

# IT@Intel: Increasing EDA Performance and Throughput with the Intel® Xeon® Processor Scalable Family

Intel IT testing of the Intel® Xeon® Gold 6300 processor Series yields the best per-core performance (up to 1.26x); that same series delivers the best throughput (up to 2.76x) with high core count CPUs compared to a four-generation-older Intel Xeon processor E5-2680 v4

## Intel IT Authors

**Shesha Krishnapura**

Intel Fellow and IT CTO

**Murty Ayyalasomayajula**

Senior Staff Engineer

**Shaji Kootaal Achuthan**

Senior Staff Engineer

**Vipul Lal**

Senior Principal Engineer

**Archana Somasekhara**

Systems Engineer

**Ty Tang**

Senior Principal Engineer

## Table of Contents

Executive Overview .....	1
Background .....	2
Evolution of the Intel® Xeon® Processor and the EDA Workflow. ....	2
Test Methodology .....	3
Results: Faster Servers Process More EDA Jobs in Less Time. ....	4
Conclusion .....	7
Related Content .....	7

## Executive Overview

Intel IT operates 56 data center modules at 16 data center sites. These sites have a total capacity of 103 megawatts, housing more than 360,000 servers that underpin the computing needs of 116,000 employees. Intel IT has four main segments of operation: Design, Office, Manufacturing and Enterprise. This paper focuses on only the Design segment.

Intel's silicon Design engineers need significant increases in computing capacity to deliver each new generation of silicon chips. To meet those requirements, Intel IT conducts ongoing throughput performance tests using real-world Intel silicon Design workloads. These tests measure Electronic Design Automation (EDA) workload throughput and help us analyze the performance improvements—and in turn, business benefit offered by newer generations of Intel® processors.

We recently tested two-socket servers based on the Intel® Xeon® Gold 6300 processor Series, running single- and multi-threaded EDA applications operating on more than 248 hours of Intel silicon Design workloads. Select results include the following:<sup>1</sup>

- **Higher frequency for per-core performance.** For critical-path EDA workloads, selecting a high-frequency CPU like the Intel Xeon Gold 6334 processor (16 cores per server) can deliver up to 1.26x higher per-core performance compared to lower-frequency CPUs in the same generation of processors.
- **Higher core counts for throughput.** For volume validation runs, selecting a higher-core-count CPU at optimal frequency like the Intel Xeon Gold 6342 processor (48 cores per server) can deliver up to 2.56x higher Register Transfer Level (RTL) Simulation throughput per server when compared to a lower core-count CPU (16 cores per server) in the same generation of processors. The Intel Xeon Gold 6342 processor (48 cores per server) completed workloads up to 2.87x faster than a previous-generation Intel Xeon 6250 processor-based server, which has only 16 cores. Even compared to a four-generation-older Intel Xeon processor E5-2680 v4 (28 cores per server), the server with the newer processor outperformed the older processor by up to 2.76x in throughput.

Based on our performance assessment and our refresh cycle, we are deploying servers based on the 3rd Gen Intel Xeon processor Scalable family in our data centers. By doing so, we have significantly increased EDA throughput performance to improve the overall EDA design cycles and optimize time to market of Intel® chips.

## 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. Design engineers run more than 213 million compute-intensive batch jobs every week. Each job takes from a few seconds to several days to complete.

As design complexity increases, so do the requirements for compute capacity, 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, helping to avoid expensive data center construction and reduce operational costs 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.

## Evolution of the Intel® Xeon® Processor and the EDA Workflow

The architectural enhancements shown in 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

As shown in Figure 1, EDA includes several phases, including front-end logic design, followed by back-end physical design and then by tape-in/tape-out. This paper discusses selected tools in the front-end and back-end design phases.

