High Performance Computing for Silicon Design

Shesha Krishnapura

Senior Principal Engineer

April 2009

Authors & Contributors:

Raju Nallapa, Doug Austin, Ananth Sankaranarayanan, Shaji Achuthan, Vipul Lal, Ty Tang, Shesha Krishnapura, Elwood Coslett

Intel Computing Environment and Computing Demand

Design Team Challenges

Design Factors

- Shorter product development time
- Nano process technology
- Reduced headcount and cost

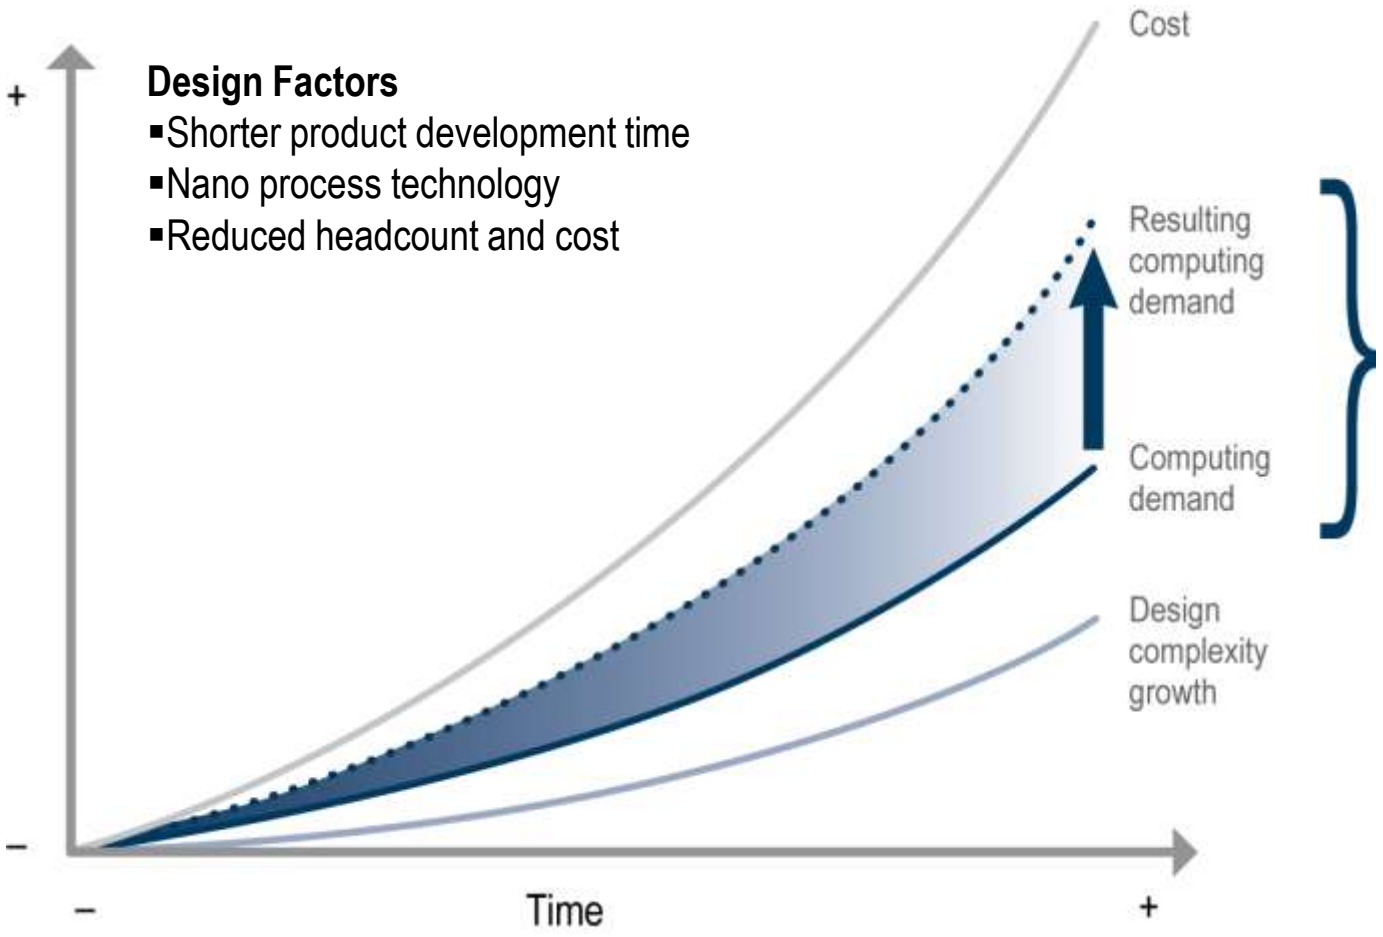

- Increased pre-silicon verification
- High degree of design automation
- Global design team and collaboration

Intel Data Center Profile

- Intel has four major groups currently driving individual data center requirements (DOME):
  - (**D**esign) Design Computing: Support the chip design community and they have most of the servers within Intel
  - (**O**ffice) General Purpose: Supporting typical IT and customer services
  - (**M**anufacturing) FAB/ATM: Manufacturing computing supporting fabrication and assembly
  - (**E**nterprise) Enterprise applications supporting eBiz

70% of servers in Intel are in “D”
30% of servers are in “OME”

Design Computing Env Overview

- Classification by server type in "D" environment
  - 65,000 Servers running Linux
  - 52% Blades (Xeon 2S) – All multi-core servers are at 4GB-8GB per core
  - 42% 1U (Xeon 2S, Few 1S) – Multi-core servers are at 8GB per core
  - 6% Rest (Xeon MP) – 128GB to 1TB per server
- Classification by use model in "D" environment
  - Batch servers (70%)
  - Interactive & large memory batch servers (30%)
- High Performance Mega Data Centers:
  - Each data center has multiple modules design to handle over 500+ watts/SF
  - 6000 sq ft/ per module with ~3MW of useable power
  - ~200 cabinets/racks per module
  - 15-22KW power allocation per rack (48-64 blades per rack)
  - Some data centers support 30KW (Up to 84 blades per rack)

Compute Demand Drivers inside Intel

- Pre-Silicon Design Computing
  - More than 100,000 simulation jobs per chip design each week for several quarters till tape-in
  - Small, Medium, Large memory workloads
  - Many chip designs in flight at a given time
  - Primarily CPU, and Physical Memory Bound – Lately Storage is of concern

- Tape-out Computing
  - 16,000 Optical Proximity Correction (OPC) jobs for each of the complex silicon layer
  - Small, Large, Very Large memory workloads
  - CPU, Network, and Storage Bound

