Winds In The East, Mist Coming In …

Rajeeb Hazra, Ph.D.
Vice President, Data Center Group
General Manager, Technical Computing

Race For Leadership … *Accelerating*

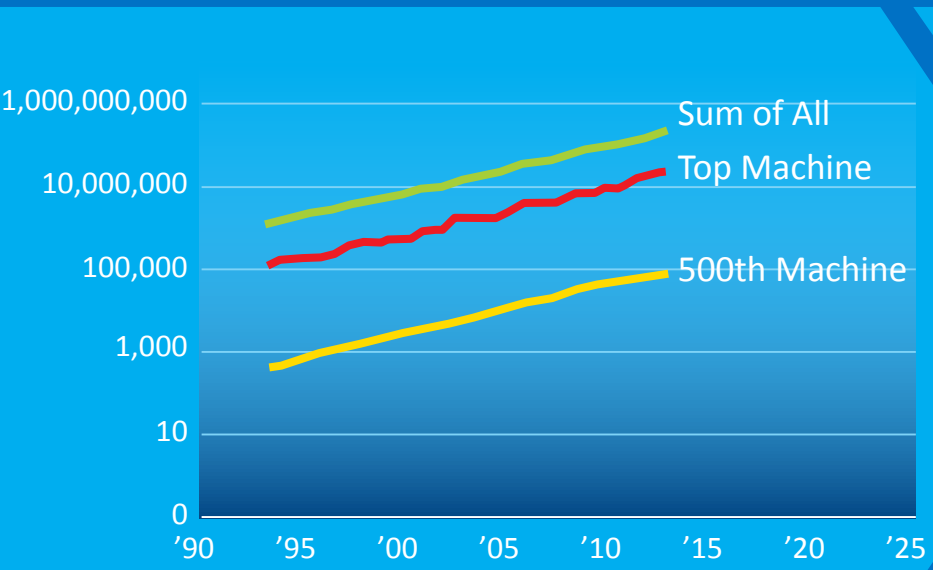

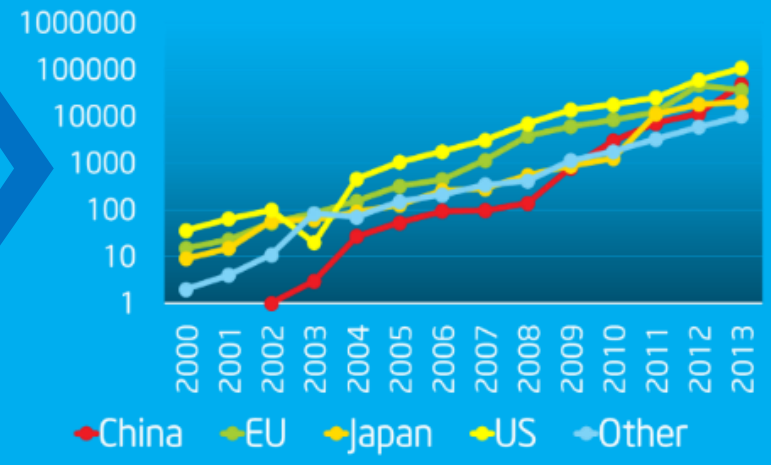

Top 500 FLOPS >50%[1] CAGR For Past Decade

Source: Top500.org

Need For HPC … *Expanding*

Physical Phenomena

Current Analanalytical Model

Real-time Data

Data Driven Model

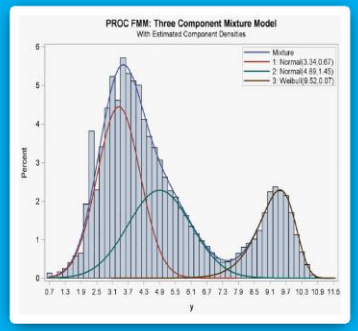

HPC

Knowledge

HPDA-Focused Servers Revenue:

2012: $743.8 Million

2017: $1.4 Billion

Source: IDC

Major Industry Shifts … *Brewing*

New contributors: ARM, Broadcom

Industry structure morphing:
OpenPower.org ... BYO

Indigenous movements: PRC, Japan, India …

HPC Modalities … *Evolving* *(again)*

The Path Forward

Proprietary	Democratized	Customized
		Vertically Designed Systems
Vertical Systems	Beowulf Cluster	