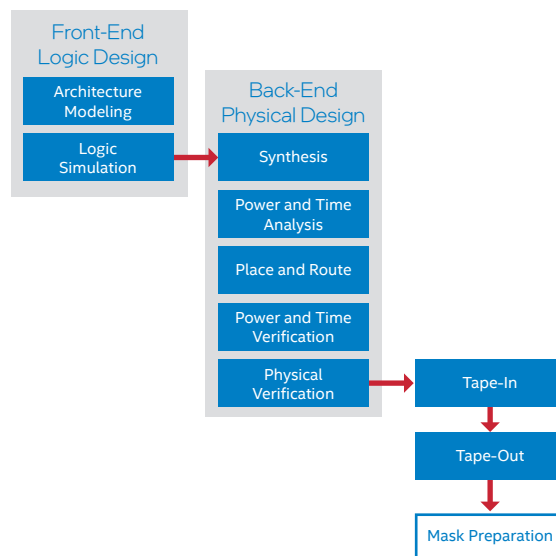


Figure 1. The EDA phases of silicon design.

Table 1. Comparison of Two-socket Servers Based on Intel® Xeon® Processors Over Time

	2004-2005	2006-2008	2009-2011	2012	2013	2014	2016	2017-2020	2021	
<b>Design</b>										
<b>Intel® Chipset</b>	E7520	5400	5520	C600		C610	C620	C620A		
<b>Process Technology</b>	90nm	65nm and 45nm	45nm and 32nm	32nm	22nm		14nm			
<b>Cores per Socket</b>	1	2 or 4	4 or 6	8	10	14	22	28	40	
<b>Interconnect Speed</b>	6.4 GB/s	21-25 GB/s	25.6 GB/s	32 GB/s		38.4 GB/s		41.6 GB/s	44.8 GB/s	
<b>Intel® QuickPath Interconnect</b>	—	—	Yes	Yes					—	
<b>Intel® UltraPath Interconnect</b>	—	—	—	—					Yes	Yes
<b>DIMMs</b>	Up to 8	Up to 16	Up to 18	Up to 24					Up to 32	
<b>Memory Type</b>	DDR2	FB-DIMM/DDR2 or FB-DIMM/DDR2	DDR3	DDR3	DDR3	DDR4	DDR4	DDR4	DDR4	
<b>Memory Speed</b>	400 MHz	667 MHz or 800 MHz	800/1066/1333 MHz	1333/1600 MHz	1333/1600/1866 MHz	1600/1866/2133 MHz	2400 MHz	2666 MHz	3200 MHz	
<b>Memory Bandwidth</b>	Up to 6.4 GB/s	Between 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	Up to 204.76 GB/s	

## Front-end Logic Design

In the silicon design process, front-end logic design includes architecture specification and functional verification design. Front-end EDA workloads are single-threaded or lightly multi-threaded and run in a highly distributed compute environment.

## Back-end Physical Design

Synthesis converts a Register Transfer Level (RTL) description to a structural gate-level netlist, which instantiates standard cells, macros and areas that compose the circuit and its connections. The synthesized netlist is verified for functionality and timing to ensure it operates as intended before the Place-and-Route stage translates the gate-level netlist into a physical design. Static-Timing-Analysis as well as other post-layout static and dynamic analyses are then performed to check all possible paths for timing violations, voltage drop analysis and more, and to deliver accurate signoff information for timing, signal integrity and power analysis for the design. Finally, the Physical-Verification stage ensures the physical design meets manufacturing constraints imposed by process technology; the verification includes Design Rule Check, Layout versus Schematic and Electrical Rule Check. Back-end EDA workloads are generally multi-threaded and consume large memory and terabytes of data.

## Tape-in/Tape-out

During tape-in, Intel chip design teams create multi-gigabyte hierarchical layout databases that specify the design to be manufactured. During tape-out, these layout databases are processed using EDA tools, which apply extremely compute-intensive resolution enhancement techniques (RET) to update layout data for mask manufacturability and verify the data for compliance to mask manufacturing rules.

## Test Methodology

We performed various tests on two-socket servers. Some tests compared several different CPUs in the 3rd Gen Intel® Xeon® Gold 6300 processor Series. Other tests compared the 3rd Gen processors to a baseline of an older Intel® Xeon® processor E5-2680 v4. We conducted front-end and back-end tests using industry-leading EDA applications to run single- and multi-threaded Intel silicon Design workloads.

