STAC Summit 2012 – New York

Low Effort, High Optimization:

STAC Benchmarks & Free Improvements, Courtesy of Moore's Law

Daryan Dehghanpisheh

June 18 - 2012

Join The Discussion

finteligent

#finteligent @finteligent

facebook.com/finteligent

finteligent

Join in the discussion!

How Intel Uses the STAC Benchmarks

- Effective in demonstrating relevant platform and performance advantages of new Intel products.
- A trusted rallying point for ecosystem collaboration
- A credible feedback mechanism for development
- 1H/12 Reports :
  - STAC-M3 (tick databases)
  - STAC-T (tick-to-trade)
  - STAC-A2 (options risk) – Earlier today

STAC-M3 (tick databases)

- STAC-M3 projects are fostering greater industry collaboration
- Major gains YTD, with more expected before EOY.
- STAC Report just released:
  - kdb+ with Intel SSD 320 Series in ION storage appliance
- **Next Phase Tests:**
  - **Sandy Bridge 4-socket**
  - **AVX enabling**

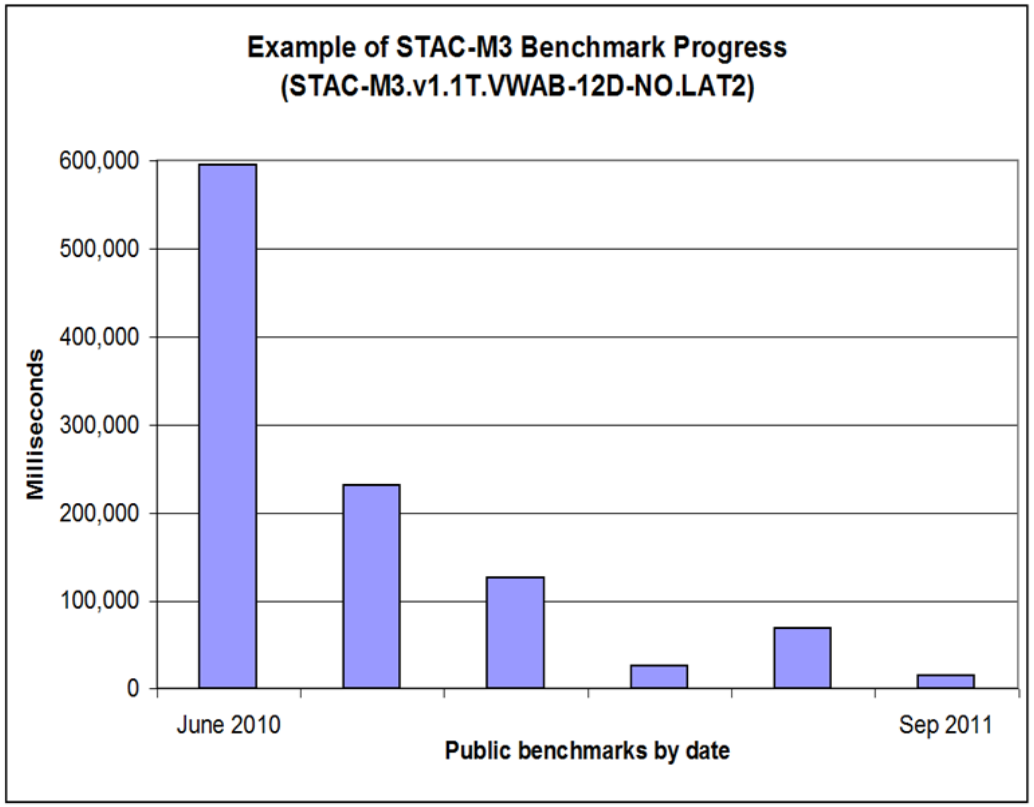

ION SR-71 with Intel SSD 320 Series:

High performance per price, power, and space.

