

HPC CHANGING LANDSCAPE

INTEL COMPUTE & MEMORY TECHNOLOGY

EMiT at Barcelona Computing Center

June 3rd, 2016

Intel EMEA

LEGAL DISCLAIMER

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.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

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.

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.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

Intel, Intel Xeon, Intel Xeon Phi™, Intel® Atom™ are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States or other countries.

Copyright © 2016, Intel Corporation

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

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase. The cost reduction scenarios described in this document are intended to enable you to get a better understanding of how the purchase of a given Intel product, combined with a number of situation-specific variables, might affect your future cost and savings. Nothing in this document should be interpreted as either a promise of or contract for a given level of costs.

Intel® Advanced Vector Extensions (Intel® AVX)* are designed to achieve higher throughput to certain integer and floating point operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you should consult your system manufacturer for more information.

*Intel® Advanced Vector Extensions refers to Intel® AVX, Intel® AVX2 or Intel® AVX-512. For more information on Intel® Turbo Boost Technology 2.0, visit <http://www.intel.com/go/turbo>
All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice

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

Copyright © 2016 Intel Corporation. All rights reserved.



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 © 2016, Intel Corporation.

RISK FACTORS

The above statements and any others in this document that refer to plans and expectations for the second quarter, the year and the future are forward looking statements that involve a number of risks and uncertainties. Words such as “anticipates,” “expects,” “intends,” “plans,” “believes,” “seeks,” “estimates,” “may,” “will,” “should” and their variations identify forward-looking statements. Statements that refer to or are based on projections, uncertain events or assumptions also identify forward-looking statements. Many factors could affect Intel’s actual results, and variances from Intel’s current expectations regarding such factors could cause actual results to differ materially from those expressed in these forward-looking statements. Intel presently considers the following to be important factors that could cause actual results to differ materially from the company’s expectations. Demand for Intel’s products is highly variable and, in recent years, Intel has experienced declining orders in the traditional PC market segment. Demand could be different from Intel’s expectations due to factors including changes in business and economic conditions; consumer confidence or income levels; customer acceptance of Intel’s and competitors’ products; competitive and pricing pressures, including actions taken by competitors; supply constraints and other disruptions affecting customers; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Intel operates in highly competitive industries and its operations have high costs that are either fixed or difficult to reduce in the short term. Intel’s gross margin percentage could vary significantly from expectations based on capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; changes in revenue levels; segment product mix; the timing and execution of the manufacturing ramp and associated costs; excess or obsolete inventory; changes in unit costs; defects or disruptions in the supply of materials or resources; and product manufacturing quality/yields. Variations in gross margin may also be caused by the timing of Intel product introductions and related expenses, including marketing expenses, and Intel’s ability to respond quickly to technological developments and to introduce new products or incorporate new features into existing products, which may result in restructuring and asset impairment charges. Intel’s results could be affected by adverse economic, social, political and physical/infrastructure conditions in countries where Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Intel’s results could be affected by the timing of closing of acquisitions, divestitures and other significant transactions. Intel’s results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust, disclosure and other issues, such as the litigation and regulatory matters described in Intel’s SEC filings. An unfavorable ruling could include monetary damages or an injunction prohibiting Intel from manufacturing or selling one or more products, precluding particular business practices, impacting Intel’s ability to design its products, or requiring other remedies such as compulsory licensing of intellectual property. A detailed discussion of these and other factors that could affect Intel’s results is included in Intel’s SEC filings, including the company’s most recent reports on Form 10-Q, Form 10-K and earnings release.

Rev. 4/15/14



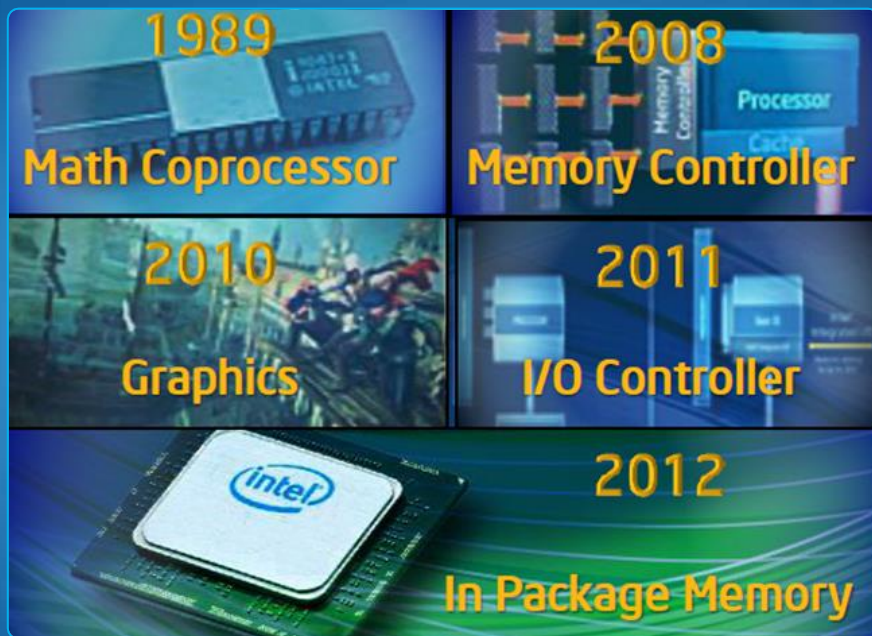
Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*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

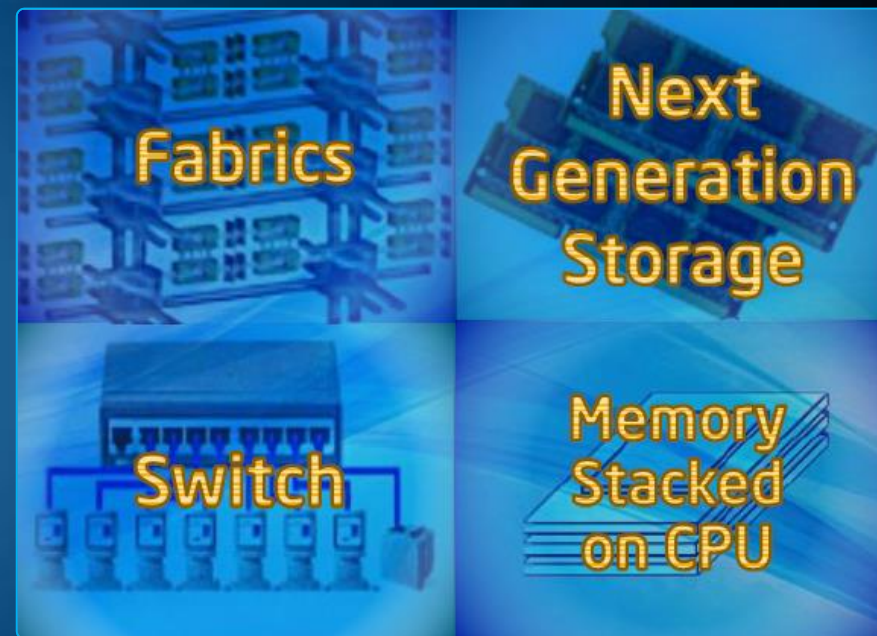
© 2016, Intel Corporation.

DRIVING INNOVATION AND INTEGRATION

Enabled by Leading Edge Process Technologies



Integrated Today



Possible Tomorrow**

SYSTEM LEVEL BENEFITS IN COST, POWER, DENSITY, SCALABILITY & PERFORMANCE

**For illustration only, not drawn to scale. Potential future options are forecasts and subject to change without notice.

“FUTUREWORLD”

HIGH-PERFORMANCE COMPUTING NEW COMPUTE PARADIGMS

PAST

Single Core
CPU



PRESENT

Multi-Core
Many-Core

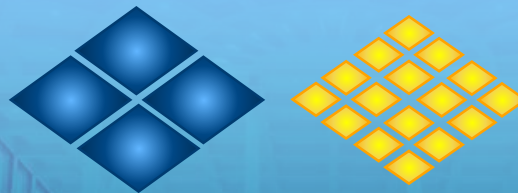


FUTURE ?

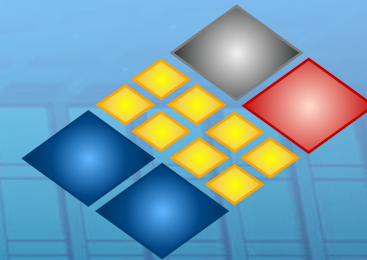
Mix of Cores, integrating
FPGA, Accelerators, ...**



transistor constraint



power constraint

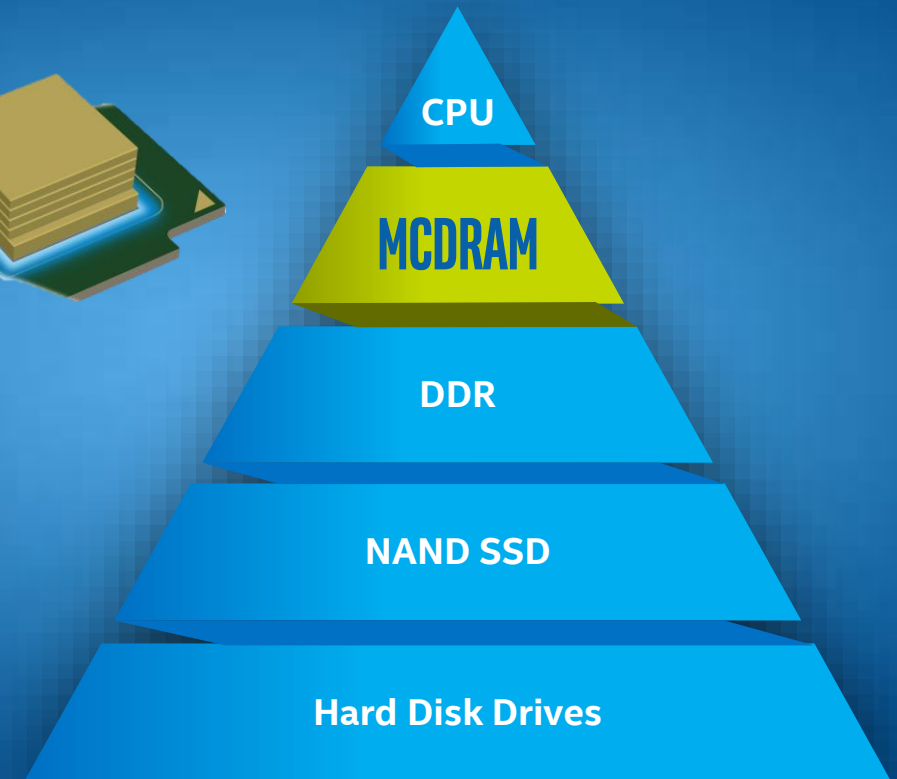
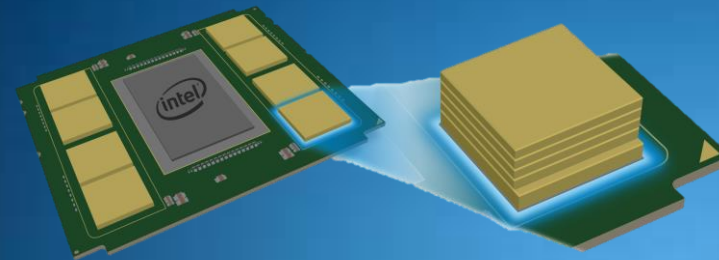