Intel HPC Environment

HPC Solution Stack

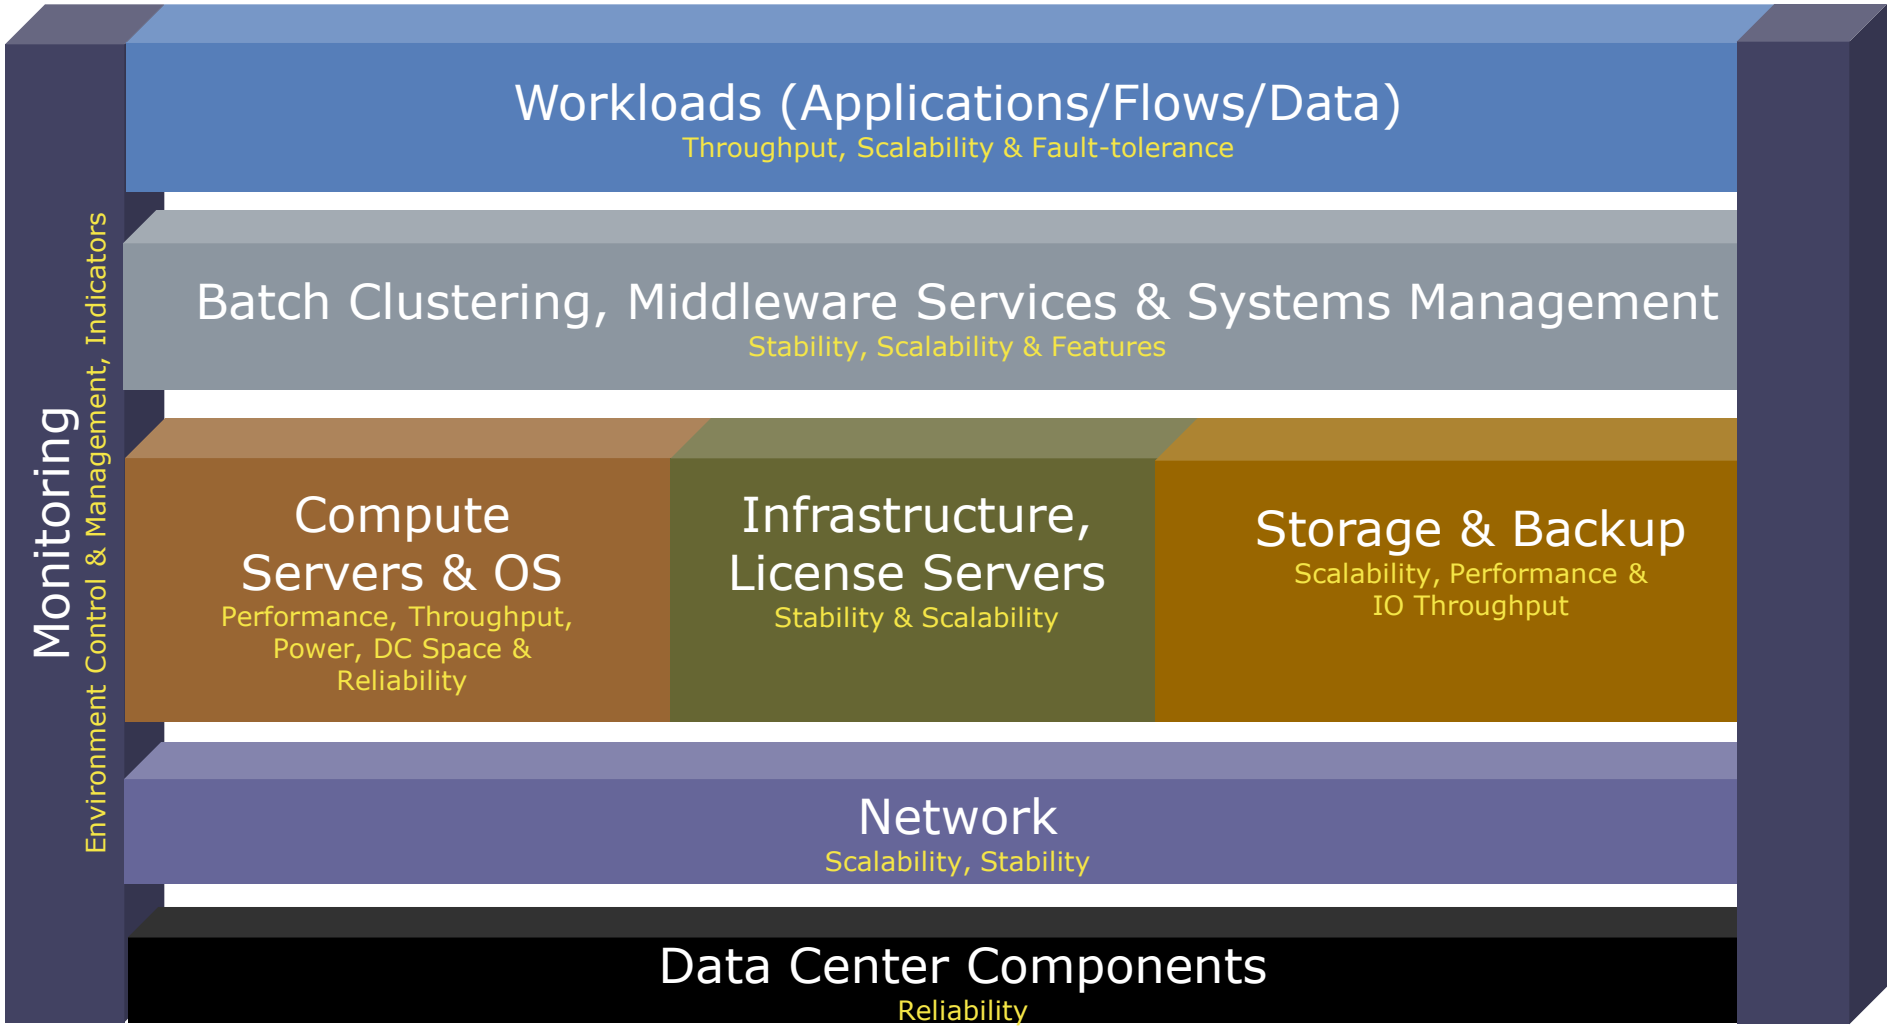

HPC Capability and Target Use Roadmap

	2006	2007	2008	2009	2010
HPC Technology Areas	**HPC-1** Optimize for 45nm, Support >=65nm		**HPC-2** Optimize for 32nm, Support >=45nm		**HPC-3 Target Additions** Optimize for 22nm, Support >=32nm
Batch Clustering Stability, Scalability, Features	Systems/Pool: 8.5K (**1.3x**)		11K (**1.3x**)		11K (**1x**)
	Jobs/Pool: 20K+ (**1.5x**)		30K+ (**1.5x**)		60K+ (**2x**)
	Scheduling: Preferential		Smart Class		Support for Virtualization
Storage & Backup Scalability, Performance IO Throughput	IO Spec TP#: 5,120 (**10x**)* HVM IO TP# : 3,200+MBps Volume Size : 3.2TB (**8x**)		5,120 3,500+ MBps (**1.1x**) 6.4TB (**2x**)		14,080 (**2.75x**) 5,300+ MBps (**1.5x**) TBD
	Single-Stream Perf*:70MBps (**1x**)		160 MBps (**2.3x**)		240 MBps (**1.5x**)
	HW/SW: Parallel-Storage-Gen1$		Parallel-Storage-Gen2$		Parallel-Storage/Open-Standard
Network Scalability, Stability	Storage: 40Gbps (**10x**)		40Gbps (**1x**)		11x1x10Gbps (**2.75x**)
	Master: 1Gbps (**10x**)		2x1Gbps (**1x**, **Redundancy**)		10Gbps (**5x**)
	Slave: 100Mbps (**1x**)		100Mbps (**1x**)		100Mbps/1Gbps (**1x/10x**)
Compute Optimized for Perf, Throughput, Capacity Power & DC Space	Large RAM: **512GB** (**4x**) (Based on Intel Architecture) Perf. TP#: 1.6-5x		**1TB** (**2x**) (Based on Intel Architecture) 1.7x		TBD
	Batch Node: 2S/**Dual**-Core/16GB Perf. TP: **2.1x** *(With Intel® Xeon® Processor 5150)*		2S/**Quad**-Core/32GB **2.3x** *(With Intel® Xeon® Processor E5450)*		2S/**Quad**-Core/48GB **1.74x** *(With Intel® Xeon® Processor X5570 – No HT)*
OS New HW Feature Support, Scalability, Stability, Perf.	Enterprise Feature: **Stable,** Inter-System **NUMA** Support		**Multi-Core** Optimized		**Virtualization** Optimized
License Servers Stability, Scalability	Platform: IA Based (**3x** over RISC)		Latest IA based solution		Latest IA based solution
Apps Tuning Throughput	Tuning: CPU Prefetch (**1.2x**) Enablement: **512GB Support**		SSE4		Hyper-Threading

(10x)* = 10x Spec Limit improvement over prior gen. solution (5120 MBps vs. 512MBps);
"Single-Stream Performance" is relevant for Backup & Vol. size; $ Proprietary Software
TP – Throughput; HWA – Hardware Acceleration

