From Milliwatts to PFLOPS

GridKa School 2013

Karlsruhe, Germany
27 August 2013

Dr. Herbert Cornelius
Intel

Legal Disclaimer

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

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

Intel product plans in this presentation do not constitute Intel plan of record product roadmaps. Please contact your Intel representative to obtain Intel's current plan of record product roadmaps.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not uniqueto Intel microprocessors. These optimizations include SSE2, SSE3, and SSE3 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 specificinstruction sets covered by this notice.Notice revision #20110804

All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

Intel processor numbers are not a measure of performance.Processor numbers differentiate features within each processor family, not across different processor families.Go to: http://www.intel.com/products/processor_number

Intel, processors, chipsets, and desktop boards 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.

Intel, Intel Xeon, Intel Xeon Phi, Intel Hadoop Distribution, Intel Cluster Ready, Intel OpenMP, Intel CilkPlus, Intel Threaded Buildiingblocks, Intel Cluster Studio, Intel Parallel Studio, Intel CoarrayFortran, Intel Math KernalLibrary, Intel Enterprise Edition for LustreSoftware, Intel Composer, the Intel Xeon Phi logo, the Intel Xeon logo and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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.

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

Legal Disclaimers: Performance

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, Go to: http://www.intel.com/performance/resources/benchmark_limitations.htm.

Intel does not control or audit the design or implementation of third party benchmarks or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmarks are reported and confirm whether the referenced benchmarks are accurate and reflect performance of systems available for purchase.

Relative performance is calculated by assigning a baseline value of 1.0 to one benchmark result, and then dividing the actual benchmark result for the baseline platform into each of the specific benchmark results of each of the other platforms, and assigning them a relative performance number that correlates with the performance improvements reported.

SPEC, SPECint, SPECfp, SPECrate. SPECpower, SPECjAppServer, SPECjEnterprise, SPECjbb, SPECompM, SPECompL, and SPEC MPI are trademarks of the Standard Performance Evaluation Corporation. See http://www.spec.org for more information.

TPC Benchmark is a trademark of the Transaction Processing Council. See http://www.tpc.org for more information.

SAP and SAP NetWeaver are the registered trademarks of SAP AG in Germany and in several other countries. See http://www.sap.com/benchmark for more information.

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

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/software/products.

Optimization Notice

Intel® compilers, associated libraries and associated development tools may include or utilize options that optimize for instruction sets that are available in both Intel® and non-Intel microprocessors (for example SIMD instruction sets), but do not optimize equally for non-Intel microprocessors. In addition, certain compiler options for Intel compilers, including some that are not specific to Intel micro-architecture, are reserved for Intel microprocessors. For a detailed description of Intel compiler options, including the instruction sets and specific microprocessors they implicate, please refer to the "Intel® Compiler User and Reference Guides" under "Compiler Options." Many library routines that are part of Intel® compiler products are more highly optimized for Intel microprocessors than for other microprocessors. While the compilers and libraries in Intel® compiler products offer optimizations for both Intel and Intel-compatible microprocessors, depending on the options you select, your code and other factors, you likely will get extra performance on Intel microprocessors.

Intel® compilers, associated libraries and associated development tools 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 Intel® Streaming SIMD Extensions 2 (Intel® SSE2), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), and Supplemental Streaming SIMD Extensions 3 (Intel® 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.

While Intel believes our compilers and libraries are excellent choices to assist in obtaining the best performance on Intel® and non-Intel microprocessors, Intel recommends that you evaluate other compilers and libraries to determine which best meet your requirements. We hope to win your business by striving to offer the best performance of any compiler or library; please let us know if you find we do not.

Notice revision #20101101

Performance

it's all about Parallelism

and

Energy Efficiency

DATACENTER AS A SYSTEM

FACILITIES NETWORKING HARDWARE SOFTWARE OPERATIONS

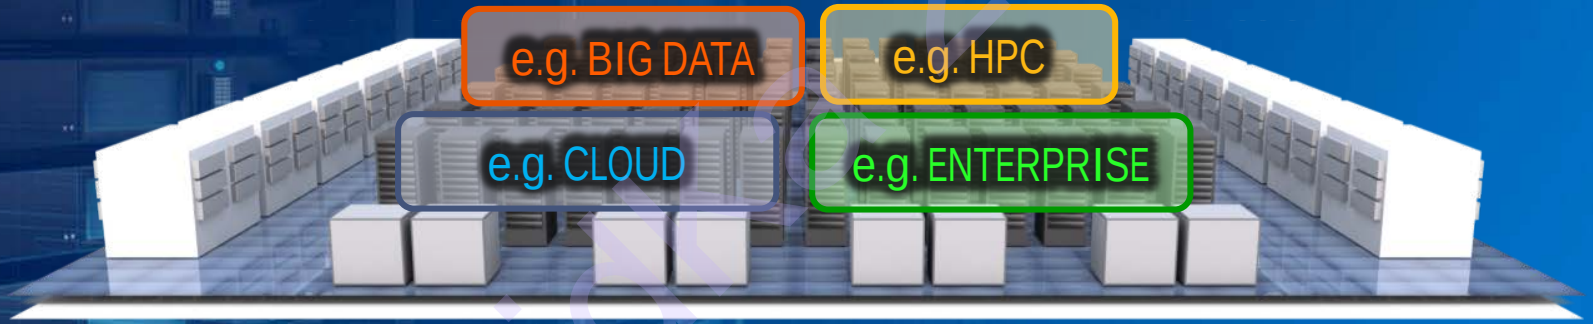

ENABLING "*IT as a SERVICE*"

seamlessly integrated system architecture for dynamically composable resources

HPC IMPERATIVES

High Performance
Capabilities & Capacity

Energy Efficiency
TCO

Ease of Use
Productivity & Sustainability

Simplicity
is the ultimate sophistication.