Our goal was to assess performance and throughput improvements by measuring the time needed 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 wherever possible.



## Intel's Latest Processors for Data Center Workloads

3rd Gen Intel Xeon Scalable processors are packed with performance- and security-enhancing features, including the following:

- 10nm process technology
- Enhanced per-core performance, with up to 40 cores in a standard socket
- Enhanced memory performance with support for up to 3200 MT/s DIMMs (2 DIMMs per channel)—32 percent more than the previous generation
- Increased memory capacity with up to eight channels
- Database compression with Intel® Vector Byte Manipulation Instructions
- Support for Intel® Optane™ persistent memory 200 series
- Built-in AI acceleration with enhanced performance of Intel® Deep Learning Boost
- Faster inter-node connections with three Intel® Ultra Path Interconnect links at 11.2 GT/s
- Significantly increased I/O bandwidth, along with PCI Express 4 support on up to 64 lanes (per socket) at 16 GT/s
- Hardware-enhanced security of Intel® Crypto Acceleration

## Results: Faster Servers Process More EDA Jobs in Less Time

Test results are shown in Figures 2 through 6; system specifications and runtimes are provided in Tables 2 and 3.

### Benefits of the 3rd Gen Intel Xeon Processor Scalable Family

We conducted tests to understand throughput and the core-to-core speed up that we are able to obtain using 3rd Gen Xeon Scalable processors instead of 2nd Gen Xeon Scalable processors. For this test, we selected four different types of SKUs in each generation:

- An 8-core very-high-frequency SKU with a total of 16 cores in a two-socket system—Intel Xeon Gold 6250 processor versus Intel Xeon Gold 6334 processor
- A 16-core high-frequency SKU with a total of 32 cores in a two-socket system —Intel Xeon Gold 6246R processor versus Intel Xeon Gold 6346 processor
- A 24-core scalable performance SKU with a total of 48 cores in a two-socket system—Intel Xeon Gold 6240R processor versus Intel Xeon Gold 6336Y processor
- A 24-core higher scalable performance SKU with a total of 48 cores in a two-socket system —Intel Xeon Gold 6248R processor versus Intel Xeon Gold 6342 processor

As shown in Figure 2, comparing the performance of same-core-count CPUs across generations results in performance increases up to 1.24x.

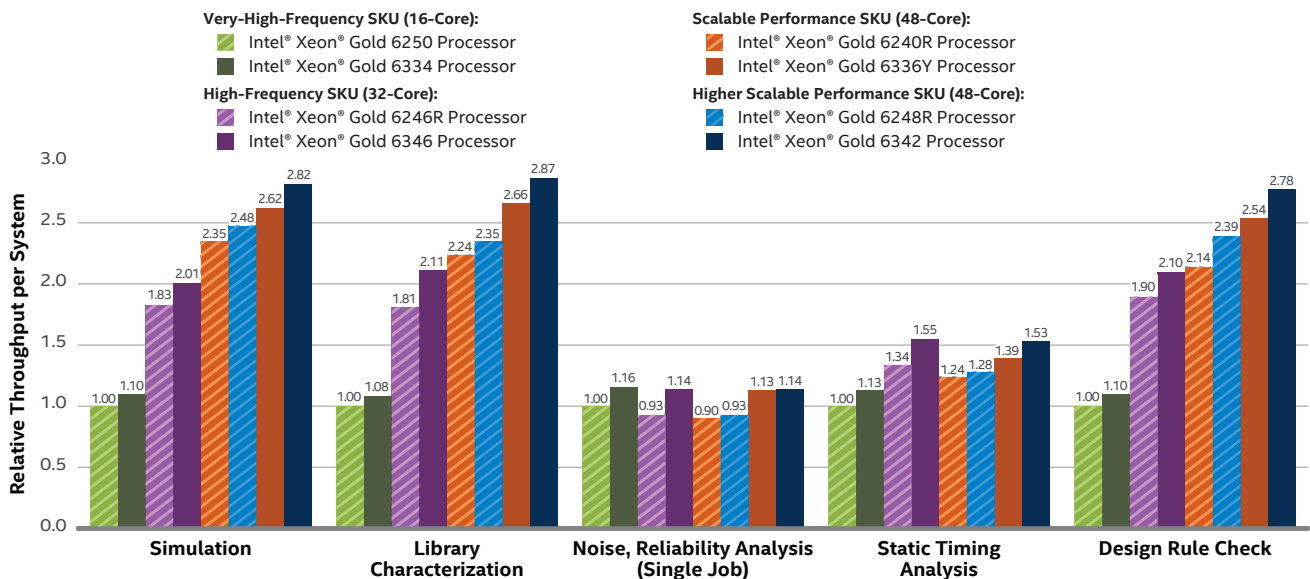
### Optimizing Platform Selection within the 3rd Gen Intel Xeon Processor Scalable Family