HPC Demand & Benefits for Intel Tapeout

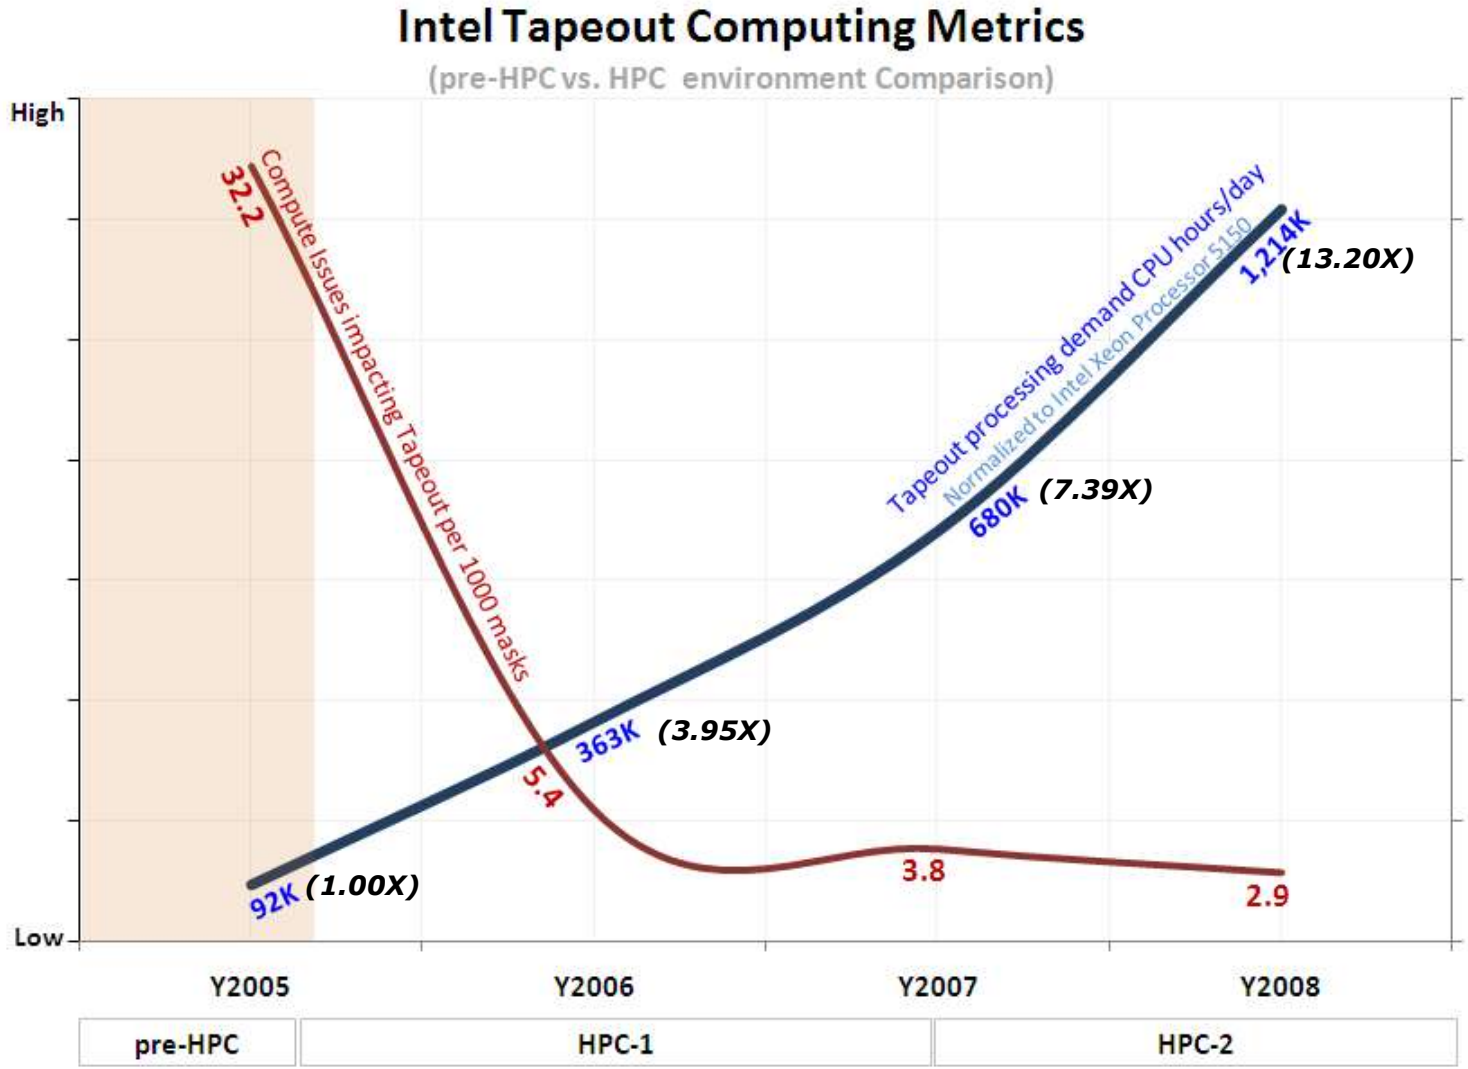

EDA Performance Improvement with Intel® Xeon® Processor Generations

Intel® Architecture Performance Improvement for OPC

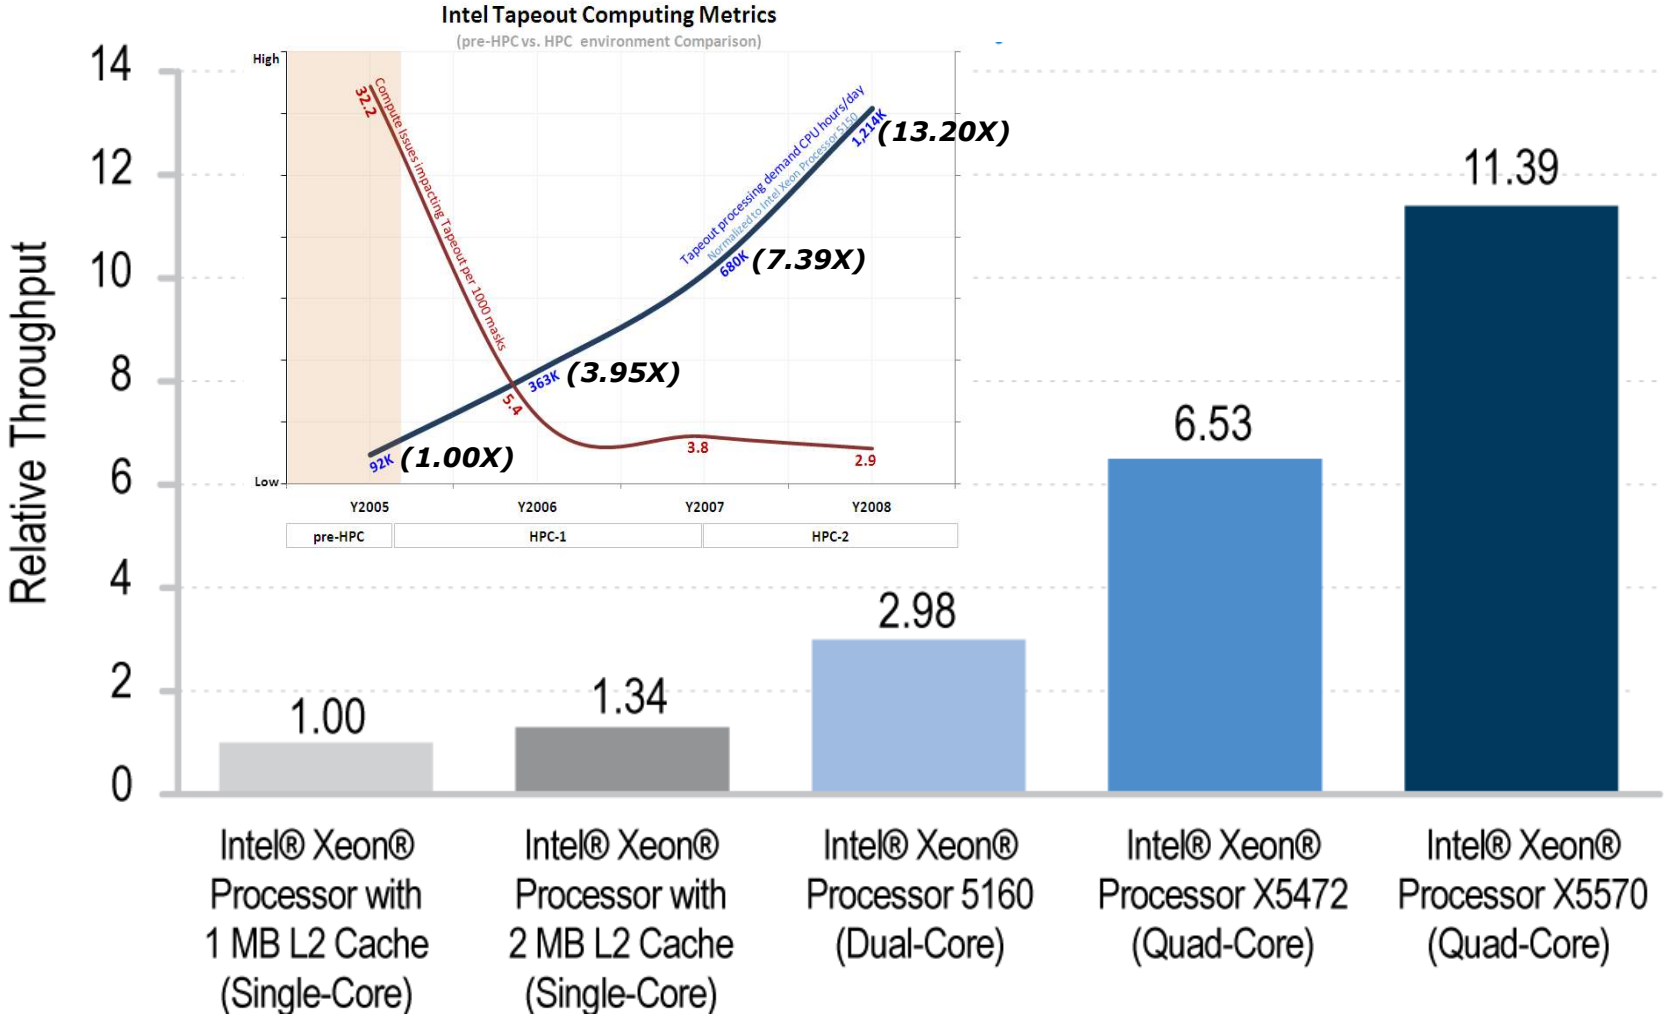

Intel internal measurements May 2007, November 2007, and February 2009.

Runtime Performance for OPC Application

Processor	OPC Jobs§	Runtime (HH:MM:SS)	Relative Throughput
64-bit Intel® Xeon® Processor with 1 MB L2 Cache (3.6 GHz)	2	10:40:12	1.00
64-bit Intel® Xeon® Processor with 2 MB L2 Cache (3.8 GHz)	2	07:58:31	1.34
Intel® Xeon® Processor 5160 (3.0 GHz)	4	03:34:39	2.98
Intel® Xeon® Processor X5472 (3.0 GHz)	8	01:37:58	6.53
Intel® Xeon® Processor X5570 (2.93 GHz)	8	00:56:11	11.39

§ One OPC job per core.

Intel internal measurements May 2007, November 2007, and February 2009.

