The Intel® Xeon® Processor E5 Family

Architecture, Power Efficiency, and Performance

Jeff Gilbert, Sr PE
Mark Rowland, PE
August 2012

Agenda

1. Architecting Performance
2. Energy Efficiency from the Load Line to the Data Center
3. Measured Performance

Legal Disclaimer - Notice

- 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.

  UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

  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 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® AES-NI requires a computer system with an AES-NI enabled processor, as well as non-Intel software to execute the instructions in the correct sequence. AES-NI is available on select Intel® processors. For availability, consult your reseller or system manufacturer. For more information, see http://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni/
- No computer system can provide absolute security under all conditions. Intel® Trusted Execution Technology (Intel® TXT) requires a computer with Intel® Virtualization Technology, an Intel TXT-enabled processor, chipset, BIOS, Authenticated Code Modules and an Intel TXT-compatible measured launched environment (MLE). Intel TXT also requires the system to contain a TPM v1.s. For more information, visit http://www.intel.com/technology/security
- Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, and virtual machine monitor (VMM). Functionality, performance or other benefits will vary depending on hardware and software configurations. Software applications may not be compatible with all operating systems. Consult your PC manufacturer. For more information, visit http://www.intel.com/go/virtualization
- Requires a system with Intel® Turbo Boost Technology. Intel Turbo Boost Technology and Intel Turbo Boost Technology 2.0 are only available on select Intel® processors. Consult your PC manufacturer. Performance varies depending on hardware, software, and system configuration. For more information, visit http://www.intel.com/go/turbo
- Intel product is manufactured on a lead-free process. Lead is below 1000 PPM per EU RoHS directive (2002/95/EC, Annex A). No exemptions required
- Halogen-free: Applies only to halogenated flame retardants and PVC in components. Halogens are below 900ppm bromine and 900ppm chlorine.
-

Legal Disclaimers - Performance

- 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.
- 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_ssj, 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.