For volume validation runs, overall cluster throughput is desirable; but for critical-path runs, the highest per-core performance is needed. Both types of workloads can be supported using a variety of Intel Xeon processor Scalable family SKUs. We compared selected 8-core, 16-core and two 24-core offerings from the 3rd Generation Intel Xeon processor Scalable family. This provides a choice for critical path versus volume validation runs.

- An 8-core very-high-frequency SKU: Intel Xeon Gold 6334 processor
- A 16-core high-frequency SKU: Intel Xeon Gold 6346 processor
- A 24-core scalable performance SKU: Intel Xeon Gold 6336Y processor
- A 24-core higher scalable performance SKU: Intel Xeon Gold 6342 processor

## Relative Two-Socket System Throughput for EDA Workloads

2nd Gen Intel® Xeon® Processors vs. 3rd Gen Intel Xeon Processors  
HIGHER IS BETTER



**Figure 2.** 3rd Generation Intel® Xeon® Scalable processor vs. 2nd Generation Intel Xeon Scalable processor: More cores provide more system throughput for EDA workloads. Note: Same application binary used across all the platforms.<sup>1</sup>

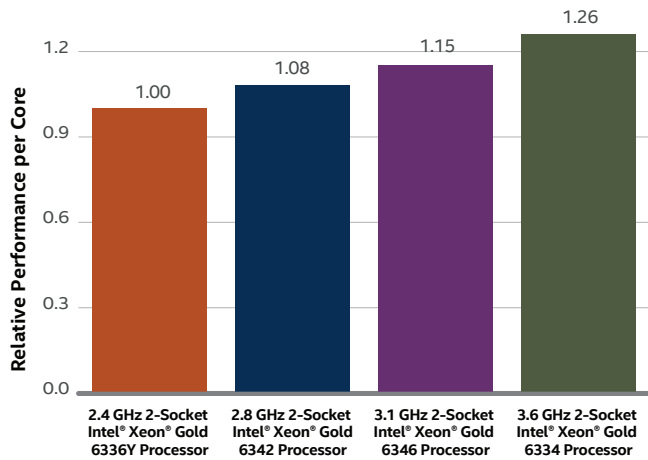
Our findings, shown in Figures 3 and 4, indicate the following best practices:

- Selecting a higher-frequency CPU can deliver up to 1.26x higher RTL Simulation per-core performance for critical-path EDA runs.
- Selecting a higher-core-count CPU can deliver up to 2.56x higher RTL Simulation throughput per server when compared to a lower core-count SKU, which is ideal for volume validation runs. The choice of optimal SKU is based on the end-user workload use model.