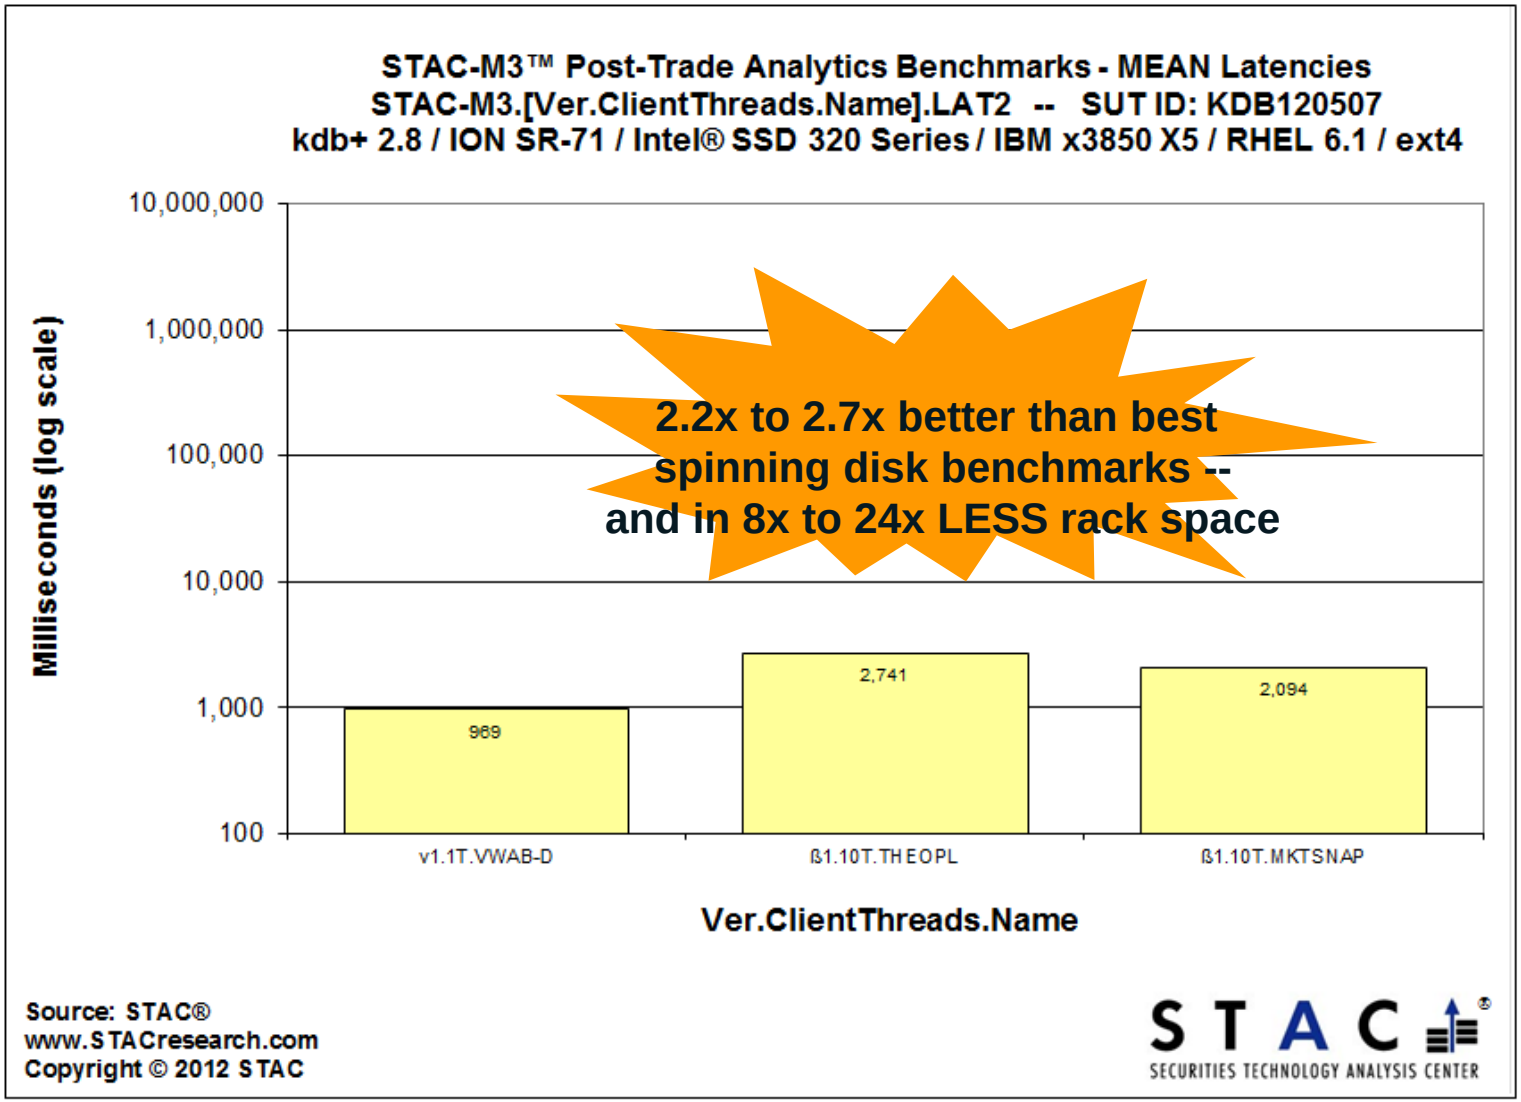

STAC-T: A New Engagement

- STAC-T is a test methodology and software
- Workload is tick-to-trade
  - Inbound: Market data and execution reports
  - Outbound: Trade messages
- Can benchmark any kind of trading system
- Flexible toolset for designing useful benchmarks
- No standard benchmark specifications yet
- Harness can integrate with hardware-based monitoring and playback for highest accuracy

Xeon Platforms tested with STAC-T

Software from Redline Trading Systems

InRush Ticker Plant for CME, version 2.12.7

Execution Gateway for CME, version 1.3.0

"Nearly no-op" algorithm

Server Configuration for Redline Solution

HP DL380p Gen8 (Sandy Bridge) and DL360 (Westmere)

Mellanox ConnectX-3 HCA (10GbE)

