WHITE PAPER
MARCH 2018

IT@INTEL

Increasing EDA Throughput with the Intel® Xeon® Processor Scalable Family

The Intel® Xeon® Platinum 8168 processor improves throughput up to 1.37x compared to a previous-generation Intel Xeon processor E5-2699 v4-based server.[Δ] Historically, it is up to 86x faster than a 64-bit Intel Xeon processor with a single core.[Δ]

Shesha Krishnapura
Intel Fellow and Intel IT CTO, Intel IT

Vipul Lal
Senior Principal Engineer, Intel IT

Ty Tang
Senior Principal Engineer, Intel IT

Shaji Kootaal Achuthan
Senior Staff Engineer, Intel IT

Murty Ayyalasomayajula
Senior Staff Engineer, Intel IT

Executive Overview

Intel's silicon design engineers need significant increases in computing capacity—both on their workstations and on data center servers—to deliver each new generation of silicon chips. To meet those requirements, Intel IT conducts ongoing throughput performance tests, using the Intel® silicon design workloads, to analyze the benefits of introducing compute servers based on new, more powerful processors in the field of electronic design automation (EDA).

We recently tested a dual-socket server based on the latest Intel® Xeon® Platinum 8168 processor running single-threaded and multi-threaded EDA applications operating on more than 200 hours of Intel design workloads. By utilizing all available cores, the server completed the workloads up to 86x faster than a server based on a 64-bit Intel Xeon processor (3.6 GHz) with a single core.[Δ] The Intel Xeon Platinum 8168 processor-based server was up to 24x faster than a server based on the Intel Xeon processor 5160 (3.0 GHz) with two cores.[Δ]

Based on our performance assessment and our refresh cycle, we plan to deploy servers based on the new Intel Xeon processor Scalable family, completing our replacement of servers based on the 8-core Intel Xeon processor E5-2600 series that are more than four years old. By doing so we expect to significantly increase EDA throughput while realizing savings, because we can avoid data center construction and reduce additional power consumption.

[Δ] For more complete information about performance and benchmark results, visit intel.com/benchmarks. Performance results based on testing details and system configuration. See the full disclaimer and system configurations on page 6.

Contents

Acronyms

DRC Design Rule Check

EDA Electronic Design Automation

NAC Node Antenna Check

Background

Silicon chip design engineers at Intel face ongoing challenges: integrating more features into ever-shrinking silicon chips, bringing products to market faster, and keeping design engineering and manufacturing costs low.

As design complexity increases, the requirements for compute capacity also increase, so refreshing servers and workstations with higher performing systems is cost-effective and offers a competitive advantage by enabling faster chip design. Refreshing older servers also enables us to realize data center cost savings. By taking advantage of the performance and power-efficiency improvements in new server generations, we can increase computing capacity within the same data center footprint, avoiding expensive data center construction and achieving operational cost savings due to reduced power consumption.

Intel IT conducts ongoing performance tests, based on the latest Intel® silicon design data, to analyze the potential performance and data center benefits of introducing servers based on new processors into our electronic design automation (EDA) computing environment. Table 1 illustrates some of the architectural enhancements.

Table 1. A Comparison of Dual-Socket Servers Based on Intel® Xeon® Processors

	2004-2005	2006-2008	2009-2011	2012-2016				2017
Introduction	2004-2005	2006-2008	2009-2011	2012	2013	2014	2016	2017
Intel® Chipset	E7520	5400	5520	C600		C610		C620
Process Technology	90nm	65nm and 45nm	45nm and 32nm	32nm	22nm		14nm	
Cores per Socket	1	2 or 4	4 or 6	8	10	14	22	28
Cache	1 MB or 2 MB[1]	4 MB or 6 MB shared between 2 cores	8 MB or 12 MB shared	20 MB shared	30 MB shared	45 MB shared	55 MB shared	38.5 MB shared
Interconnect Speed	6.4 GB/s	21-25 GB/s	25.6 GB/s per Intel® QuickPath Interconnect	32 GB/s per Intel® QuickPath Interconnect		38.4 GB/s per Intel® QuickPath Interconnect		41.6 GB/s per Intel® UltraPath Interconnect
DIMMs	Up to 8	Up to 16	Up to 18	Up to 24				
Memory Type	DDR2-400 MHz	FB-DIMM/DDR2-667 MHz or FB-DIMM/DDR2-800 MHz	DDR3-800/1066/1333 MHz	DDR3-1333/1600 MHz	DDR3-1333/1600/1866 MHz	DDR4-1600/1866/2133 MHz	DDR4-2400 MHz	DDR4-2666 MHz
Memory Bandwidth	Up to 6.4 GB/s	21-25 GB/s	Up to 32 GB/s	Up to 51.2 GB/s	Up to 59.7 GB/s	Up to 68 GB/s	Up to 76.8 GB/s	Up to 128 GB/s
Maximum Memory	16 GB	64 GB or 128 GB[2]	144 GB or 288 GB[3]	Up to 768 GB[4]	Up to 1536 GB[5]			Up to 3072 GB[6]