even more power constraint

**For illustration only, not drawn to scale. Potential future options are forecasts and no indication of product plans.

HIGHLY PARALLEL PROCESSING

KNIGHTS LANDING NEXT-GEN INTEL® XEON PHI™ PROCESSOR



For illustration only. All dates, product descriptions, features, availability, and plans are forecasts and subject to change without notice.

A BW-HUNGRY FUNCTION

“STREAM” addition

```
results_t run_bench(data_t *A, data_t *B, data_t *C, const char* id)
{
    // ... begin timing
    ///!! START WORKLOAD
    for (size_t i = 0; i < NUM_ITERATIONS; ++i) {
        #pragma omp parallel for simd
        for (size_t i = 0; i < NUM_ELEMENTS; ++i) {
            C[i] = A[i] + B[i];
        }
    }
    ///!! END WORKLOAD
    // ... end timing
}
```

One function will operate on different memory operands:

- one input set in regular DDR
- other input set in MCDRAM

The input sets contain the same data.

A BW-HUNGRY CODE (MAIN0)

```
int main()
```

```
{
```

```
    data_t *A_reg = (data_t*)      malloc(sizeof(data_t) * NUM_ELEMENTS);  
    data_t *B_reg = (data_t*)      malloc(sizeof(data_t) * NUM_ELEMENTS);  
    data_t *C_reg = (data_t*)      malloc(sizeof(data_t) * NUM_ELEMENTS);
```

ALLOCATING
REGULAR MEM

```
    data_t *A_hbw = (data_t*) hbw_malloc(sizeof(data_t) * NUM_ELEMENTS);  
    data_t *B_hbw = (data_t*) hbw_malloc(sizeof(data_t) * NUM_ELEMENTS);  
    data_t *C_hbw = (data_t*) hbw_malloc(sizeof(data_t) * NUM_ELEMENTS);
```

ALLOCATING
HIGH-BW MEM

```
    init(A_reg, B_reg, C_reg, A_hbw, B_hbw, C_hbw);
```

```
    auto res_reg = run_bench(A_reg, B_reg, C_reg, "[REG]");  
    auto res_hbw = run_bench(A_hbw, B_hbw, C_hbw, "[HBW]");
```

RUNNING
BENCHMARK

```
    std::cout << "Computations happened " << res_reg/res_hbw  
               << "x times faster in high-bandwidth memory.\n";
```

RESULTS
OUTPUT

```
    free(A_reg);  
    free(B_reg);  
    free(C_reg);
```

FREEING
REGULAR MEM

```
    hbw_free(A_hbw);  
    hbw_free(B_hbw);  
    hbw_free(C_hbw);
```

FREEING HIGH-
BW MEM

```
}
```

A BW-HUNGRY WORKLOAD RESULTS (KNIGHTS LANDING)

NOTE: those are *some* units, only here to show *relative* performance difference.

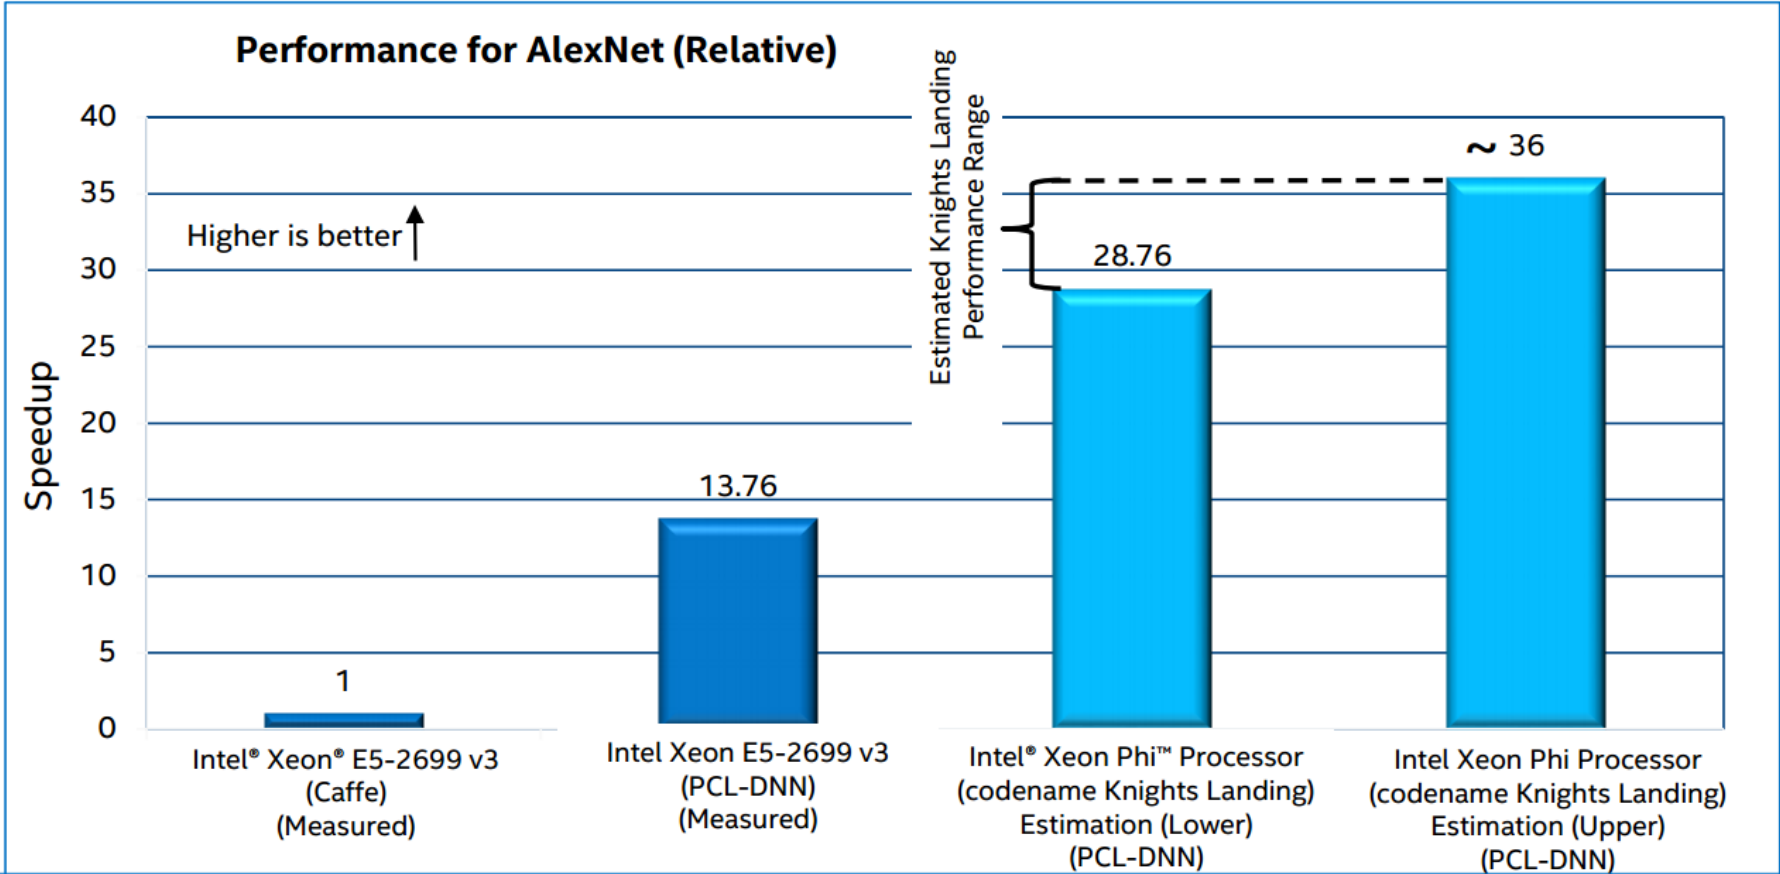
```
$ ./run.sh  
[REG] Calculations took 15376.5 [units].  
[HBW] Calculations took 3056.19 [units].  
Computations happened 5.03125x times  
faster in high-bandwidth memory.
```

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.
Notice Revision #20110804

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

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/performance>.