Vertically Designed = System ➔ Components

Power

Reliability

Scalability

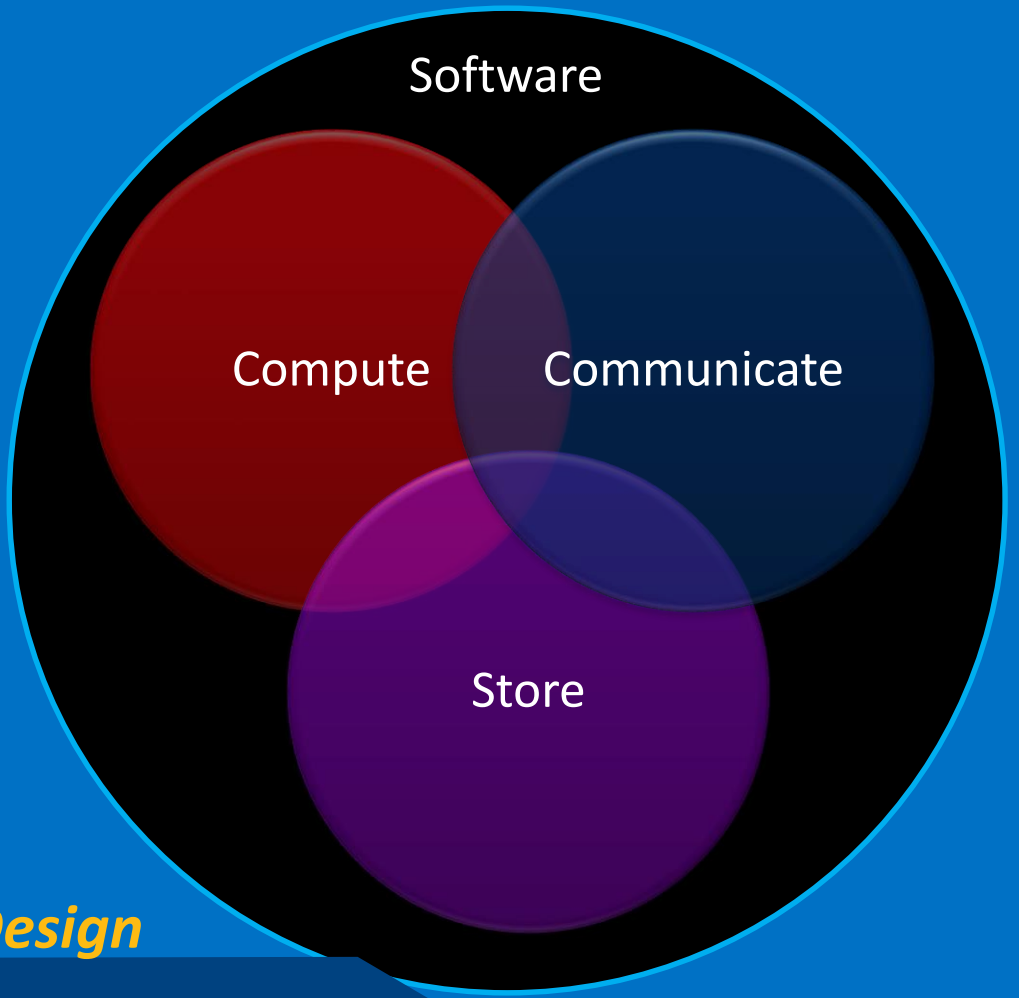

Only Way To Address Is To Design At The System Level

The System View

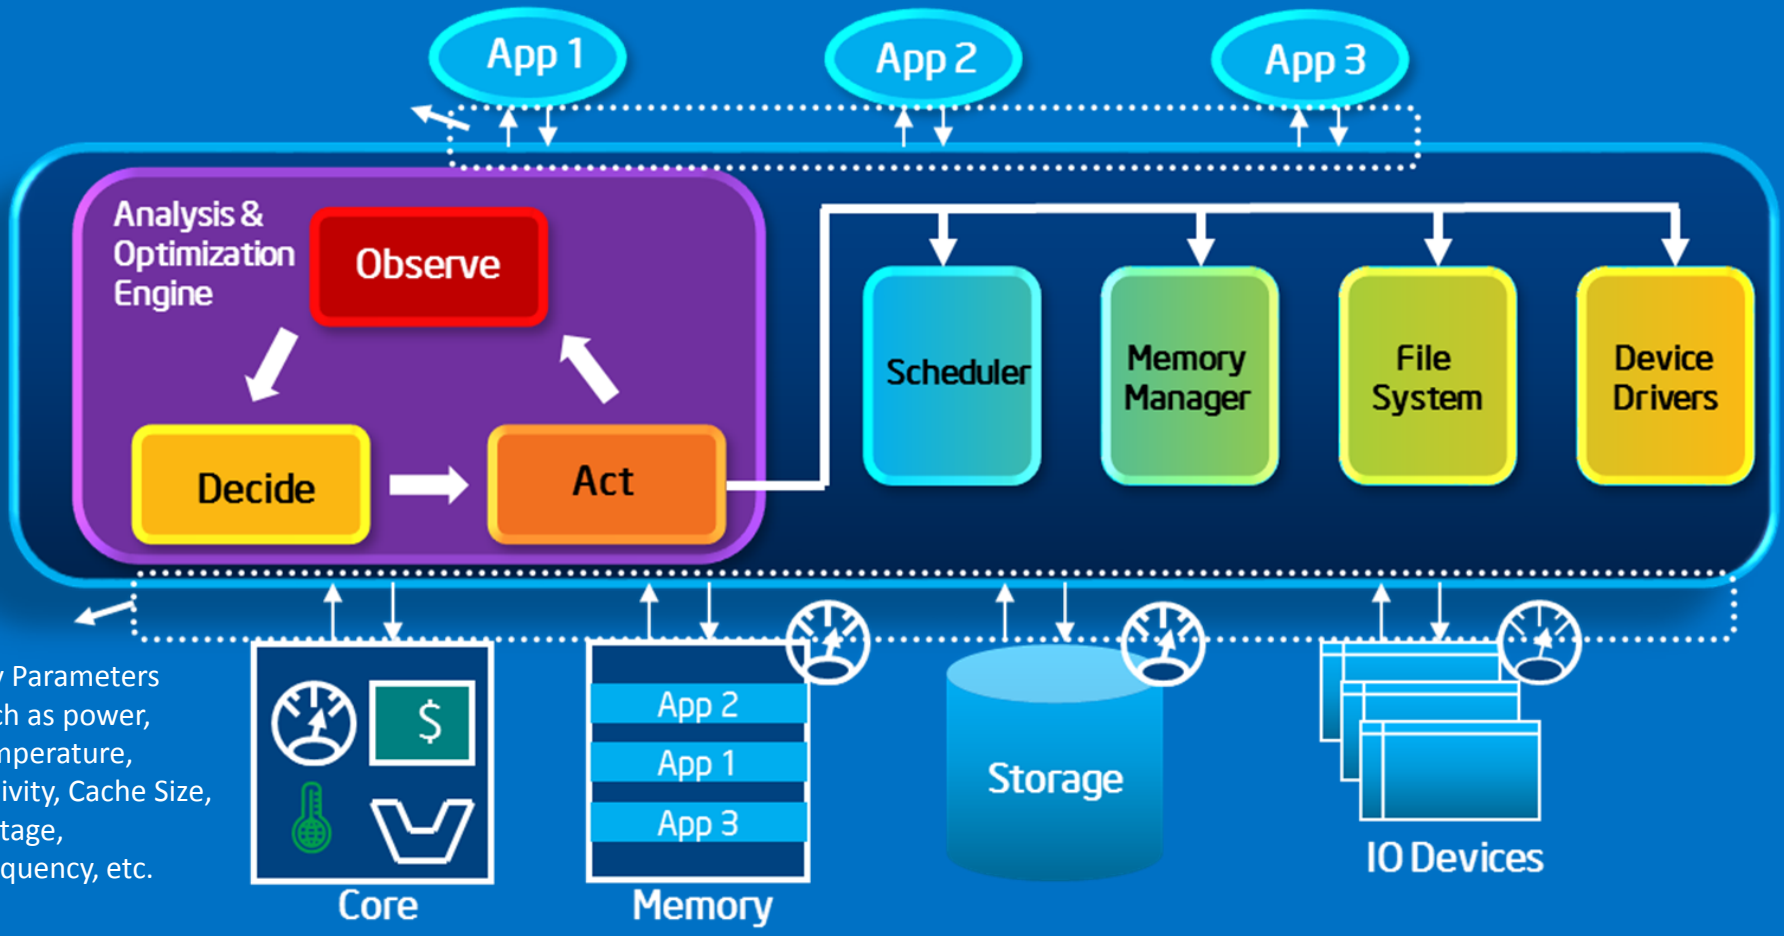

The Next Click Down: The "Rack" View

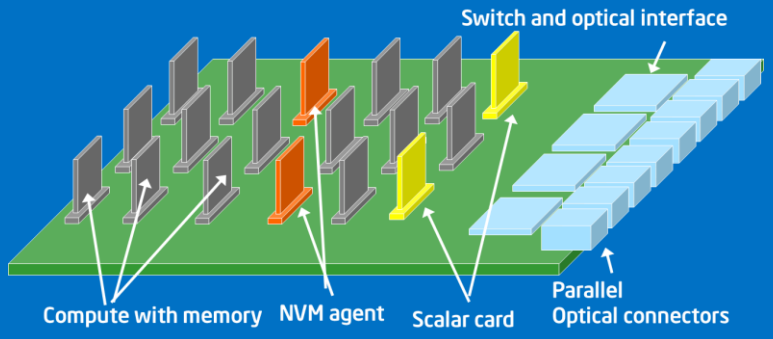

Datacenter "Rack" of the Future

- Simpler design: No DIMMs, common fabric