- Leonardo da Vinci

Transforming the Economics of HPC

Executing to Moore's Law

Predictable Silicon Track Record – well and alive at Intel.
Enabling new devices with higher performance and functionality while controlling power, cost, and size

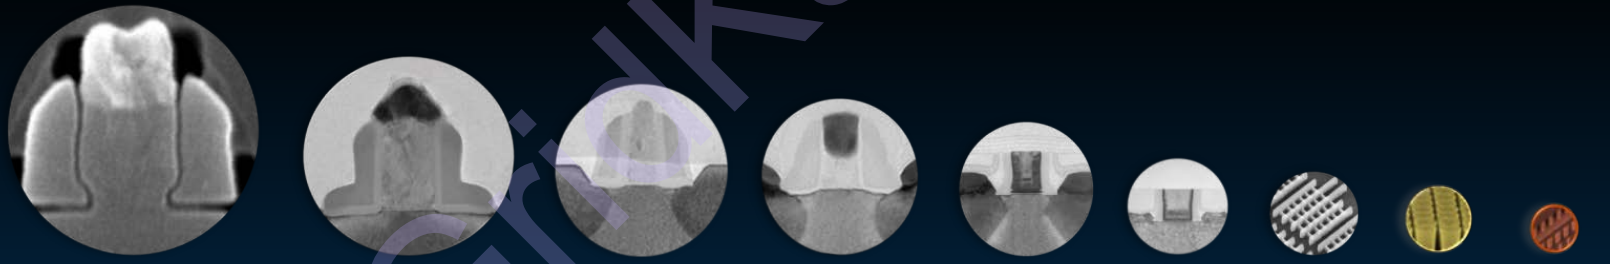

Future options subject to change without notice.