DEEP LEARNING: SINGLE NODE PERFORMANCE



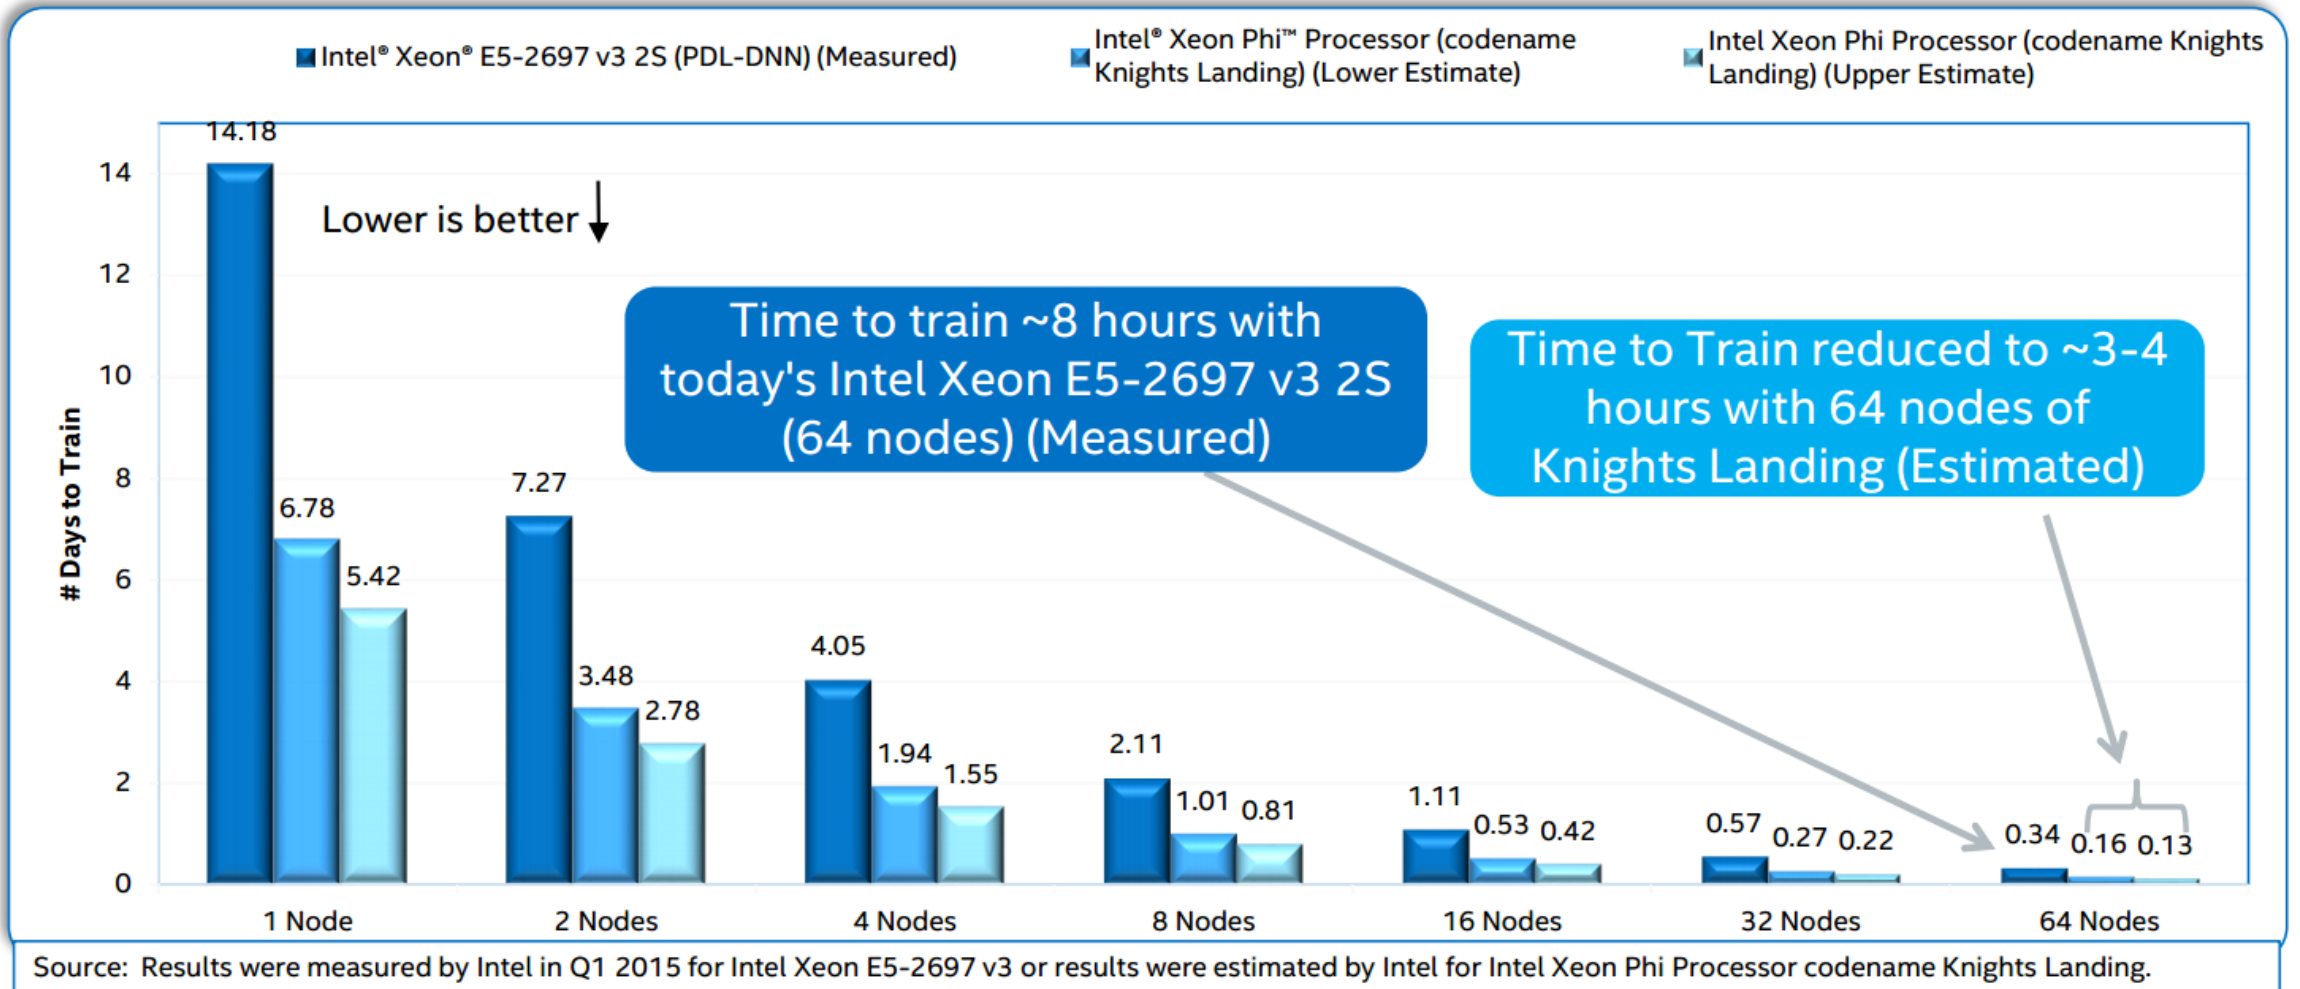
Source: Results were measured by Intel in Q1 2015 for Intel Xeon E5-2697 v3 or results were estimated by Intel for Intel® Xeon Phi™ processor codename Knights Landing.

Intel Xeon processor E5-2699v3 2S measured: 8 x 8GB DDR4 2133, AlexNet on randomly generated inputs (32,000 images) Intel® C Compiler: 15.0.2, OS: CentOS 7.0.1406

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. Configuration Details: See System Configurations slide Results are for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. For more information go to <http://www.intel.com/performance>



DEEP LEARNING: TIME TO TRAIN SCALABILITY



2 x Intel Xeon processor E5-2697 v3 @ 2.60GHz, DDR4, 2133GHz, 64 GB; RHEL 6.5, Network interface: InfiniBand® FDR, Intel® C Compiler 15.0.2 with Intel® Advanced Vector Extensions 2 (Intel® AVX2), OpenMP®, Intel® MPI library, DNN Library: PCL-DNN Library, PCL-DNN Harness & PCL-CML Library, randomly generated inputs (64000 images), training on 1.3M images of ImageNet-1k

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> Results are for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