Optimization Notice

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

Foundations of SNB-EP Performance

Start with the Sandy Bridge Core

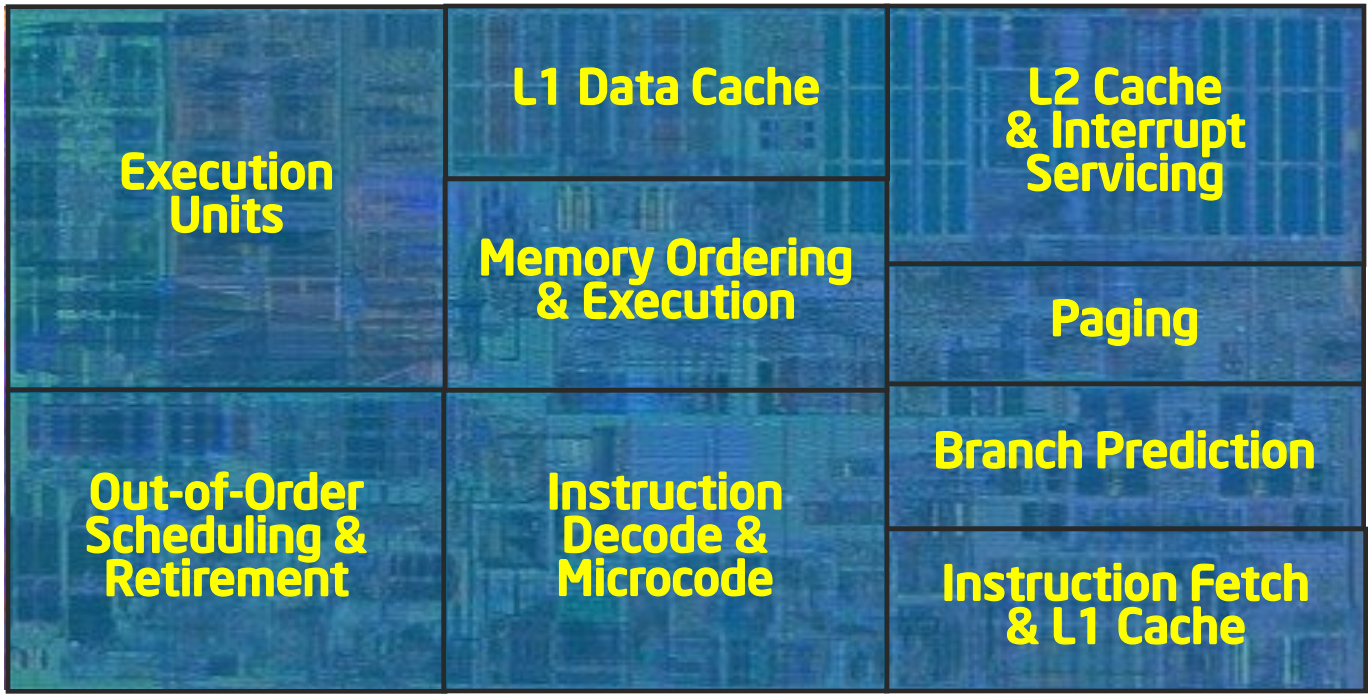

Foundations of SNB-EP Performance

Put Eight Cores on a High BW Interconnect: The Ring

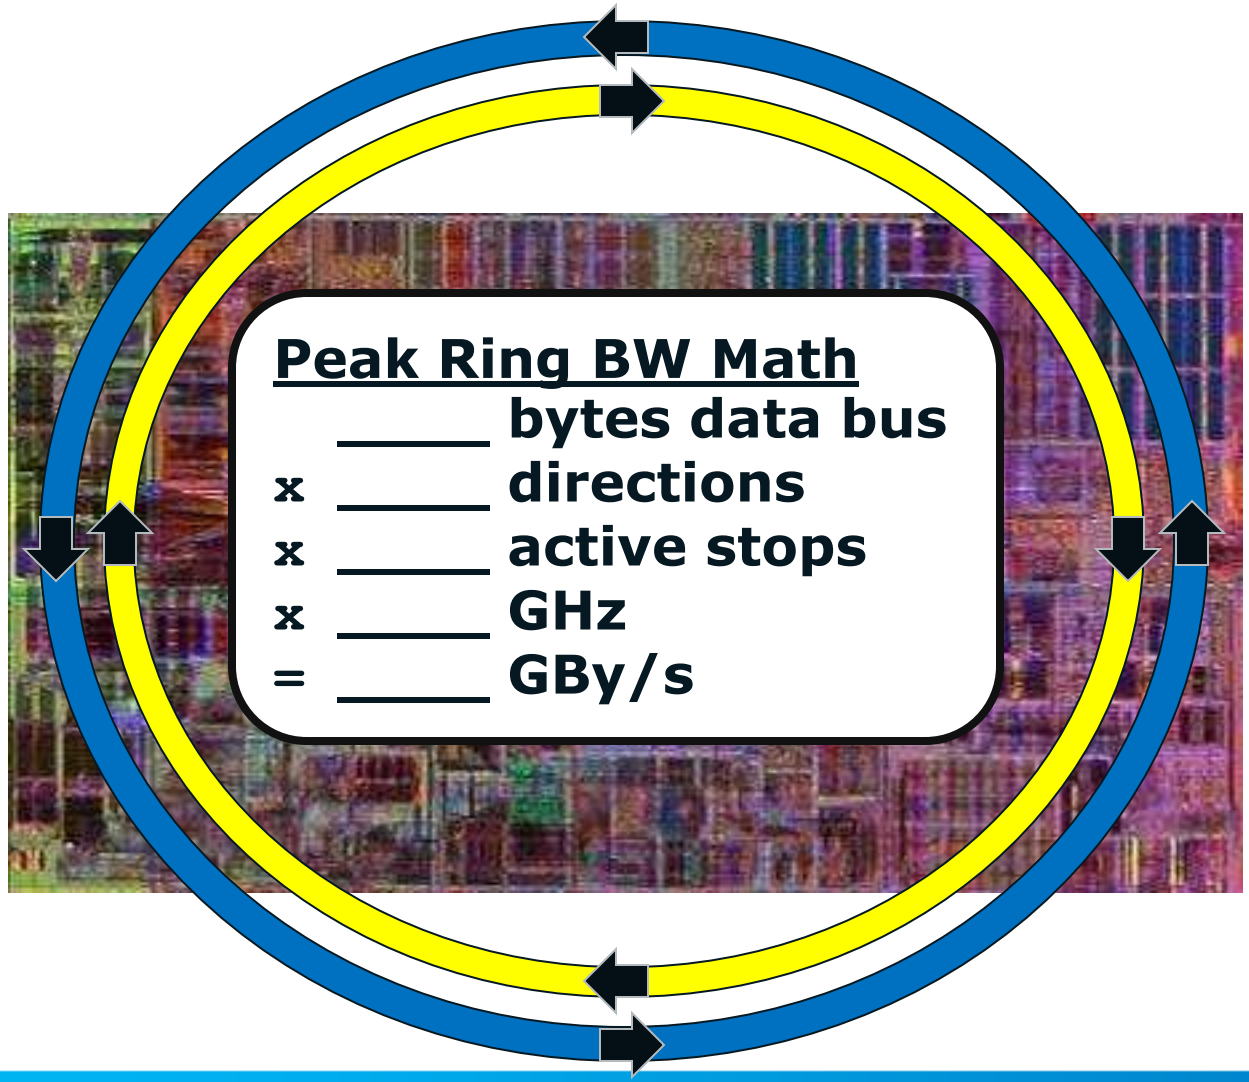

Foundations of SNB-EP Performance

Add an LLC, System Agents, and Power Management