Intel Xeon processor 5500 series offers Higher Density, Superior Performance, and Lower Power for OPC

30 Dual-Core Intel® Xeon® 5150 based servers

16GB RAM per server (4GB RAM/core)

Total Power Consumed **9360** Watts

120 Distributed OPC jobs

11 minutes and **41** seconds

15 Quad-Core Intel® Xeon® X5355 based servers

32GB RAM per server (4GB RAM/core)

Total power consumed **6450** Watts

120 Distributed OPC jobs

11 minutes and **04** seconds

10 Quad-Core Intel® Xeon® E5450 based servers

32GB RAM per server (4GB RAM/core)

Total power consumed **3200** Watts

80 Distributed OPC jobs

11 minutes and **36** seconds

6 Quad-Core Intel® Xeon® X5570 based servers

24GB RAM per server (3 GB RAM/core)

Total power consumed **2250** Watts

78 Distributed OPC jobs

10 minutes and **25** seconds

Intel quad-core server solution shows OPC throughput advantages

Intel internal measurements May 2006, November 2006, November 2007, and February 2009.

Profile: Intel® Xeon® Processor 5500 Series

- Up to 13.14x improved performance over single-core processors for simulation workloads
- Up to 11.39x improved performance over single-core processors for OPC workloads
- Up to 13:1 server consolidation ratio for simulation workloads and 11:1 for OPC workloads

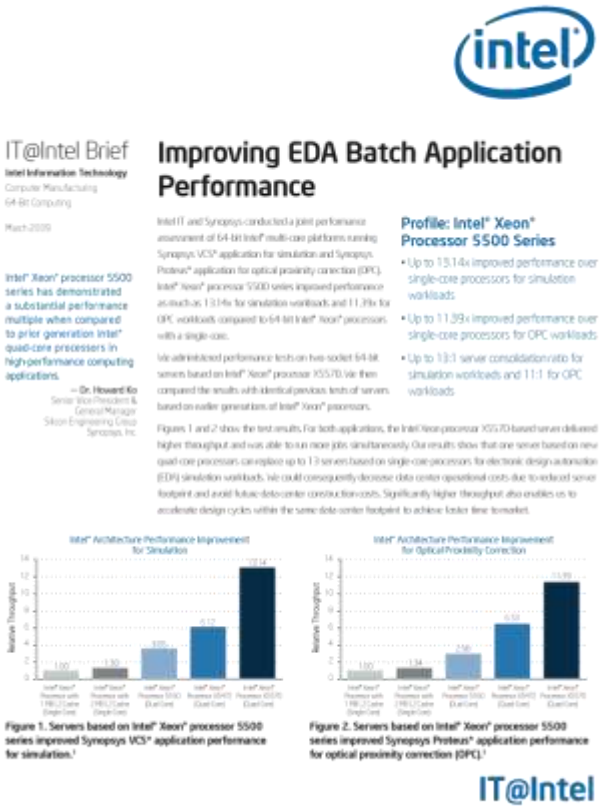

IT@Intel Brief

Intel Information Technology

Computer Manufacturing

64-bit Computing

March 2009

Intel® Xeon® processor 5500 series has demonstrated a substantial performance multiple when compared to prior generation Intel® quad-core processors in high-performance computing applications.

— Dr. Howard Ko
Senior Vice President & General Manager
Silicon Engineering Group
Synopsys, Inc.

Improving EDA Batch Application Performance

Intel IT and Synopsys conducted a joint performance assessment of 64-bit Intel® multi-core platforms running Synopsys VCS® application for simulation and Synopsys Proteus® application for optical proximity correction (OPC). Intel® Xeon® processor 5500 series improved performance as much as 13.14x for simulation workloads and 11.39x for OPC workloads compared to 64-bit Intel® Xeon® processors with a single core.

We administered performance tests on two-socket 64-bit servers based on Intel® Xeon® processor X5570. We then compared the results with identical previous tests of servers based on earlier generations of Intel® Xeon® processors.

Profile: Intel® Xeon® Processor 5500 Series

- Up to 13.14x improved performance over single-core processors for simulation workloads
- Up to 11.39x improved performance over single-core processors for OPC workloads
- Up to 13:1 server consolidation ratio for simulation workloads and 11:1 for OPC workloads

Figures 1 and 2 show the test results. For both applications, the Intel Xeon processor X5570-based server delivered higher throughput and was able to run more jobs simultaneously. Our results show that one server based on new quad-core processors can replace up to 13 servers based on single-core processors for electronic design automation (EDA) simulation workloads. We could consequently decrease data center operational costs due to reduced server footprint and avoid future data center construction costs. Significantly higher throughput also enables us to accelerate design cycles within the same data center footprint to achieve faster time to market.

Intel® Architecture Performance Improvement for Simulation

Intel® Architecture Performance Improvement for Optical Proximity Correction

Figure 1. Servers based on Intel® Xeon® processor 5500 series improved Synopsys VCS® application performance for simulation.[1]

Figure 2. Servers based on Intel® Xeon® processor 5500 series improved Synopsys Proteus® application performance for optical proximity correction (OPC).[1]

