Hewlett Packard
Enterprise

High Performance Computing Solutions for Key FSI Workloads

Lacee McGee, WW Sr. Product Manager

Growing Number of HPC Use Cases

Traditional HPC

- Modeling & Simulation
- More iterative methods (stochastic, parametric, ensemble)
- More SMEs

High Performance Data Analytics

- Today: Knowledge Discovery, BI/BA, Anomaly Detection, Marketing
- Emerging: Precision Medicine, Cognitive, AI, IoT

HPC Anywhere

- On-Premise
- Cloud (Public, Private, Hybrid)
- Private Hosted

IDC Market Analysis Perspective: Worldwide Technical Computing, 2016 New Forces at Play in the High-Performance Computing and High-Performance Data Analysis Sectors

High Performance Data Analytics

3x Growth of HPC Market

- 2019 TAM:$4.9B
- 63% from server systems

Competition = War of Algorithm's

- Divide between machine learning and HPDA algorithms
- Enterprise algorithm's lack parallelism

HPC Meets Big Data

- Shift from extreme compute centric
- Data friendly configurations

In-Memory Solutions

- Dominant by 2019
- Two available strategies
- Energy consumption driving need

HPC Parallelism

- Improving solution times and accuracy
- Dynamic pattern discovery
- Complex problem solving

Dilemma of What to Store

- Data volumes double every 2-3 years
- Bite the bullet
- Monetize risk and compliance

IDC Report: #259921

How is HPE addressing these needs?

Innovation

Purpose Built Infrastructure

Open Source Contribution

Roadmap Alignment

Strategic Partnerships

intel®

CoolIT systems™

End-2-End Solutions

Center of Excellence

Workload Optimized

Hewlett Packard Enterprise

HPE Apollo platforms and solutions optimized for HPC / Big Data

	HPE Apollo 8000 Supercomputing	**HPE Apollo 6500** Rack-Scale GPU computing	**HPE Apollo 6000** Rack-Scale HPC	**HPE Apollo 2000** Enterprise bridge to scale-out compute	**HPE Moonshot** Optimized for workspace mobility and media	**HPE Apollo 4000** Server solutions purpose built for Big Data
Platforms						
Workloads	High performance computing				Specialized	Big Data
	Energy / Oil and gas	**Health / life Sciences**	**Financial services**	**Manufacturing CAD/CAE**	**Academia / Research**	**Mobility / Media**	**Object storage**	**Data analytics**
Solutions and ISVs	Halliburton Paradigm Schlumberger	BIOVIA Gaussian	Altimesh Murex	ANSYS Simulia Synopsys	Custom apps	Citrix Mobile workplace Mobility	Ceph Scality	Cloudera Hortonworks
Tech partners	Intel		Mellanox		NVIDIA		Seagate	
	HPE Technology Services							

Hewlett Packard Enterprise

Deliver Automated Intelligence in Real-time for Deep Learning

Unprecedented Performance and Scale with HPE Apollo 6500 High density Accelerator solution

Use Cases

Automated Intelligence
delivered by HPE Apollo 6500 and Deep Learning software solutions

Video, Image, Text, Audio, time series pattern recognition	Large, highly complex, unstructured simulation & modeling	Real-time, near real-time analytics

Faster Model Training Time, Better Fusion of Data*

Customer Benefits

HPE Apollo 6500 is an ideal HPC and Deep Learning platform providing **unprecedented performance with 8 accelerators, high bandwidth fabric and a configurable accelerator topology to match deep learning workloads**

- Up to 8 high powered accelerators per tray (node), 2P Intel E5-2600 v4 support
- Choice of high-speed, low latency fabrics with 2x IO expansion
- Workload optimized using flexible configuration capabilities

Hewlett Packard Enterprise

* **Benchmarking results provided at or shortly after announcement**

Optimized performance targeting Financial Services Industry

HPE Trade and Match Server Solution

Best in class speed with Leadership Reliability

- Maximum Frequency for HFT Order Execution
- Minimize cache coherent memory operations
- + 20% overclocking Speedups
- Impressive real world benchmark results

Altimesh Hybridizer on HPE Apollo 6000

Code Modernization on HPE Apollo 6000

- Code modernization to help code take advantage of new micro architectures
- Lower customer TCO
- Service oriented transformation project

HPE Risk Compliant Archive Solution

Meet grueling regulations while lowering TCO

- Enterprise Wide Storage Architecture
- Achieve Lowest $/GB
- Verified by Cohasset Association

HPE Moonshot Trader Workstation

Maximize trader productivity

- Match and exceed existing end user experience
- Reduce square meter cooling cost
- Superior compute and graphics performance

Hewlett Packard Enterprise

Performance optimized solutions that address today's FSI challenges and fuel innovation

Fraud Detection

Potential Next Generation Solution

Customer Challenges

Customer experience considerations becoming a driving force

Instant data capture to maximize financial returns / minimize financial

Complex graphs with high-degree of connections

Fraud Detection Solution

- High memory-to-processor ratio to handle the demands of in-memory database applications
- Built in Reliability to help protect applications from down time
- Rich I/O capabilities and flexibility

HPE Integrity MC990 X Server

Specifications:

- 8 Socket Intel Xeon E7
- Up to 192 cores
- 45 MB of L3 Cache
- Memory: Up to 12 TB
- Expansion: Up to 20 slots

Hewlett Packard Enterprise

Hewlett Packard Enterprise

CoolIT systems

Liquid Cooling for HFT Trading

Pat McGinn | VP of Product Marketing
pat.mcginn@coolitsystems.com

September 20, 2016

The world leading manufacturer of energy efficient data center, server and desktop liquid cooling solutions for the HPC, Cloud and Enterprise markets.