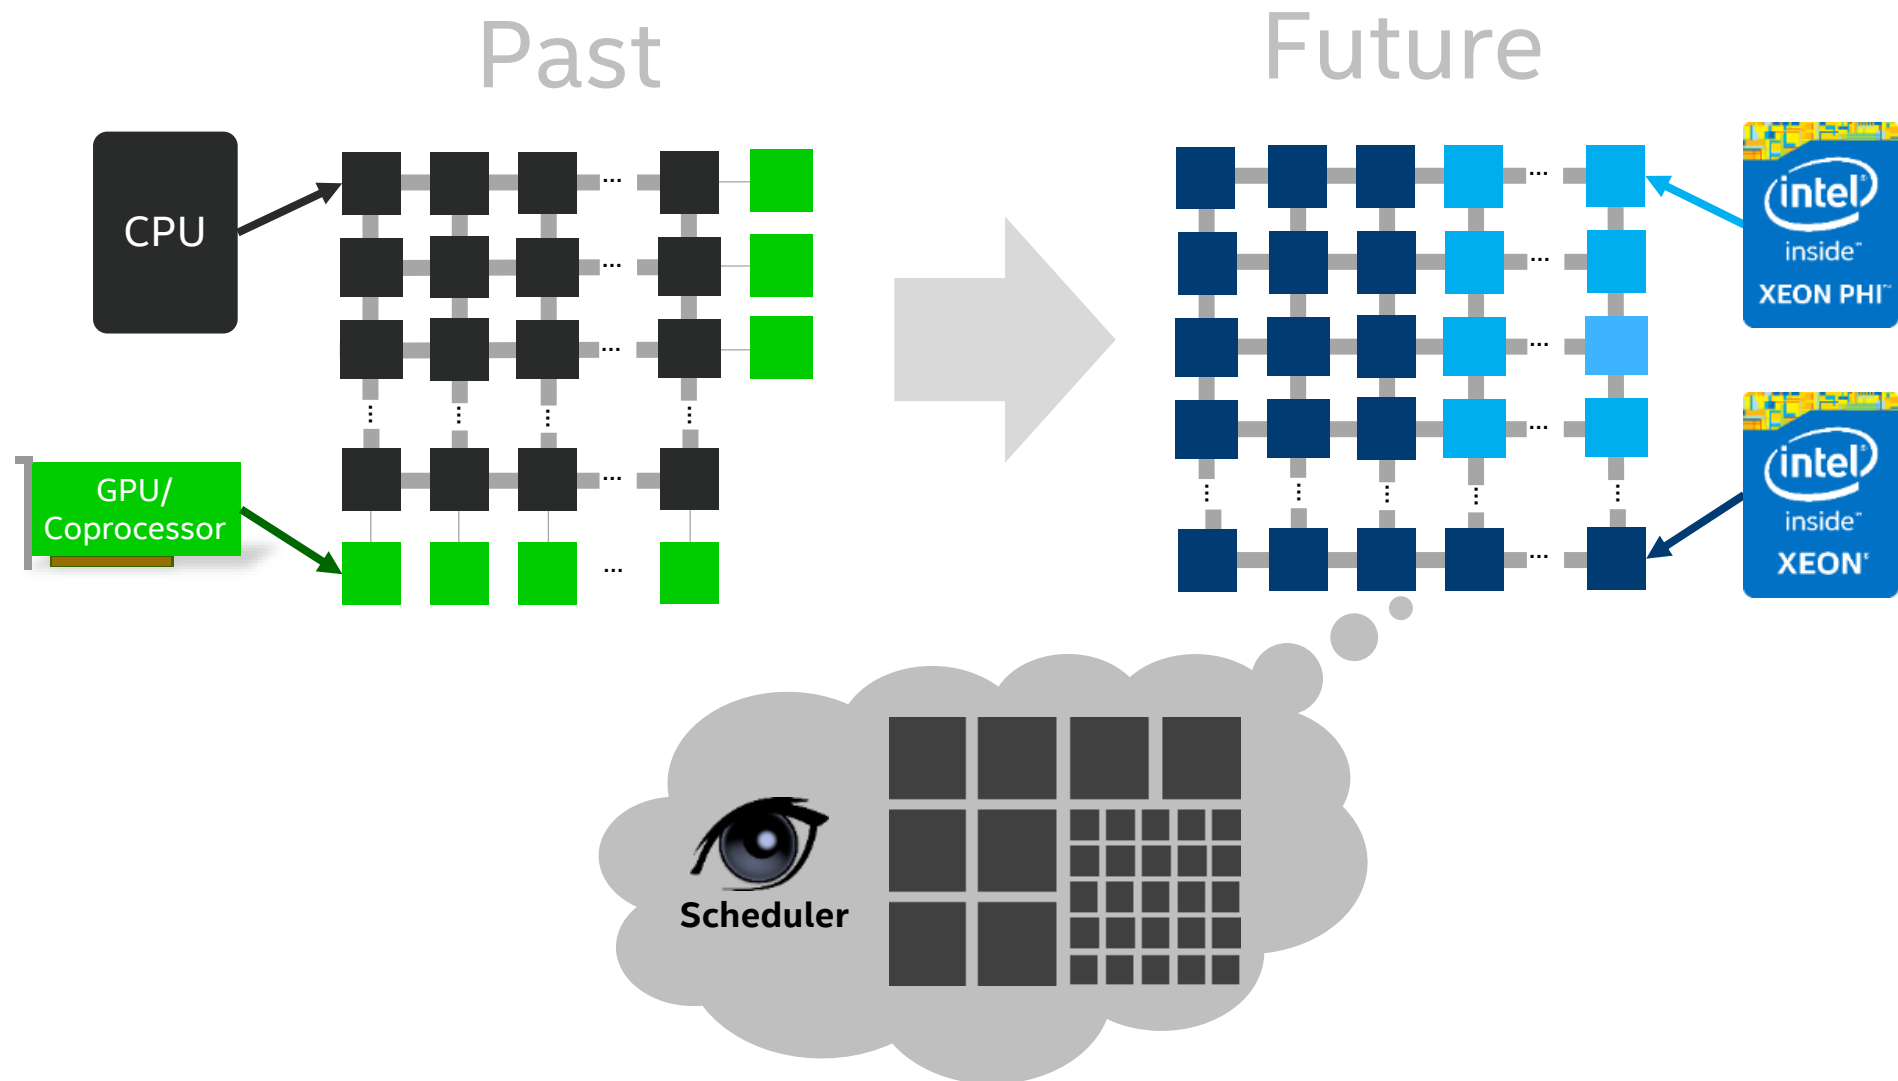


Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*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

© 2016, Intel Corporation.

FUTURE CLUSTERS...



FPGA: PERFORMANCE & ENERGY EFFICIENCY EXAMPLE

Application performs post processing of **3D textures for analyzing rock samples**

- Code Labels 128x128x128 pixel textures in memory and calculates overlap
- Large data size
 - >20GB input file
 - >8M textures

Kernel code:

```
for(iz=-tw/2; iz < tw/2 ; iz++) {  
    for(iy=-tw/2; iy < tw/2 ; iy++) {  
        for(ix=-tw/2; ix < tw/2 ; ix++) {  
            /* Copy texture into buffer */  
            buf[(iz+tw/2)*tw*tw + (iy+tw/2)*tw + ix + tw/2] =  
                image[(z+iz)*dimx*dimy + (y+iy)*dimx + (x+ix)];  
  
            /* Label the texture */  
            label_texture(buf, tw, tw, tw, val_min, val_max, nl, nb,  
                bodies, matrix_a, matrix_a_b, matrix_a_c, matrix_a_d);  
        }  
    }  
}
```

Offload kernel to an accelerator to meet the goal of labeling **8 million textures in 30 minutes**.



438x K80 cards
4 racks
116KW



162x Stratix V A7 cards
2 racks
13.7KW

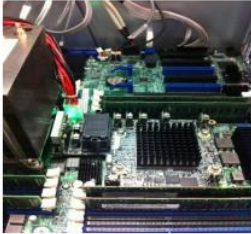
LESS ENERGY
SAME PERFORMANCE

8.5X

MORE PERFORMANCE
SAME ENERGY

Source: M. Hilgeman, Dell Accelerating Understanding Summit 2015

SHIPPING TODAY: IN 2ND XEON[®] SOCKET FPGA



Pactron FPGA Accelerated Computing Solutions

“Intel[®] Xeon + Altera FPGA”
Software Development Platforms

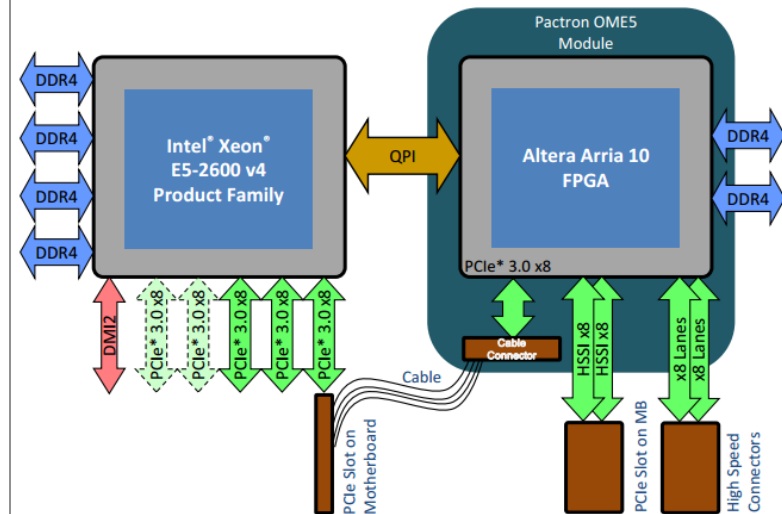
10/21/2015

© 2015 Pactron HJPC Corporation

6

Pactron's Grantley “HSX/BSX” SDP Platforms

Software Development for Accelerating Workloads using Xeon and coherently attached FPGA in-socket



Processor	Intel® Xeon® E5-26xx v4 Processor
FPGA Module	Pactron OME5
QPI Speed	6.4 GT/s full width (target 8.0 GT/s at full width)
Memory to FPGA Module	2 channels of DDR4 (up to 64 GB)
Expansion connector to FPGA Module	Two x8 HSSI lanes (PCIe electrical) - maybe used for direct I/O e.g. Ethernet
Control Port	PCIe 3.0 x8 connected to cable connector on module. For discovery/enumeration only, not for bulk data.
Features	Configuration Agent, Caching Agent,, (optional) Memory Controller
Software	Accelerator Abstraction Layer (AAL) runtime, drivers, sample applications