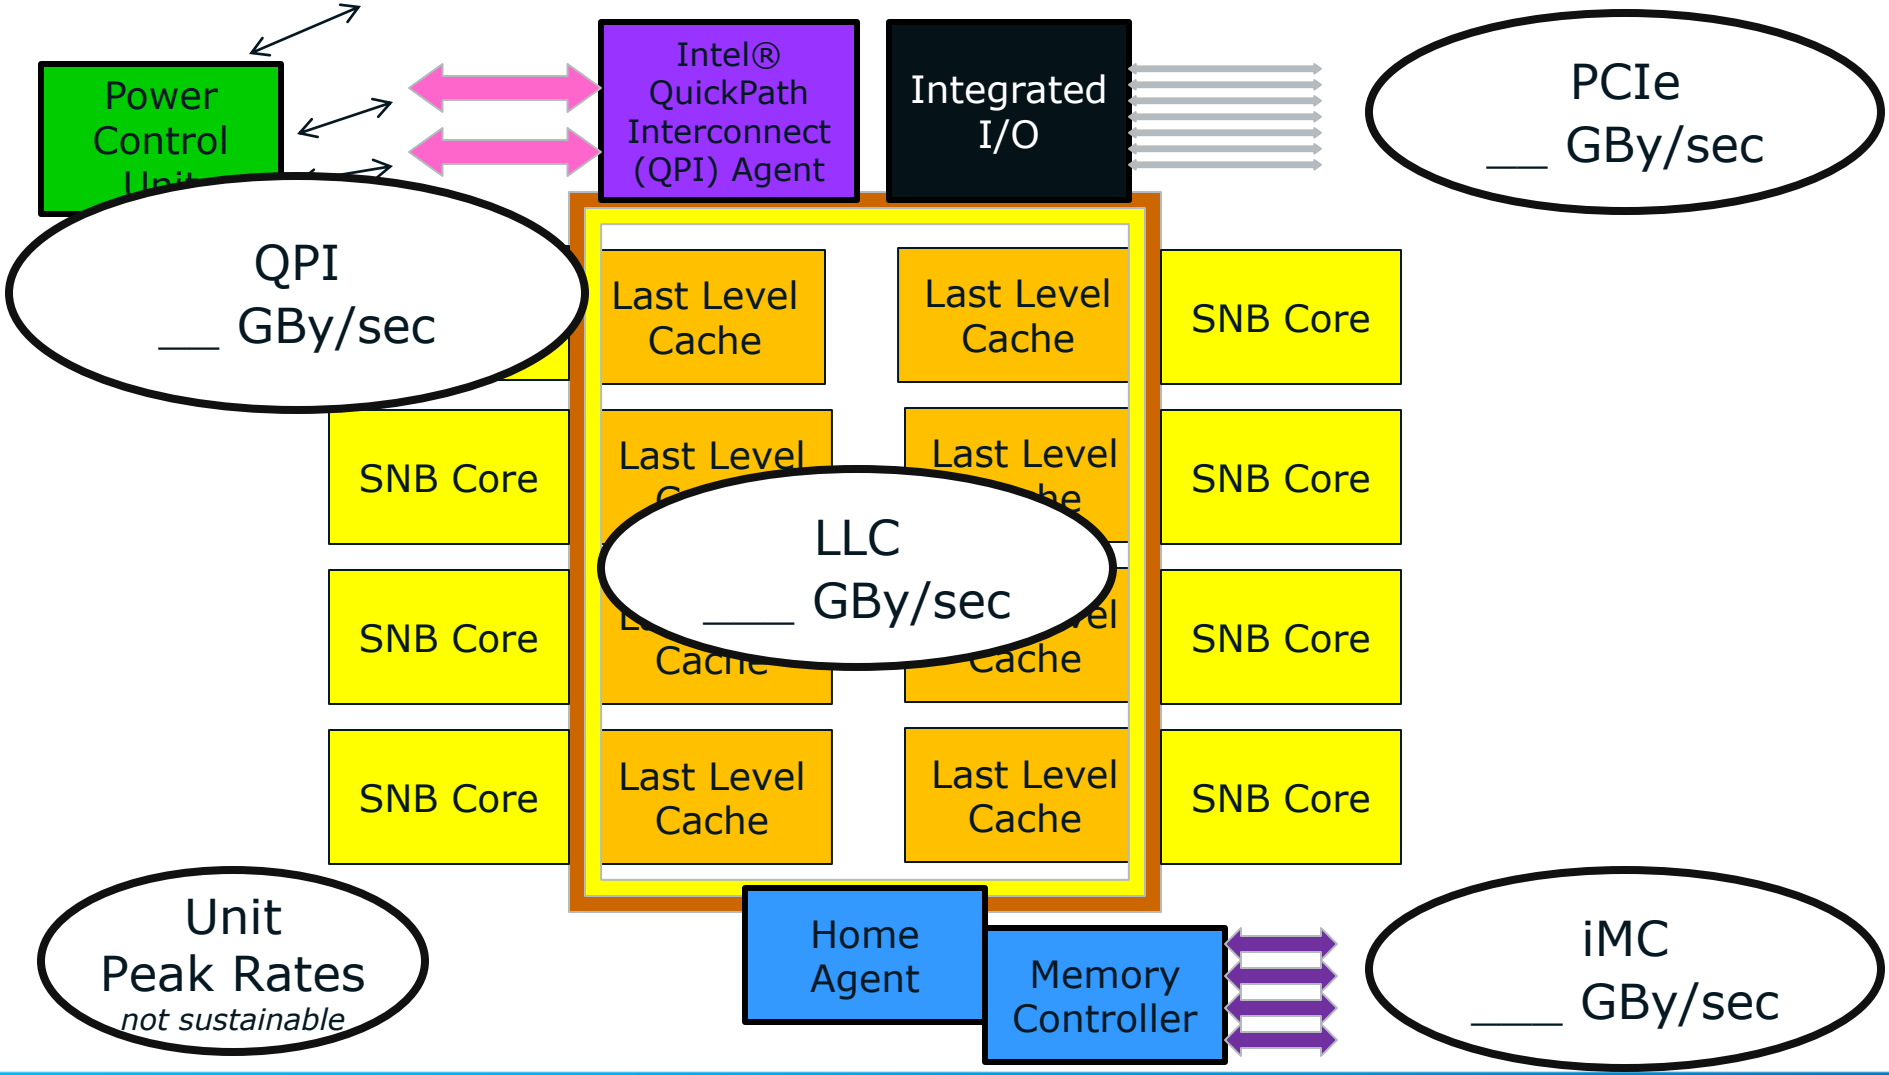

Thurley Platform Review

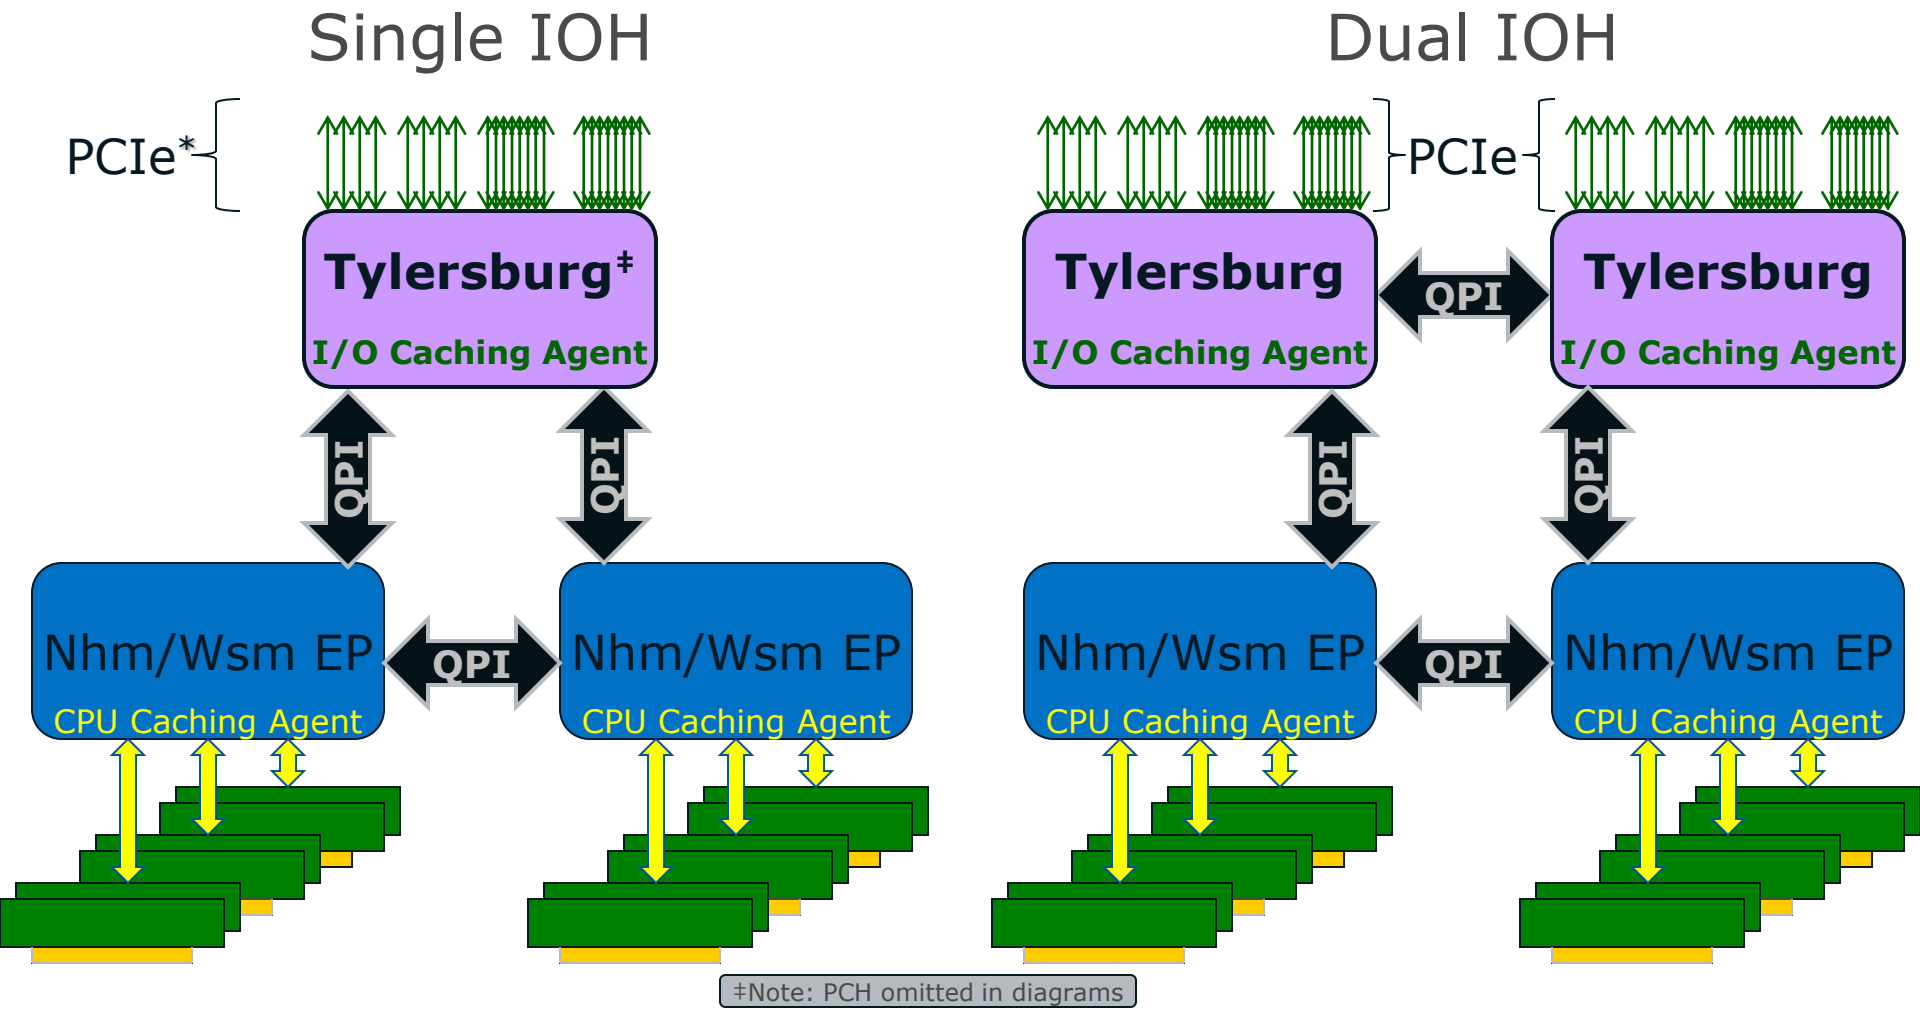

Foundations of SNB-EP Performance

Higher Performance Platform

Romley

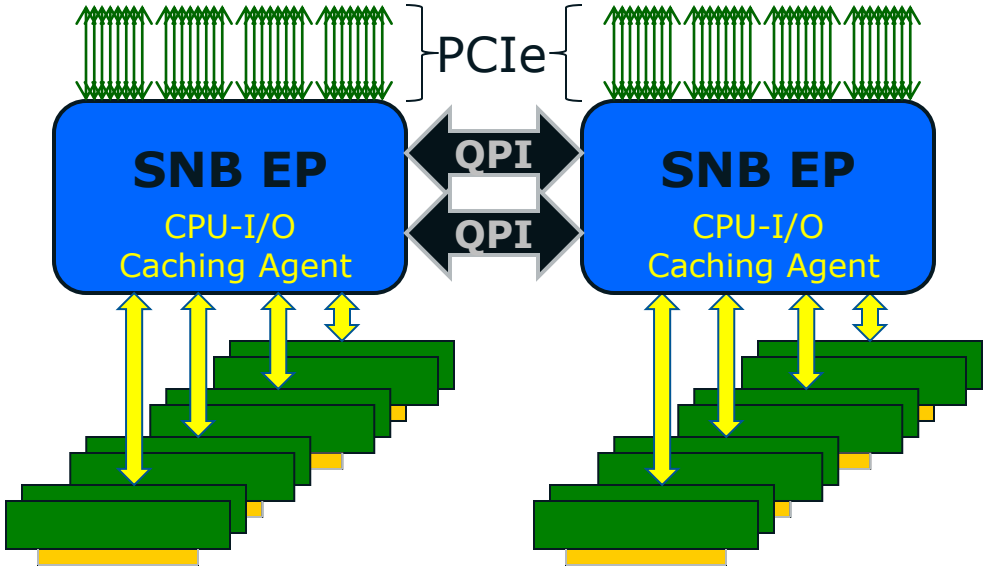

Thurley

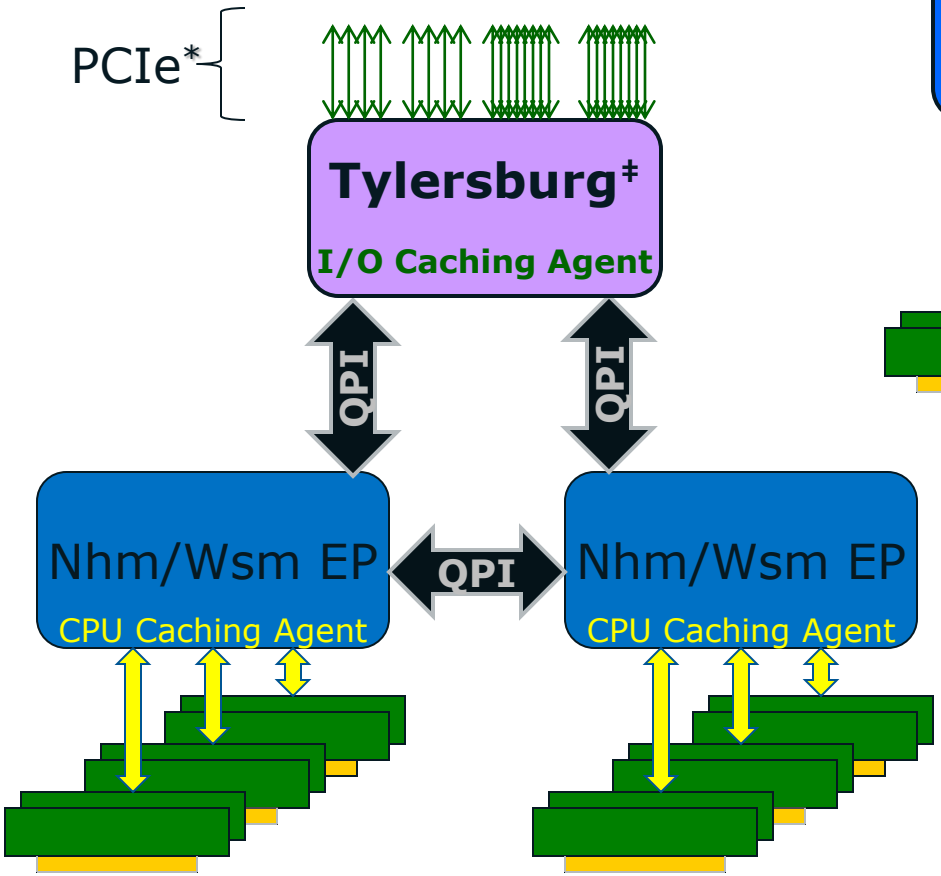

Topology Performance Changes

- 40 Lanes of 8 GT/s Integrated PCIe
- Dual Inter-processor QPI links
- Four higher speed memory channels