The World Leading Liquid Cooling Supplier

15 years in the market

- HQ in Calgary, Canada
- 50 staff
- Taipei, Shenzhen (manufacturing), Rotterdam, Gothenburg, Austin
- Steady growth rate in last five years
- Alberta Exporter of the Year 2015

Industry Awards

Proven technology

- Selling 30-40,000 units/month
- >2M units sold worldwide
- 99.998% leak free and improving
- Intellectual Property: 53 issued patents, 19 pending
- Products offered by major OEMs

Proud to Support

Heat is problematic in Data Centers

Data center managers run out of options:

- Efficiency obstacles, environmental concerns & cost issues
- Server density increases pushing boundaries of traditional air cooling
- CPU performance & longevity reduced
- New “hot” chips push conventional heat boundaries
- Intensive computing increases while power reduction at the chip-level stalls

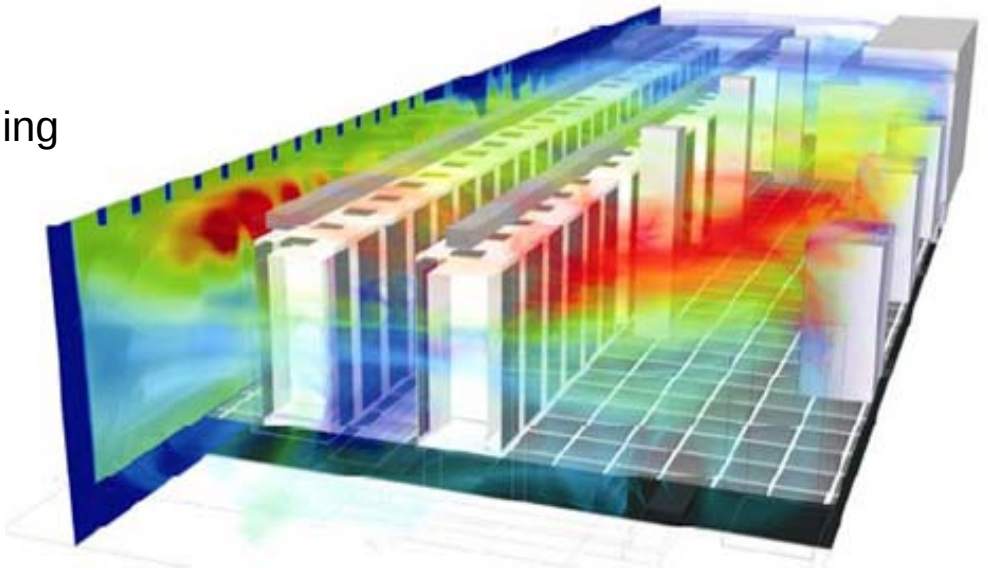

The Solution

Enthusiast

- Desktop
- Overclocking
- Acoustics, Reliability & High Performance

Closed Loop DCLC™

- 1U & 2U rack-mount Servers
- Big Data, HPC/HFC
- Performance & Density

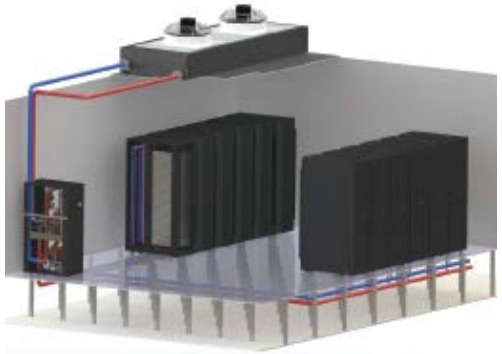

Rack DCLC™

- Rack-level cooling with/without facility water
- Data Centers
- Performance, Efficiency & Density

INCREASING PERFORMANCE, EFFICIENCY & DENSITY

DCLC™ Advantage

Performance

Facilitates peak performance for higher powered or overclocked processors

Efficiency

Provides a significant reduction in total data center energy consumed

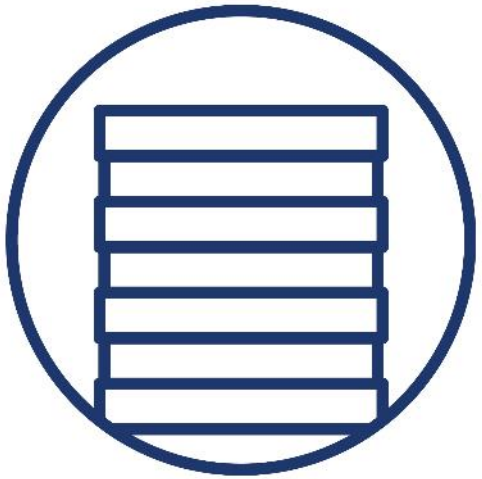

Density

Enables 100% utilization of rack & data center spaces

Direct Contact Liquid Cooling for Servers and Desktops

CoolIT Systems has supplied **Closed-Loop DCLC™** cooling solutions to distributors, system integrators and OEM's for over 10 years.

Anchored by the best-in-class E3 active coldplate assembly, the following components can be used to develop the ideal solution for your application.

CoolIT Systems E3
Active Coldplate Assembly

Features:

- Patented Split-flow technology
- Extremely quiet
- Very low power
- Available Intel, AMD and custom retention

Benefits:

- Thermal Resistance of 0.037 C/W maintains CPU well below specification
- MTTF validated to 80000 hours @ 60C for a long operating life

CoolIT Systems EP2
Active Coldplate Assembly

Specifically designed and optimized for the unique power distribution of the **Intel® Xeon Phi™ X200 Processor family** (previously codenamed Knights Landing or KNL).