Driving Innovation and Integration

Enabled by Leading Edge Process Technologies

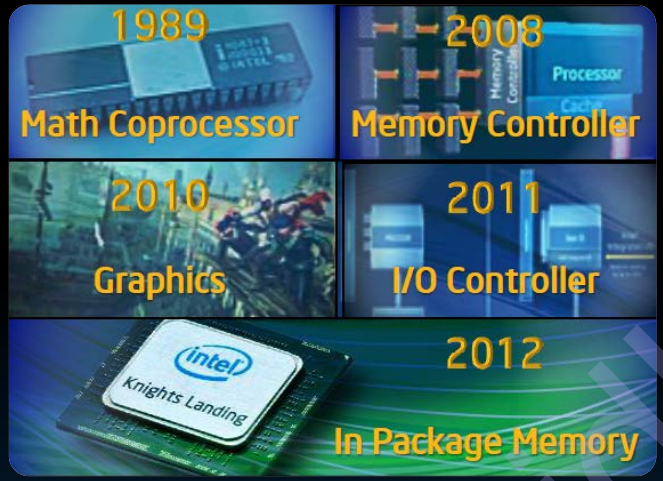

Integrated Today

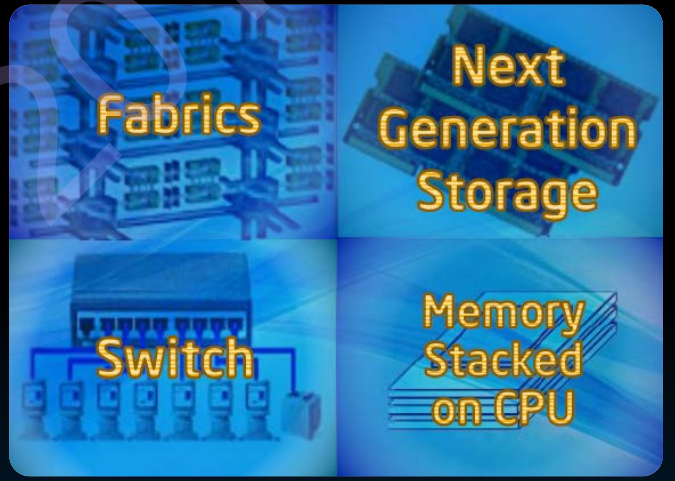

Coming in the Future

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

From MILLIWATTS to TERAFLOPS

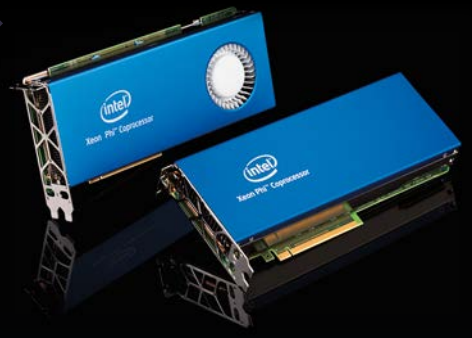

Smartphones with Intel® Inside

Intel® Xeon® Processors

Intel® Many Integrated Core Architecture

Energy Efficient

#1 TOP500 June 2013

33 PFLOPS HPL

54 PFLOPS Peak

32000 Intel® Xeon® E5v2 Processors

48000 Intel® Xeon Phi™ Coprocessors

Intel's Assets for HPC

Processors
Intel® Xeon® Processor

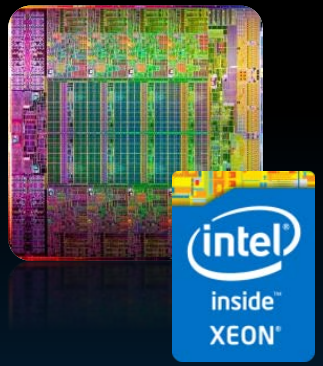

Coprocessor
Intel® Many Integrated Core

Network & Fabrics

Storage

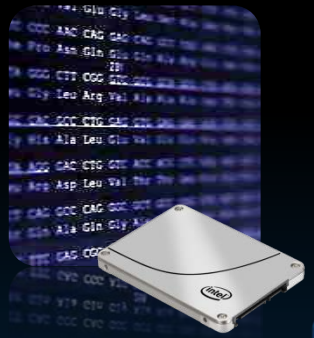

Software & Services

Increasing Processor Performance

Through Many-Core Technologies for Highly Parallel Workloads

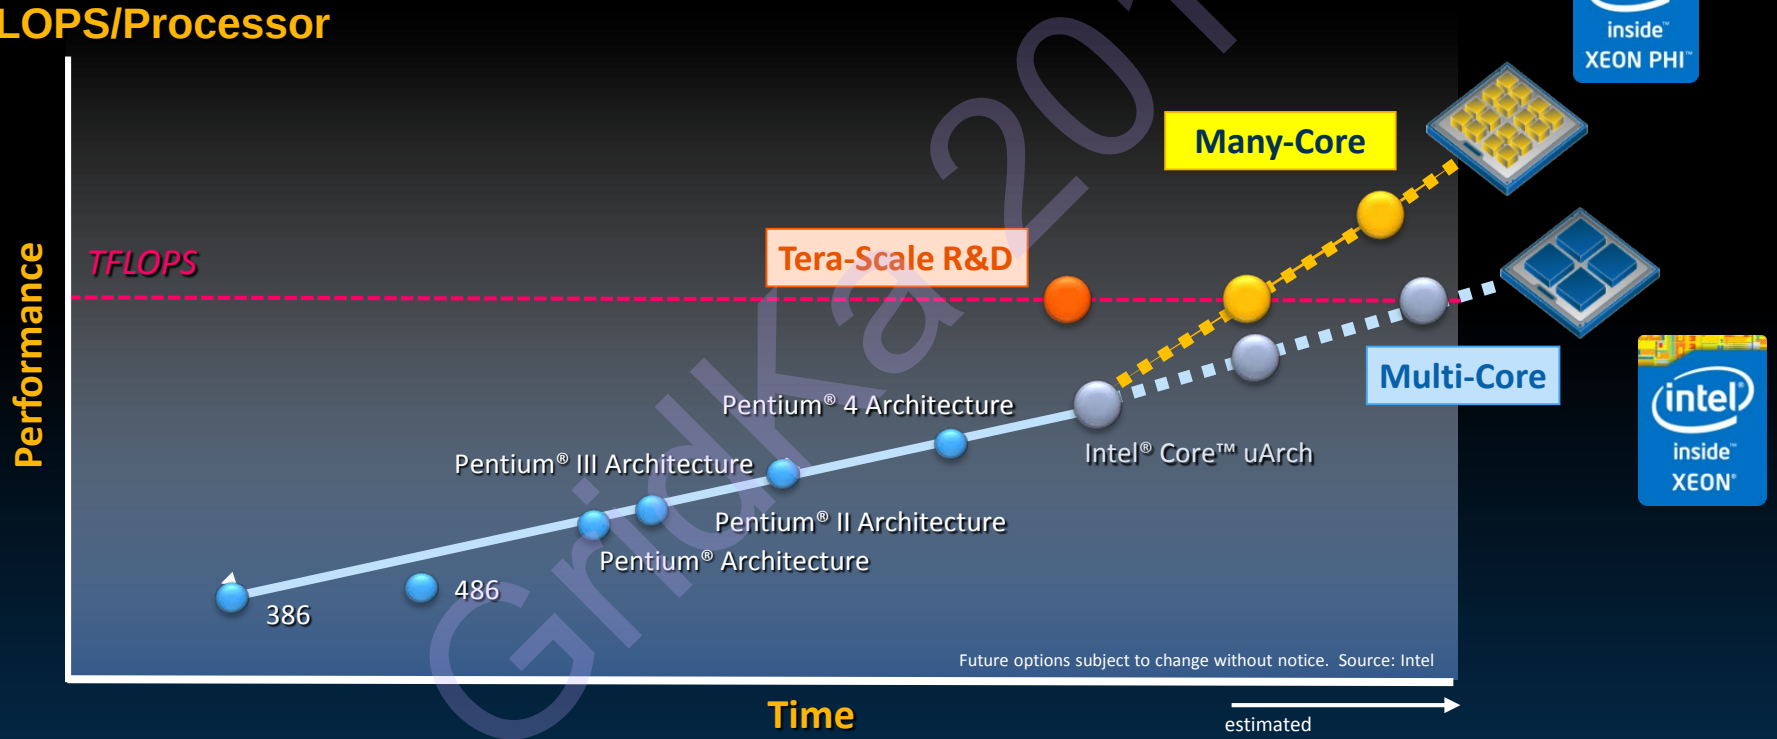

For illustration only. All dates, product descriptions, features, availability, and plans are forecasts and subject to change without notice.

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

„Big Core“ – „Small Core“

Different Optimization Points
Common Programming Models
and Architectural Elements

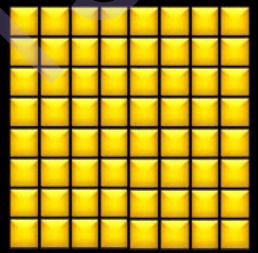

Intel® Xeon® Processor

Simply aggregating more cores generation after generation is not sufficient

Performance per core/thread must increase each generation, be as fast as possible

Power envelopes should stay flat or go down each generation

Balanced platform (Memory, I/O, Compute)

Cores, Threads, Caches, SIMD

Intel® Xeon Phi™ Coprocessor

Optimized for highest compute per watt

Willing to trade performance per core/thread for aggregate performance

Power envelopes should also stay flat or go down every generation

Optimized for highly parallel workloads

Cores, Threads, Caches, SIMD

For illustration only

Intel Roadmap to Exascale

1.000.000.000.000.000.000

Intel's Exascale Goal:
Reach Exascale by ~2020 with Intel technologies including Intel® Xeon Phi™ Coprocessors

Intel® Xeon Phi™ Product Family
Key ingredient in Intel Exascale Roadmap:

- Programmability
- Power efficiency
- Scalability
- Resiliency

Future options subject to change without notice.

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

Common Programming Models & Software Tools

Common Intel® architecture enables applications to run across the full spectrum of Intel® Xeon® family based servers so programmers don't have to "start over".

Use the same development tools you used for Intel® Xeon® processors, such as Intel® Cluster Studio XE and Intel® Parallel Studio XE

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

Intel® Xeon® E5 Processor Family

Foundation of HPC Performance
suited for full scope of workloads

Industry leading performance and
performance/watt
for serial & parallel workloads

General purpose with focus on
fast single core/thread performance
with "moderate" number of cores

www.intel.com/xeon

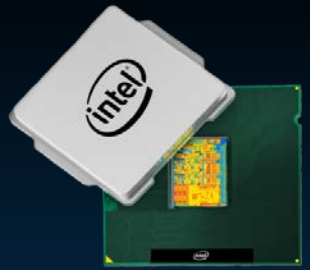

Tick-Tock Development Cycles

Integrate. Innovate.