IT@Intel

Intel® Architecture Performance Improvement for Simulation

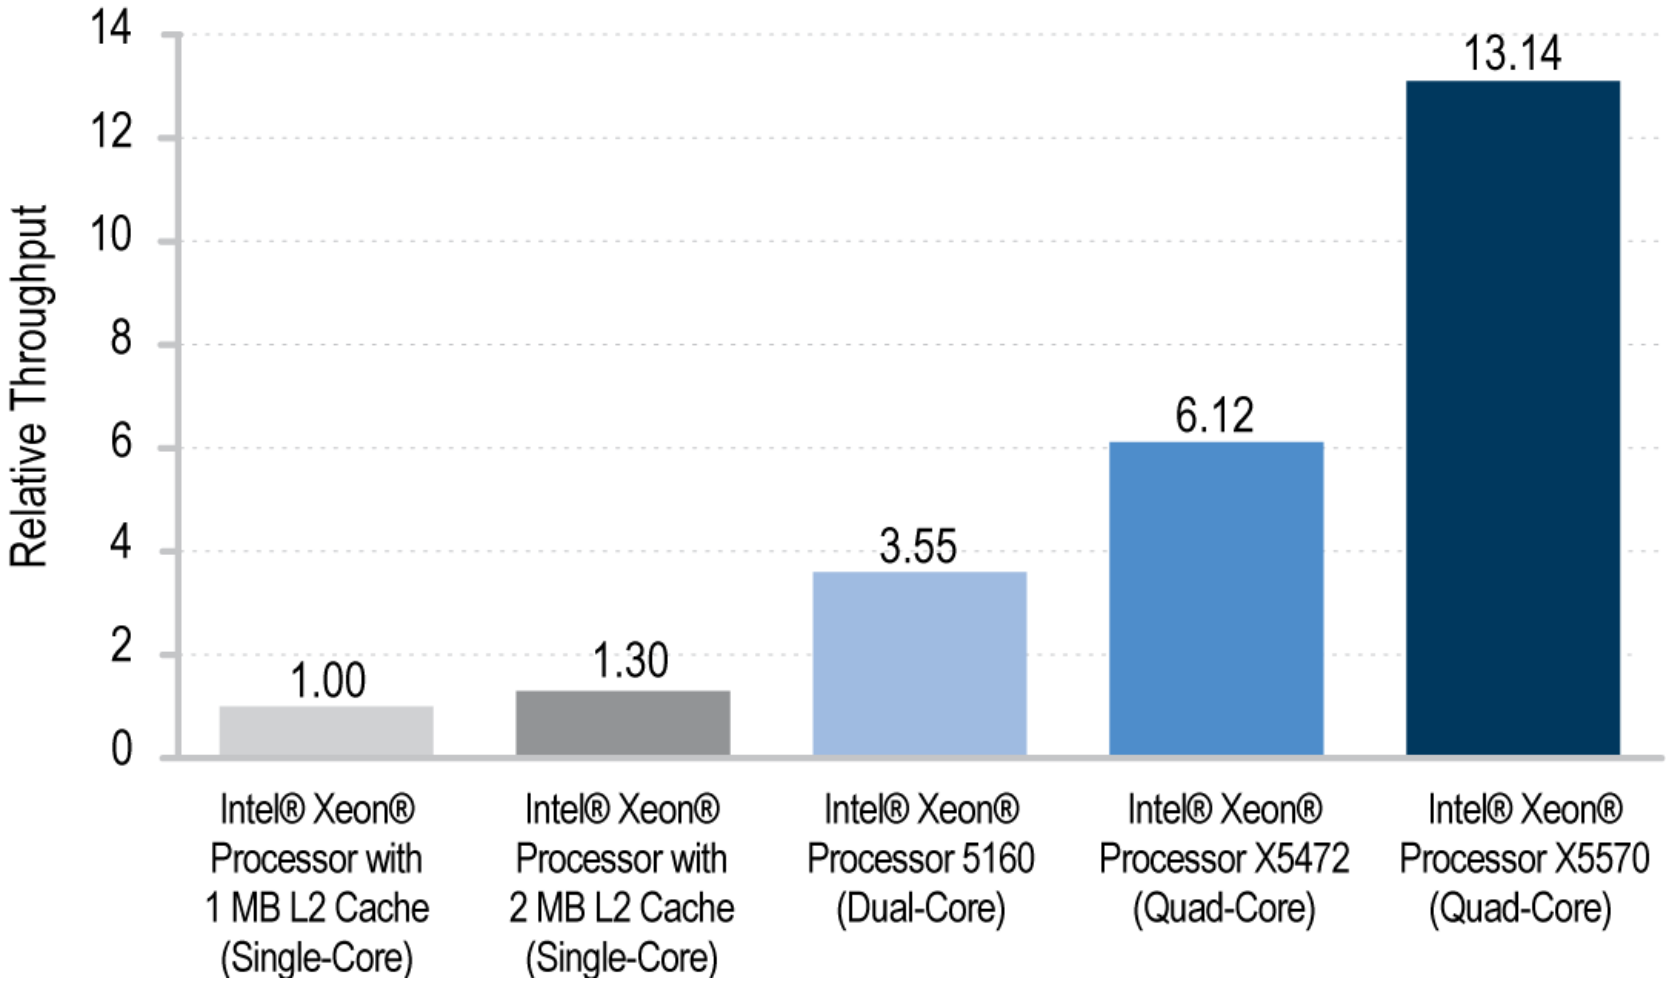

Intel internal measurements May 2007, November 2007, and February 2009.

Runtime Performance for Simulation Application

Processor	Simultaneous Simulation Jobs	Runtime (HH:MM:SS)	Relative Throughput
64-bit Intel® Xeon® Processor with 1 MB L2 Cache (3.6 GHz)	2	93:51:07	1.00
64-bit Intel® Xeon® Processor with 2 MB L2 Cache (3.8 GHz)	2	72:23:11	1.30
Intel® Xeon® Processor 5160 (3.0 GHz)	4	26:26:16	3.55
Intel® Xeon® Processor X5472 (3.0 GHz)	8	15:20:01	6.12
Intel® Xeon® Processor X5570 (2.93 GHz)[Δ]	16	07:08:36	13.14

[Δ] Tests run on Intel Xeon Processor X5570 series had Intel® Hyper-Threading Technology and Intel® Turbo Boost Technology enabled.

Intel internal measurements May 2007, November 2007, and February 2009.

Mainstream Intel® Xeon® Processor 5500 Series Segments

	Basic	Standard	Advanced
CPU Frequency	2.0 GHz to 2.13 GHz	2.26 GHz to 2.53 GHz	2.66 GHz to 2.93 GHz
CPU Power	80 W	80 W	95 W
QPI	4.8 GT/S	5.86 GT/S	6.4 GT/S
CPU Cache Size	4 MB	8 MB	8 MB
Memory Speed	800 MHz	800/1066 MHz	800/1066/1333 MHz
Intel® Turbo Boost Technology	No	Yes	Yes
Intel® HT Technology	No	Yes	Yes

GT/S – Gigatransfers/Second; Intel® HT – Intel® Hyper-Threading Technology; QPI – Intel® QuickPath Interconnect

EDA Throughput and Total Cost of Ownership

- In tests with real Intel EDA workloads, we required fewer servers based on high-end processors to achieve the same performance. This resulted fewer EDA application licenses; reduced data center power, space, and connectivity requirements; and substantially lower estimated TCO.