- Leverages E3 pump technology and Split-Flow design theory
- Ensures appropriate cooling for both CPU and MCDRAM
- Thermal Resistance of 0.050 C/W
- 1U chassis compatible
- Reference retention scheme
- Includes CPU carrier
- Can be spec’d to supply optimized flow rates for varying radiator requirements

Speed is the essence of HFT Market Trends

Business Trends

- Trading exchanges are providing deterministic behavior and **value added services**
- Improve proximity to trading exchange to **minimize latency**
- **Cost efficiencies** in trading operations to improve ROI

Technology Trends

- Shift from program trading to a more **automated process**
- Leverage **machine learning, big data, and analytics**
- Incorporate **non traditional data points** in decision making process e.g. Twitter feed
- Faster processing through **high performance computing architectures**

CoolIT Systems Liquid Cooling for HFT

- CoolIT is approaching 15,000 HFT systems supported in the market
- Multiple radiators and coldplates to match different server configurations
- Liquid loops can be retrofitted into servers
- Power savings due to improved leakage current and fan reduction
- No impact to server management and serviceability

Enhanced Configuration for HPE Apollo r2000 System

- Four 1U HPE ProLiant XL170r Gen 9 Servers per system
- One E5-1680 v3 processor per server (140 watts)
- Single **CoolIT Systems Closed-Loop DCLC™** device per Server
- Up to four 8GB DIMMs per Server
- Up to twelve LFF or 24 SFF HDDs
- Two 1400W Power Supplies

Speeding up HFT Order Execution

Save Time

- Overclocking capability - optimized for improved frequency
- Improved Price/Performance

Improve ROI

- Optimized for applications that perform better at high frequency and with lower core count
- Save power with lower fan count and less leakage current

Reliability

- Solution utilizes work station processors with ECC Memory
- Run processors cleaner and at lower temperatures

Ease of Deployment

- Plug & Play solution optimized for Co-location data center deployments
- Significantly reduce noise in the compute environment

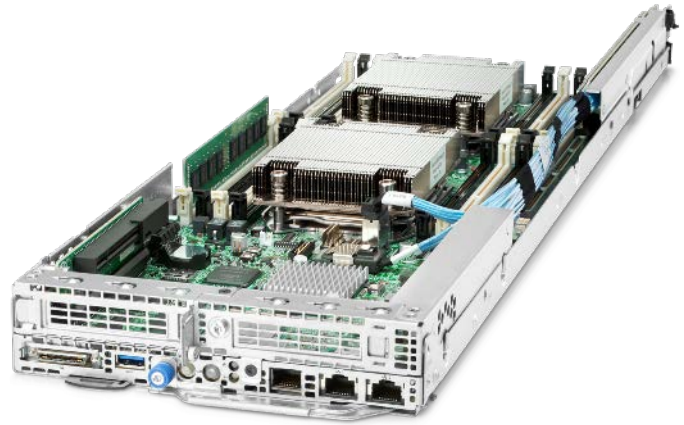

Gain competitive advantage for **High Frequency Trading workloads**

Enhance Performance

Facilitate +18% overclock speedup of four HPE ProLiant XL170r Servers, speeding up HFT order execution.

Optimize Density

Pack more compute into less space and enable a smaller IT footprint. Reduce the need for expansion in existing facilities.

Reduce TCO

Reduce energy consumption and overall TCO. Increase ROI by improving trade operations and reducing time latencies.