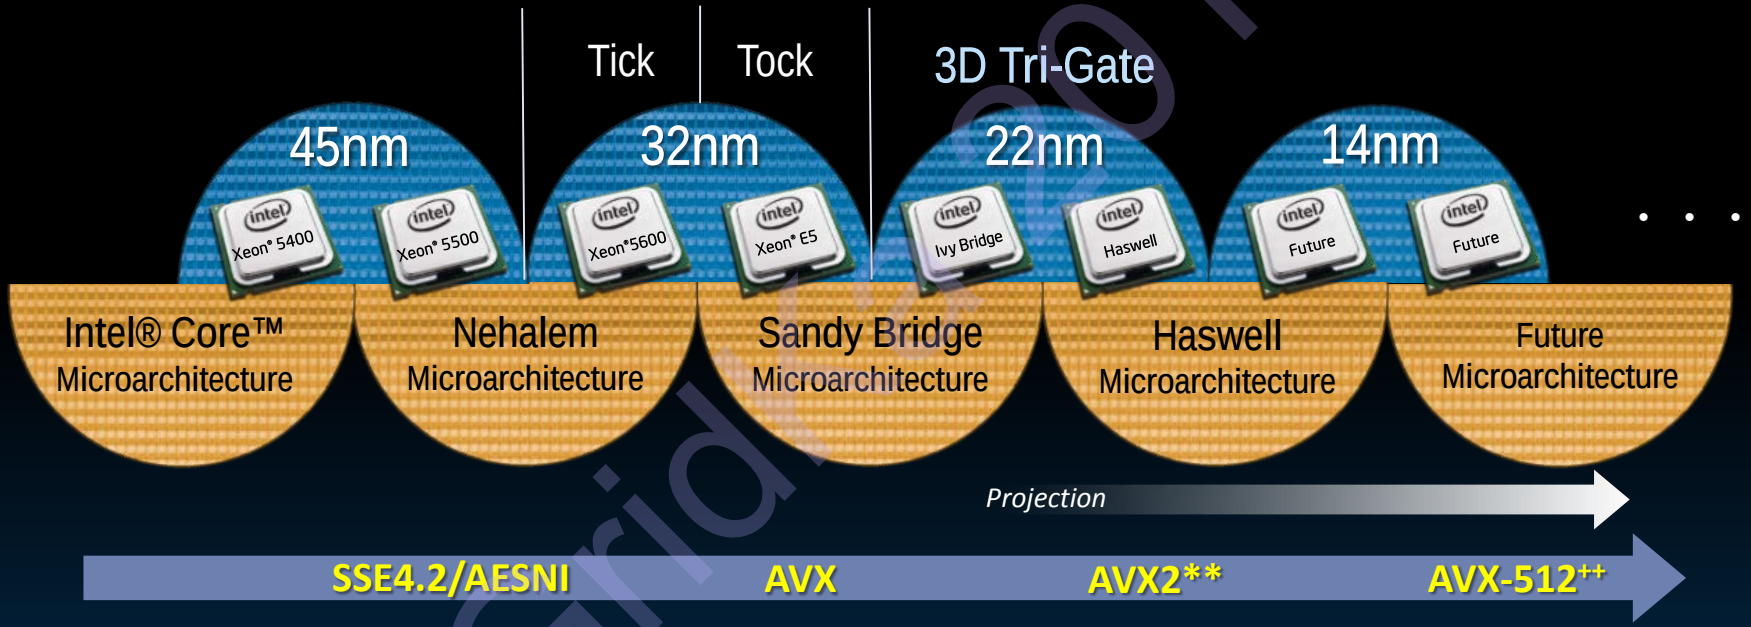

**Intel® Architecture Instruction Set Extensions Programming Reference, #319433-012A, FEBRUARY 2012

++Intel® Architecture Instruction Set Extensions Programming Reference, #319433-015, JULY 2013

Potential future options, subject to change without notice.

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

www.intel.com/xeonphi

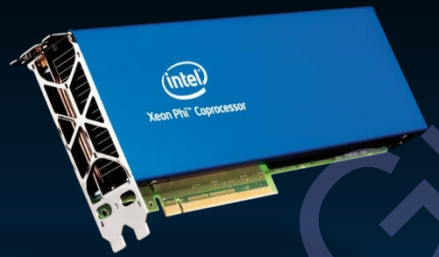

Intel® Xeon Phi™ Coprocessor

Up to 61 Cores, 244 Threads
512-bit SIMD instructions
>1TFLOPS DP-F.P. peak
Up to 16GB GDDR5 Memory, 352 GB/s
PCIe* x16
Up to 300W TDP (card)

22nm with the world's first
3-D Tri-Gate transistors
Linux* operating system
IP addressable native node
Common x86/IA
Programming Models and SW-Tools

Intel® Xeon Phi™ Coprocessor

Codename: Knights Corner - It is so much more

Restricted Architectures

DSP
FPGA
GPU
ASIC

Run restricted code

Custom HW Acceleration

Supercomputer on a chip

Operate as a compute node
Run a full OS (Linux*)
Run MPI
Run OpenMP*
Run x86 native code
Run offloaded code

Intel® Xeon Phi™ Coprocessor

Restrictive architectures limit the ability for applications to use arbitrary nested parallelism, functions calls and threading models