- Coherent multi-slot cards
- Broad configurability:
  - Compute intensive agents
  - Single thread focused nodes
  - NVM nodes
  - DRAM cards
  - Bridges for traditional interfaces

The Node: Now A Building Block

The “Right Sized” Building Block

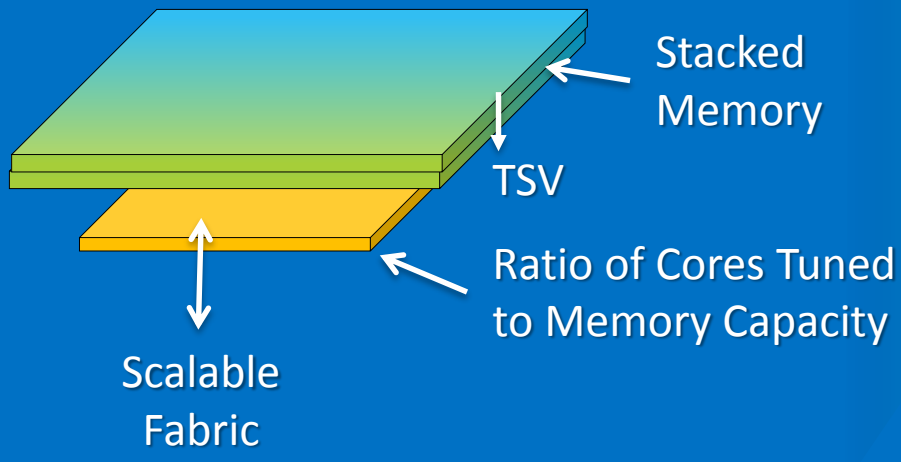

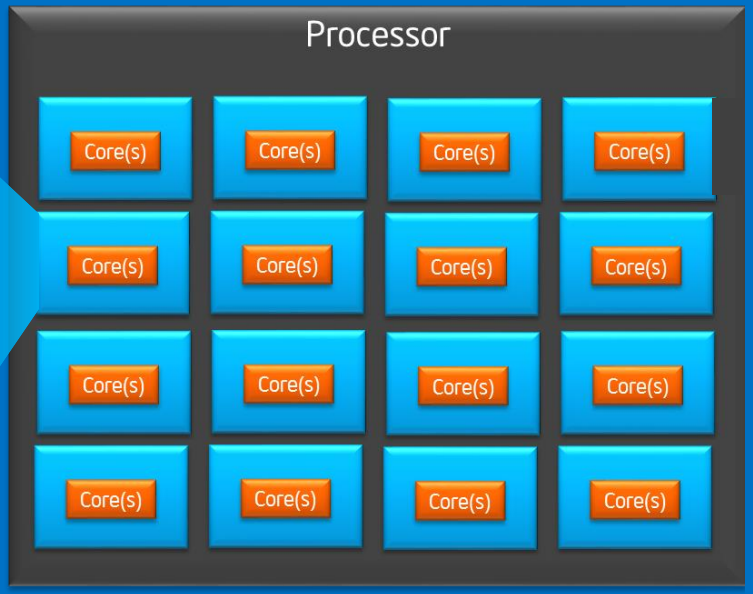

System Performance ≠ Max Node Performance

Interconnect: No Longer Just A “Bus”