### Relative Performance for RTL Simulation

Across the Intel® Xeon® Gold 6300 Processor Series

HIGHER IS BETTER

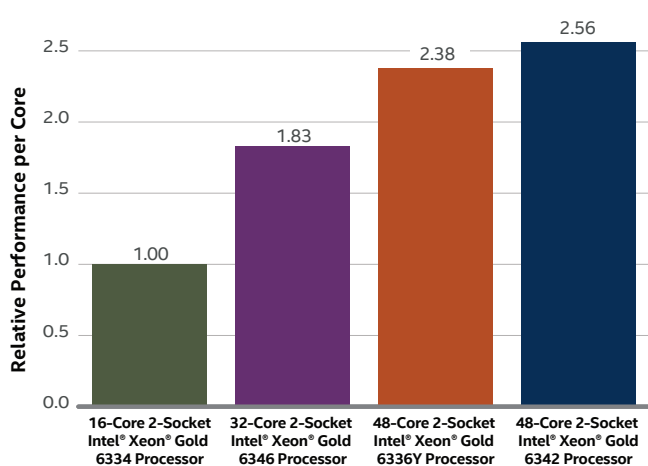


**Figure 3.** 3rd Generation Intel® Xeon® Scalable processors: A higher frequency results in better per-core performance for critical-path EDA workloads. Note: Same application binary used across all the platforms.<sup>1</sup>

### Relative System Throughput per Server for RTL Simulation

Across the Intel® Xeon® Gold 6300 Processor Series

HIGHER IS BETTER



**Figure 4.** 3rd Generation Intel® Xeon® Scalable processor: A higher core count results in better per system throughput for volume validation runs. Note: Same application binary used across all the platforms.<sup>1</sup>

### EDA Per-Core Performance Across Four Generations of Intel Xeon Processors

We tested four generations of Intel Xeon processors to compare the per-core performance for critical-path EDA workloads.

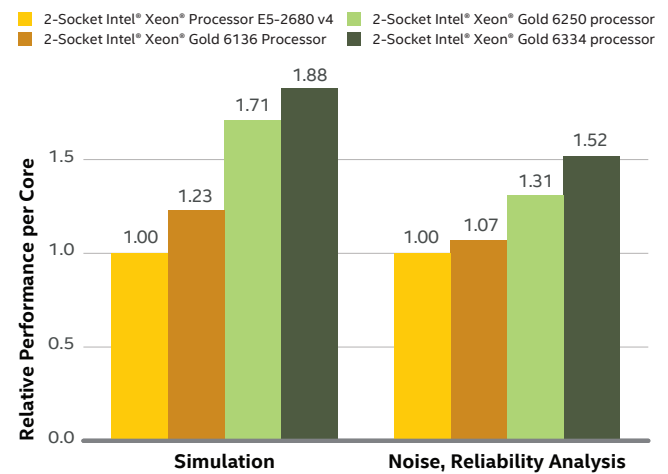
- A 14-core Intel Xeon processor E5-2680 v4
- A 12-core Intel Xeon Gold 6136 processor (1st Gen)
- An 8-core Intel Xeon Gold 6250 processor (2nd Gen)
- An 8-core Intel Xeon Gold 6334 processor (3rd Gen)

Based on the test results, for performance-critical workloads, an 8-core 3rd Gen Intel Xeon Gold 6334 processor-based server can reduce license consumption time and provide a performance boost of 1.52x to 1.88x across all the tested workloads (see Figure 5).

### Per-Core EDA Relative Performance

Across Four Generations of Intel® Xeon® Processors

HIGHER IS BETTER



**Figure 5.** Per-core EDA performance of select Intel® Xeon® processors across four generations. Note: Same application binary used across all the platforms.<sup>1</sup>