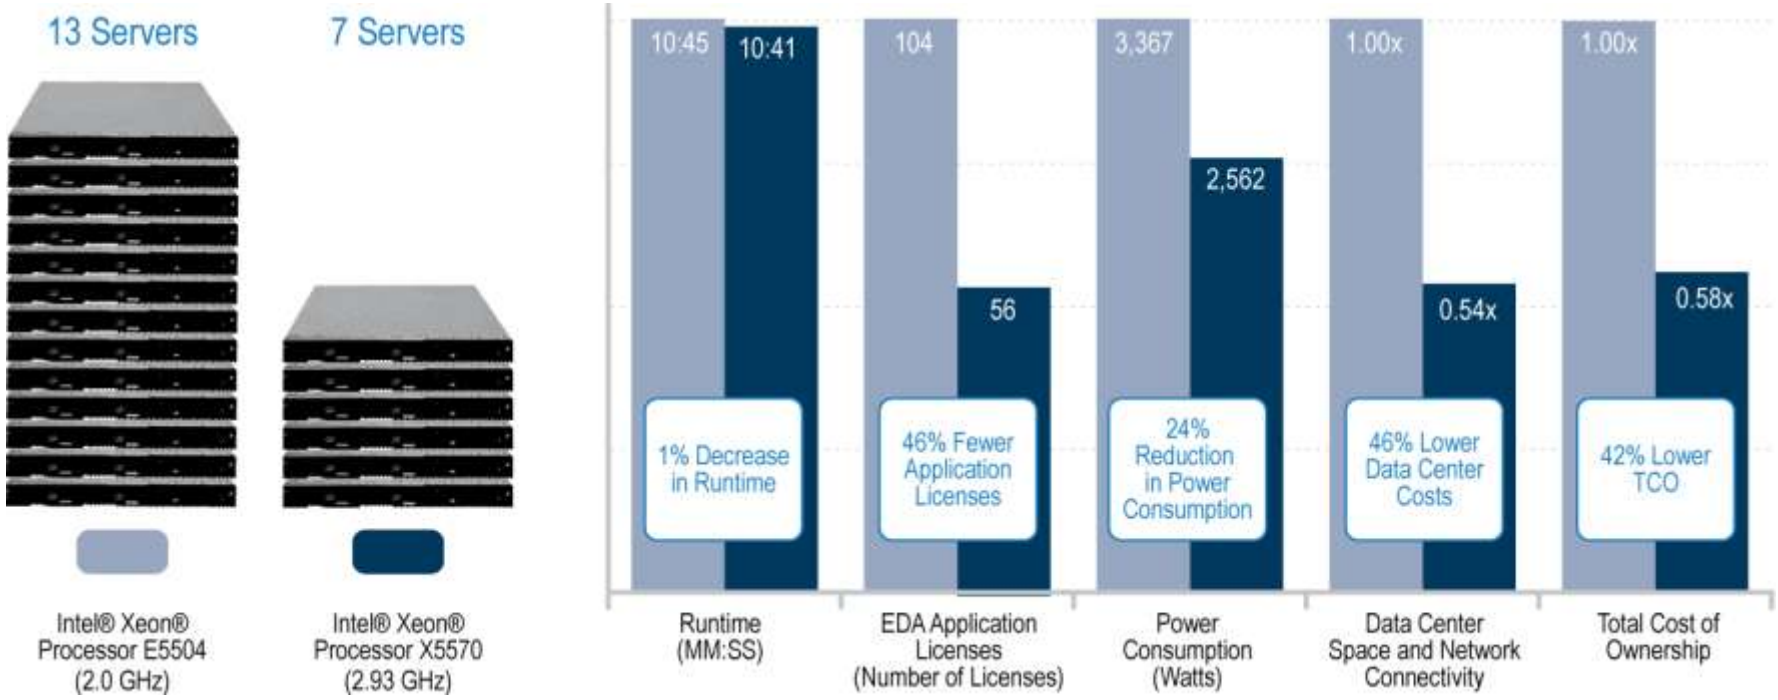

Intel internal measurements, February 2009.

Server TCO

- The hardware platform accounts for a small proportion of server total cost of ownership (TCO). TCO calculations based on Intel® Xeon® processor X5570 (2.93 GHz).

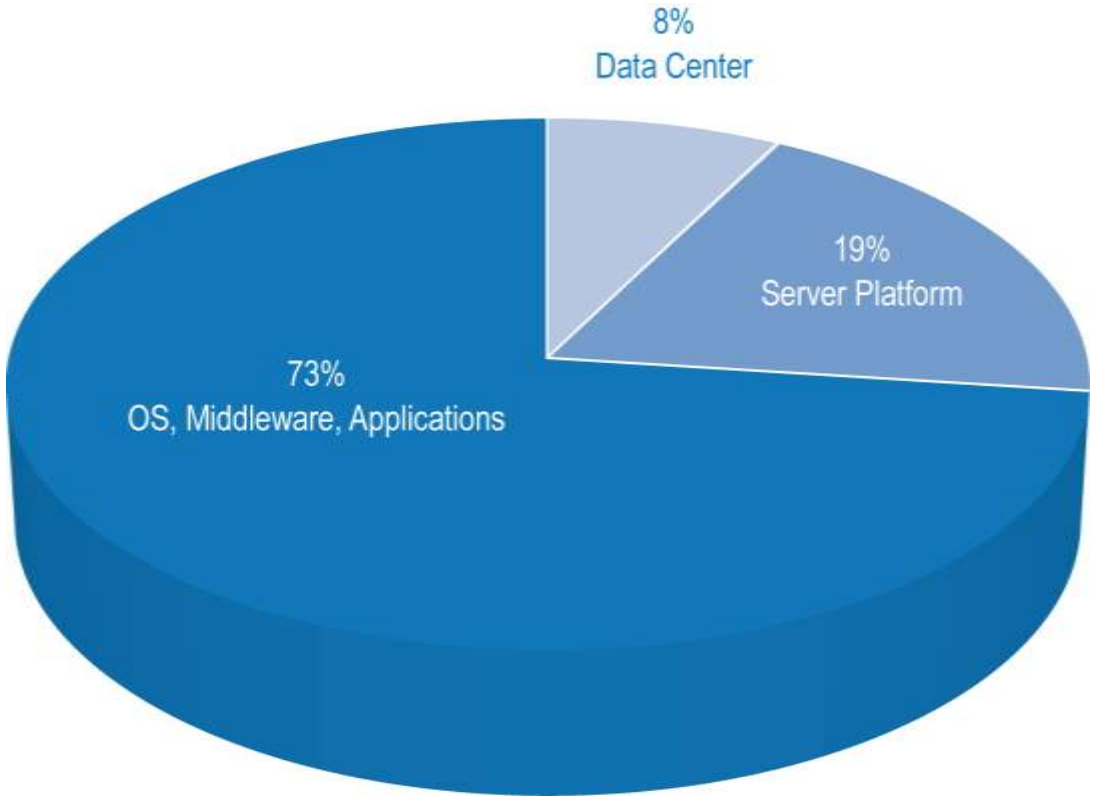

Intel internal measurements, February 2009.

Profile: Intel® Xeon® Processor 5500 Series

- High-end processors reduce server TCO by 42 percent compared to low-end processors
- High-end processors deliver up to 87 percent faster performance

IT@Intel Brief

Intel Information Technology

Computer Manufacturing

Server TCO

March 2009

With the introduction of the Intel® Xeon® processor 5500 series-based platforms, the benefits we are seeing from our IT strategy to standardize on higher-end processors for our servers purchases is even more compelling and results in a significantly lower TCO.

— Diane Bryant
Chief Information Officer
Intel Corporation

Selecting Server Processors to Reduce Total Cost

Intel IT is standardizing on Intel® Xeon® processor X5570 (2.93 GHz) for two-socket servers for design computing and enterprise server virtualization. Our testing and analysis demonstrates that the newest high-end Intel® Xeon® processors based on Next-Generation Intel® Microarchitecture (Nehalem) can significantly enhance server performance, providing an opportunity for Intel IT to reduce total cost of ownership (TCO) by 42 percent.

Profile: Intel® Xeon® Processor 5500 Series

- High-end processors reduce server TCO by 42 percent compared to low-end processors
- High-end processors deliver up to 87 percent faster performance

We compared the high-end Intel Xeon processor X5570 (2.93 GHz) with the low-end Intel® Xeon® processor E5504 (2.0 GHz) for two Intel IT computing environments: design and enterprise. In real-world application testing with Intel's electronic design automation (EDA) workloads, two-socket servers based on the 2.93-GHz processor delivered the same performance for an estimated 42 percent lower TCO over four years, as shown in Figure 1. For enterprise computing, analysis confirmed a similar relationship: A two-socket server based on the 2.93-GHz processor delivered 87 percent better performance for an estimated 8 percent increase in TCO.