‡Note: PCH omitted in diagrams

Foundations of SNB-EP Performance

Focus on I/O Performance

- **PCIe G3: 8 GT/s vs. 4 GT/s**
  - DMI2 (4 GT/s) vs. DMI1 (2 GT/s) (not shown in diagram)
- **I/O capacity scales with sockets (memory BW)**
- **Inherent benefit from Integration:**
  QPI link to I/O controller replaced with direct ring interconnect reducing latency and increasing BW
- **CPUs and PCIe are a unified Caching Agent**
  - Less resource partitioning
    - More scalable, higher performance
  - Reduces the latency of cacheable traffic
  - PCIe acts under the auspices of and uses the LLC (more later)

Foundations of SNB-EP Performance

Focus on I/O Performance (cont'd)

- **I/O-related Optimizations**
  - Double width data buses in the I/O unit
  - ReadCurrent semantics rather the Code Read
    - Potentially reduces memory write traffic – maybe a lot
  - Inbound writes
    - Cache line pre-allocated but ownership can be preempted
    - Prefetch of data (for write merging)
- **40 lanes vs. 36 lanes**
- **Physical address range (46b vs. 41b)**

Foundations of SNB-EP Performance

Focus on I/O Performance (cont'd)

- **Intel® Direct Data I/O Technology (Intel® DDIO): IIO allocates and transfers directly into LLC**
  - IIO cache allocating is generally limited to 2 (of 20) ways
    - Can use a line that's already been allocated by, say, a core
  - Circular buffers of reasonable size (a few to ten MBy) can reside in the LLC *and, in practice, almost never be written*.
  - Making use of this can effectively double the achievable I/O bandwidth of a core and of a socket.
  - Permits practically linear scaling as multiple high bandwidth I/O devices are added (e.g., 10 GbE adapters) with achieving nearly zero read and write bandwidth to memory
    - Saves power, too

Mid-Game Summary

- **Improved performance by improving the parts**
  - Sandy Bridge core
  - On-die interconnect (“Ring”)
  - More and faster memory channels with improved scheduling
  - Faster inter-socket communication (Intel® QPI)
  - Integrating and accelerating I/O

- **Coming Up in the Next Half:**

  Performance with Power Efficiency

Energy Efficient Load Line

Energy Efficient Load Line

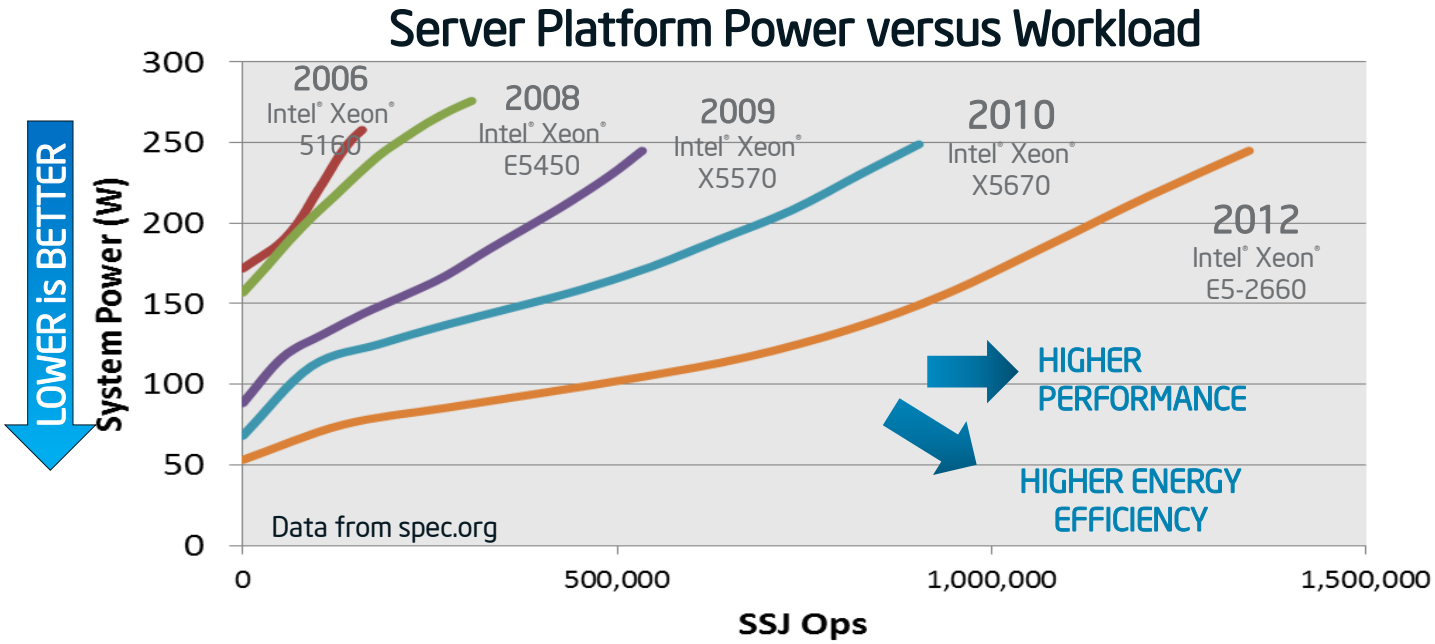

Performance:
45% CAGR

Energy Efficiency:
60% CAGR

- **Platform efficiency at low Power**
  - CPU and DRAM VR Phase shedding
- **Scalable Uncore Power**
  - Uncore voltage frequency scaling
- **Scalable Memory Power**
  - Multi-rank slow CKE
- **Processor Power**
  - Energy Perf BIAS, Dynamic Switching
- **I/O Power management**
  - QPI L0p/L1, PCIE ASPM L1

Significant Improvement to Proportional Energy

Dynamic Performance Load Line

PCU dynamically adjusts to OS Power Management Policy

- OS communicates Policy through EPB (Energy Perf BIAS)
- PCU monitors and adjusts autonomous on die power saving engines