**** Available Dec 2015 ****

10/21/2015

© 2015 Pactron HJPC Corporation

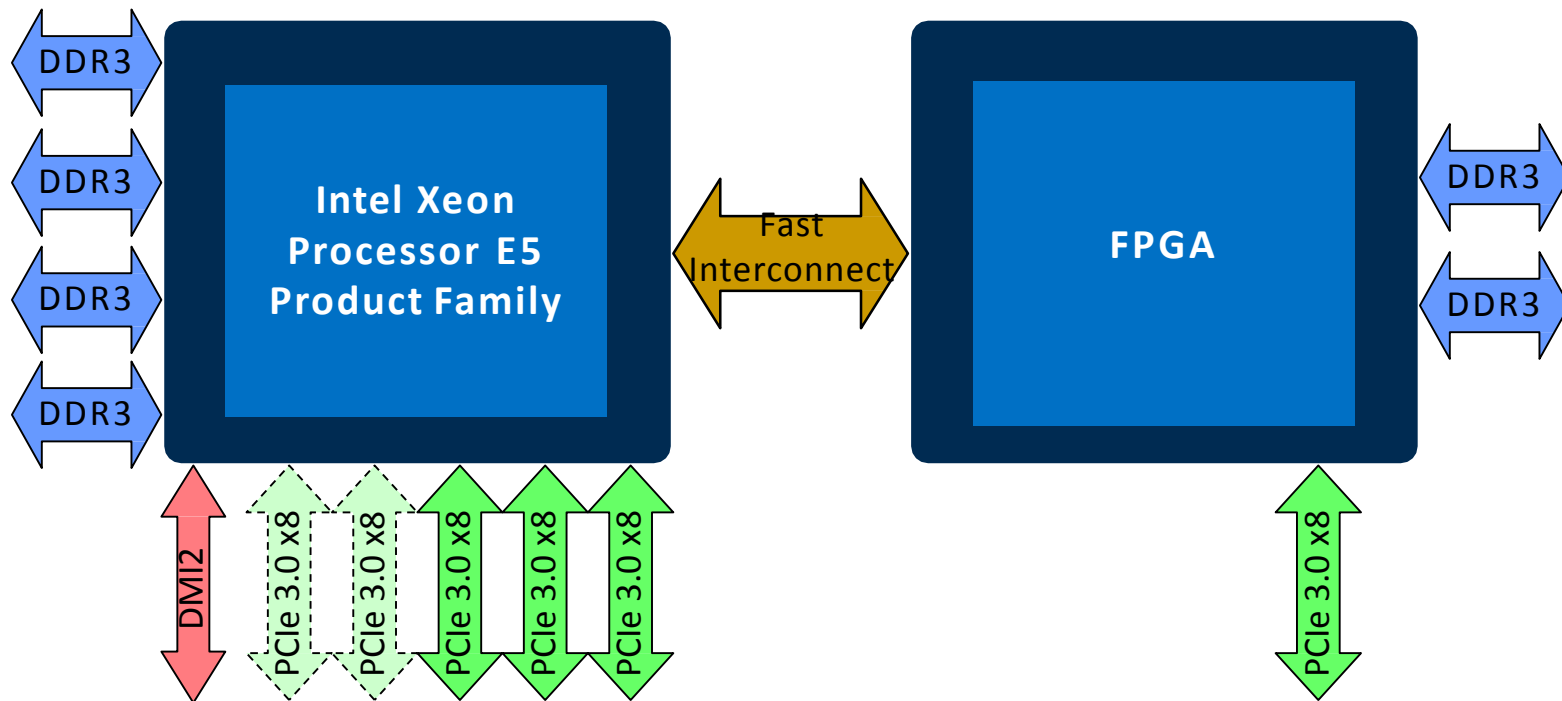
10

Pactron: <http://www.pactroninc.com/> → <http://www.pactroninc.com/qpi/>

INTEL® XEON® E5 + FIELD PROGRAMMABLE GATE ARRAY

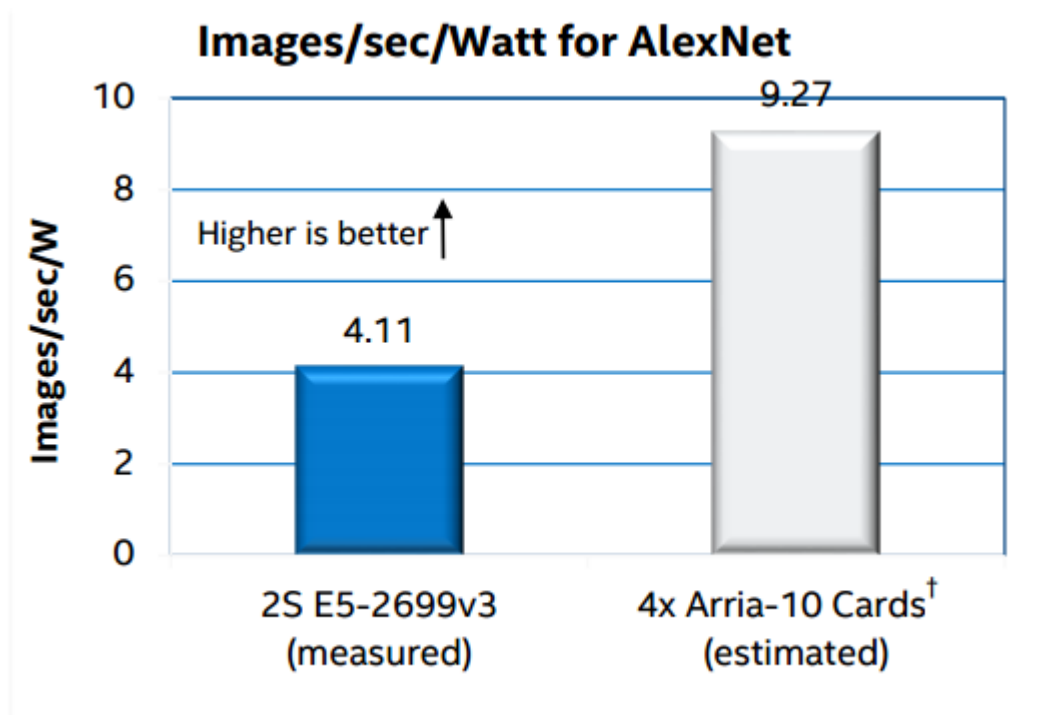
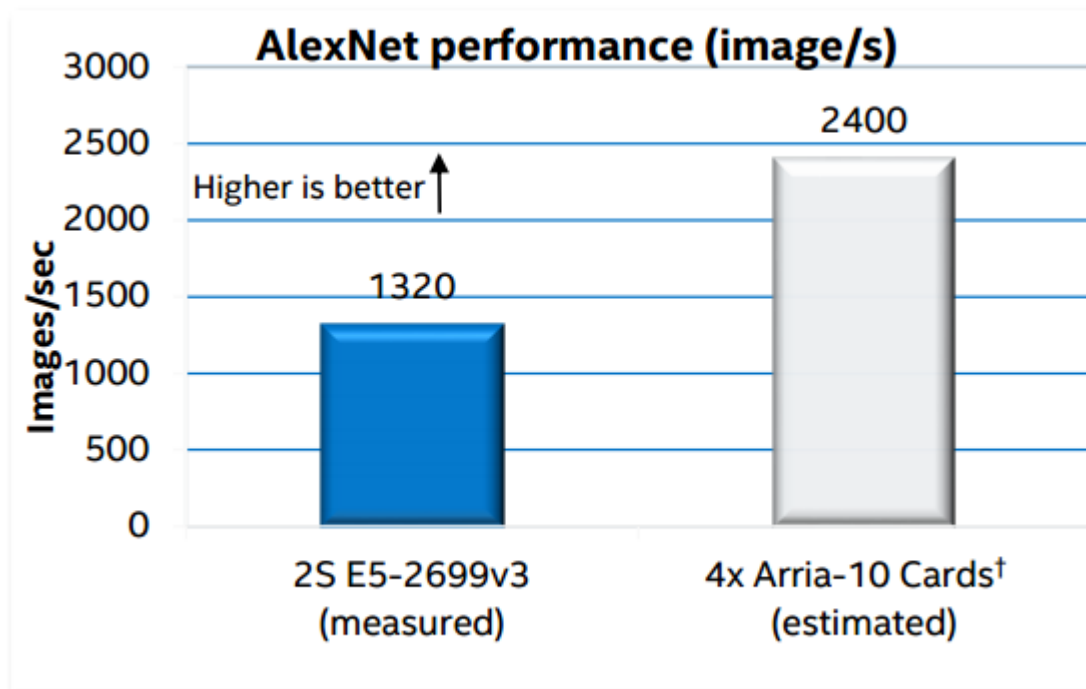
SOFTWARE DEVELOPMENT PLATFORM (SDP) SHIPPING TODAY

Software Development for Accelerating Workloads using Intel® Xeon® processors and coherently attached FPGA in-socket



Processor	Intel Xeon Processor E5
FPGA Module	Altera® Stratix® V
Fast Interconnect Speed	6.4 GT/s full width (target 8.0 GT/s at full width)
Memory to FPGA Module	2 channels of DDR3 (up to 64 GB)
Expansion connector to FPGA Module	PCI Express® (PCIe) 3.0 x8 lanes - maybe used for direct I/O e.g. Ethernet
Features	Configuration Agent, Caching Agent, (optional) Memory Controller
Software	Accelerator Abstraction Layer (AAL) runtime, drivers, sample applications

MACHINE LEARNING ACCELERATION WITH INTEL® XEON®+FPGA



Power-performance of CNN classification boosted up to 2.2X

<http://www.intel.com/content/www/us/en/benchmarks/intel-product-performance.html>

Source: Intel Measured (E5-2699v3 results); Altera* Estimated (4x Arria* 10 results) †2S E5-2699v3 + 4x GX1150PCIe cards. Most computations executed on Arria-10 FPGA's, 2S E5-2699v3 host assumed to be near idle, doing misc. networking/housekeeping functions. Arria-10 results estimated by Altera with Altera custom classification network. 2x E5-2699v3 power estimated @ 139W while doing "housekeeping" for GX1150 cards based on Intel measured microbenchmark. In order to sustain ~2400 img/s we need a I/O bandwidth of ~500 MB/s, which can be supported by a 10GigE link and software stack 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. Configuration Details: See System Configurations slide For more information go to <http://www.intel.com/performance> 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.



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

© 2016, Intel Corporation.