Flexible Execution Models

Optimized Performance for all Workloads

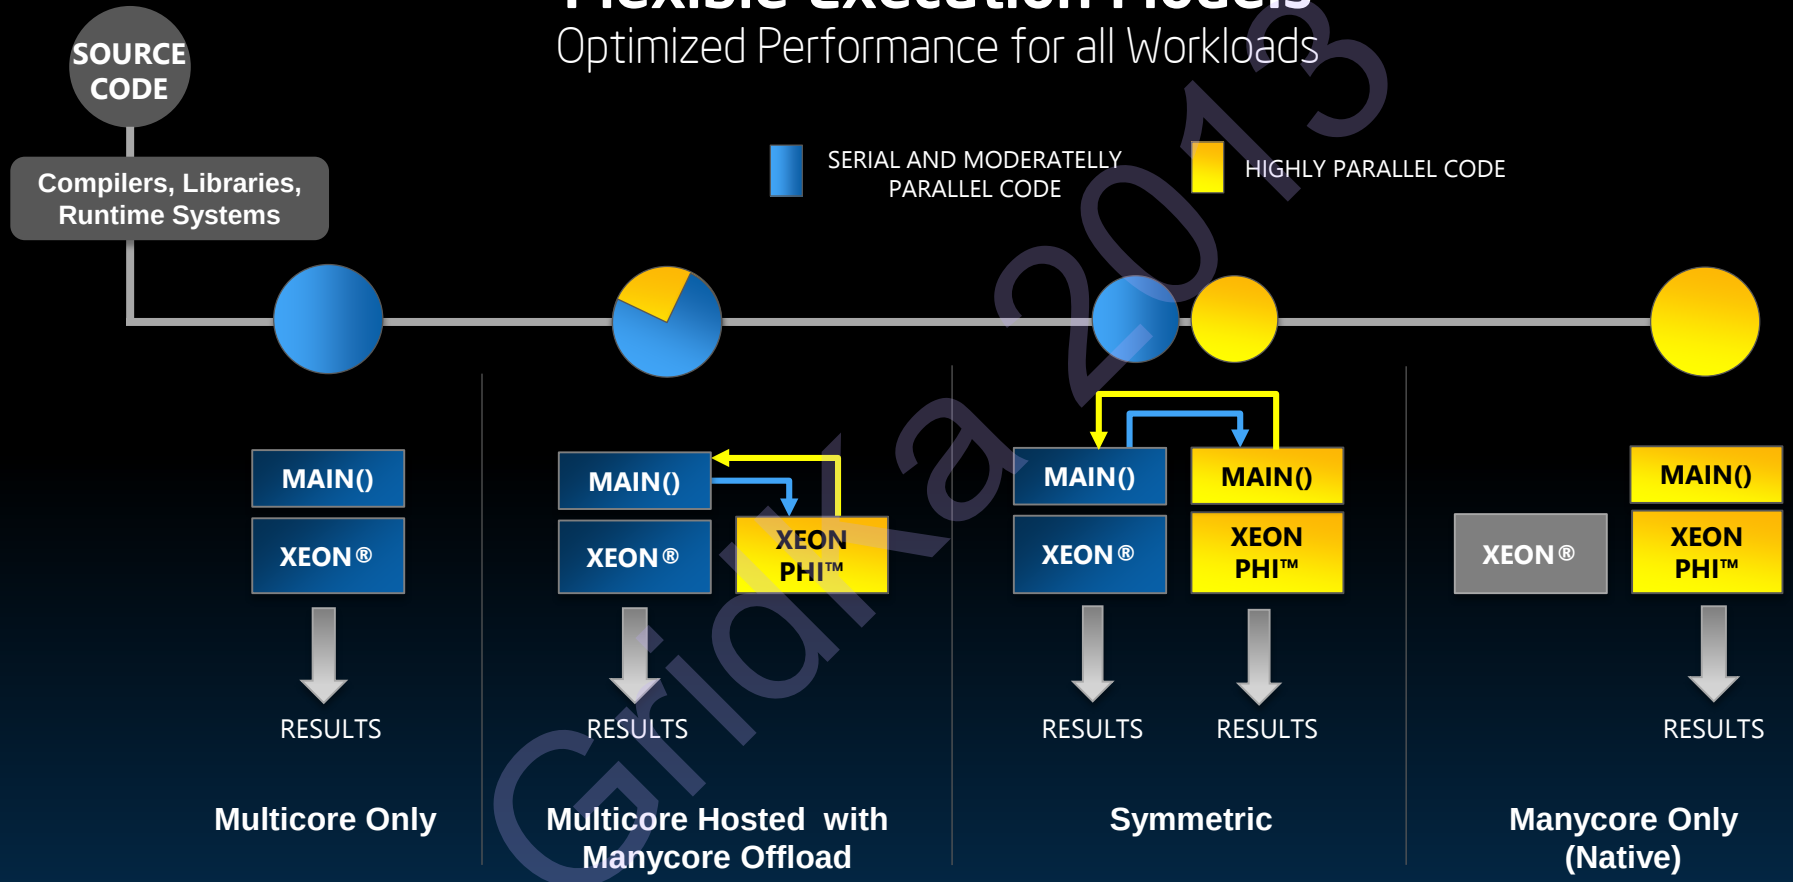

Manycore Processors – Example HPC Use Cases

1 MPI process
64 Threads each

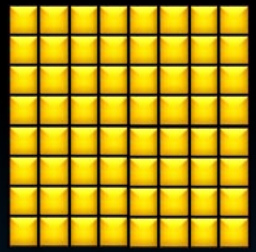

4 MPI processes
16 Threads each

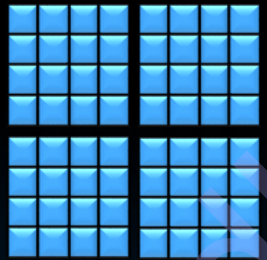

16 MPI processes
4 Threads each

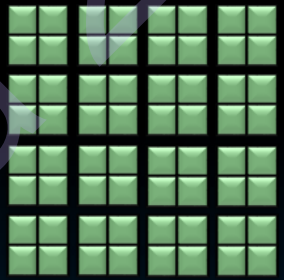

64 MPI processes
1 Threads each

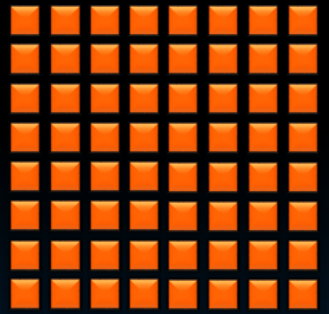

For illustration only.