PCU automatically adjusts for Performance at high utilization

- Leverages EPB to switch into performance mode when necessary

Optimized across a range of workloads

- Single-threaded workloads
- Multi-threaded workloads

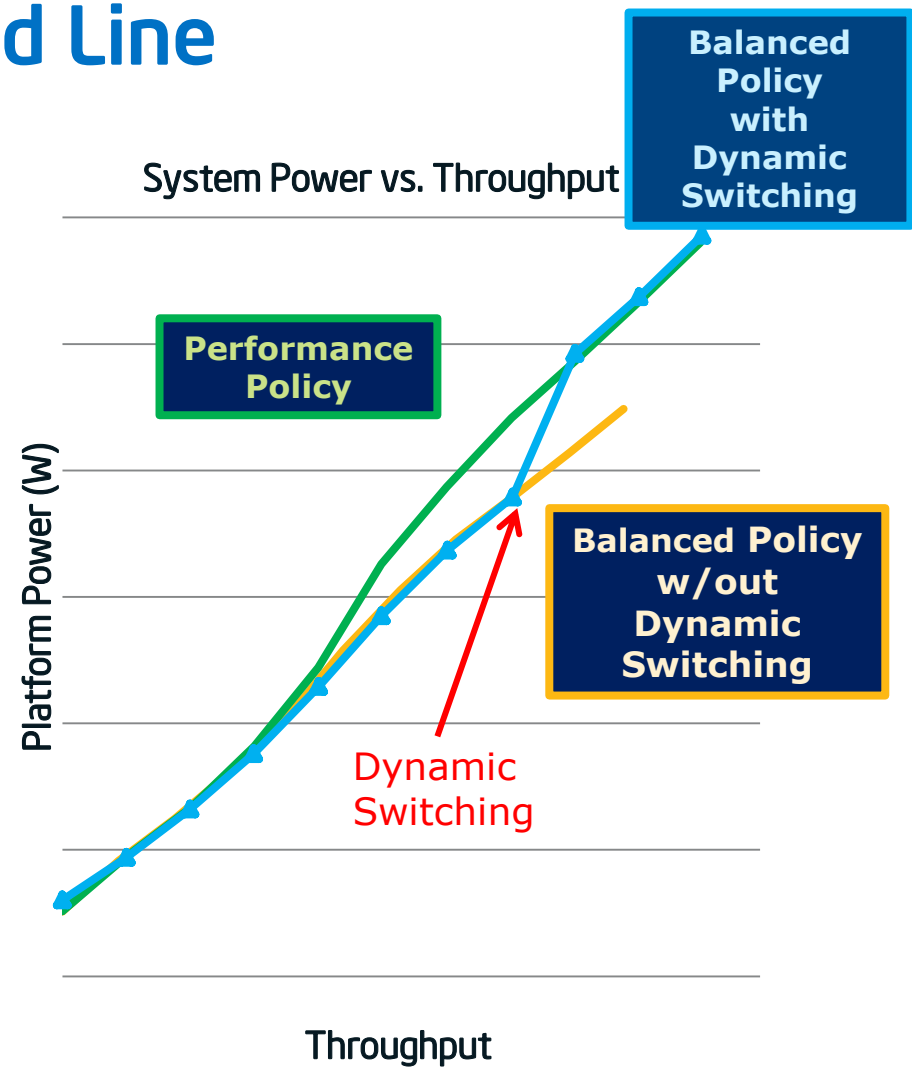

PCU works synergistically with OS Power Policy


Energy Efficiency in the Data Center

Running Average Power Limiting (RAPL)