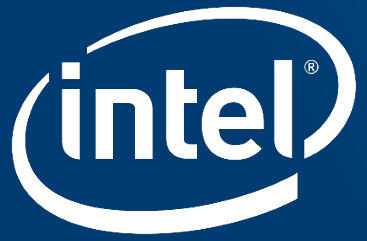

Financial Applications on Intel® Architecture

Nimisha S. Raut
nimisha.s.raut@intel.com

Code Modernization

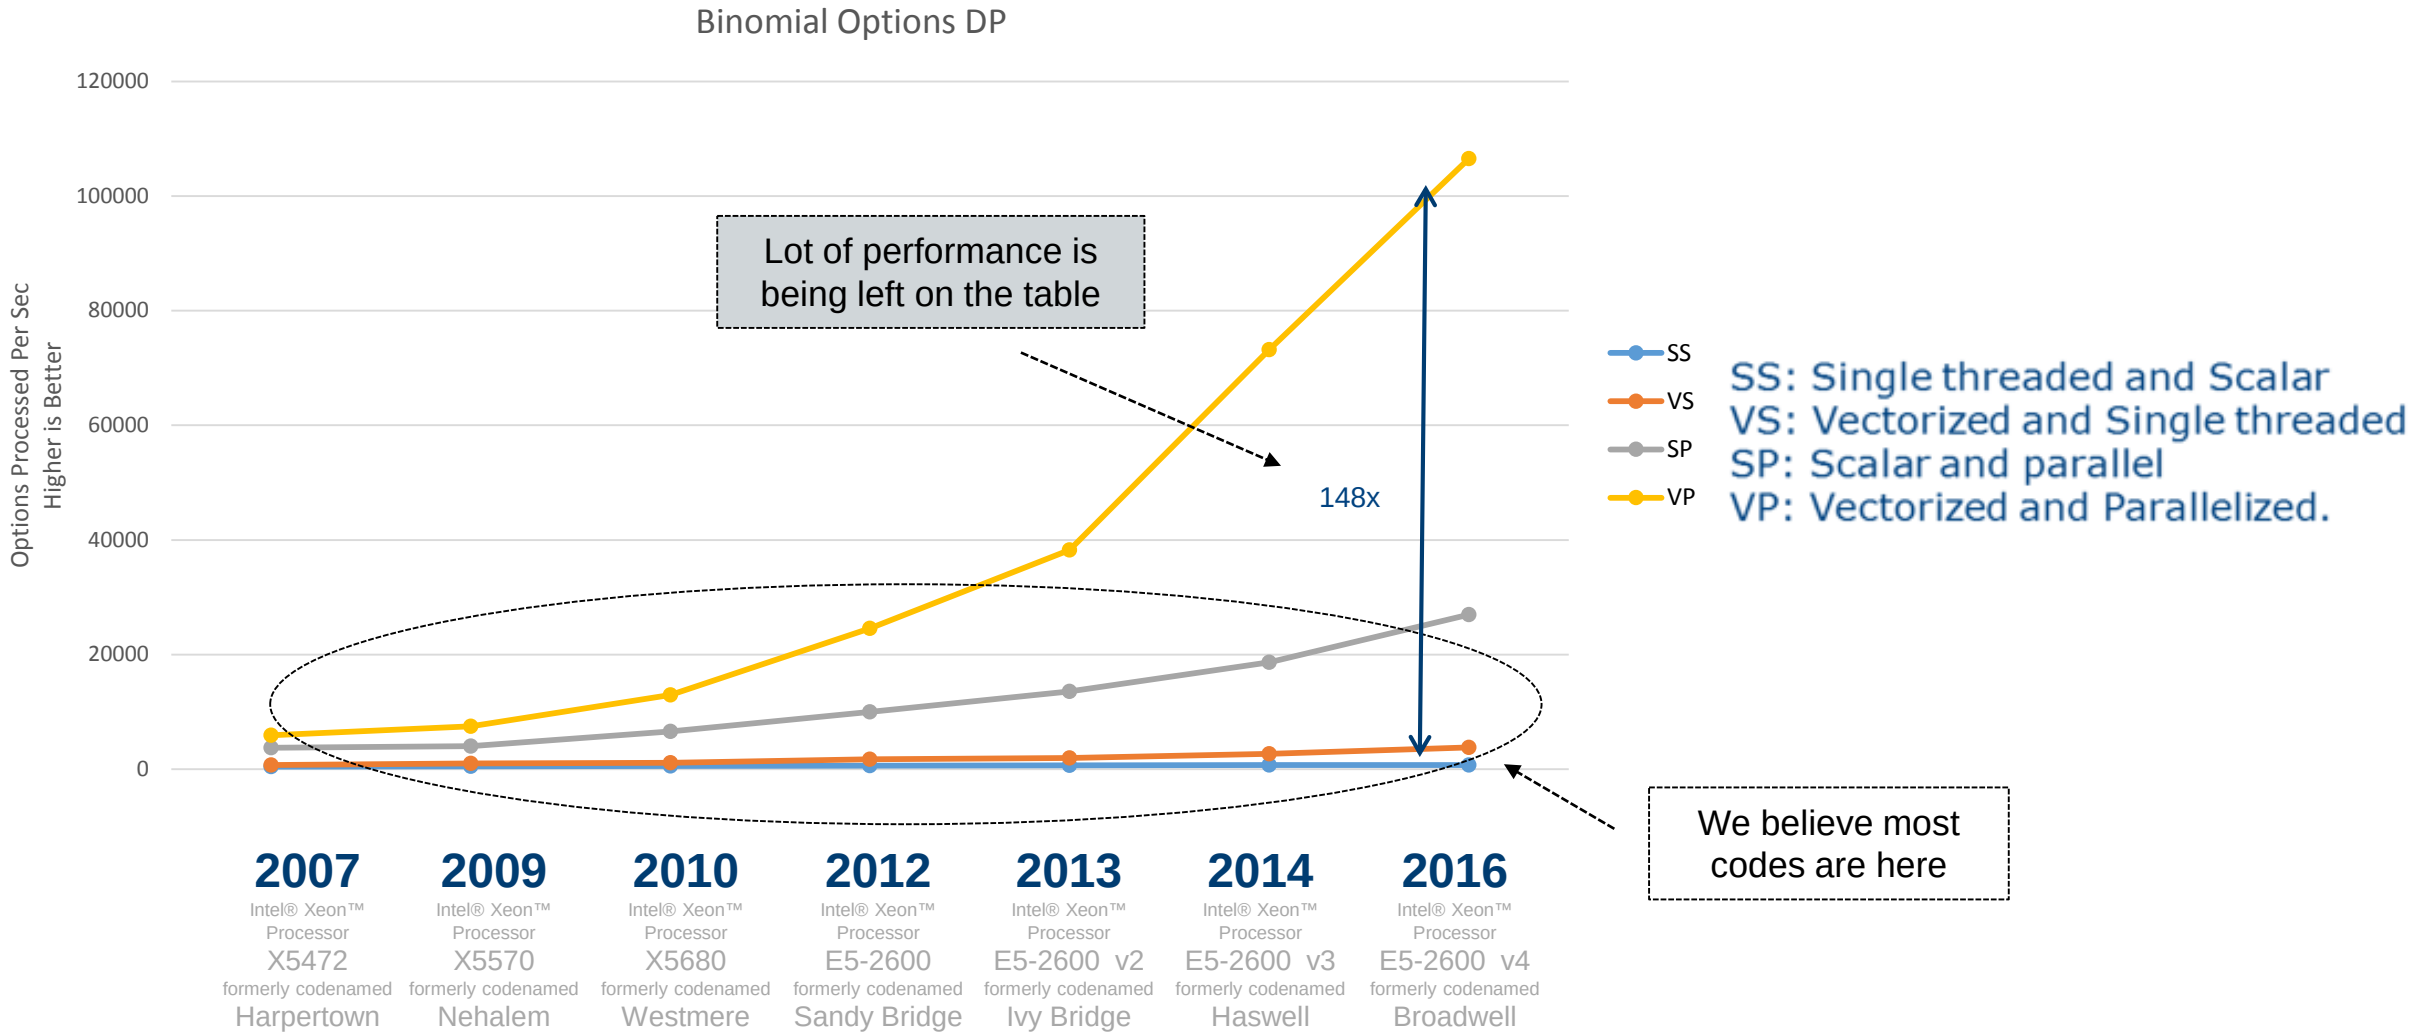

Parallelization and vectorization of your code will maximize your ROI

Intel® Xeon® and Intel® Xeon Phi™ Product Families

More cores → More Threads → Wider vectors

	Intel® Xeon® processor 5400 series 2007	Intel® Xeon® processor 5500 series 2009	Intel® Xeon® processor 5600 series 2010	Intel® Xeon® processor E5-2600 series 2012	Intel® Xeon® processor E5-2600 v2 series 2013	Intel® Xeon® processor E5-2600 v3 series 2014	Intel® Xeon® processor E5-2600 v4 series 2016	Intel® Xeon Phi™ x100 coprocessor	Intel® Xeon Phi™ x200 processor & coprocessor
Up to Core(s)	4	4	6	8	12	18	22	57-61	Up to 72
Up to Threads	4	8	12	16	24	36	44	228-244	Up to 288
SIMD Width	128	128	128	256	256	256	256	512	512
Vector ISA	Intel® SSE4.1	Intel® SSE4.2	Intel® SSE4.2	Intel® AVX	Intel® AVX	Intel® AVX2	Intel® AVX2	Intel® MIC-512	Intel® AVX-512