Highly Parallel Applications

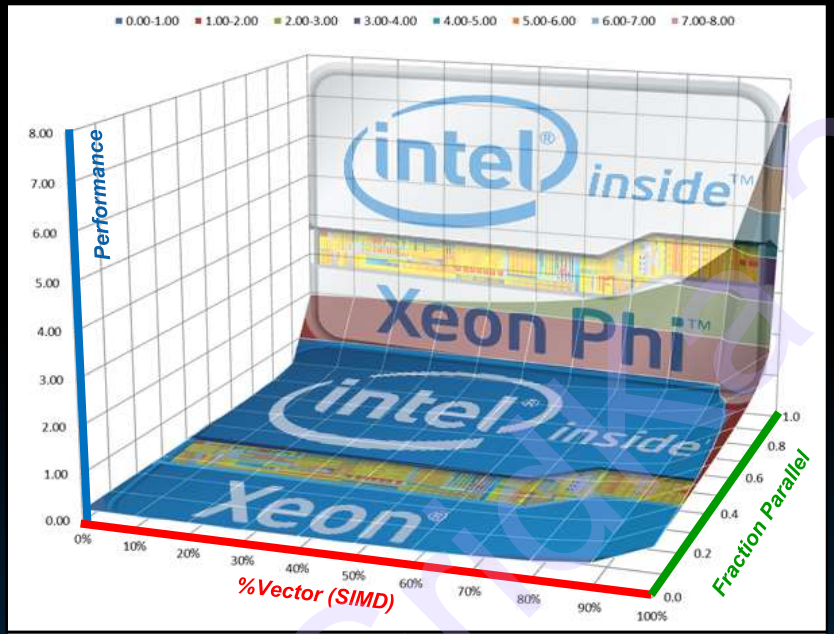

Theoretical acceleration of a highly parallel processor over a Intel® Xeon® parallel processor (<1: Intel® Xeon® faster) – For illustration only

Efficient vectorization, threading, and parallel execution drives higher performance for suitable scalable applications

Parallel Programming for Intel® Architecture (IA)

CORES	Use threads directly or e.g. via OpenMP*, pthreads Use tasking, Intel® TBB / Cilk™ Plus
VECTORS	Intrinsics, auto-vectorization, vector-libraries Language extensions for vector programming (SIMD)
BLOCKING	Use caches to hide memory latency Organize memory access for data reuse
DATA LAYOUT	Structure of arrays facilitates vector loads / stores, unit stride Align data for vector accesses

Parallel programming to utilize the hardware resources, in an abstracted and portable way

More Cores. Wider Vectors. Performance Delivered.

Intel® Parallel Studio XE 2013 and Intel® Cluster Studio XE 2013

More Cores

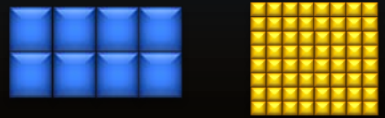

Wider SIMD/Vector

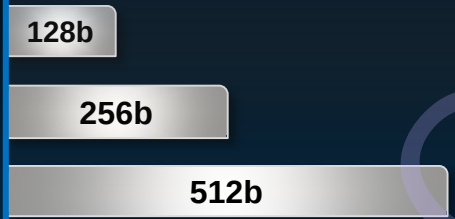

Scaling Performance Efficiently

Industry Leading Software Tools

High-Performance from advanced compilers

Comprehensive libraries

Parallel programming models

Insightful analysis tools

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

- 512-bit SIMD Instructions
- 8x 64-bit F.P./INT or 16x 32-bit F.P./INT per 512-bit register
- 32x 512-bit registers (ZMM0-ZMM31)
- 8x mask registers
- First implemented in the future Intel® Xeon Phi™ processor and coprocessor known by the code name Knights Landing
- Also supported by some future Intel® Xeon® processors

For testing the Intel® Software Development Emulator (Intel® SDE) has been extended for Intel AVX-512 and is available at http://www.intel.com/software/sde.

http://software.intel.com/en-us/blogs/2013/07/10/avx-512-instructions

http://download-software.intel.com/sites/default/files/319433-015.pdf

Intel SIMD Evolution

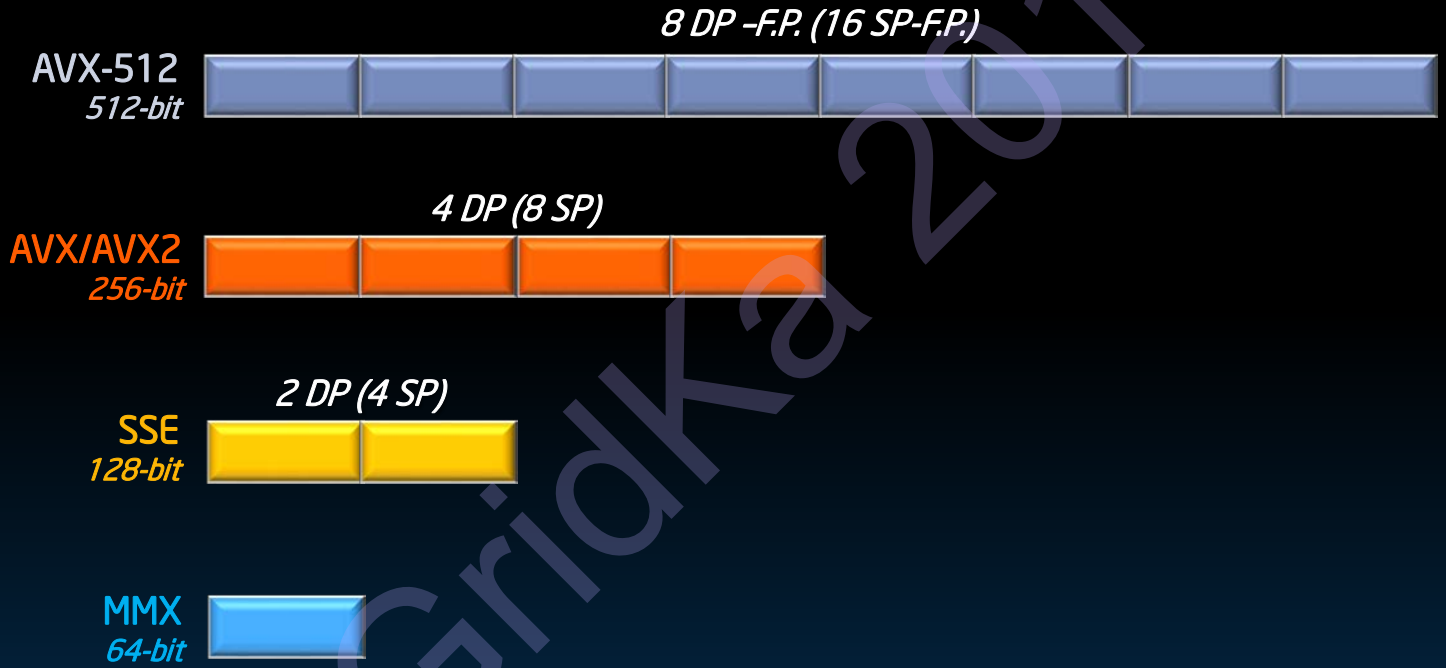

Potential future options and features subject to change without notice.