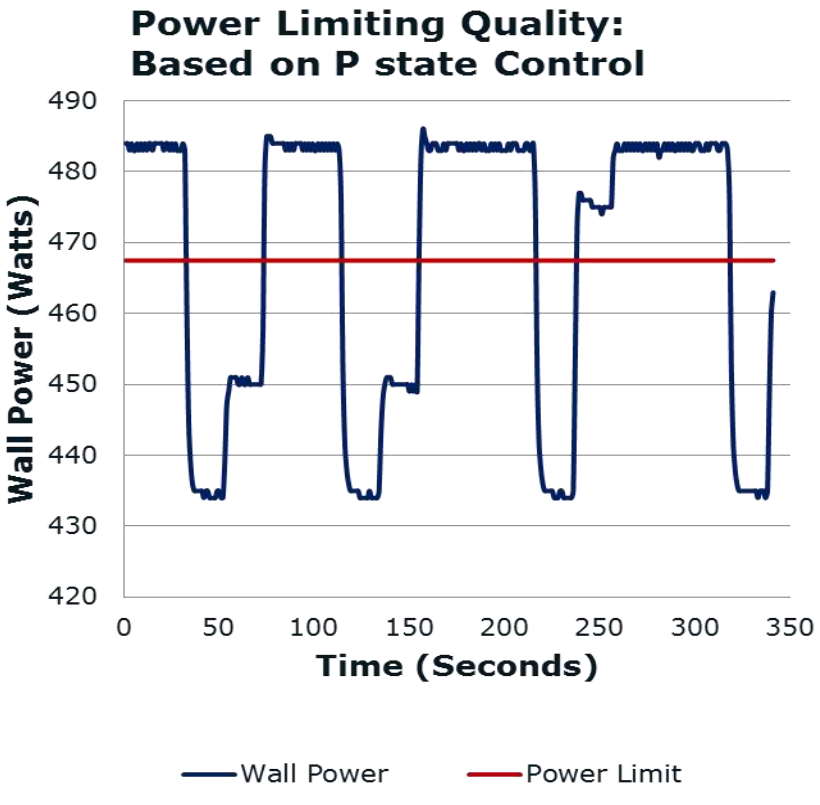

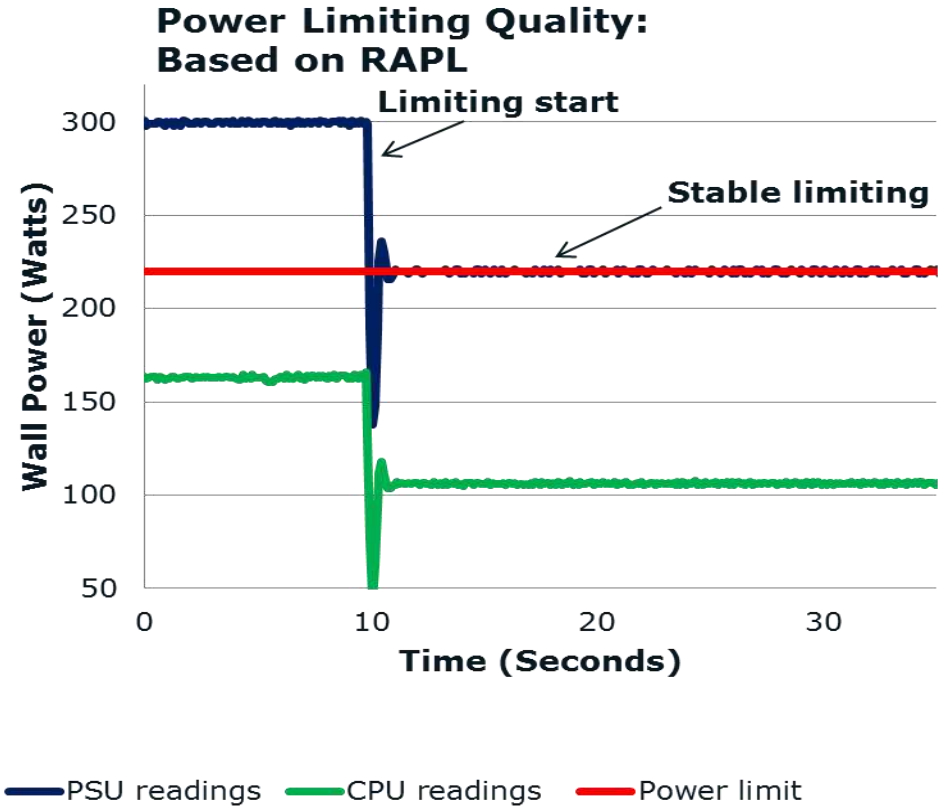

RAPL gives accurate and stable power limiting than P state control

Improved Efficiency with RAPL

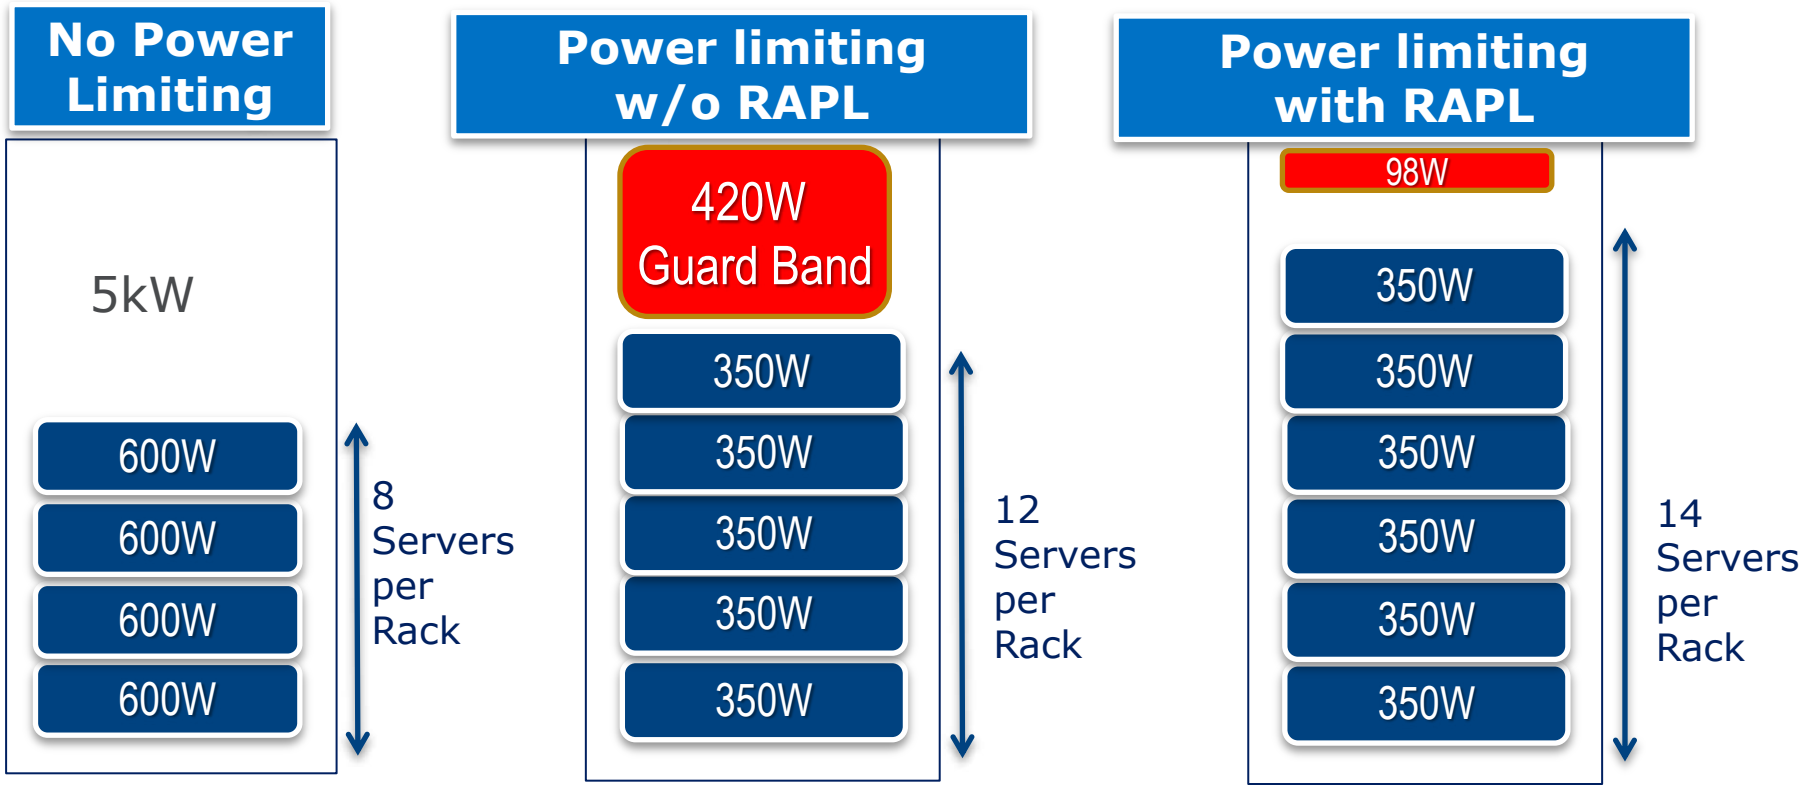

Max System Power: 600W
Typical System Power: 350W
Rack Power: 5kW

Improved Power Limiting Accuracy Allows for Smaller Guard bands and Increased Rack Density.

Socket RAPL & the Power/Performance Load Line

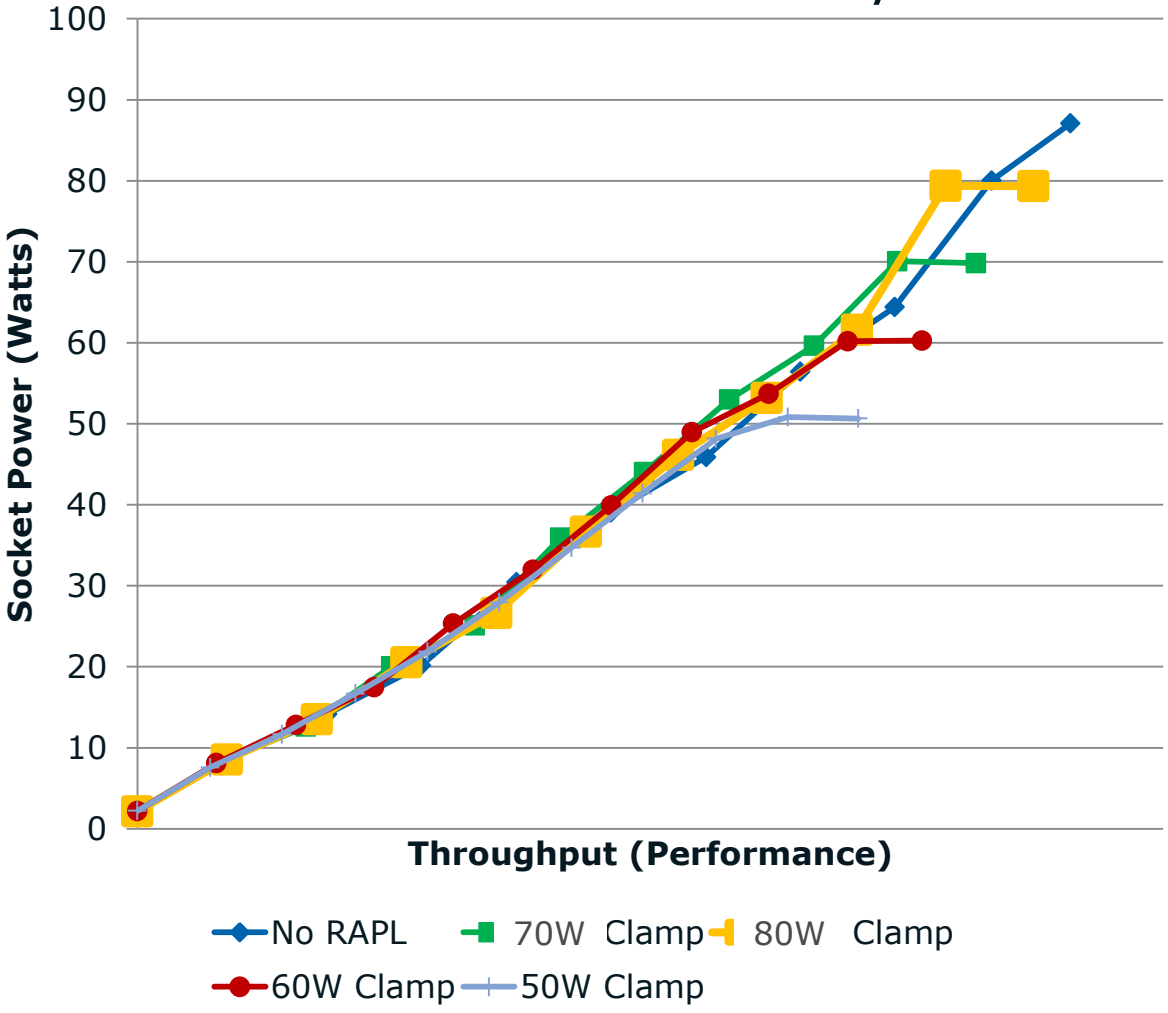

- **Power Clamps effectively limit socket power**
- **Max possible performance achieved for given power limit**
- **Efficient Power Scaling across the load line**

Internal Estimates
Q2 2011
SNB EP, C0 stepping

Maximum Performance

Memory Latency Optimizations

- Early Snoop
  Dynamic Direct 2 Core
- Uncore Frequency Change
- Dynamic Memory CKE Disable
- New LLC Prefetcher
- **Distributed L3**
  - Theoretical Peak: ~844GB/s (1s @ 3.3GHz w/ 8 cores)
  - Core->L3 Read Throughput: >250GB/s (1s @ 3.3GHz w/ 8 cores)
- **Dual Load Ports on L1 D-Cache**
- **SandyBridge Turbo 2.0**

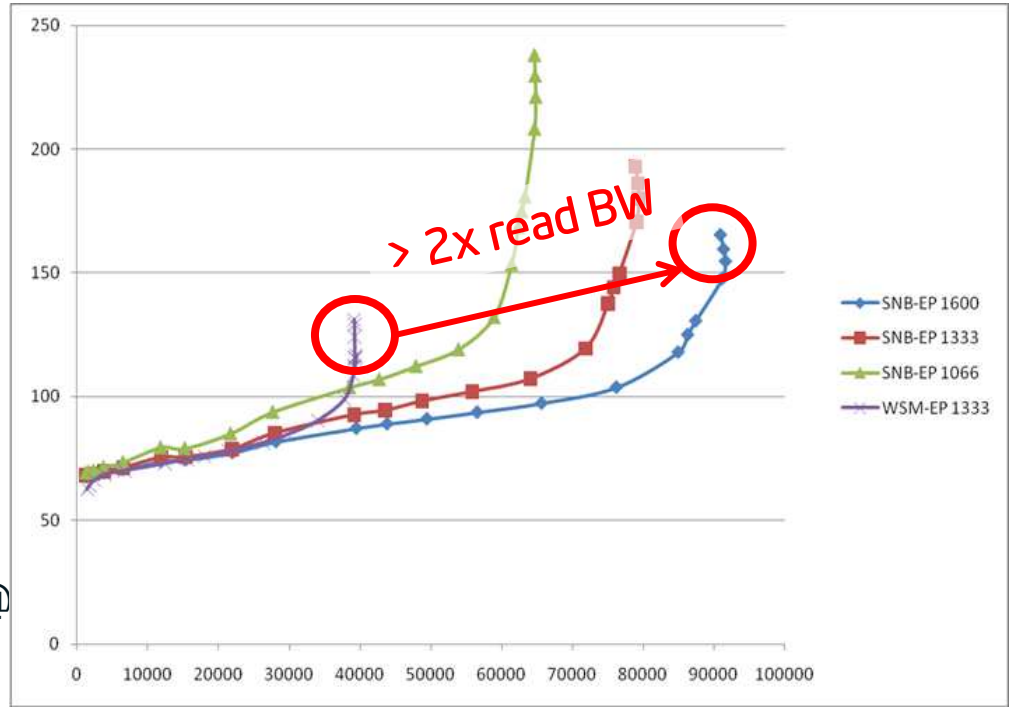

>2x max bandwidth from Xeon 5600 on read BW
- 3->4 channels (+33%)
- 1333->1600 (+20%)
- Improved Efficiency (+~40%)