Knights Landing Overview

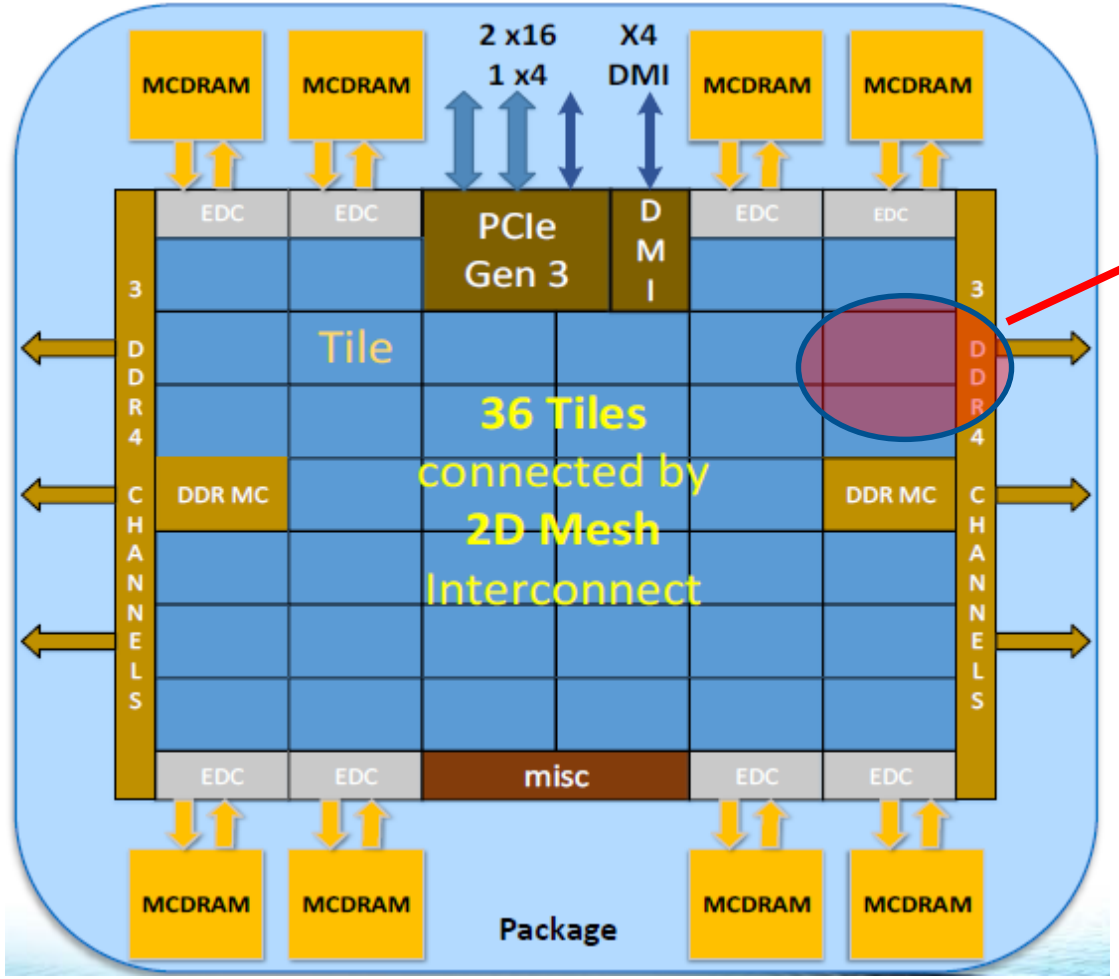

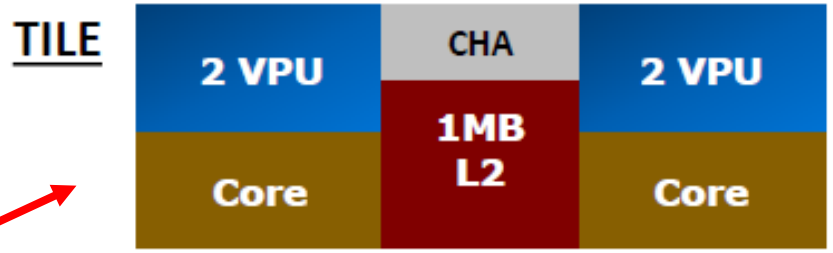

Chip: Up to 36 Tiles interconnected by **2D Mesh**
Tile: 2 Cores + 2 VPU/core + 1 MB L2

Memory: MCDRAM: 16 GB on-package; High BW
DDR4: 6 channels @ 2400 up to 384 GB

IO: 36 lanes PCIe* Gen3 + 4 lanes of DMI for chipset
Node: 1-Socket only
Fabric: Intel® Omni-Path Architecture on-package (not shown)

Vector Peak Perf: 3+TF DP and 6+TF SP Flops
Scalar Perf: ~3x over Knights Corner
Streams Triad (GB/s): MCDRAM : 400+; DDR: 90+

Source Intel: All products, computer systems, dates, and figures specified are preliminary based on current expectations and are subject to change without notice. KNL data are preliminary based on current expectations and are subject to change without notice. 1. Binary Compatible with Intel Xeon processors using Haswell Instructions Set (except TSX), 2 Bandwidth numbers are based on STREAM-like memory access pattern when MCDRAM used as flat memory. Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware & software design or configuration may affect actual performance.

Integrated On-Package Memory Usage Modes

Mode configurable at boot time and software exposed through NUMA[1]

Platform Memory (DDR4) only available for bootable KNL host processor

	Cache Mode	Flat Mode	Hybrid Mode
	64B cache lines direct-mapped 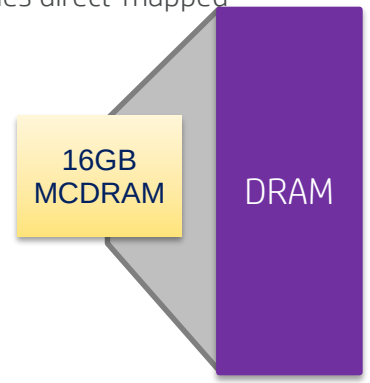	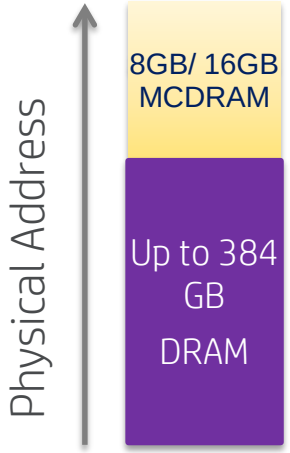	Split Options[2]: 25/75% or 50/50% 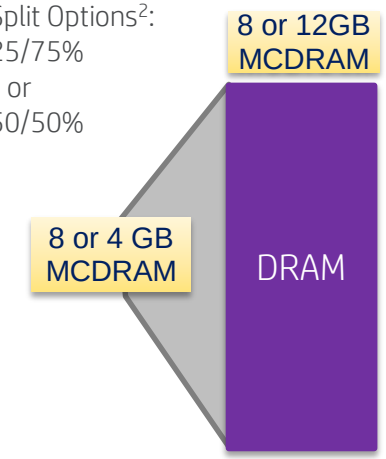
Description	Hardware automatically manages the MCDRAM as a "L3 cache" between CPU and external DDR memory	Manually manage how the app uses the integrated on-package memory and external DDR for peak perf	Harness the benefits of both Cache and Flat modes by segmenting the integrated on-package memory
Usage Model	▪ App and/or data set is very large and will not fit into MCDRAM ▪ Unknown or unstructured memory access behavior	▪ App or portion of an app or data set that can be, or is needed to be "locked" into MCDRAM so it doesn't get flushed out	▪ Need to "lock" in a relatively small portion of an app or data set via the Flat mode ▪ Remaining MCDRAM can then be configured as Cache