RHEL 6.1

Server Platform Comparison: Then & Now

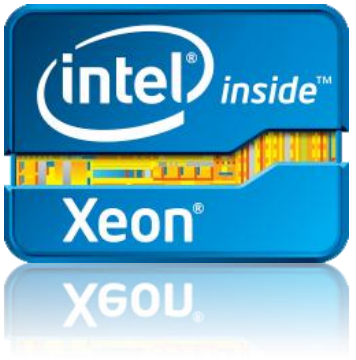

Sandybridge: 16 Cores at 2.9 Ghz

Westmere Platform: 12 cores at 2.93 GHz

Block Level View of STAC-T Harness & Testing

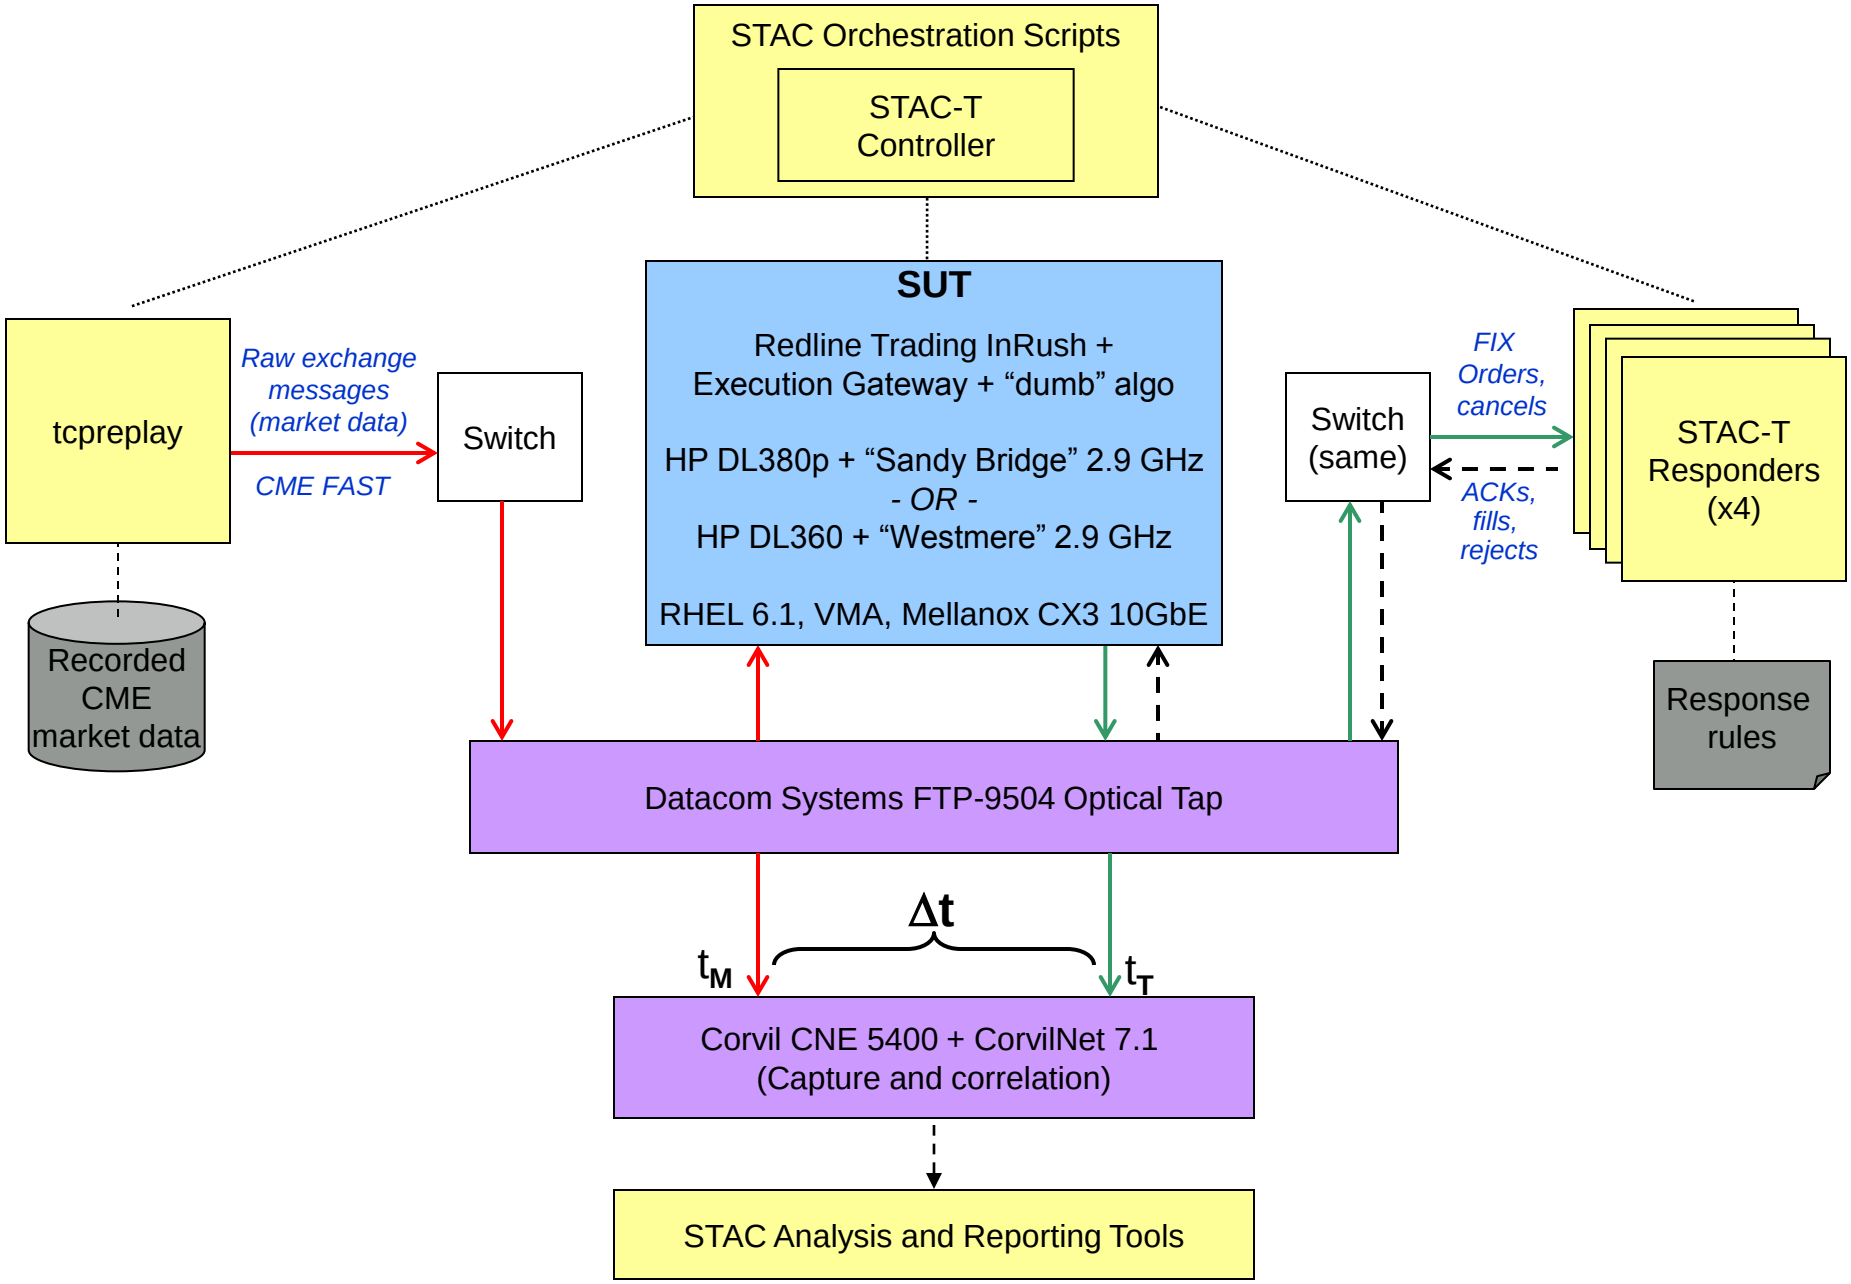

The Result: GHz for GHz, 23% Faster

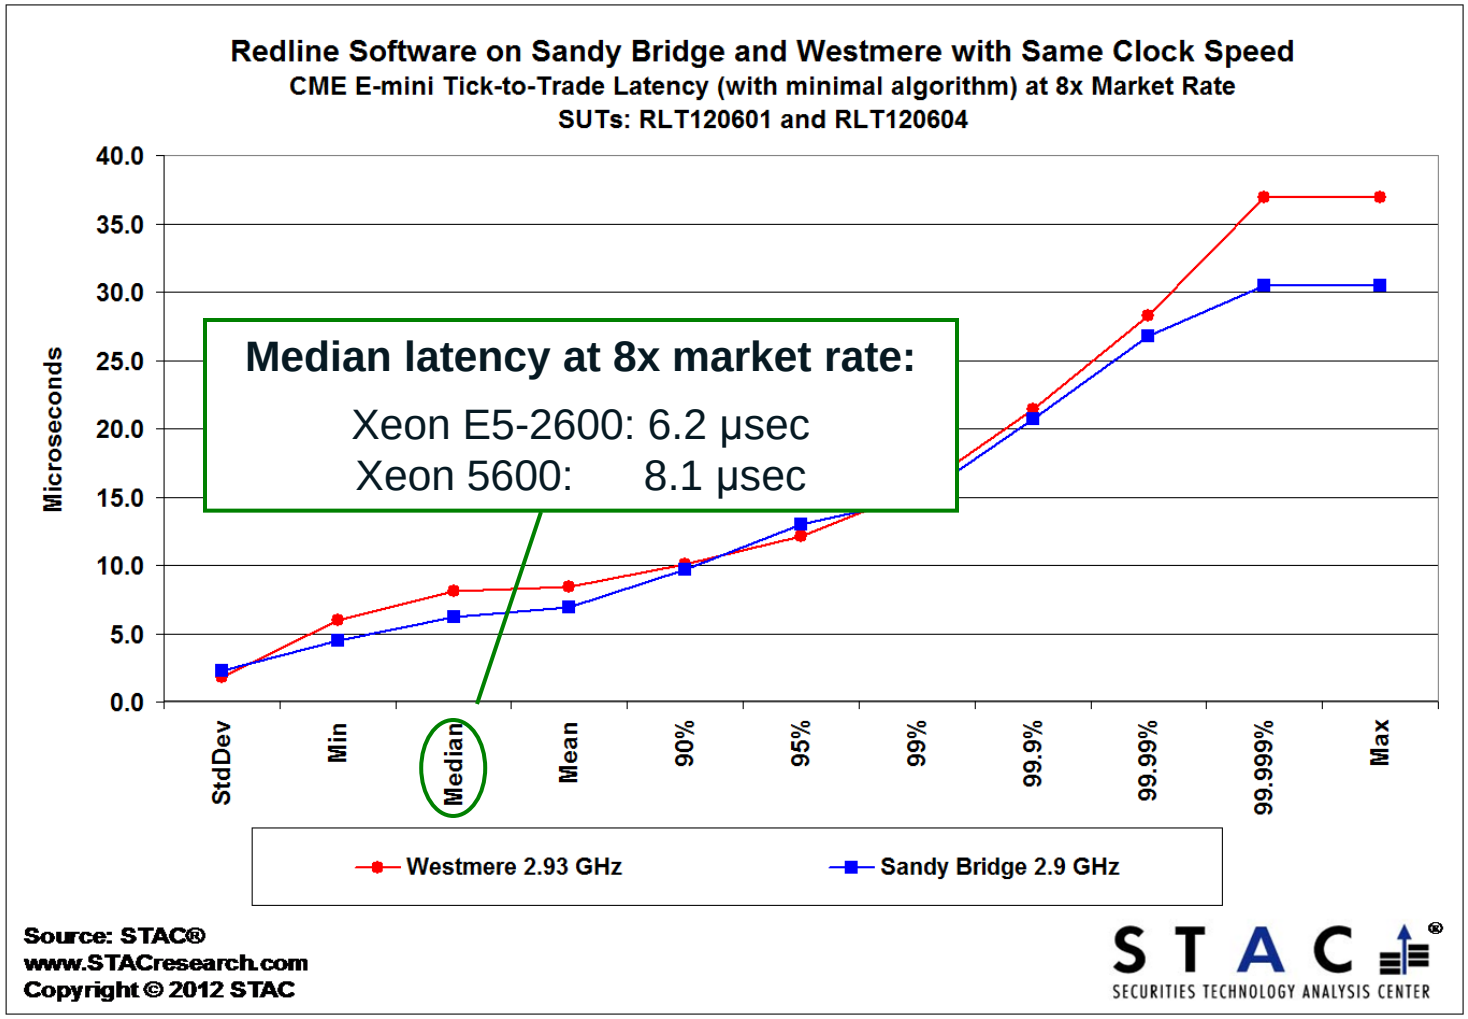

The Heart of a Next-Generation Data Center

Leading Performance