Benchmark Notes:
- Intel internal tool for BW and Latency

Comparison of 4 core to 8 core Scaling @ 3.3GHz

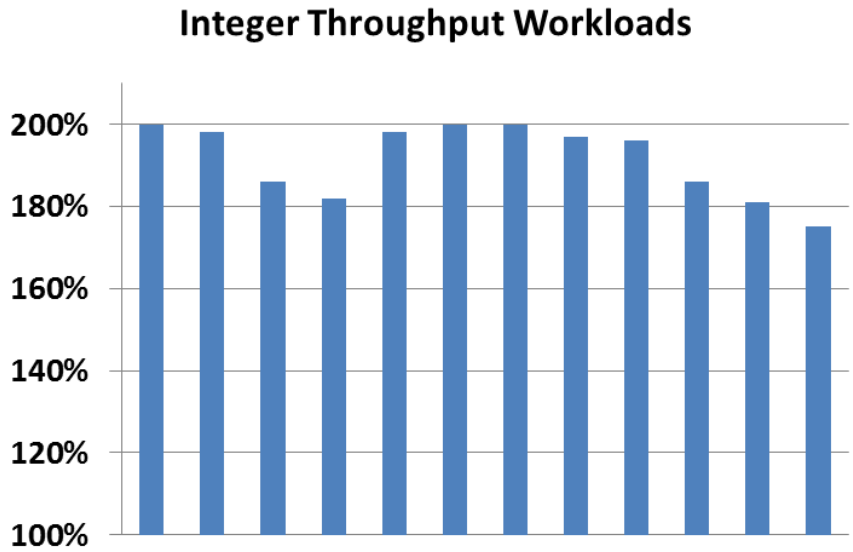

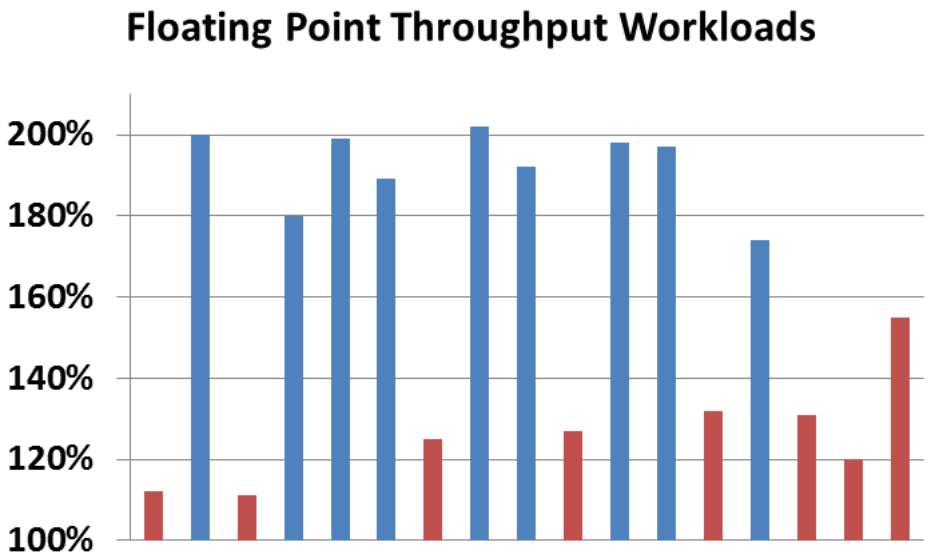

- Core sensitive apps in both INT and FP show excellent performance scaling
- Memory sensitive apps show less scaling (as expected shown in red)

Internal Testing – Estimate
4c: SNB E5-2643 w/out Turbo (1 DPC, DDR 1600)
8c: SNB E5-2690 w/ Turbo (2 DPC, DDR 1600)
ICC 12.1 / RHEL 6.1 / 2.6.32.131

Apps highlighted in Red are Memory Bandwidth sensitive

Intel® Xeon® E5 uncore provides significant core Scaling

Configuration Details for Foil #25

For the SPEC benchmarks, please see http://www.spec.org for more information
Configuration Details: As of 31 May 2012
SAP* SD 2-tier
2x Intel Xeon processor X5690 (12M cache, 3.46GHz, 6.40GT/s Intel QPI) score 5220 SD users. Certification #2011005. Source:
http://download.sap.com/download.epd?context=40E2D9D5E00EEF7C4B299992CE278ECED5166ED278FF20DF78759DC5B1E5FE79;
2x Intel Xeon processor E5-2690 (20M cache, 2.9GHz, 8.0GT/s Intel QPI) score 7865 SD users. Source:
http://download.sap.com/download.epd?context=40E2D9D5E00EEF7C5DDB3927818D671E00ECF023B5CE29EE68B565E9F19F1254
SPECvirt_sc*2011
2x Intel® Xeon® processor X5690 (6C,12M, 3.06GHz) score 1367 @ 84 VMs. Source: http://www.spec.org/virt_sc2010/results/res2011q1/virt_sc2010-20110209-00022-perf.html;
2x Intel® Xeon® processor E5-2690 (8C, 2.9GHz, C0) score 2,388 @ 150 VMs. Source: http://www.spec.org/virt_sc2010/results/res2012q2/virt_sc2010-20120403-00045-perf.html
SPECpower_ssj*2008
metrics for SPECpower are efficiency based and expressed as ssj_ops/watt.
2x Intel Xeon processor X5675 (12M cache, 3.45GHz, 6.40GT/s Intel QPI) score 3,329. Source: http://www.spec.org/power_ssj2008/results/res2011q4/power_ssj2008-20110713-00386.html;
2x Intel Xeon processor E5-2660 (20M cache, 2.2GHz, 8.0GT/s Intel QPI, C1) score 5,088. Source: http://www.spec.org/power_ssj2008/results/res2012q2/power_ssj2008-20120427-00454.html
TPC-E*
2x Intel Xeon processor X5690 (12M Cache, 3.46GHz, 2P/12C/24T) referenced as published at 1,284.14 tpsE, $250 USD/tpsE, available 5/4/11. Source: http://www.tpc.org/tpce/results/tpce_result_detail.asp?id=111050403;
Intel: 2x Intel Xeon processor E5-2690 (20M cache, 2.9GHz, 2P/16C/32T) referenced as published at 1,863.23 tpsE, $207.85 USD/tpsE, available 3/6/12. Source: http://www.tpc.org/tpce/results/tpce_result_detail.asp?id=112030601
VMmark* 2
2x Intel Xeon processor X5690 (12M cache, 3.46GHz, 6.40GT/s Intel QPI) score 7.59 @ 7 Tiles. Source: http://www.vmware.com/a/assets/vmmark/pdf/2011-10-18-Fujitsu-RX300S6.pdf;
2x Intel Xeon processor E5-2690 (20M cache, 2.9GHz, 8.0GT/s Intel QPI, C1) score 11.13 @ 10 Tiles. Source: http://www.vmware.com/a/assets/vmmark/pdf/2012-05-15-HP-DL360pG8.pdf
TPC-C*
2x Intel Xeon processor X5690 (12M Cache, 3.46GHz, 2P/12C/24T) referenced as published at 1,053,100 tpmC, $0.57 USD/tpmC, available 6/20/11. Source: http://www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=111120802;
2x Intel Xeon processor E5-2690 (20M cache, 2.9GHz, 8.0GT/s Intel QPI) referenced as published at 1,503,544 tpmC, $0.53 USD/tpmC, available 4/11/12. Source: http://www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=112041101
SPECjbb*2005
2x Intel Xeon processor X5690 (12M cache, 3.46GHz, 6.40GT/s Intel QPI) score 975,257 bops, 487,629 bops/JVM. Source: http://www.spec.org/osg/jbb2005/results/res2011q1/jbb2005-20110215-00950.html;
2x Intel Xeon processor E5-2690 (2.9GHz, 8C) score 1,584,567 bops. Source: http://www.spec.org/osg/jbb2005/results/res2012q1/jbb2005-20120306-01056.html
SPECint*_rate_base2006
2x Intel Xeon processor X5690 (12M cache, 3.45GHz, 6.40GT/s Intel QPI) baseline score 425. Source: http://www.spec.org/cpu2006/results/res2012q2/cpu2006-20120322-20154.html
2x Intel Xeon processor E5-2690 (20M cache, 2.9GHz, 8.0GT/s Intel QPI) baseline score 671. Source: http://www.spec.org/cpu2006/results/res2012q1/cpu2006-20120307-19618.html
SPECjEnterprise*2010
2x Intel Xeon processor X5690 (12M cache, 3.46GHz, 6.40GT/s Intel QPI) score 5,427 EjOPS. Source: http://www.spec.org/jEnterprise2010/results/jEnterprise2010.html;
2x Intel Xeon processor E5-2690 (20M cache, 2.9GHz, 8.0GT/s Intel QPI) score 8,310.19 EjOPS. Source: http://www.spec.org/jEnterprise2010/results/jEnterprise2010.html
SPECfp*_rate_base2006
2x Intel Xeon processor X5690 (12M cache, 3.45GHz, 6.40GT/s Intel QPI) baseline score 271. Source: http://www.spec.org/cpu2006/results/res2012q1/cpu2006-20111219-19195.html
2x Intel Xeon processor E5-2690 (20M cache, 2.9GHz, 8.0GT/s Intel QPI) baseline score 496. Source: http://www.spec.org/cpu2006/results/res2012q1/cpu2006-20120307-19617.html
STREAM*_MP Triad (NTW)
2x Intel Xeon processor X5690 (12M cache, 3.45GHz, 6.40GT/s Intel QPI) TRIAD score 42GB/s. Source: Intel TR#1241
2x Intel Xeon processor E5-2690 (20M cache, 2.9GHz, 8.0GT/s Intel QPI, C1) score 79.5 GB/s. Source: Intel TR#1241
Linpack
2x Intel Xeon X5690 (12M cache, 3.45GHz, 6.40GT/s Intel QPI) score 159.4. Source: Intel TR#1236
2x Intel Xeon processor E5-2690 (20M cache, 2.9GHz, 8.0GT/s Intel QPI, C1) score 347.7. Source: Intel TR#1236