### EDA Throughput Across Four Generations of Intel Xeon Processors

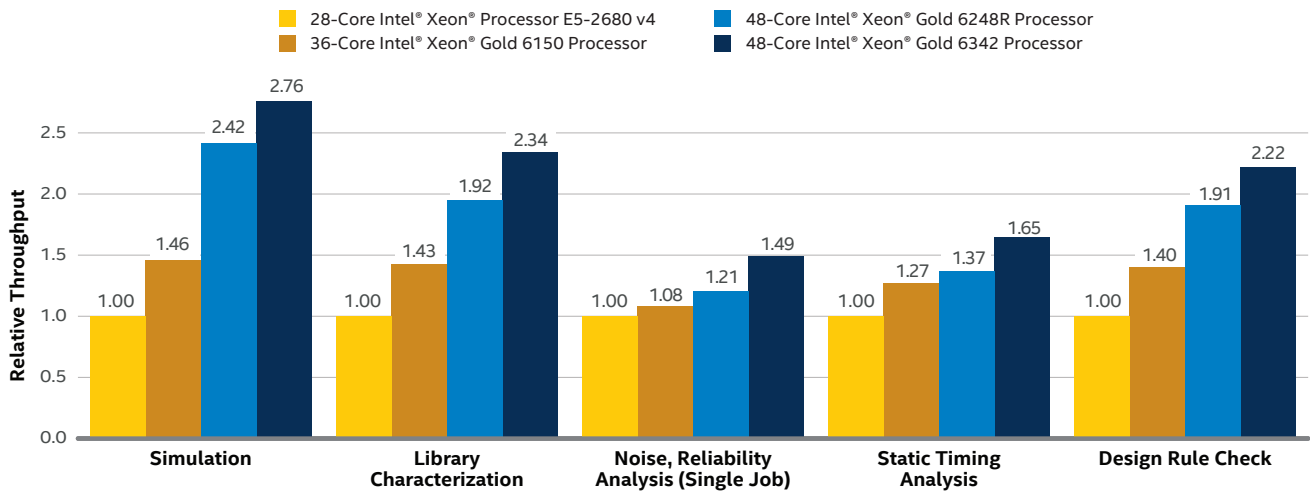
We tested four generations of Intel Xeon processors to compare the throughput:

- A 14-core Intel Xeon processor E5-2680 v4
- An 18-core Intel Xeon Gold 6150 processor
- A 24-core Intel Xeon Gold 6248R processor
- A 24-core Intel Xeon Gold 6342 processor

Based on the test results, a 3rd Gen Intel Xeon Gold 6342 processor-based server provides up to a 2.76x increase in throughput per server, compared to a four-generation-older processor-based server (see Figure 6). In other words, 11 older Intel Xeon processor E5-2680 v4-based servers can be replaced with only four servers based on the latest generation of Intel Xeon Scalable processor.

## Relative Two-Socket System Throughput for EDA Workloads Across Four Generations of Intel® Xeon® Processors

HIGHER IS BETTER



**Figure 6.** EDA throughput of select Intel® Xeon® processors across four generations: Consolidate servers by up to 2.76x to reduce data center footprint, power and cooling costs and software license costs.

**Table 2.** Specification of the Systems Used for Testing

	Intel® Xeon® Processor E5 v4 Family	1st Generation Intel Xeon Processor Scalable Family		2nd Generation Intel Xeon Processor Scalable Family				3rd Generation Intel Xeon Processor Scalable Family			
Processor	Intel Xeon E5-2680 v4	Intel Xeon 6136	Intel Xeon 6150	Intel Xeon 6250	Intel Xeon 6246R	Intel Xeon 6240R	Intel Xeon 6248R	Intel Xeon 6334	Intel Xeon 6346	Intel Xeon 6336Y	Intel Xeon 6342
Cores per Socket	14	12	18	8	16	24	24	8	16	24	24
Frequency	2.4 GHz	3.0 GHz	2.7 GHz	3.9 GHz	3.4 GHz	2.4 GHz	3.0 GHz	3.6 GHz	3.1 GHz	2.4 GHz	2.8 GHz
Cache per CPU	35 MB	24.75 MB	24.75 MB	35.75 MB	35.75 MB	35.75 MB	35.75 MB	18 MB	36 MB	36 MB	36 MB
Bus Speed	9.6 GT/s	10.4 GT/s	10.4 GT/s	10.4 GT/s	10.4 GT/s	10.4 GT/s	10.4 GT/s	11.2 GT/s	11.2 GT/s	11.2 GT/s	11.2 GT/s
RAM	512 GB	768 GB	768 GB	768 GB	768 GB	768 GB	768 GB	1 TB	1 TB	1 TB	1 TB
Memory Type	DDR4-2400 MHz	DDR4-2666 MHz	DDR4-2666 MHz	DDR4-2933 MHz	DDR4-2933 MHz	DDR4-2933 MHz	DDR4-2933 MHz	DDR4-3200 MHz	DDR4-3200 MHz	DDR4-3200 MHz	DDR4-3200 MHz

Testing by Intel IT as of April 2021 through January 2022.