Intel® Xeon Phi™ Coprocessor Microarchitecture Overview

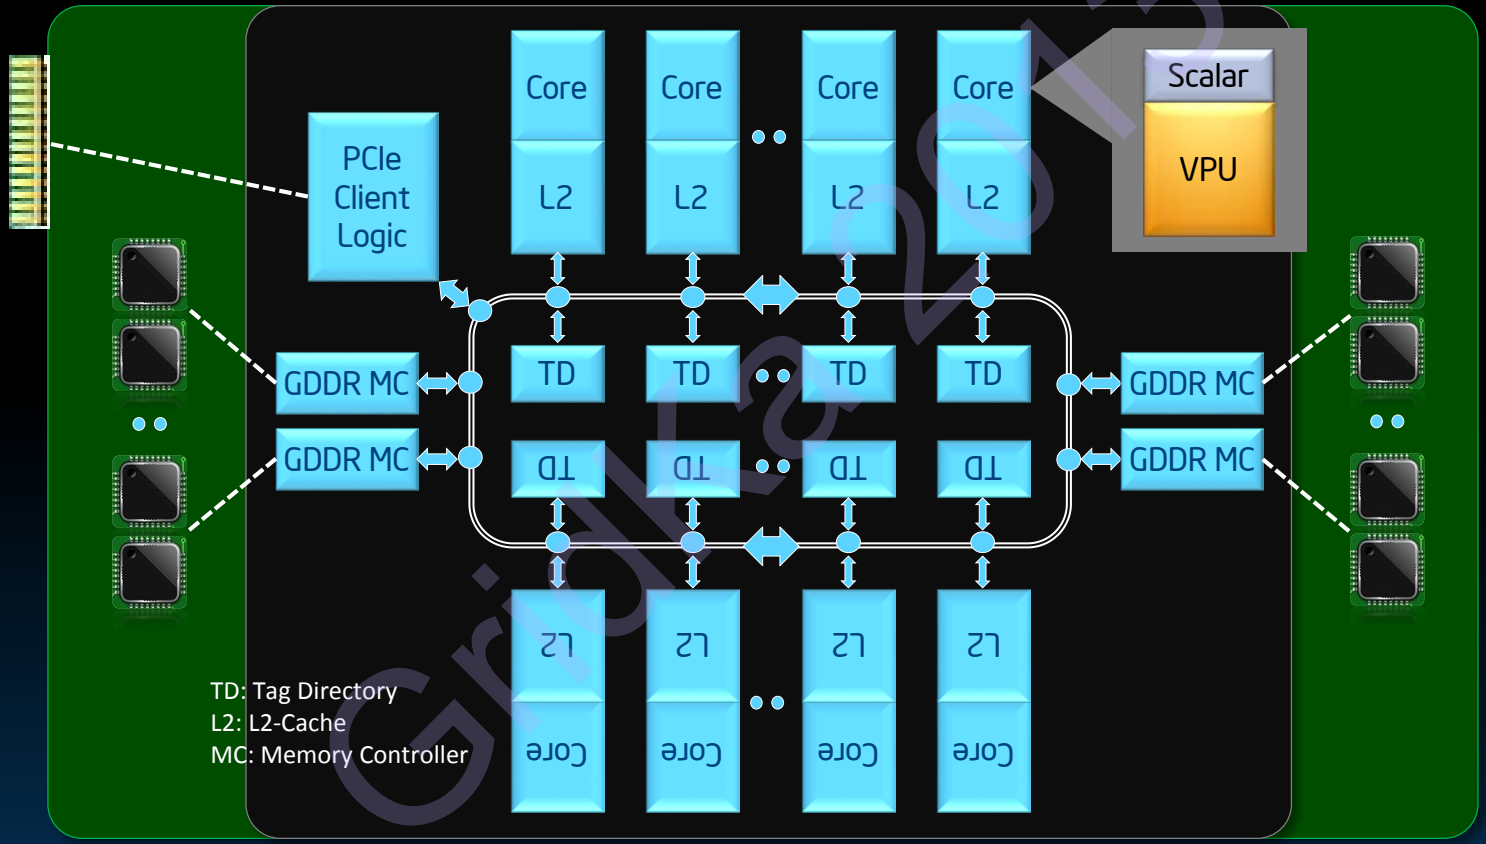

For illustration only.

Intel® Xeon Phi™ Coprocessor: SKUs H2'2013

7100 : Best Performance (highest level of features)

Premium Offering for Most Demanding Users
Passively Cooled and No Thermal Solution Enabling Large D

16GB GDDR5
352GB/s
61Cores
>1.2TF DP

7120X

7120P

5100 : Best Performance/Watt (optimized for high density er

Ideal for Memory BW Bound (STREAM, Energy) & Memory C
Energy), Innovative Dense Form Factor, Lowest TDP Pass

8GB GDDR5
320+GB/s
60Cores
>1TF DP

5120D

5110P

3100 : Best Value (outstanding parallel computing s

Ideal for Compute Bound Workloads (Monte Carlo, B
etc.), Active and Passive Cooling for Wide Range of

6GB GDDR5
240GB/s
57Cores
>1TF DP

3120A

3120P

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

Intel® Xeon Phi™ Coprocessor x100 Family Reference Table

Processor Brand Name	Codename	Process	SKU #	Form Factor, Thermal	Board TDP (Watts)	Max # of Cores	Clock Speed (GHz)	Peak Double Precision (GFLOP)	GDDR5 Memory Speeds (GT/s)	Peak Memory BW	Memory Capacity (GB)	Total Cache (MB)	Production Si Stepping	Enabled Turbo	Turbo Clock Speed (GHz)
Intel® Xeon Phi™ Coprocessor x100	Knights Corner	22nm	SE10P	PCIe Card, Passively Cooled	300	61	1.1	1073.6	5.5	352	8	30.5	B	N	N/A
			SE10X	PCIe Card, No Thermal Solution	300	61	1.1	1073.6	5.5	352	8	30.5	B	N	N/A
			7120P	PCIe Card, Passively Cooled	300	61	1.238	1208	5.5	352	16	30.5	C	Y	1.333
			7120X	PCIe Card, No Thermal Solution	300	61	1.238	1208	5.5	352	16	30.5	C	Y	1.333
			5120D	Dense Form, No Thermal Solution	245	60	1.053	1011	5.5	352	8	30	C	N	N/A
			5110P	PCIe Card, Passively Cooled	225	60	1.053	1011	5.0	320	8	30	B, C	N	N/A
			3120P	PCIe Card, Passively Cooled	300	57	1.1	1003	5.0	240	6	28.5	C	N	N/A
			3120A	PCIe Card, Actively Cooled	300	57	1.1	1003	5.0	240	6	28.5	C	N	N/A