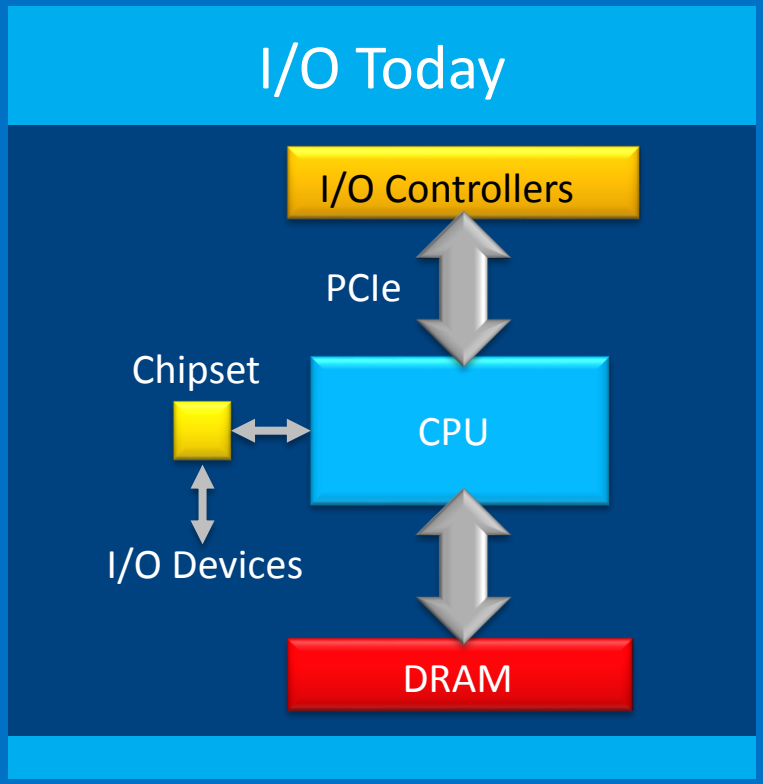

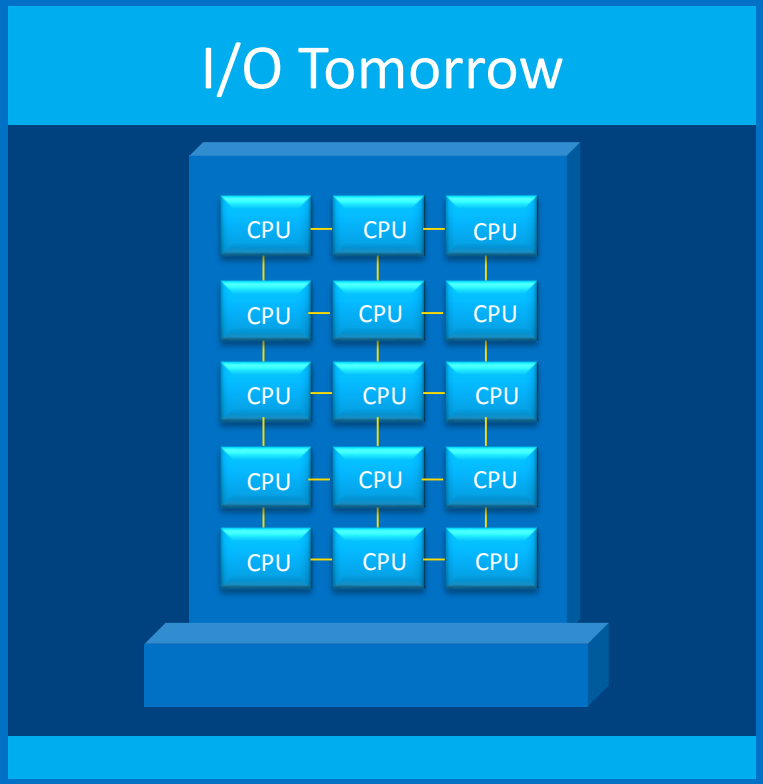

- Off-node → Predominantly for memory access
- Bisection bandwidth/latency → Hops
- Re-configurability: Adaptive + User-Defined

System Level Reliability

With Unreliability Everywhere

Component level design needed… but, at scale, some will fail

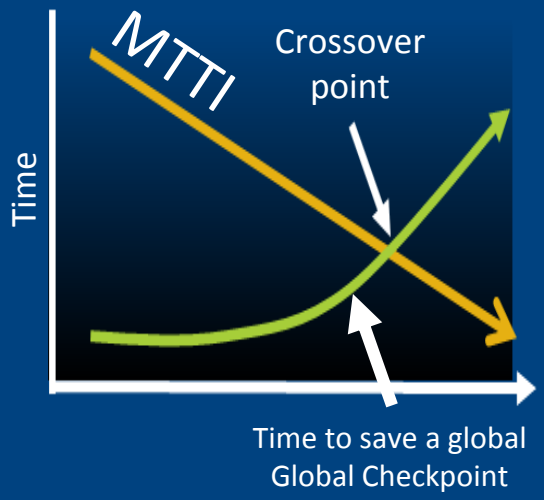

Hardware: Integrated approach for system level fault tolerance

Protect:
Harden all system ingredients

Detect & Reporting:
System Level MCA, Coordinated global insightful rack level fault reporting