Up to 80% performance boost over Intel® Xeon® processor 5600 series-based servers[1]

Flexible & Efficient
Advanced features automate power consumption across the platform

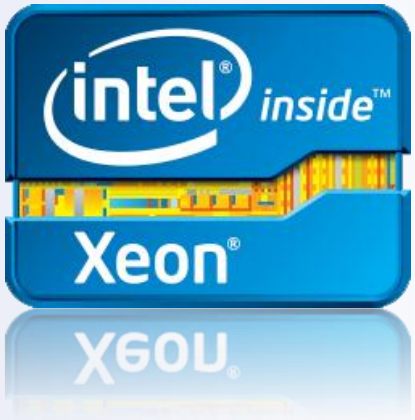

Best combination of performance, power efficiency, and cost

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 intel.com/performance"

1 Performance comparison using best submitted/published 2-socket server results on the SPECfp*_rate_base2006 benchmark as of 6 March 2012. Configuration details in backup

Reduce Bottlenecks With Intel® Integrated I/O

Would you put a
racecar engine in this…

…or this?

Intel® Integrated I/O

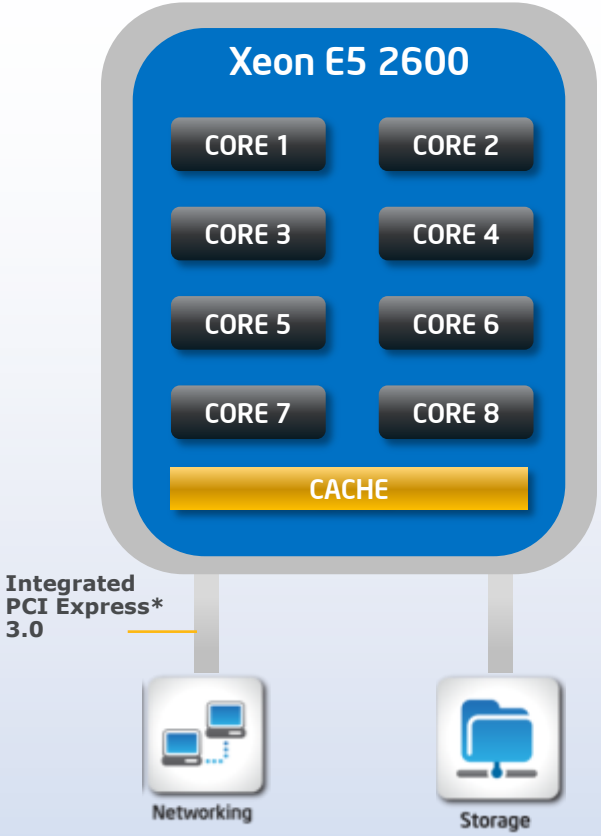

New Intel® Integrated I/O

1st server processor with Intel® Integrated I/O

Reduces I/O latency by as much as 30%[1]