[1] Data provided only for 1 MB cache. [2] 128 GB support with Intel® 5400 Chipset introduced in 2007. [3] 144 GB assumes 18 memory slots populated with 8-GB DIMMs; 288 GB assumes 18 memory slots populated with 16-GB DIMMs, and validated only with Intel® Xeon® processor 5600 series. [4] 768 GB assumes 24 memory slots populated with 32-GB DIMMs. [5] 1536 GB assumes 24 memory slots populated with 64-GB DIMMs. [6] 3072 GB assumes 24 memory slots populated with 128-GB DIMMs

Faster Servers Process More EDA Jobs in Less Time

The architectural enhancements shown Table 1 illustrate how the Intel® Xeon® processor has evolved over the last few years. We have found that refreshing data center servers to use the latest processor technology substantially improves EDA throughput.

While our assessments focus on EDA applications, throughput improvements may also be achieved with other applications used in high-performance computing environments where simulation and verification are large parts of the workflow, including:

- Computational fluid dynamics and simulation in the aeronautical and automobile industries
- Synthesis and simulation applications in the life sciences industry
- Simulation in the oil and gas industries

Test Methodology

We ran tests on dual-socket servers based on the Intel® Xeon® Platinum 8168 processor. This processor includes new features designed to increase throughput compared with previous processor generations, including 14nm process technology, 24 cores, and 33 MB L3 cache.

We ran several tests using industry-leading EDA single-threaded and multi-threaded EDA applications comprising Intel Xeon processor and chipset design workloads.

Our goal was to assess throughput improvement by measuring the time taken to complete a specific number of design workloads. To maximize throughput, we configured each application to utilize all available cores, resulting in one job or process per core. The test configuration is shown in Table 2. We then compared our results with previous tests conducted using the same approach on servers based on the processors.

Maximizing Throughput with Intel® HT Technology

The Intel® Xeon® Platinum 8168 processor with Intel® Hyper-Threading Technology (Intel® HT Technology) enabled can support up to 96 concurrent software threads in a single two-socket platform and deliver higher performance throughput compared to HT Technology being disabled. Intel HT Technology increased performance by up to 1.25x when completing the same number of jobs using two times the application licenses.