Recover:
Recovery actions based on global information @ the transaction level

Software: Reliability designed in and coordinated at all layers

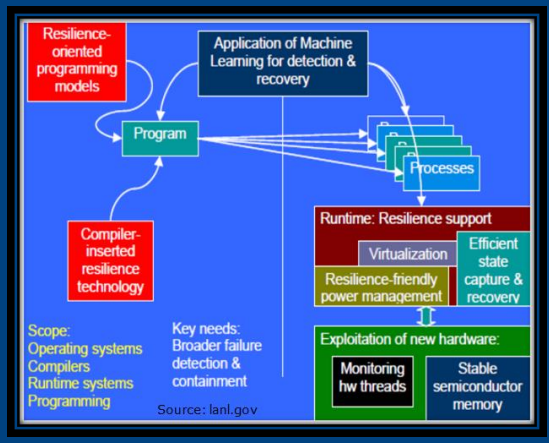

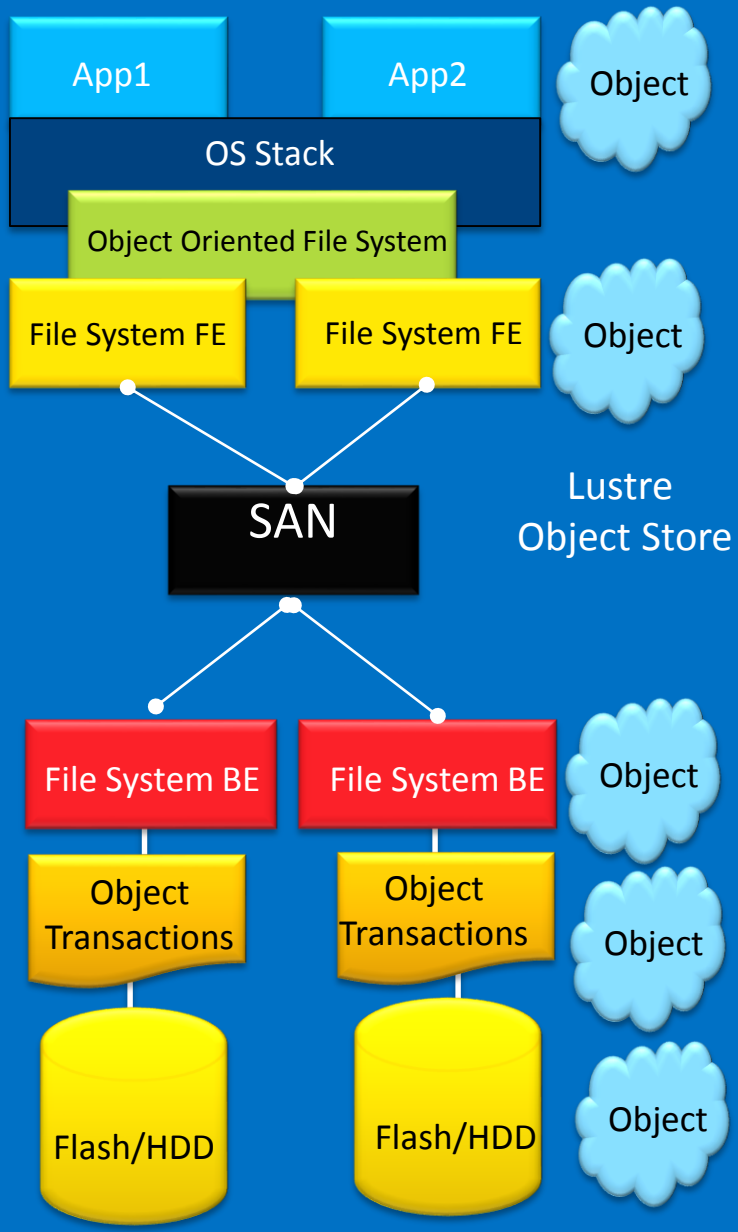

Object-based Storage Model

- Storage of objects is abstracted
- Metadata is accreted during object creation and IO
- Enables distributed data intensive computing model
- Enables analytics