Improves IO bandwidth by as much as 2x with PCI Express* 3.0 support[2]

Intel® **Integrated I/O**

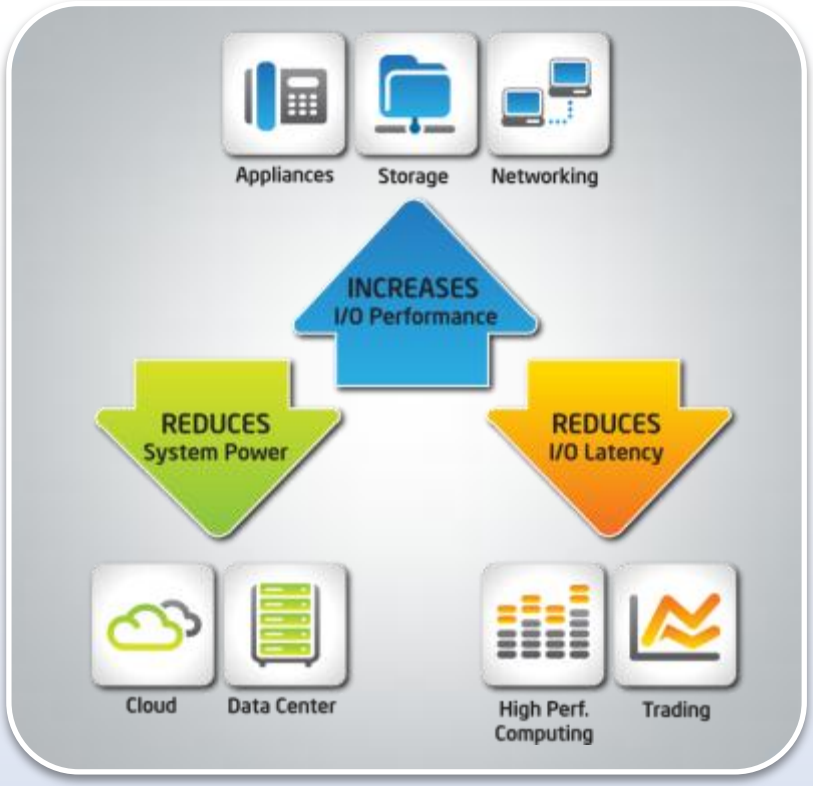

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may 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.

- 1 Source: Intel internal measurements of average time for an I/O device read to local system memory under idle conditions comparing Intel® Xeon® processor E5-2600 product family (230 ns) vs. Intel® Xeon® processor 5500 series (340 ns). See notes in backup for configuration details
- 2 Source: 8 GT/s and 128b/130b encoding in PCIe* 3.0 specification enables double the interconnect bandwidth over the PCIe* 2.0 specification (www.pcisig.com/news_room/November_18_2010_Press_Release/).

New Intel® Data Direct I/O Technology (Intel® DDIO)

Send I/O directly **to and from processor cache** for all I/O traffic types

Can allow system memory to **remain in low power state**

Reduce latency by eliminating unneeded trips to memory

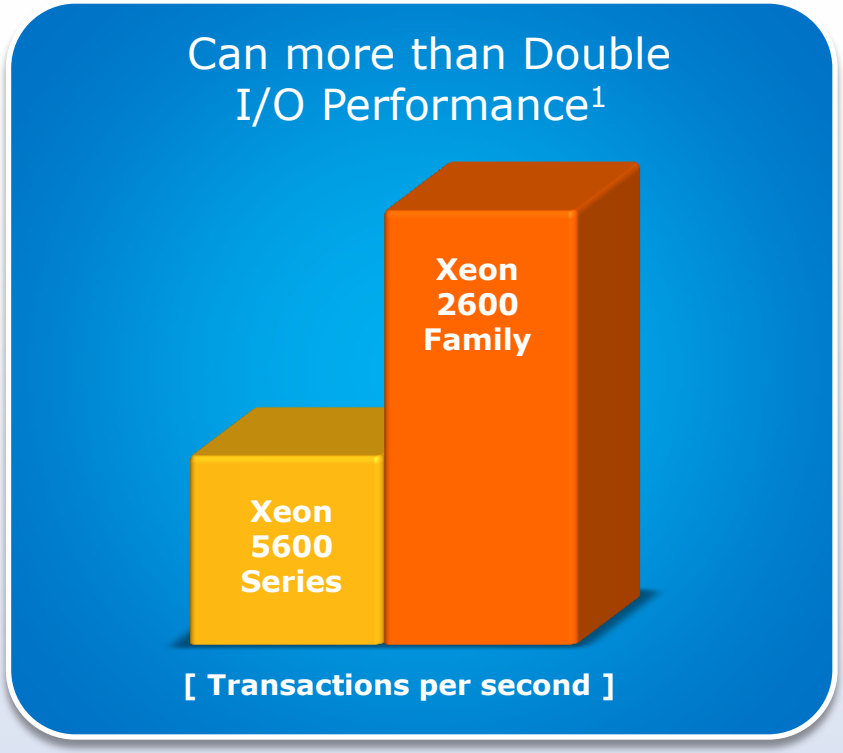

- **1** Up to 2.3x I/O performance is 1S with a Xeon processor 5600 series vs. 1S Xeon Processor E5-2600 data for L2 forwarding test using 8x10GbE ports .See notes in backup for configuration details

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.