1. NUMA = non-uniform memory access
2. As projected based on early product definition

KNL ISA

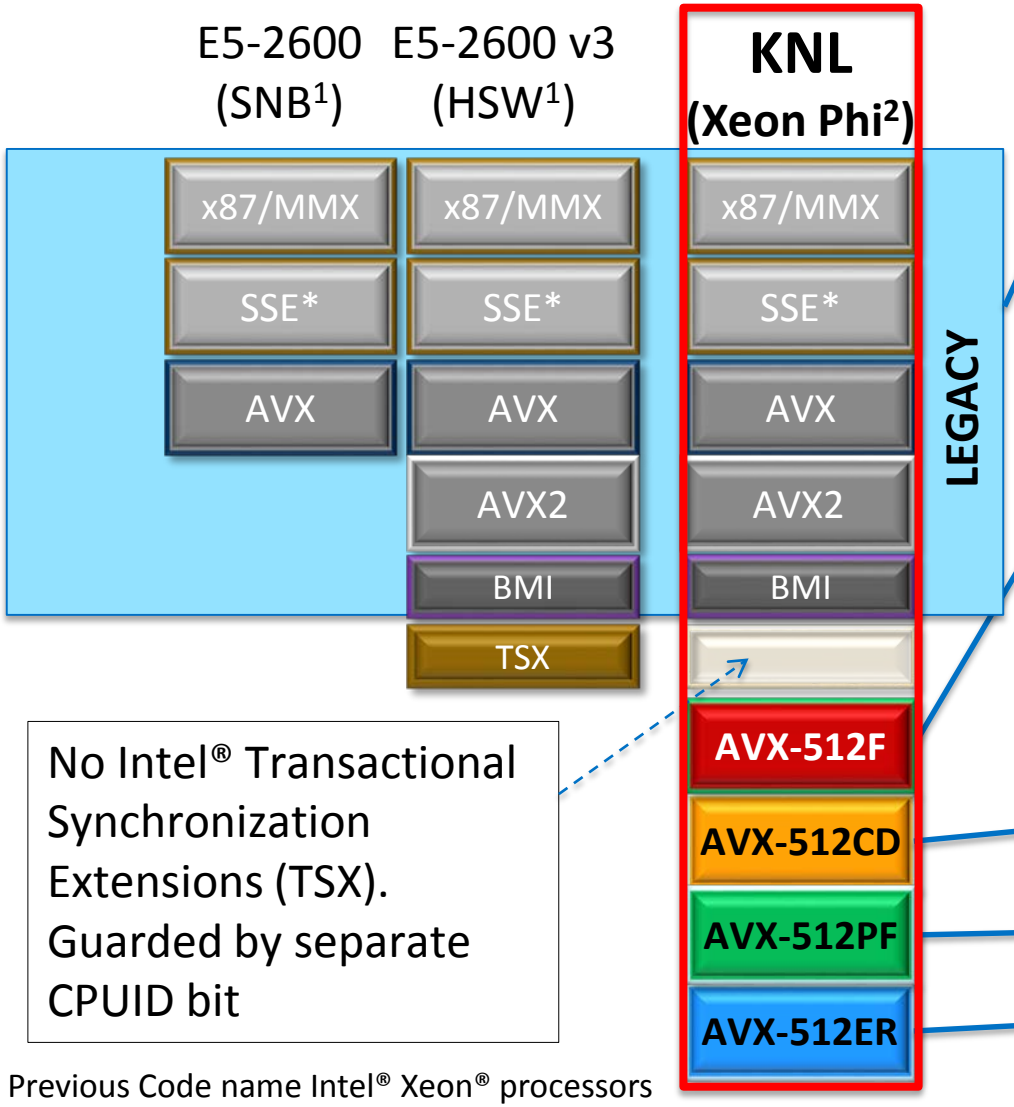

No Intel® Transactional Synchronization Extensions (TSX). Guarded by separate CPUID bit