BURROWS-WHEELER ALIGNER* (BWA*) ON FPGA

BWA-MEM is made of three main kernels

SMEM generation

Find likely mapping locations (or seeds) on the reference genome

45% total application performance increase compared to CPU implementation



Seed extension

Chain and extend seeds together with a dynamic programming algorithm



Output generation

Sort and perform general mapping



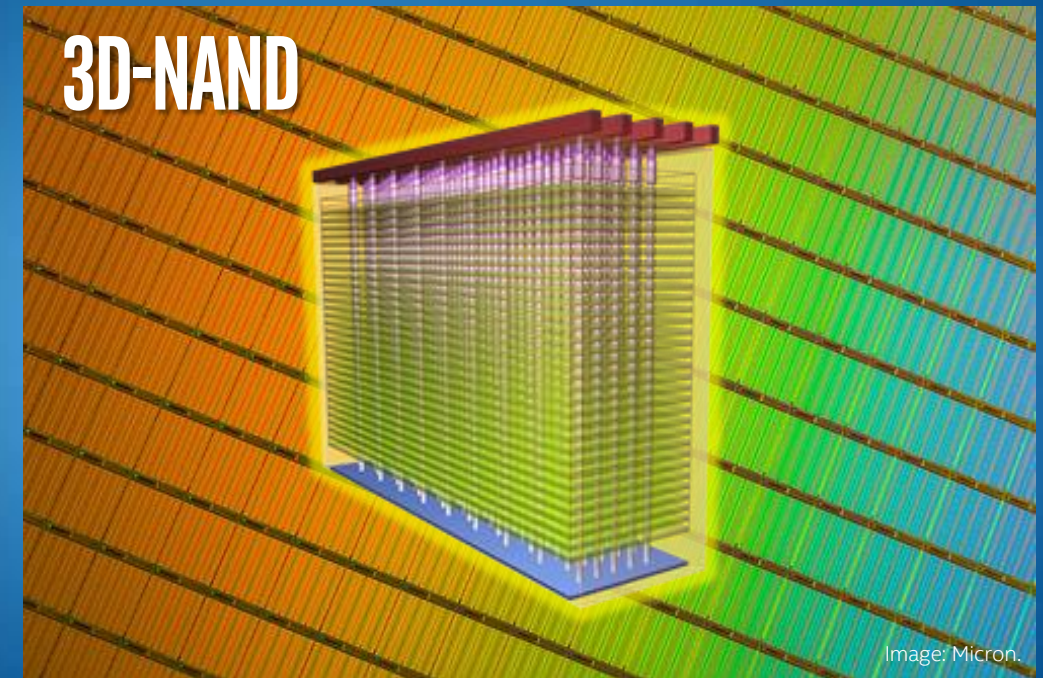
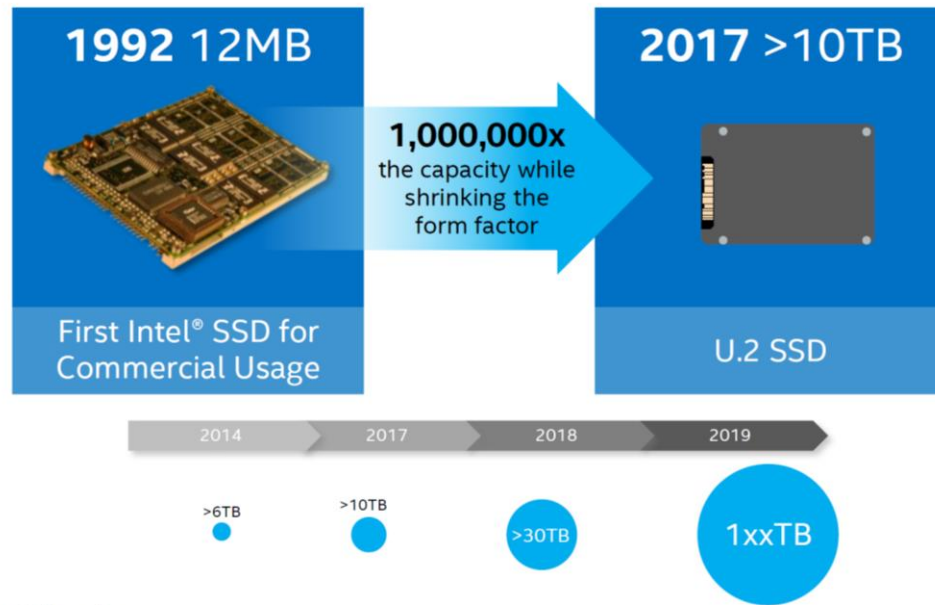
x3 kernel speed-up compared to CPU implementation

An FPGA-Based Systolic Array to Accelerate the BWA-MEM Genomic Mapping Algorithm

<http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7363679>

Ernst Joachim Houtgast, Vlad-Mihai Sima, Koen Bertels and Zaid Al-Ars
Faculty of EEMCS, Delft University of Technology, Delft, The Netherlands

SSD CAPACITY (R)EVOLUTION



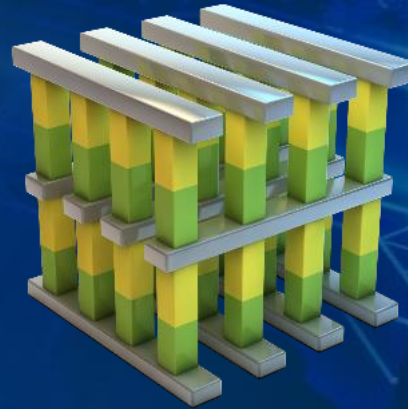
Future options are forecasts and subject to change without notice.

*Other brands and names are the property of their respective owners.



NEW MEMORY AND STORAGE

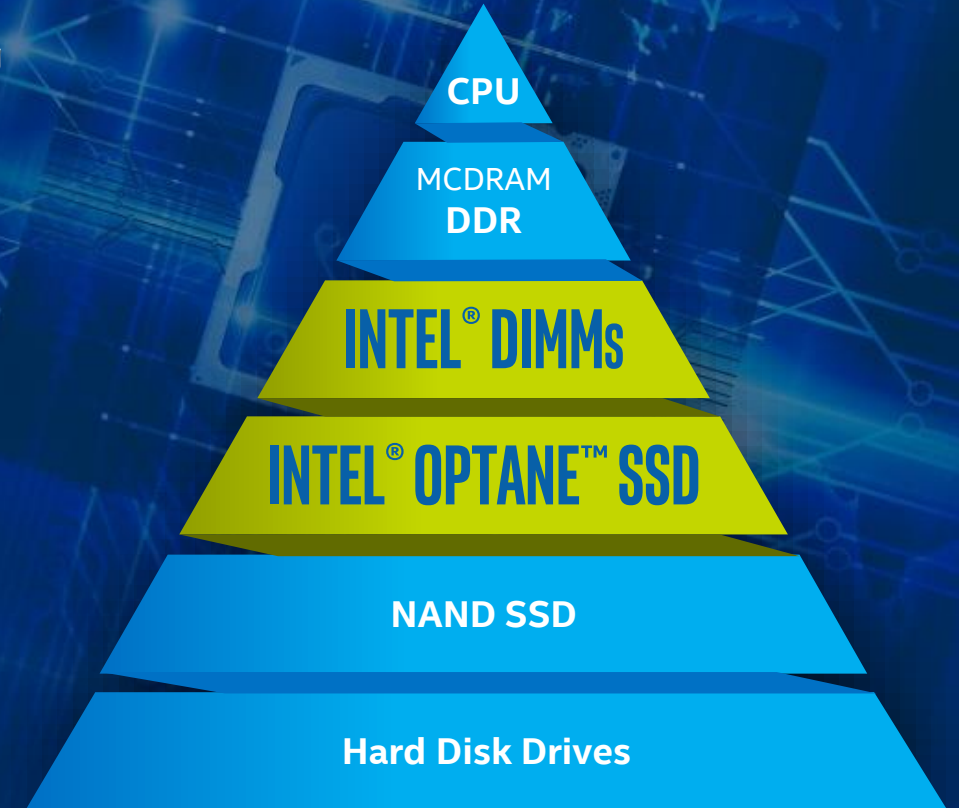
3D XPOINT™ TECHNOLOGY



1000X
FASTER
THAN NAND

1000X
ENDURANCE
OF NAND

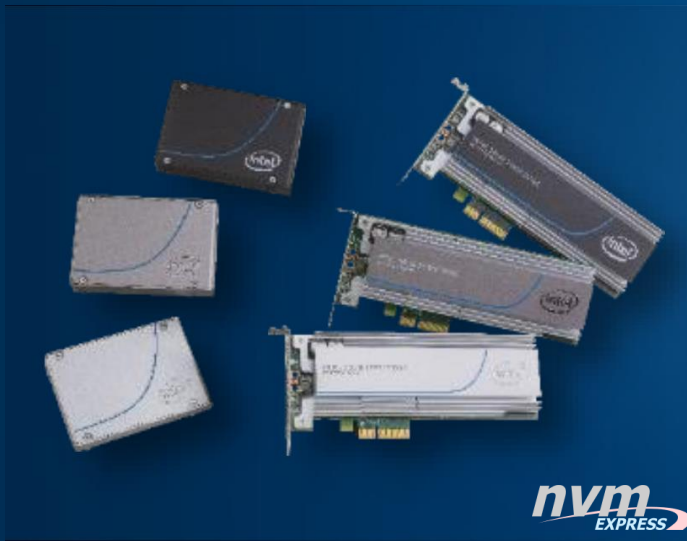
10X
DENSER
THAN DRAM



For illustration only. All dates, product descriptions, features, availability, and plans are forecasts and subject to change without notice.

3D XPOINT™ TECHNOLOGY (NVM)

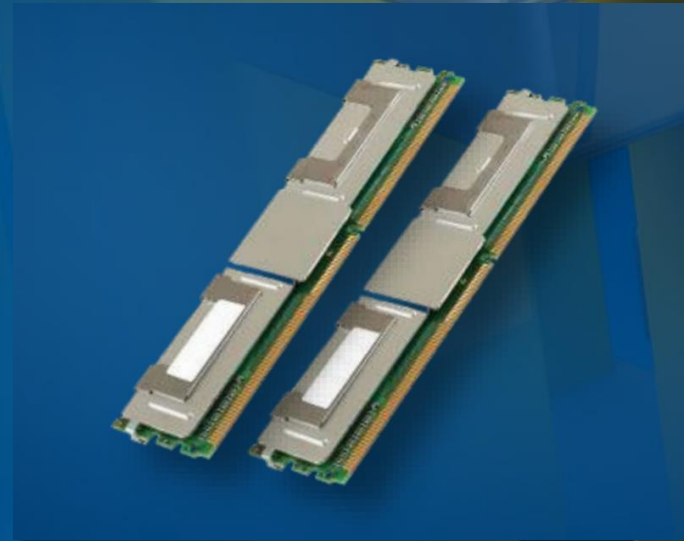
INTEL® OPTANE™ SSD



NEW CLASS OF NON-VOLATILE STORAGE

1000x faster than NAND
1000x endurance of NAND

DIMMs** BASED ON 3D XPOINT™



NEW CLASS OF NON-VOLATILE MEMORY**

4x more memory capacity
1/2 cost of DRAM

**For illustration only, future potential options are forecasts and subject to change without notice.

*Other brands and names are the property of their respective owners.