**Table 3.** Workload Run Times

Workload (Cores Per Server)	Intel® Xeon® Processor E5 v4 Family	1st Gen Intel Xeon Processor Scalable Family		2nd Generation Intel Xeon Processor Scalable Family				3rd Generation Intel Xeon Processor Scalable Family			
	Intel Xeon E5-2680 v4 (28)	Intel Xeon 6136 (24)	Intel Xeon 6150 (36)	Intel Xeon 6250 (16)	Intel Xeon 6246R (32)	Intel Xeon 6240R (48)	Intel Xeon 6248R (48)	Intel Xeon 6334 (16)	Intel Xeon 6346 (32)	Intel Xeon 6336Y (48)	Intel Xeon 6342 (48)
Simulation	8:35:01	8:08:26	5:53:44	8:46:10	4:48:16	3:43:52	3:32:24	7:58:34	4:21:43	3:20:40	3:06:37
Library Characterization	1:45:31	1:41:35	1:14:01	2:09:01	1:11:07	0:57:43	0:54:53	2:00:00	1:01:01	0:48:31	0:45:00
Power, Noise Signal Integrity	9:15:15	8:38:26	8:31:48	7:05:01	7:36:40	7:50:06	7:38:39	6:06:09	6:12:57	6:17:37	6:11:43
Static Timing Analysis	4:55:48	4:18:54	3:52:20	4:35:18	3:25:30	3:42:33	3:35:50	4:03:19	2:58:04	3:17:55	2:59:40
Design Rule Check	5:49:02	5:37:24	4:10:05	7:17:46	3:50:27	3:24:29	3:02:53	6:36:39	3:28:27	2:52:16	2:37:33

NOTE: Benchmarks were conducted with all the cores loaded; static timing analysis workload is limited to a maximum of 32 threads. Results are dependent on tool type, version, and data set.



## Conclusion

The new Intel Xeon processor Scalable family delivers significant improvements in throughput and per-core performance for Intel silicon Design workloads across a range of EDA applications in the data center.

For performance-centric workloads, selecting a higher-frequency Intel Xeon 6300 processor Series SKU can increase performance by up to 1.26x, compared to lower-frequency 6300 Series SKUs. For throughput-centric workloads, an Intel Xeon 6300 processor Series SKU with balanced frequency and core count can deliver up to 2.56x better throughput than other 6300 Series SKUs and 2.87x better throughput than a 2nd Gen Intel Xeon Gold processor. And compared to a four-generation-older Intel Xeon processor E5-2680 v4 with 24 cores, a server with a newer processor can outperform the older processor by up to 2.76x.

Using a weighted performance measure of end-to-end EDA applications based on Intel silicon Design tests, we found that we can replace 11 older-generation Intel Xeon processor E5-2600 v4-based servers with just four of the latest Intel Xeon Gold 6300 processor Series-based servers. Based on our performance assessment and our [refresh cycle](#), we have deployed servers based on the new Intel Xeon Gold 6300 processor Series, enabling 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.

## Related Content

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

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

For more information on Intel IT best practices, visit [www.intel.com/IT](http://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 on [Twitter](#) or at [#IntelIT](#).

Visit us today at [intel.com/IT](http://intel.com/IT) or contact your local Intel representative if you would like to learn more.

### Intel IT Contributors

**Matt Ammann**, Principal Engineer, for providing server network operations support

**Juan J. Sanchez and Chandra Sudireddy**, IT Operations, for providing multiple server configurations and operational support



<sup>1</sup> Testing by Intel IT as of April 2021 through January 2022. See [Tables 2 and 3](#) for configuration details and workload run times.

Performance varies by use, configuration and other factors. Learn more at [intel.com/PerformanceIndex](http://intel.com/PerformanceIndex). Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure. Your costs and results may vary. Intel technologies may require enabled hardware, software or service activation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation in the U.S. and other countries. Other names and brands may be claimed as the property of others. Intel Corporation. 0122/WWES/KC/PDF