Integration Is The Game Changer

Unprecedented Innovations Only Enabled by the Leading Edge Process Technology

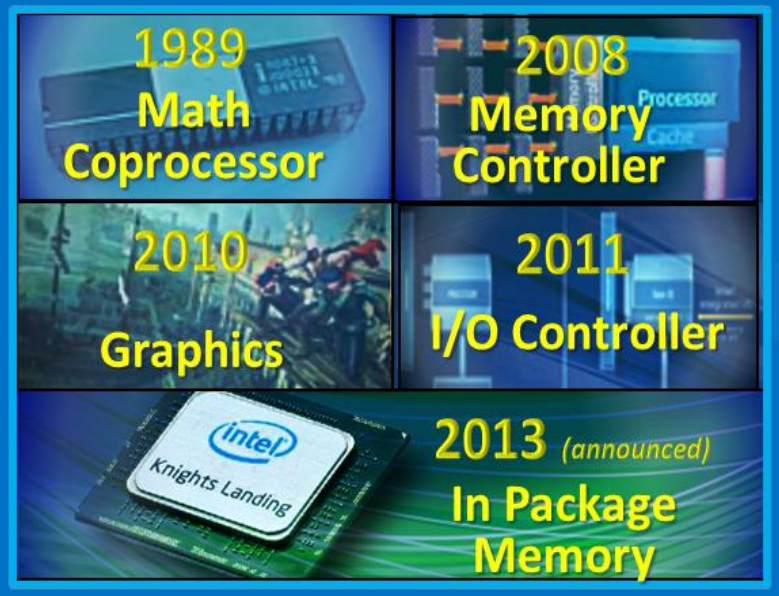

Today

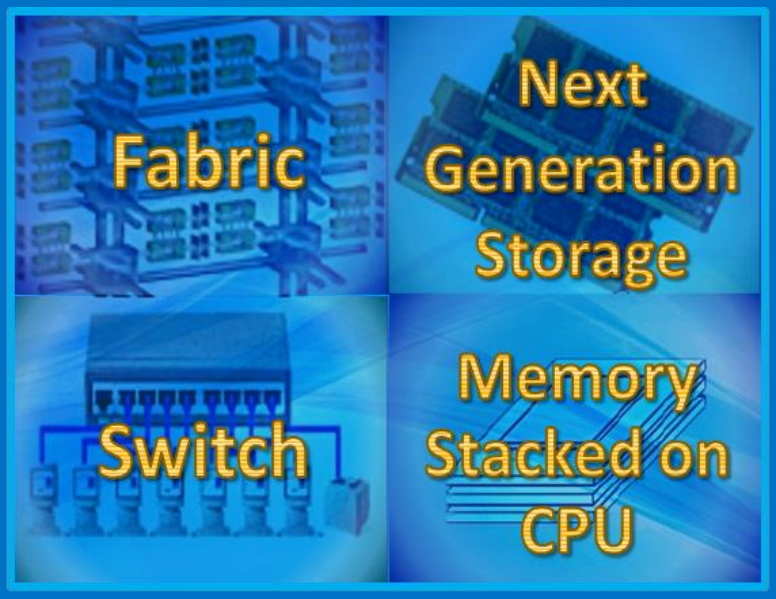

Possibilities For Tomorrow

But … Will It Run My Program(s)??

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

The Future Of Programming Models

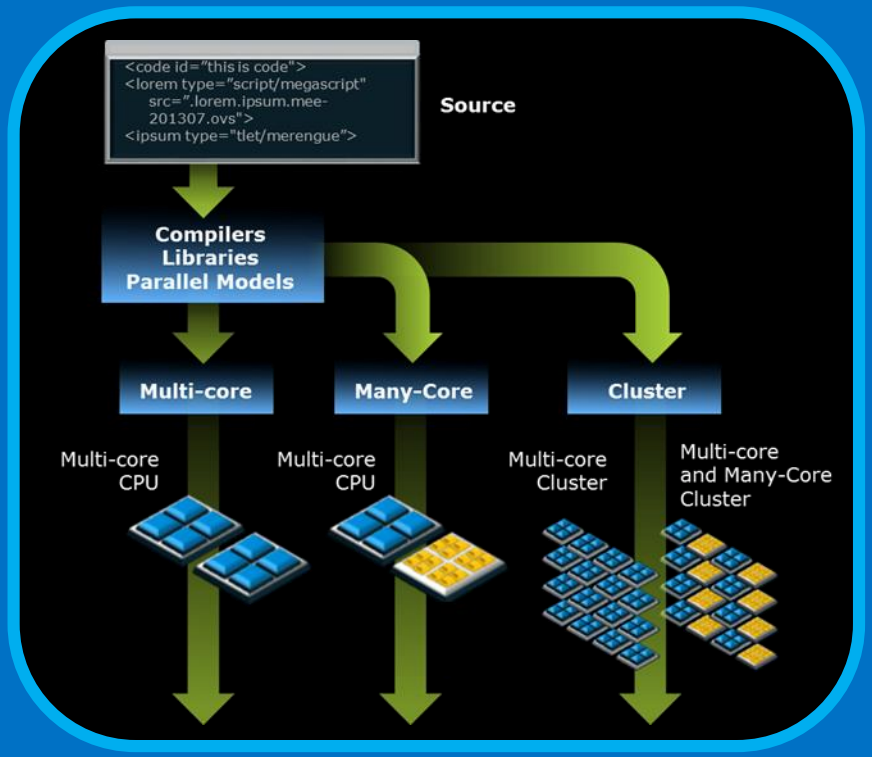

Instruction Parallelism	Data Parallelism	Thread Parallelism	Cluster / Process Parallelism

Serial Code

Fast Scalar performance, Optimized C/C++,FORTRAN, Threading and Performance Libraries, Debug / Analysis Tools

Parallel Node Level

Multi-core, Multi-Socket, SSE and AVX instructions, OpenMP, Threading and Performance Libraries, Thread Checker , Ct