Simulation Jobs Comparison

Time Needed to Complete 113 Jobs on Intel® Xeon® Platinum 8168 Processor

HIGHER IS BETTER

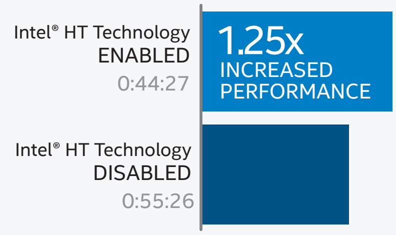

Table 2. Test Configuration for Dual-Socket Servers

	64-bit Intel Xeon® Processor	Intel® Xeon® Processor 5160	Intel® Xeon® Processor X5365	Intel® Xeon® Processor X5460	Intel® Xeon® Processor X5570	Intel® Xeon® Processor X5675	Intel® Xeon® Processor E5-2680	Intel® Xeon® Processor E5-2680 v2	Intel® Xeon® Processor E5-2697 v3	Intel® Xeon® Processor E5-2699 v4	Intel® Xeon® Platinum 8168 Processor
Cores	1	2	4	4	4	6	8	10	14	22	24
Frequency	3.6 GHz	3.0 GHz	3.0 GHz	3.16 GHz	2.93 GHz	3.06 GHz	2.7 GHz	2.8 GHz	2.6 GHz	2.2 GHz	2.7 GHz
Cache	1 MB	4 MB	8 MB	12 MB	8 MB	12 MB	20 MB	25 MB	35 MB	55 MB	33 MB
Interconnect	800 MHz Shared FSB	1333 Dual Independent FSB	1333 Dual Independent FSB	1333 Dual Independent FSB	25.6 GB/s per Intel® QPI link	25.6 GB/s per Intel® QPI link	32.0 GB/s per Intel® QPI link	32.0 GB/s per Intel® QPI link	38.4 GB/s per Intel® QPI link	38.4 GB/s per Intel® QPI link	41.6 GB/s per Intel® UPI link
RAM	16 GB	16 GB	32 GB	32 GB	48 GB	96 GB	128 GB	256 GB	256 GB	256 GB	768 GB
Memory Type	DDR2-400 MHz	FB-DIMM/ DDR2-667 MHz	FB-DIMM/ DDR2-667 MHz	FB-DIMM/ DDR2-667 MHz	DDR3-1333 MHz[7]	DDR3-1333 MHz	DDR3-1333 MHz	DDR3-1600 MHz	DDR4-2133 MHz[8]	DDR4-2400 MHz	DDR4-2666 MHz

DDR – double data rate; FB-DIMM – fully buffered dual in-line memory module; FSB – front side bus; Intel® QPI – Intel® QuickPath Interconnect; Intel® UPI – Intel® UltraPath Interconnect

[7] DDR3-1333 RAM running at 1066 MHz. [8] DDR4-2133 RAM running at 1866 MHz.

Results

Results are shown in Figure 1; actual runtimes are on the following page in Table 3. The Intel Xeon Platinum 8168 processor-based server completed the tests up to 1.37x faster than a previous-generation Intel Xeon processor E5-2699 v4-based server[Δ]. For historical purposes, we also show that the latest processor-based server is up to 24x faster than a server based on the Intel Xeon processor 5160[Δ] and up to 86x faster than a server based on a single-core 64-bit Intel Xeon processor[Δ].

24X FASTER[Δ]

2006-2017

Higher is Better

All cores loaded using Intel Xeon processor 5160 as baseline

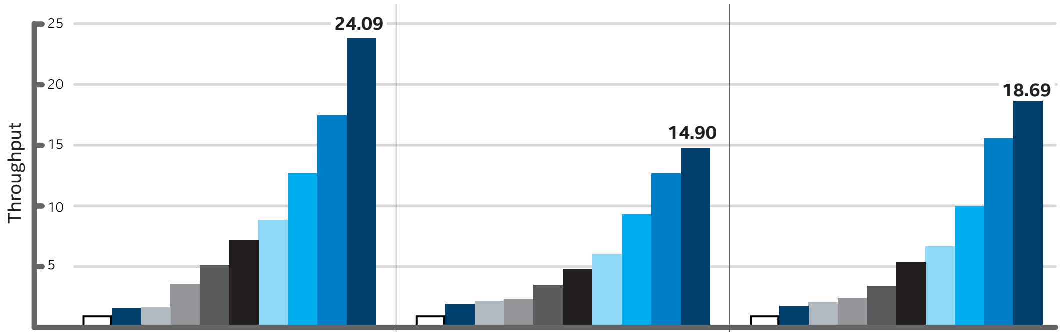

RELATIVE THROUGHPUT	Simulation (113 jobs)	Physical Verification (DRC)	Physical Verification (NAC)
Intel® Xeon® processor 5160	1.00	1.00	1.00
Intel® Xeon® processor X5365	1.58	1.94	1.79
Intel® Xeon® processor X5460	1.65	2.20	2.06
Intel® Xeon® processor X5570	3.63	2.33	2.43
Intel® Xeon® processor X5675	5.20	3.53	3.46
Intel® Xeon® processor E5-2680	7.22	4.88	5.40
Intel® Xeon® processor E5-2680 v2	8.94	6.09	6.75
Intel® Xeon® processor E5-2697 v3	12.92	9.20	10.04
Intel® Xeon® processor E5-2699 v4	17.63	12.44	15.70
Intel® Xeon® Platinum 8168 processor	24.09	14.90	18.69

86X FASTER[Δ]