Intel® Xeon® Processor E5-2600 Product Family Generational Performance Summary

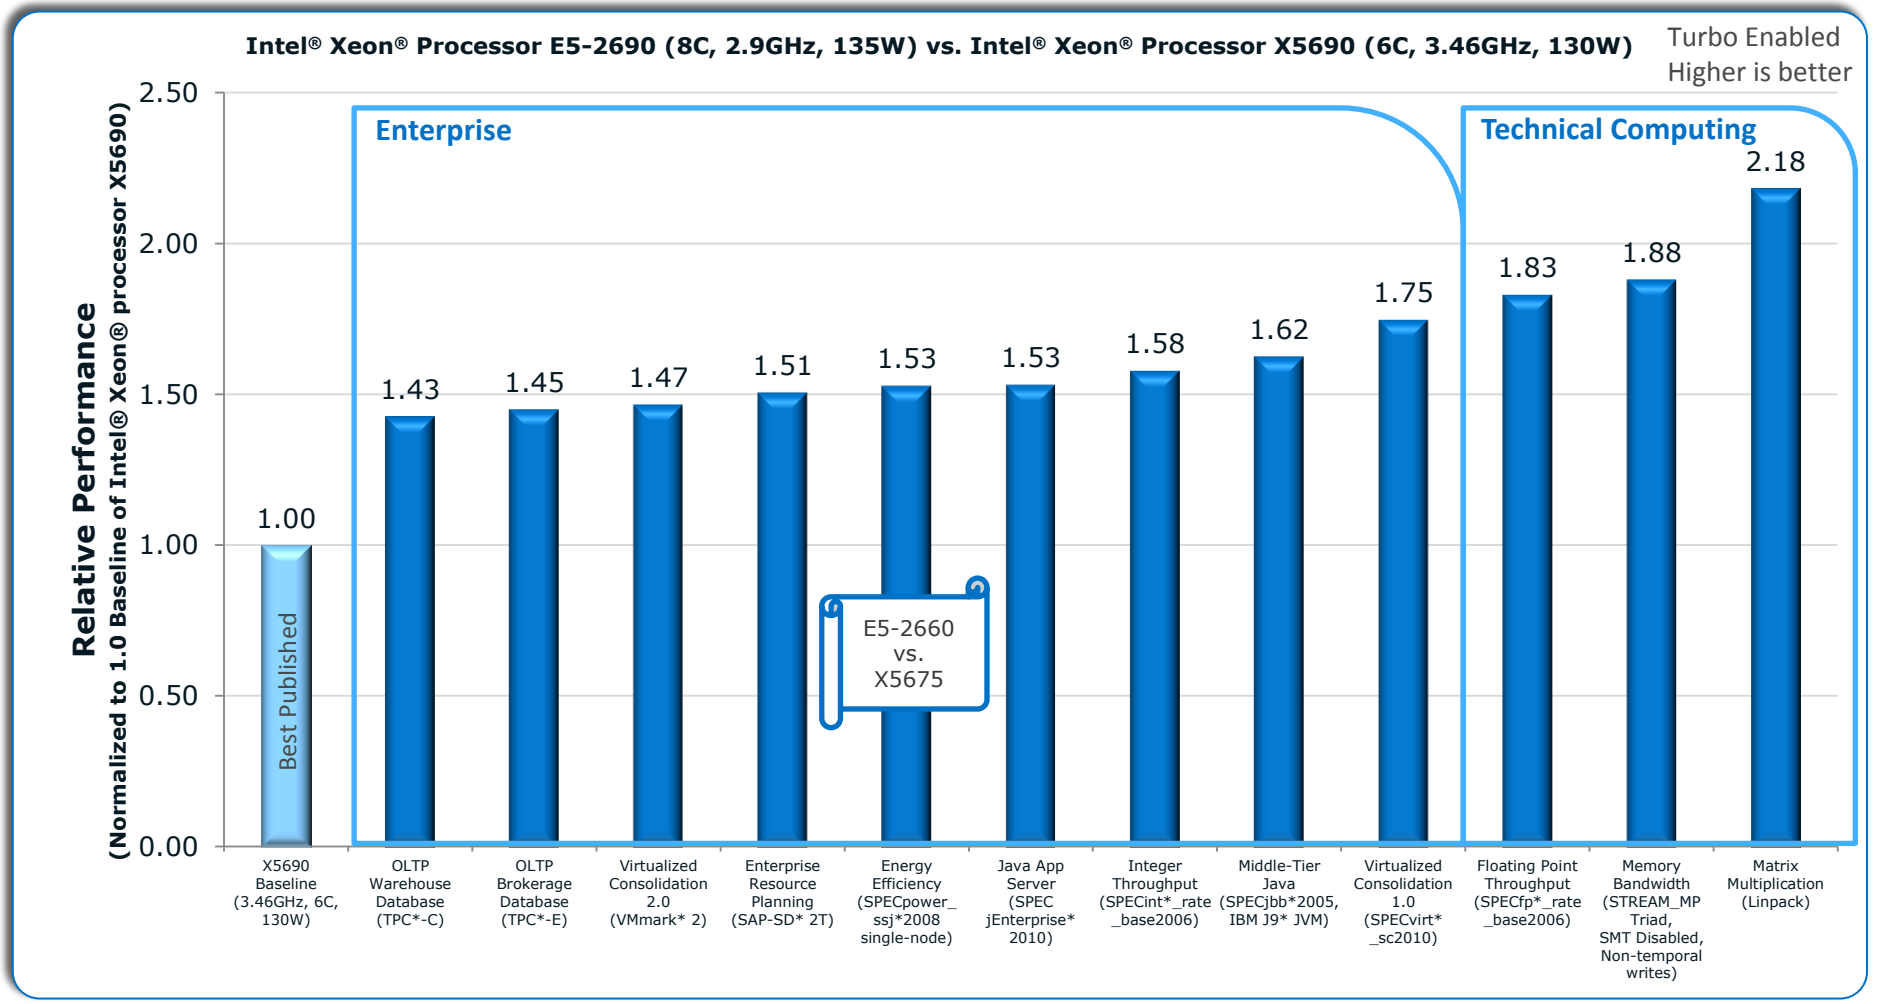

Intel® Xeon® processor E5-2690 delivers performance gains up to 2X

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.

Linpack performance may vary based on thermal solution.
Configuration Details: Please reference foil 24 for details.
For more information go to http://www.intel.com/performance

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