Multi-Node / Cluster Level

Cluster Tools, Cluster OpenMP, MPI Checker

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

The Knights Landing Processor

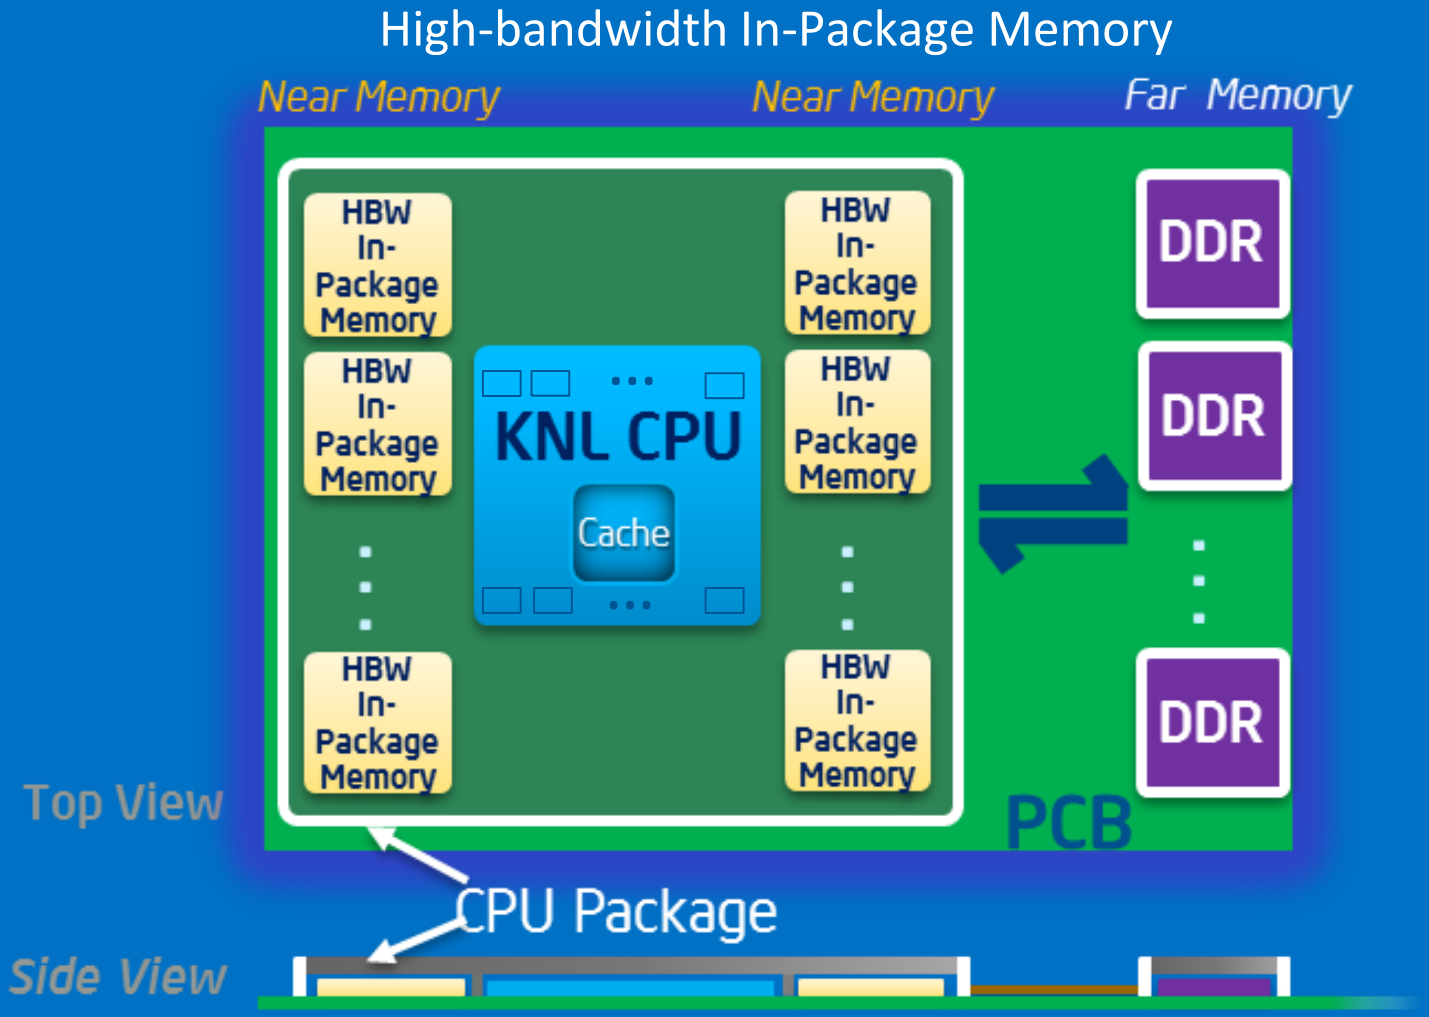

Why Is Acceleration Critical?

Welcome To The New Future Of High Performance Computing!

The Digital Economy: HPC's Next Stage

New Users, New Business Models

3D Printing … New Usages

Economic Activities … Personalized & Democratized

HPC Is intrinsic to the creation and delivery of the Digital Economy economic value

Winds In the East, Mist Coming In …

*"Sized applications of the **Internet of Things** could have direct economic impact of **$2.7 trillion to $6.2 trillion per year** in 2025." – McKinsey*

*World demand for **3D printing** is projected to rise more than 20 percent per year to **$5 billion** in 2017 – Marketwatch.com*

***The total size of digital economy** , including global business-to-business, consumer transactions and the global market for digital products and services is estimated **at $20.4 trillion** – IDC & IDate*

Will HPC Be The Same Ever Again?