2004-2017

Higher is Better

All cores loaded using 64-bit Intel Xeon processor with 1 MB L2 cache as baseline

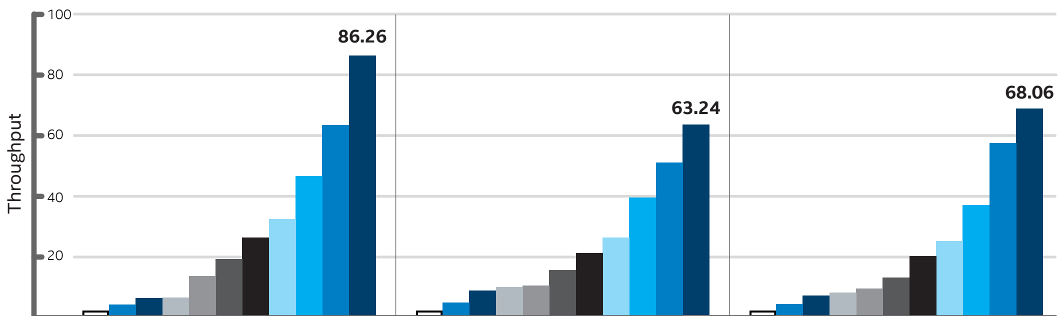

RELATIVE THROUGHPUT	Simulation (113 jobs)	Physical Verification (DRC)	Physical Verification (NAC)
Intel® Xeon® processor 1MB L2	1.00	1.00	1.00
Intel® Xeon® processor 5160	3.58	4.24	3.64
Intel® Xeon® processor X5365	5.65	8.22	6.50
Intel® Xeon® processor X5460	5.91	9.32	7.50
Intel® Xeon® processor X5570	12.98	9.89	8.84
Intel® Xeon® processor X5675	18.63	14.98	12.59
Intel® Xeon® processor E5-2680	25.87	20.70	19.66
Intel® Xeon® processor E5-2680 v2	32.01	25.86	24.59
Intel® Xeon® processor E5-2697 v3	46.28	39.04	36.55
Intel® Xeon® processor E5-2699 v4	63.14	52.79	57.18
Intel® Xeon® Platinum 8168 processor	86.26	63.24	68.06

Figure 1. Electronic Design Automation (EDA) summary test results showing relative throughput of 64-bit Intel® Xeon® processors. Note: Same application binary used across all the platforms.

[Δ] For more complete information about performance and benchmark results, visit intel.com/benchmarks. Performance results based on testing details and system configuration. See the full disclaimer and system configurations on page 6.

Share:

Table 3. Electronic Design Automation (EDA) Test Results Showing Runtimes and Workload Configurations

	64-bit Intel® Xeon® Processor with 1 MB L2 Cache	Intel® Xeon® Processor 5160	Intel® Xeon® Processor X5365	Intel® Xeon® Processor X5460	Intel® Xeon® Processor X5570	Intel® Xeon® Processor X5675	Intel® Xeon® Processor E5-2680	Intel® Xeon® Processor E5-2680 v2	Intel® Xeon® Processor E5-2697 v3	Intel® Xeon® Processor E5-2699 v4	Intel® Xeon® Platinum 8168 Processor
SIMULATION (113 CPU MODEL TESTS)											
Number of Simultaneous Jobs	2	4	8	8	8	12	16	20	28	44	48
Total Runtime (hh:mm:ss)	79:41:46	22:15:24	14:06:54	13:28:57	6:08:23	4:16:36	3:04:52	2:29:23	1:43:20	1:15:44	0:55:26
Relative Throughput	1.00	3.58	5.65	5.91	12.98	18.63	25.87	32.01	46.28	63.14	86.26
PHYSICAL VERIFICATION (DESIGN RULE CHECK [DRC])											
Simultaneous 2-Threaded Jobs	1	2	4	4	4	6	8	10	14	22	24
Total Number of Iterations	9240	4620	2310	2310	2310	1540	1155	924	660	420	385
Total Number of Jobs	9240	9240	9240	9240	9240	9240	9240	9240	9240	9240	9240
Total Runtime (hh:mm:ss)	60052:18:00	14151:19:00	7308:35:00	6443:37:00	6070:10:00	4008:16:40	2900:58:30	2321:48:24	1538:21:00	1137:37:00	949:40:40
Relative Throughput	1.00	4.24	8.22	9.32	9.89	14.98	20.70	25.86	39.04	52.79	63.24
PHYSICAL VERIFICATION (NODE ANTENNA CHECK [NAC])											
Simultaneous 2-Threaded Jobs	1	2	4	4	4	6	8	10	14	22	24
Total Number of Iterations	9240	4620	2310	2310	2310	1540	1155	924	660	420	385
Total Number of Jobs	9240	9240	9240	9240	9240	9240	9240	9240	9240	9240	9240
Total Runtime (hh:mm:ss)	16390:44:00	4500:39:00	2520:28:00	2186:09:30	1853:46:30	1302:09:20	833:31:30	666:33:48	448:26:00	286:39:00	240:50:20
Relative Throughput	1.00	3.64	6.50	7.50	8.84	12.59	19.66	24.59	36.55	57.18	68.06

