 threads); 256 GB ; 10x SSDs; 10Gbps NIC
- 19x for HDFS Erasure Coding in micro workload (RawErasureCoderBenchmark) and 1.25x in Terasort, plus 50+% storage capacity saving and higher failure tolerance level.
 - RawErasureCoderBenchmark (same for before and after): single node with Intel® Xeon® processor E5-2699 v4 @ 2.20GHz *2 (22 cores, 44 threads); 256GB; 8x HDDs; 10Gbps NIC
 - Terasort (same for before and after): 10 nodes each with Intel® Xeon® processor E5-2699 v4 @ 2.20GHz *2 (22 cores, 44 threads); 256GB; 8x HDDs; 10Gbps NIC
- 5.6x for HBase off heaping read in micro workload (PE) and 1.3x in real Alibaba production workload
 - PE (same for before and after): Intel® Xeon® Processor X5670 @ 2.93Hz *2 (6 cores, 12 threads); RAM: 150 GB; 1Gbps NIC
 - Alibaba (same for before and after): 400 nodes cluster with Intel® Xeon® processors
- 1.22x Spark Shuffle File Encryption performance for TeraSort and 1.28x for BigBench
 - Terasort (same for before and after): Single node with Intel® Xeon® Processor E5-2699 v3 @ 2.30GHz *2 (18 cores, 36 threads); 128GB; 4x SSD; 10Gbps NIC
 - BigBench (same for before and after): 6 nodes each with Intel® Xeon® Processor E5-2699 v3 @ 2.30GHz *2 (18 cores, 36 threads); 256GB; 1x SSD; 8x SATA HDD 3TB, 10Gbps NIC
- 1.35X Spark Shuffle RPC encryption performance for TeraSort and 1.18x for BigBench
 - Terasort (same for before and after): 3 nodes each with Intel® Xeon® Processor E5-2699 v3 @ 2.30GHz *2 (18 cores, 36 threads); 128GB; 4x SSD; 10Gbps NIC
 - BigBench (same for before and after): 5 nodes. 1x head node: Intel® Xeon® Processor E5-2699 v3 @ 2.30GHz *2 (18 cores, 36 threads); 384GB; 1x SSD; 8x SATA HDD 3TB, 10Gbps NIC. 4x worker nodes: each with Intel® Xeon® processor E5-2699 v4 @ 2.20GHz *2 (22 cores, 44 threads); 384GB; 1x SSD; 8x SATA HDD 3TB, 10Gbps NIC.

Intel, the Intel logo, Xeon, Xeon phi, Lake Crest, etc. are trademarks of Intel Corporation in the U.S. and/or other countries.

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

© 2017 Intel Corporation

LEGAL NOTICES & DISCLAIMERS

Configurations:

- 2.7X performance for big data decision support workload (MKL)

Decision Support Workload Configuration Summary

*** The brief query descriptions are derived from TPC-DS specification**

Query Number	Brief Description *
Q19	Select the top revenue generating products bought by out of zip code customers for a given year, month and manager.
Q42	For each item and a specific year and month calculate the sum of the extended sales price of store transactions
Q43	Report the sum of all sales from Sunday to Saturday for stores in a given data range by stores.
Q52	Report the total of extended sales price for all items of a specific brand in a specific year and month.
Q55	For a given year, month and store manager calculate the total store sales of any combination all brands
Q63	For a given year calculate the monthly sales of items of specific categories, classes and brands that were sold in stores and group the results by store manager. Additionally, for every month and manager print the yearly average sales of those items.
Q68	Compute the per customer extended sales price, extended list price and extended tax for "out of town" shoppers buying from stores located in two cities in the first two days of each month of three consecutive years. Only consider customers with specific dependent and vehicle counts.
Q73	Count the number of customers with specific buy potentials and whose dependent count to vehicle count ratio is larger than 1 and who in three consecutive years bought in stores located in 4 counties between 1 and 5 items in one purchase. Only purchases in the first 2 days of the months are considered.
Q98	Report on items sold in a given 30 day period, belonging to the specified category.

Parallel Stream Number	Executed Queries
Stream #1	Q19, Q19, Q19, Q19, Q19
Stream #2	Q42, Q42, Q42, Q42, Q42
Stream #3	Q43, Q43, Q43, Q43, Q43
Stream #4	Q52, Q52, Q52, Q52, Q52
Stream #5	Q55, Q55, Q55, Q55, Q55
Stream #6	Q63, Q63, Q63, Q63, Q63
Stream #7	Q68, Q68, Q68, Q68, Q68
Stream #8	Q73, Q73, Q73, Q73, Q73
Stream #9	Q98, Q98, Q98, Q98, Q98



Спасибо!

nikolay.mester@intel.com