HPC: *Explicit* Today, *Implicit* Tomorrow

"HPC for users"

User on client workstation

Workload files ↓ ↑ Return files

HPC Applications

OS + Middleware

On premise or virtual in the cloud

"HPC services for services"

Web APIs

HPC enabled consumer & biz services

Web APIs

HPC Service

OS + Middleware

Cloud based HPC platform

HPC expands beyond …

… highly skilled users (explicit) ***to everyone (implicit)***

… on premise high-end computing ***to services & customized solutions***

… high-end applications ***to new embedded applications & workflows***

Cray & Intel:
Partnering To Create The New HPC Frontier

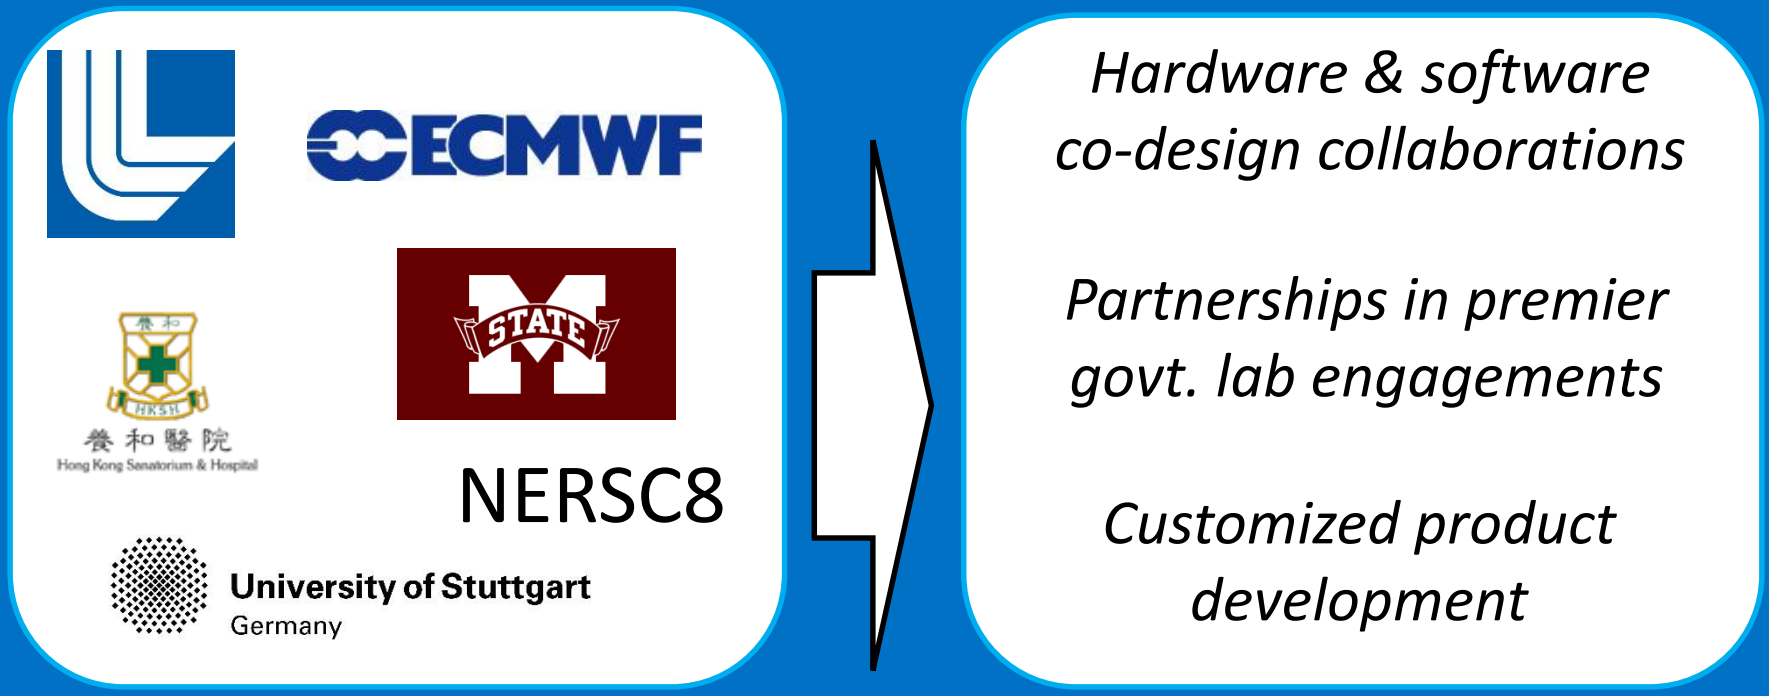

TODAY

TOMORROW

Complementary Capabilities, Unified Vision

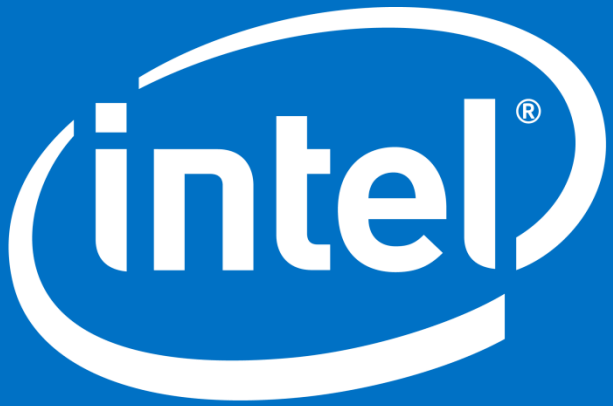

Look Inside.™

Optimization 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 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 specific instruction sets covered by this notice.

Notice revision #20110804

Legal Disclaimers

- 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.
- 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.
- 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.
- Intel® Hyper-Threading Technology Available on select Intel® Xeon® processors. Requires an Intel® HT Technology-enabled system. Consult your PC manufacturer. Performance will vary depending on the specific hardware and software used. For more information including details on which processors support HT Technology, visit http://www.intel.com/info/hyperthreading.
- Intel® Turbo Boost Technology requires a Platform with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your platform manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see http://www.intel.com/technology/turboboost
- Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor series, not across different processor sequences. See http://www.intel.com/products/processor_number for details. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications. All dates and products specified are for planning purposes only and are subject to change without notice
- 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. Product plans, dates, and specifications are preliminary and subject to change without notice
-
- *Other names and brands may be claimed as the property of others.
-