All SKUs, pricing and features are subject to change without notice

Next Intel® Xeon Phi™ Processor

Codename: Knights Landing

Designed using Intel's cutting-edge

14nm process

Not bound by "offloading" bottlenecks

Standalone CPU

or PCIe Coprocessor

Leadership compute & memory bandwidth

Integrated

On-Package Memory

Assume Exascale Computing at 20MW …

New Forms of Energy

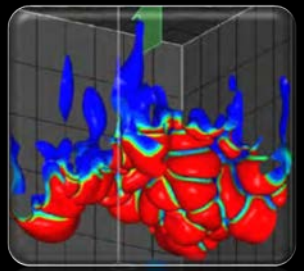

Ecological Sustainability

Space Exploration

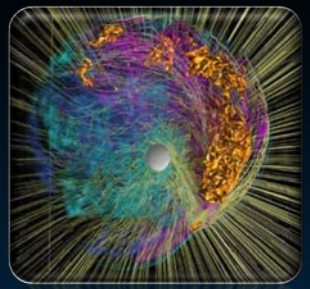

Medical Innovation

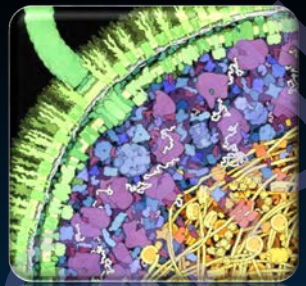

And many others….

1 Exaflop	20MW
1 Petaflop	20KW
1 Teraflop	20W
100 Gigaflop	2W
1 Gigaflop	20mW

Today's #25 system in a rack!

100x the performance of today's phone at the same power

For illustration and concept only.

HPC:
The Path to Exascale

Processors
Intel® Xeon® Processor

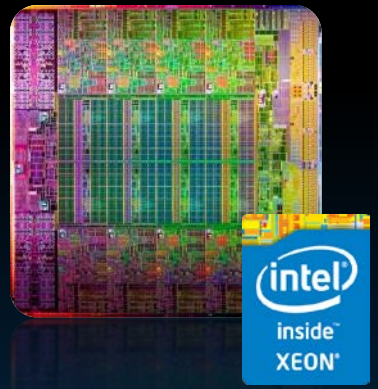

Coprocessor
Intel® Many Integrated Core

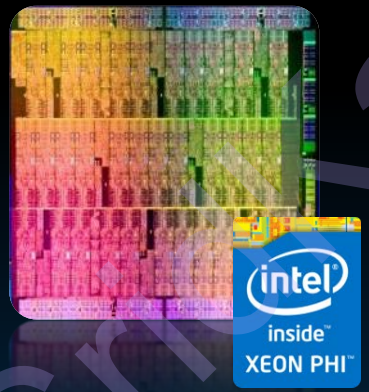

Fabrics

Software

HPC:
The Path to Exascale (cont.)

Memory & Storage

Networking

Reliability & Resiliency

Power Management

Intel TeraScale Research Areas

MANY-CORE COMPUTING

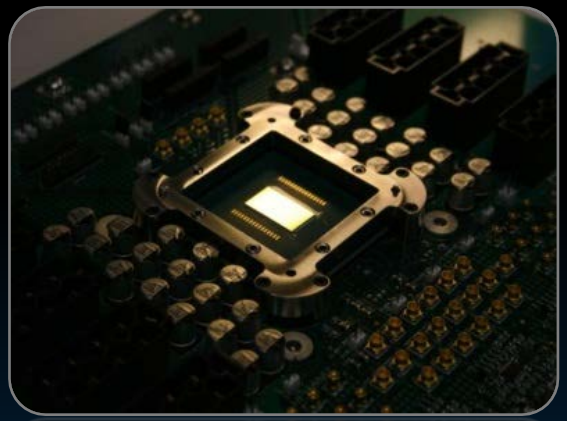

Teraflops
of computing power

STACKED MEMORY

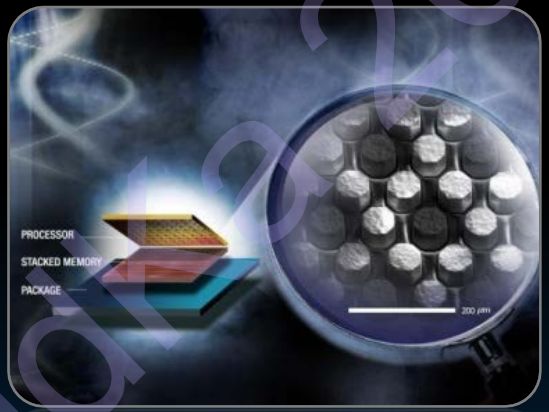

Terabytes
of memory bandwidth

SILICON PHOTONICS

Terabits
of I/O throughput

Future vision, does not represent real products.

The Power of Solutions: Big Data Example

Sort 1TB of Data:

>4 Hours

Sort 1TB of Data:

7 MINUTES

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

The Power of Platform Solutions

TeraSort for 1TB sort:
>4 hour process time

Xeon 5600
HDD
1GbE

Upgrade processor ~50% reduction

Upgrade to SSD ~80% reduction

Upgrade to 10GbE ~50% reduction

Intel distribution ~40% reduction

<7 minutes

intel


PERFORMANCE
it's all about Scalable Parallelism
INTEL
INNOVATION & LEADERSHIP
FOR THE ROAD AHEAD

Thank You.