Conclusion

The new Intel Xeon processor Scalable family delivers significant improvements in throughput performance for Intel design workloads across a range of EDA applications in the data center.

Using a weighted performance measure of end-to-end EDA applications based on Intel silicon design tests, we found that the effective refresh ratio to replace servers based on the 8-core Intel Xeon processor E5-2600 series with servers based on the Intel Xeon processor Scalable family is approximately 3.2:1. Based on our performance assessment and our refresh cycle, we plan to deploy servers based on the new Intel Xeon processor Scalable family, which will enable us to achieve greater throughput while realizing operational benefits such as cost avoidance of data center construction and reduced power consumption.

Our test results suggest that other technical applications with large memory requirements — such as simulation and verification applications in the auto, aeronautical, oil and gas, and life sciences industries — could see similar throughput improvements, depending on their workload characteristics.

For more information on Intel IT best practices, visit **www.intel.com/IT**.

IT@Intel

We connect IT professionals with their IT peers inside Intel. Our IT department solves some of today's most demanding and complex technology issues, and we want to share these lessons directly with our fellow IT professionals in an open peer-to-peer forum.

Our goal is simple: improve efficiency throughout the organization and enhance the business value of IT investments.

Follow us and join the conversation:

- Twitter
- #IntelIT
- LinkedIn
- IT Center Community

Visit us today at **intel.com/IT** or contact your local Intel representative if you would like to learn more.

Related Content

If you liked this paper, you may also be interested in these related stories:

- Disaggregated Servers Drive Data Center Efficiency and Innovation paper
- Data Center Strategy Leading Intel's Business Transformation paper
- High-Performance Computing for Silicon Design paper
- Extremely Energy-Efficient, High-Density Data Centers paper
- How Software-Defined Infrastructure is Evolving at Intel paper

The following system configurations and performance tests are discussed in this paper. For more information go to intel.com/performance.

Intel® Xeon® Platinum 8168 processor improves throughput up to 1.37x compared to a previous-generation Intel Xeon processor E5-2699 v4-based server. Intel Xeon Platinum 8168 Processor (24 cores, 2.7 GHz, 33 MB cache, 768 GB RAM, DDR4-2666 MHz) vs. Intel® Xeon® Processor E5-2699 v4 (2.2 GHz, 55 MB cache, 256 GB RAM, DDR4-2400 MHz).

Intel Xeon Platinum 8168 processor completed the workloads up to 86x faster than a server based on a 64-bit Intel Xeon processor. Intel Xeon Platinum 8168 Processor (24 cores, 2.7 GHz, 33 MB cache, 768 GB RAM, DDR4-2666 MHz) vs. 64-bit Intel® Xeon® Processor with 1 MB L2 cache (1 core, 3.6 GHz, 16 GB RAM, DDR2-400 MHz).

Intel Xeon Platinum 8168 processor-based server was up to 24x faster than a server based on the Intel Xeon processor 5160. Intel Xeon Platinum 8168 Processor (24 cores, 2.7 GHz, 33 MB cache, 768 GB RAM, DDR4-2666 MHz) vs. Intel® Xeon® Processor 5160 (2 cores, 3.0 GHz, 4 MB cache, 16 GB RAM, FB-DIMM/DDR2-667 MHz).

Printed in USA ♻ Please Recycle 0318/DPEA/KC/PDF