KNL implements all legacy instructions

- Existing binaries run w/o recompilation

KNL introduces AVX-512 Extensions

- 512-bit FP/Integer Vectors
- 32 registers, 8 mask registers
- Gather/Scatter

Conflict Detection: Improves Vectorization

Prefetch: Gather and Scatter Prefetch

Exponential and Reciprocal Instructions

1. Previous Code name Intel® Xeon® processors
2. Xeon Phi = Intel® Xeon Phi™ processor

STAC-A2* BENCHMARK

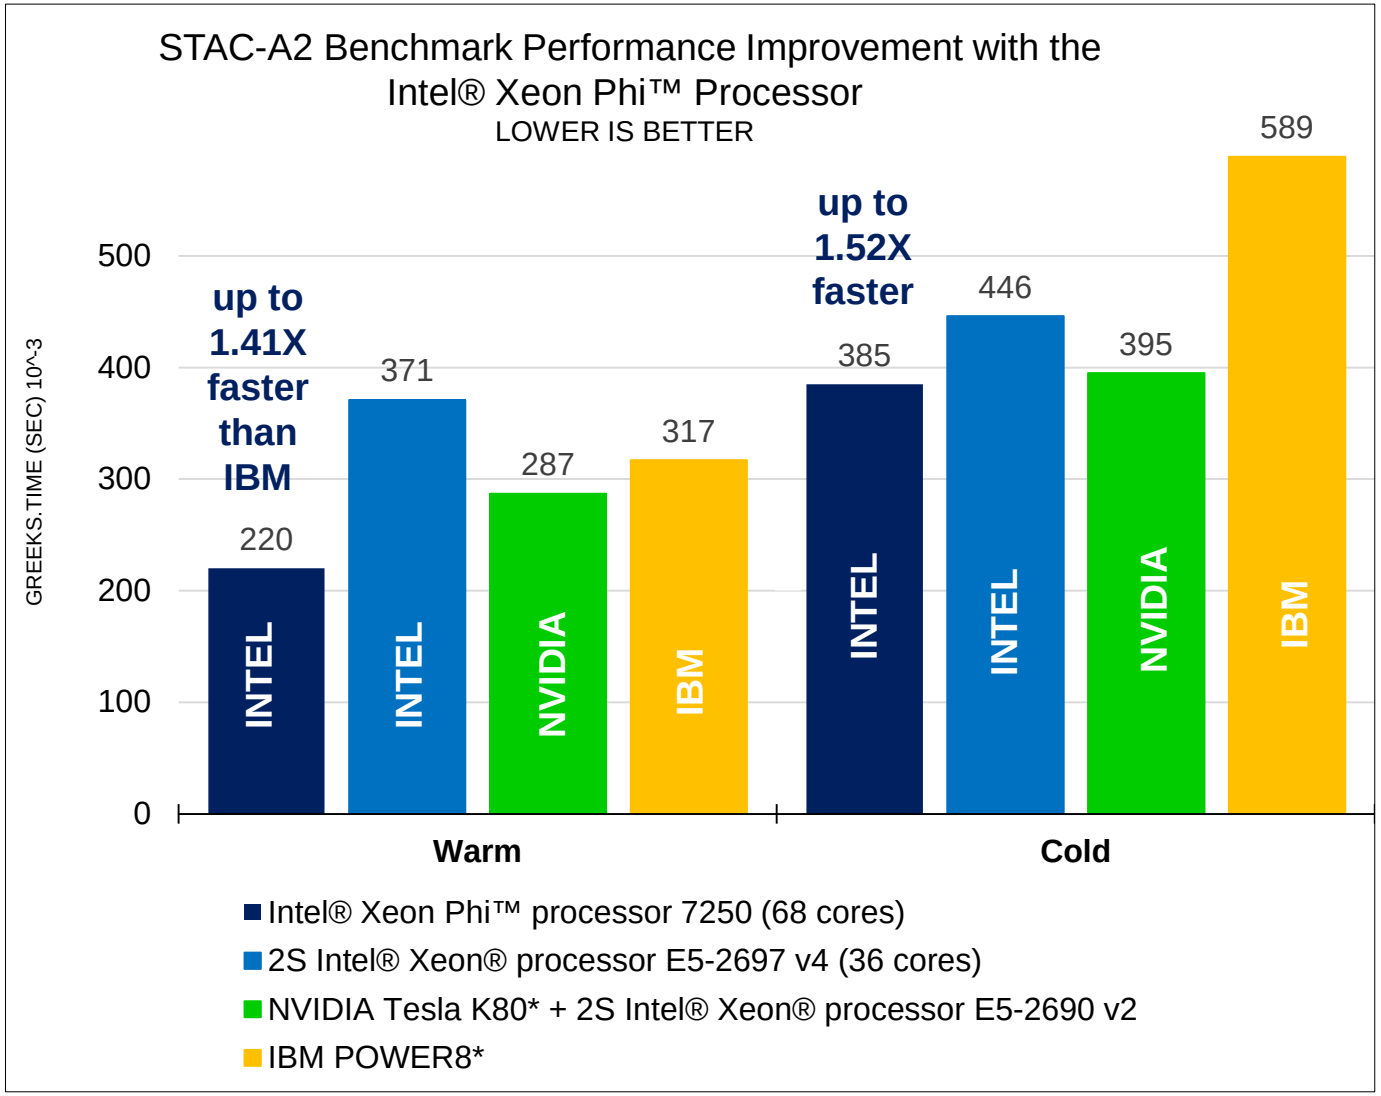

See next slide for configuration details.

SOURCE: STAC* AUDITED RESULTS AS OF MAY 2016

The STAC-A2 Benchmark suite is the industry standard for testing technology stacks used for compute-intensive analytic workloads involved in pricing and risk management.

Application: Intel Composer XE STAC Pack Rev. H

Value Proposition:

- **The Intel Xeon Phi processor based-system takes up to 5.7X less space than the IBM Power8* based-system**
- Created by the Financial community to evaluate SW/HW stacks
- Performance enhanced by Intel® AVX512 and MCDRAM

Results on baseline problem size: The Intel® Xeon Phi™ 7250 processor system is up to 1.2X faster than next competitor (NVIDIA K80* system) in warm runs, was 2X more power efficient than the IBM Power8 system, and had > 4X better space efficiency than competitor systems.

Configuration details: STAC-A2

STAC SUT ID INTC160428 - Intel® Xeon Phi™ processor 7250: Intel® Xeon Phi™ processor 7250 68 core, 272 threads, 1400 MHz core freq. (Turbo ON), MCDRAM 16 GB 7.2 GT/s, DDR4 96GB 2400 MHz, CentOS 7.2, quadrant cluster mode, flat memory mode. See www.STACresearch.com/INTC160428.

STAC SUT ID INTC160314 - Intel® Xeon® processor E5-2699 v4: Supermicro* Superserver SYS-1028GR-TR, Intel® Xeon® Dual Socket ® processor E5-2699 v4 2.2 GHz (Turbo ON) , 22 (HT on) Cores/Socket, 44 Cores, 88 Threads, DDR4 256GB, 2133 MHz, Red Hat 7.2. See www.STACresearch.com/INTC160314.

STAC SUT ID IBM150305 - IBM POWER8™ : IBM Power System* sever, 2x 12-core POWER8* @ 3.52 GHz, 24 cores / 192 Threads (only 96 used), 1 TB DDR3, RH 7.0, IBM XL C/C++ for Linux v13.1. See www.STACresearch.com/IBM150305.

STAC SUT ID NVDA141116 - NVidia® Tesla® K80 : Supermicro* SYS-2027GR-TRHF, Intel Xeon E5-2690 v2, 3.00GHz, 128GB DDR3, 2XGK210B PCI Express GEN3 Dual GPU 2496 Processor cores Base Clock 560MHz Boost Range 562-875MHz 12GB GDDR5 Memory Clock 2.5GHz. NVIDIA CUDA* 6.5 (Driver 340.58), CentOS 6.6 **+** Intel® Xeon® processor E5-2690 v2: 10 Cores/Socket, 20 Cores (HT off), DDR3 128GB. See www.STACresearch.com/NVDA141116.

Legal Notices and 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.

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**.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

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.

Intel, the Intel logo and others 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.

Thank You