The Heart of a Next-Generation Data Center

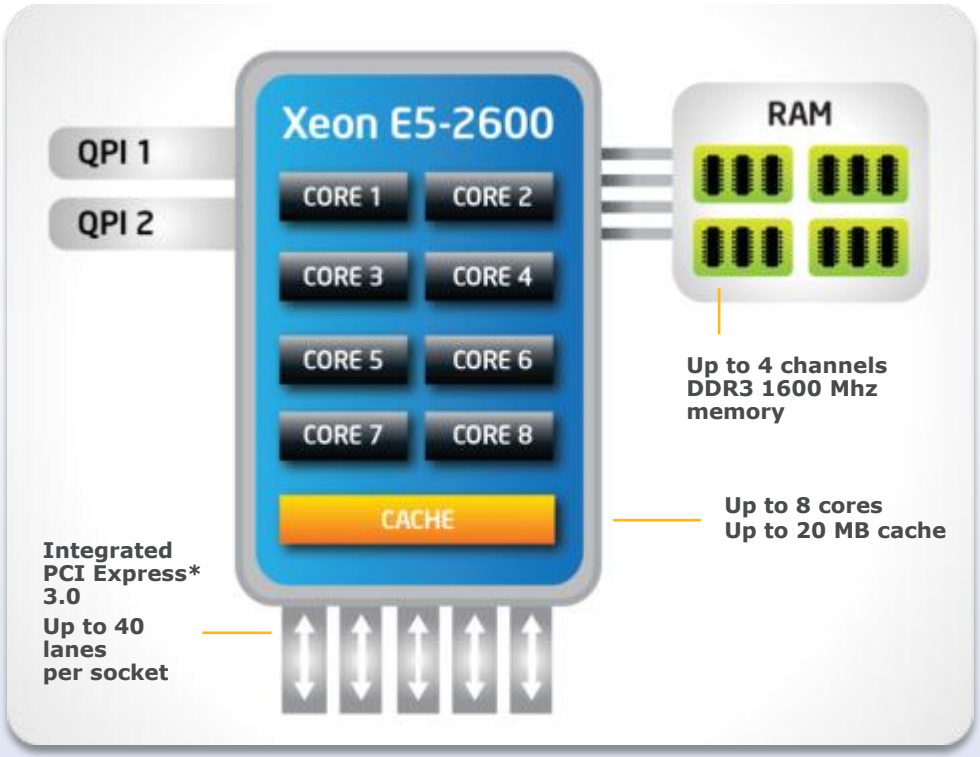

Up to **80% performance boost** vs. prior gen[1]

Dramatically reduce compute time with Intel® Advanced Vector Extensions

Performance when you need it with Intel® Turbo Boost Technology 2.0

Intel® Integrated I/O with Intel® Data Direct I/O **cuts latency**[2] while adding capacity & bandwidth

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 intel.com/performance

1 Performance comparison using best submitted/published 2-socket server results on the SPECfp*_rate_base2006 benchmark as of 6 March 2012.

2 Source: Intel internal measurements of average time for an I/O device read to local system memory under idle conditions comparing Intel® Xeon® processor E5-2600 product family (230 ns) vs. Intel® Xeon® processor 5500 series (340 ns). See notes in backup for configuration details

How to achieve the best performance

Maximizing Performance involves "System Level" optimizations

- OEM BIOS Settings: SMI, HyperThreading, C-States- All Off
  - Experiment with EIST & Turbo On/Off

- On the application: Maximize your resources by…
  1. Pin Threads, Interrupts, and Processes to individual cores using CPU_ID
  2. Place "communication" functions threads on adjacent cores
  3. Use PCM to determine L3 Cache Misses & Keep data in L3 Cache
     - ✓ http://software.intel.com/file/41604
  4. Compile w/Performance Settings, Use PGO, Evaluate IPP / SSE 4.2 Strings
     - ✓ http://software.intel.com/en-us/articles/using-avx-without-writing-avx-code/

- Determine how many cores your trading strategy requires
  1. Can it run on 8 cores? If so, match up CPU+NIC per strategy
     - ✓ https://access.redhat.com/knowledge/solutions/53031