Our analysis demonstrated to Intel IT management and purchasing groups that software acquisition and licensing costs—which represent 3x to 6x the cost of the hardware platform—are the largest components of overall TCO for servers deployed at Intel. Because of this, standardizing on high-end processors is a cost-effective way for Intel IT to maximize server return on investment (ROI).

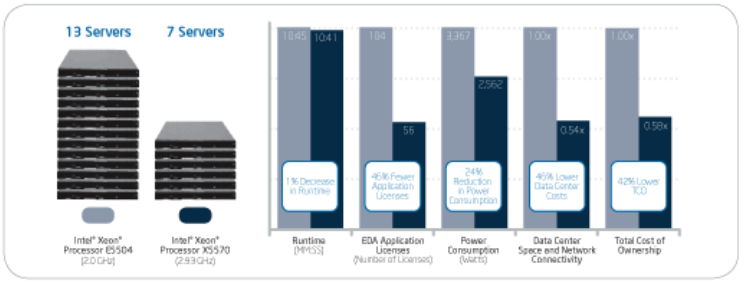

Figure 1. Electronic design automation (EDA) throughput and total cost of ownership (TCO). In tests with real Intel EDA workloads, we required fewer servers based on high-end processors to achieve the same performance. This resulted fewer EDA application licenses, reduced data center power, space, and connectivity requirements, and substantially lower estimated TCO. Intel internal measurements, February 2009.

IT@Intel

HPC-1 to HPC-2 Storage Performance

HPC-2 Generation Limit: ~**15K** OPC jobs accessing one Parallel-Storage

Category	Parallel-Storage-Gen1	Parallel-Storage-Gen2
Meta Data Server Load	**~100%**	**~80%**
Interactive Latency	**Unacceptable**	**Acceptable (no impact)**
Write (sec)	**25.00**	**4.00 (6.25x)**
Read (sec)	**25.00**	**0.47 (53x)**
File listing (sec)	**17.00**	**0.7 (24x)**
File removal (sec)	**25.00**	**0.36 (69x)**
Event: Storage Vol Offline	**Yes**	**No**

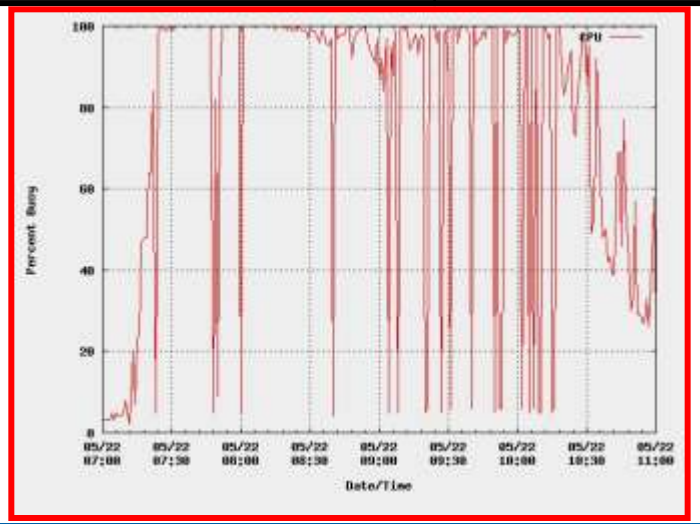

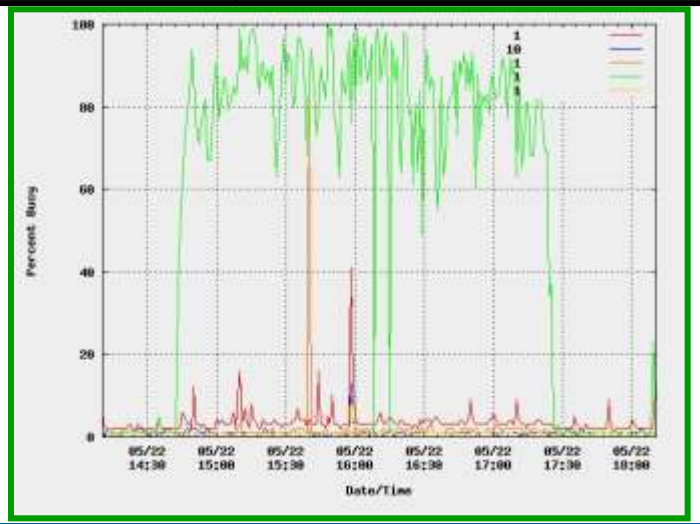

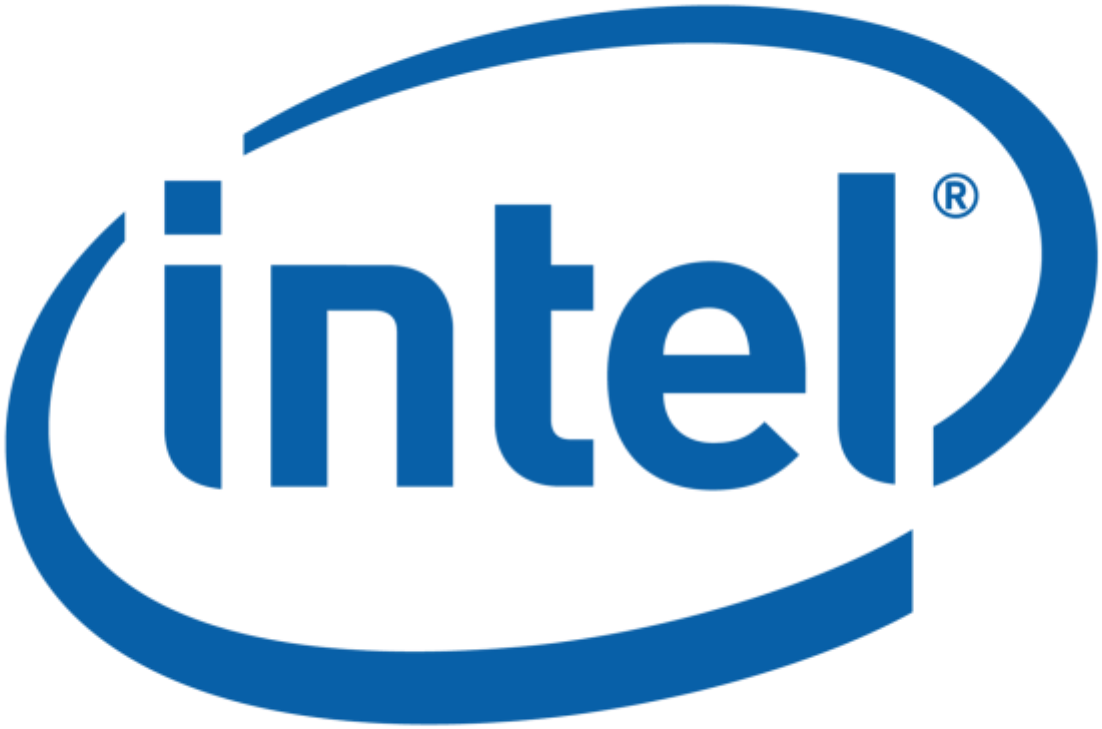
intel