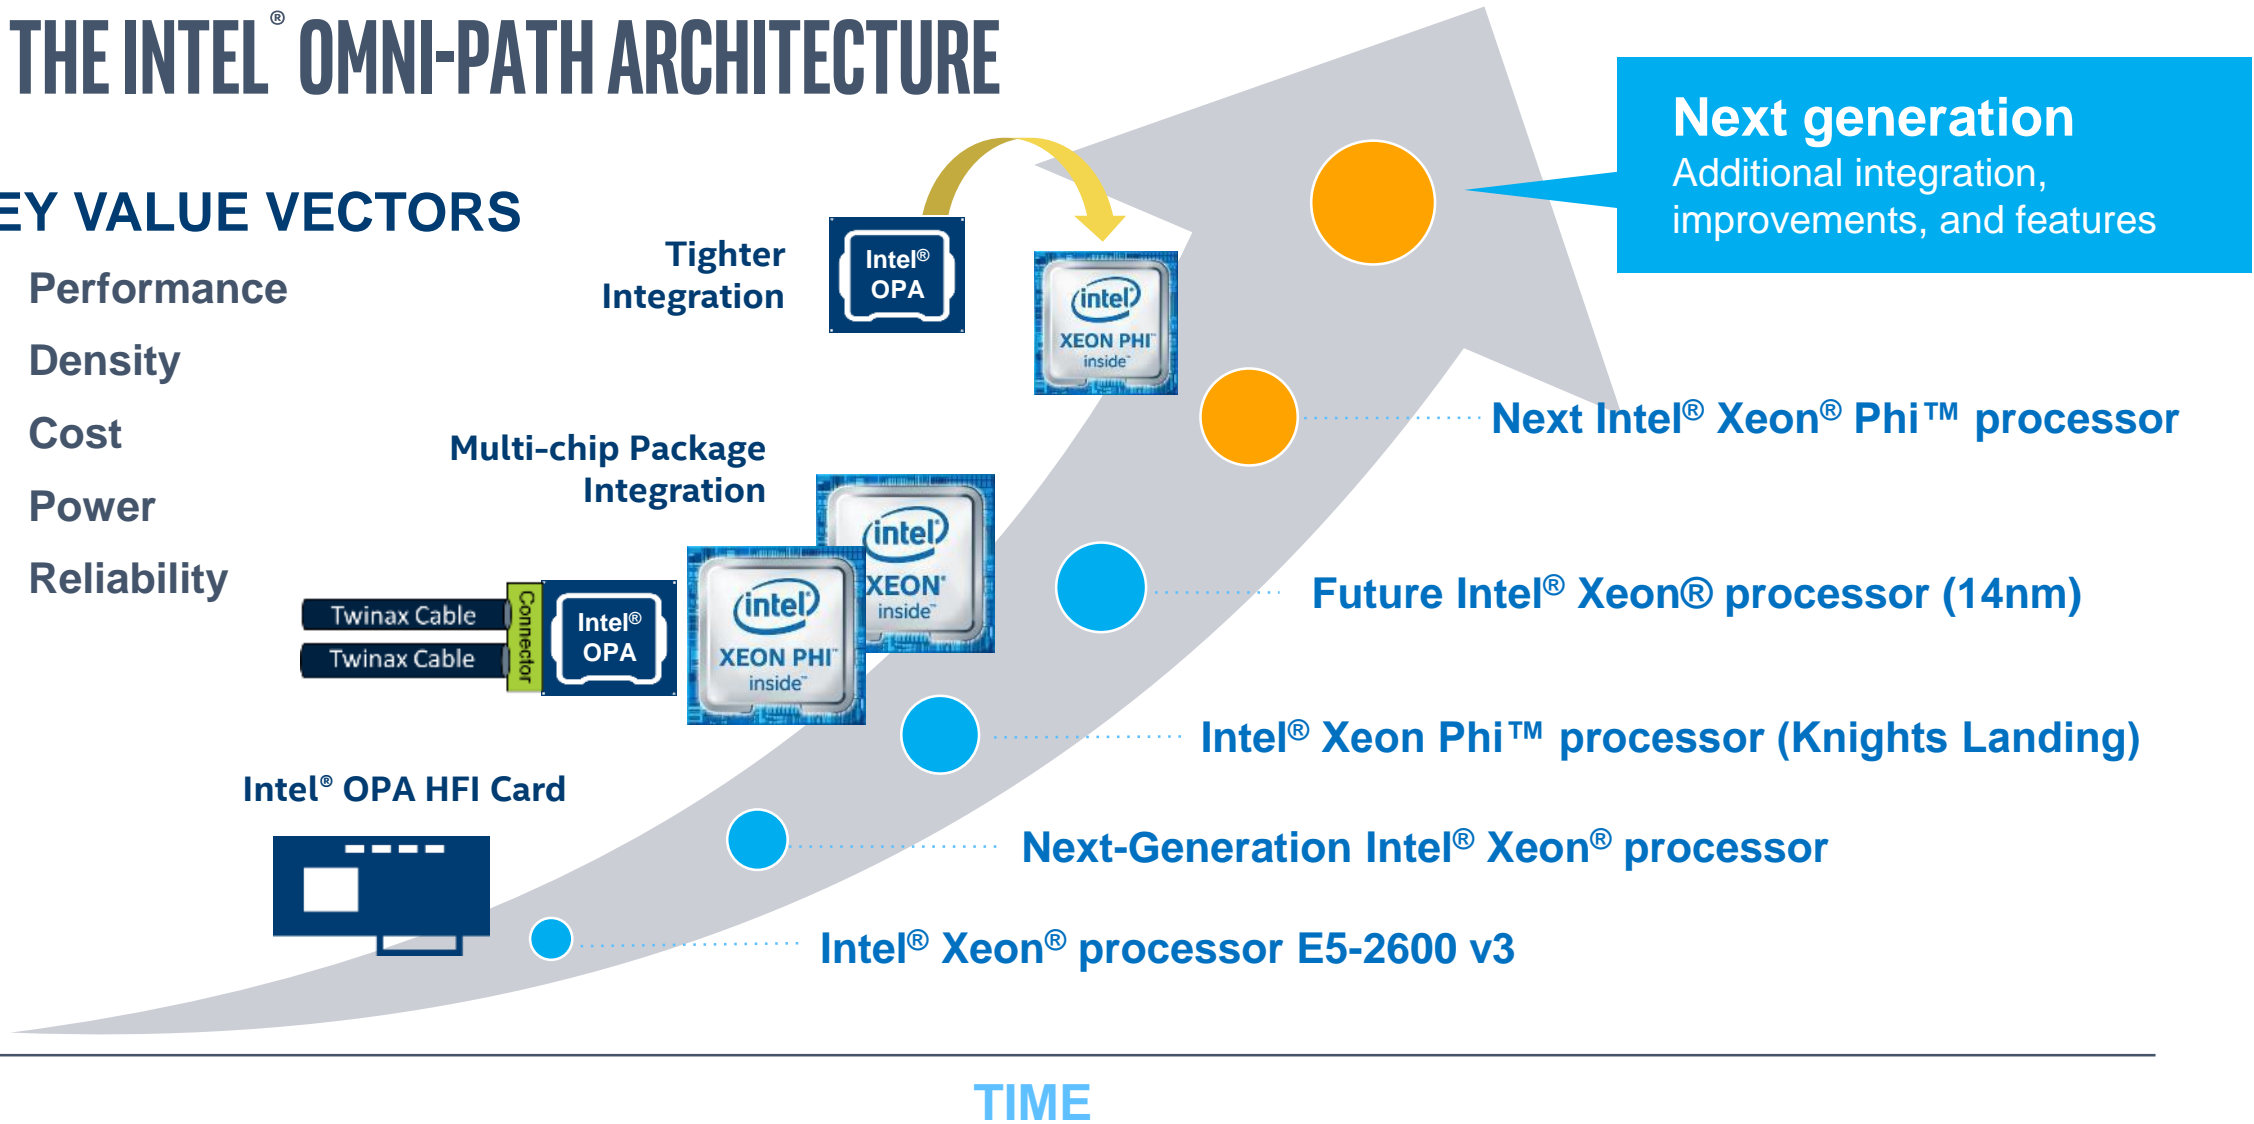
CPU-FABRIC INTEGRATION

WITH THE INTEL® OMNI-PATH ARCHITECTURE

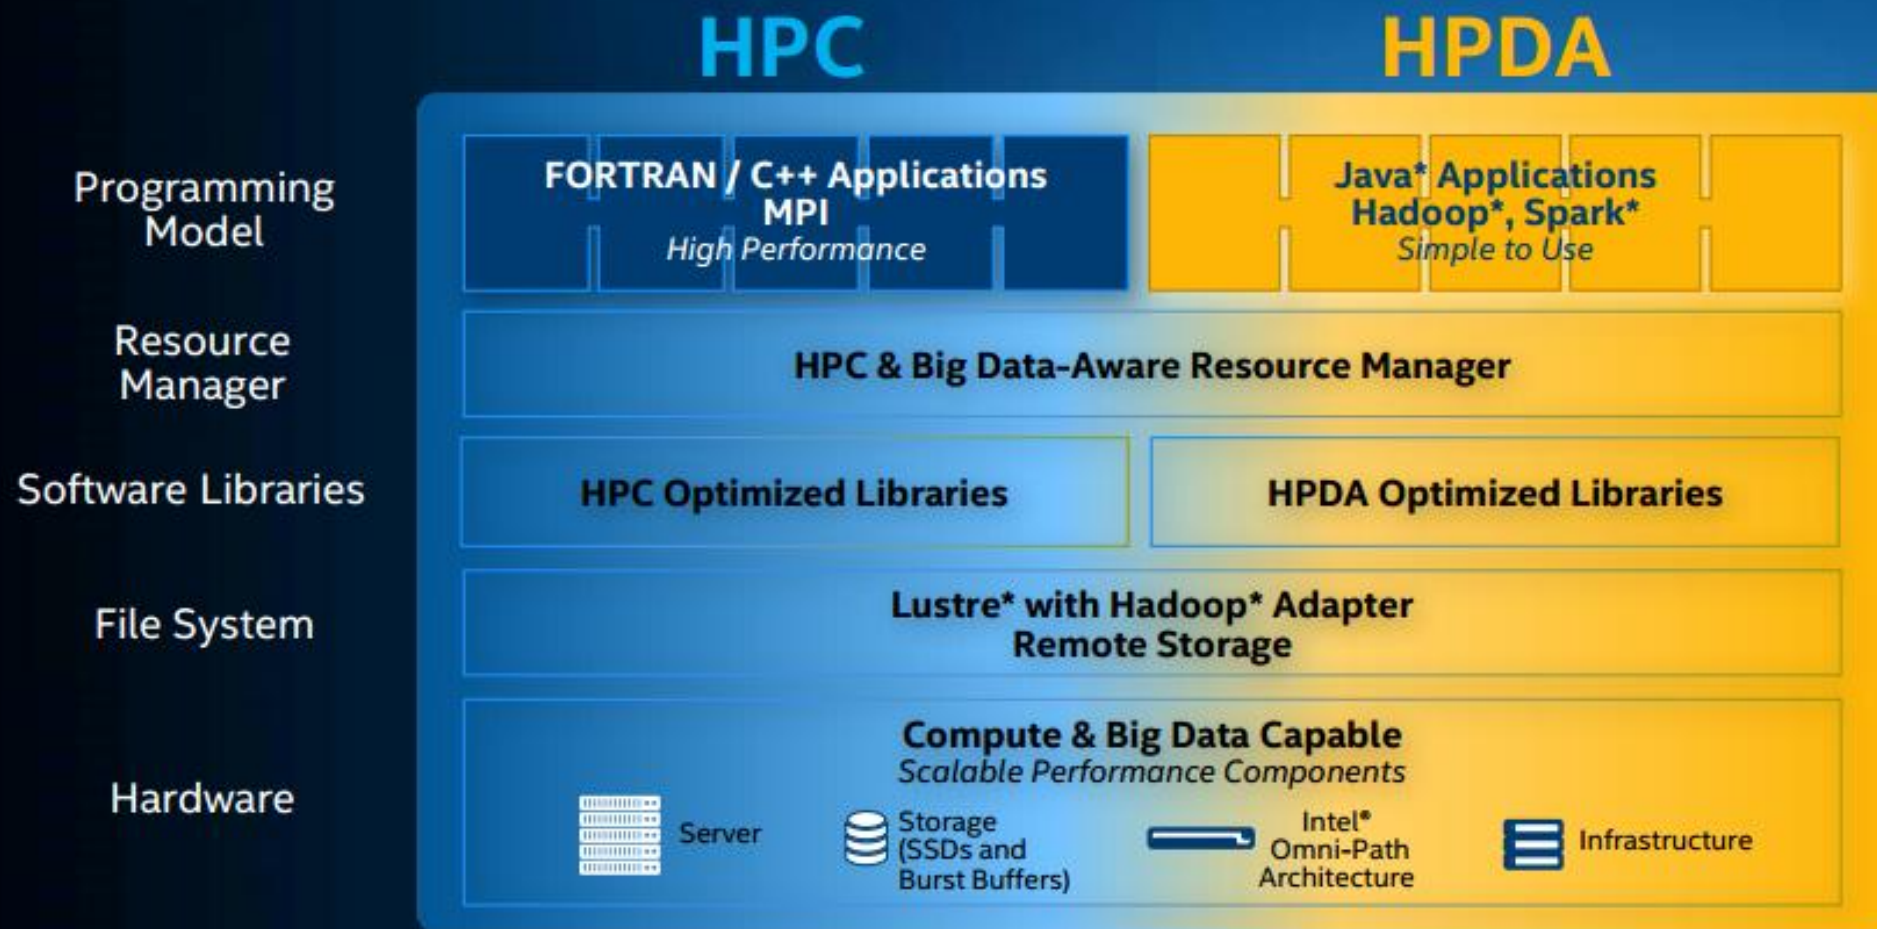
KEY VALUE VECTORS

- ✓ Performance
- ✓ Density
- ✓ Cost
- ✓ Power
- ✓ Reliability

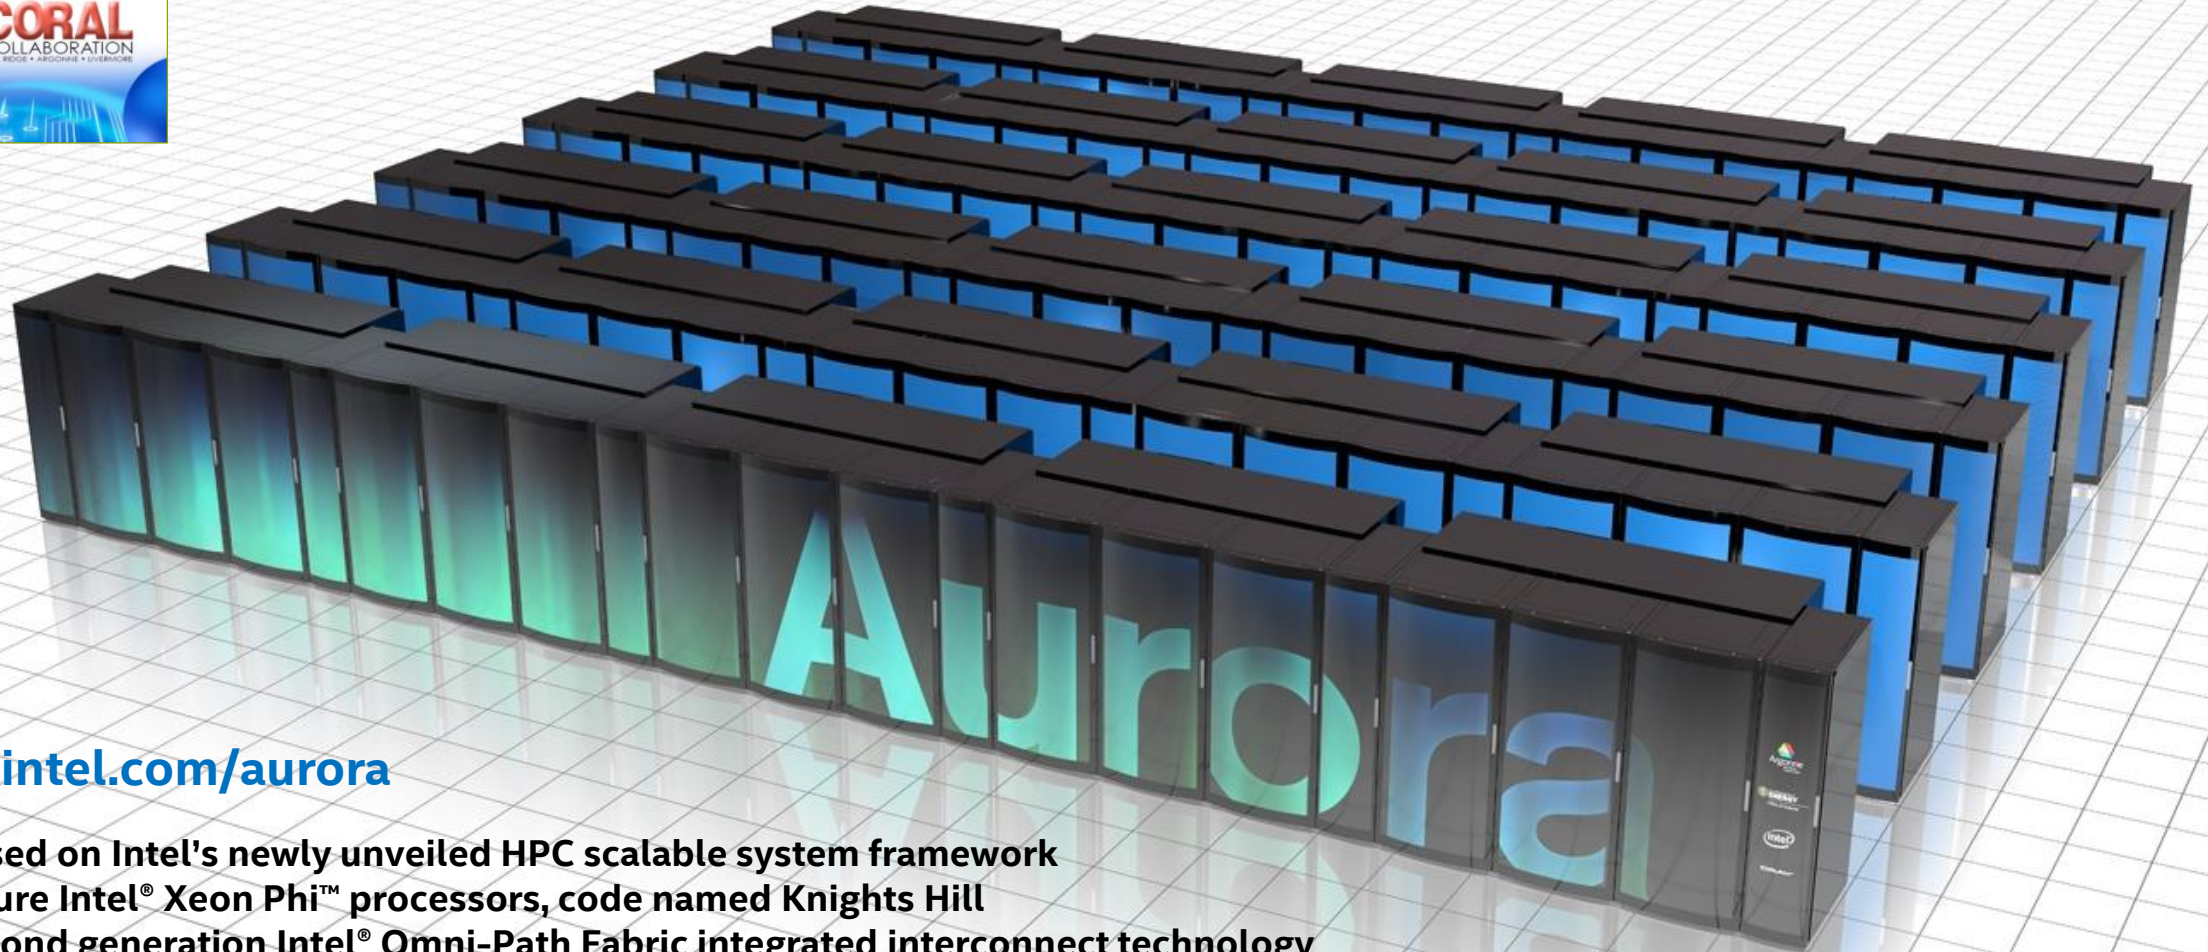
INTEGRATION



UNIFIED ARCHITECTURE FOR HPC & HPDA



Next-Generation 180 PFLOPS Peak Performance Supercomputer at Argonne National Laboratory



www.intel.com/aurora

- based on Intel's newly unveiled HPC scalable system framework
- future Intel® Xeon Phi™ processors, code named Knights Hill
- second generation Intel® Omni-Path Fabric integrated interconnect technology
- a new non-volatile memory-based storage hierarchy
- a standards-based application development environment using Intel tools
- an advanced parallel file system using Intel® Lustre* software



THANK YOU.