Announcing on June 18th 2012

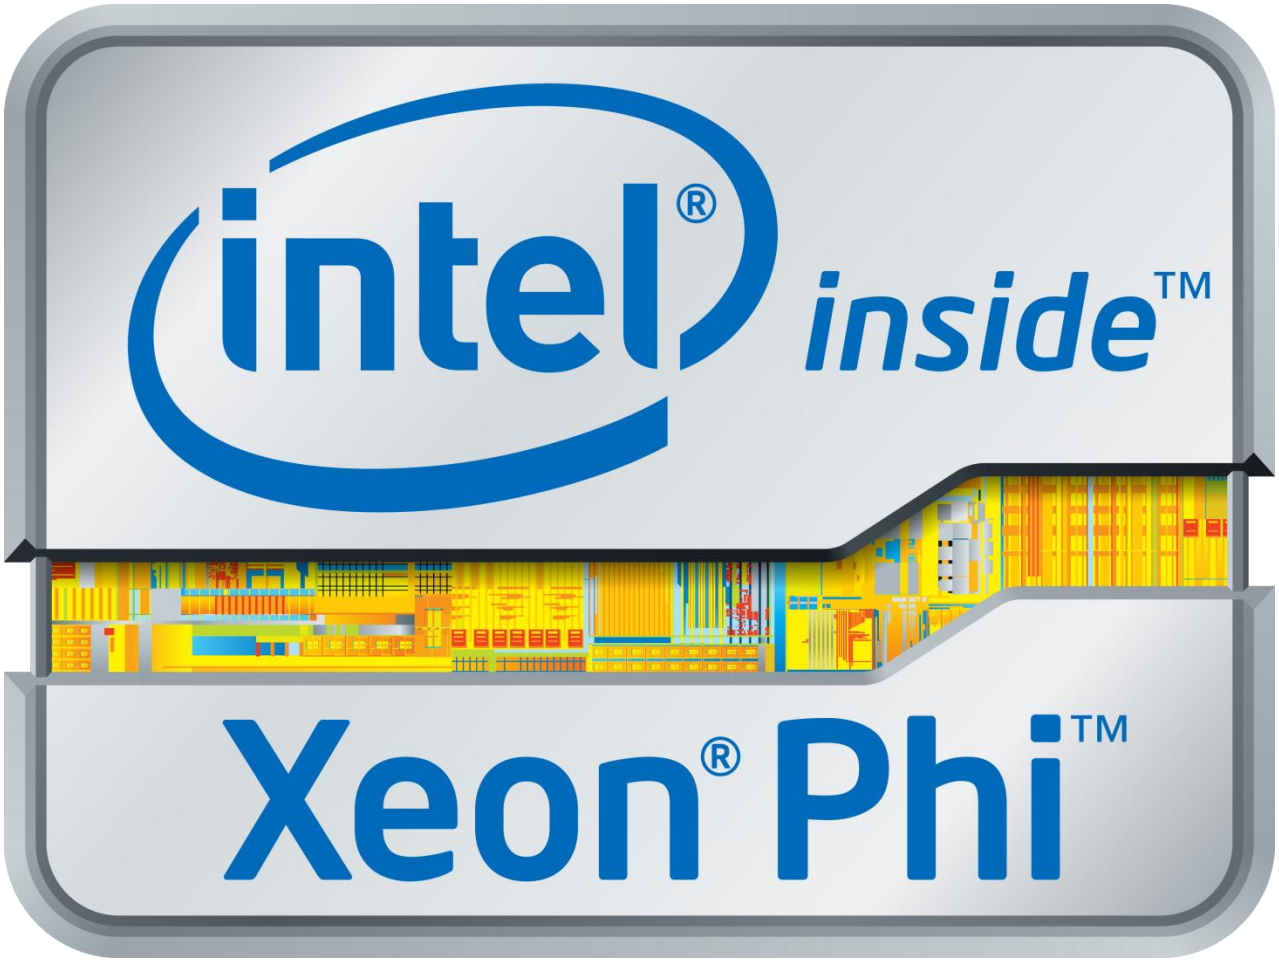

Intel® Xeon® Phi™ Product Family

Intel® Xeon® Phi™ Product Overview

- Product family based on Intel® Many Integrated Core Architecture (Intel® MIC Architecture)
  - Smaller, low power, IA cores with 512-bit wide vector engine
- First product is a coprocessor codenamed Knights Corner
  - PCIe card requiring an IA host processor;
    Next gen Knights Landing is processor and coprocessor (card)
  - Availability: In production in 2012 (Public), product availability in Q4'12 (NDA)
  - Sold as card via OEMs, Expect to price relative to similarly capable competitive alternatives
- >1 Teraflop sustained DGEMM (double precision) perf on early silicon at Nov'11
- Standards-based C/C++/FORTRAN with full support by Intel® Cluster Studio XE
- Use in offload or cluster model as autonomous compute engine
  - "More than an accelerator"
- Supports Data Parallel, Thread Parallel, Process Parallel

Performance and Programmability for Highly Parallel Applications

Next steps

- STAC-T1
  - Extract latency
  - Test higher-throughput scenarios (e.g., TVITCH)
- STAC-A2
  - Increase use of vectorized code
  - Test on Xeon® E5 series in 4-socket configuration
  - Test on Xeon® Phi™ (~50 Cores)
- STAC-M3
  - Test with Xeon® E5 series in 4-socket configuration
  - More Storage Alternatives

Thank You!

STAC

Redline

HP

Corvil

Intel MKL Team

Intel Labs