REALIZING A MORE PRODUCTIVE EDA ENVIRONMENT

Improving the economics of semiconductor design with HPE Systems and AMD EPYC processors
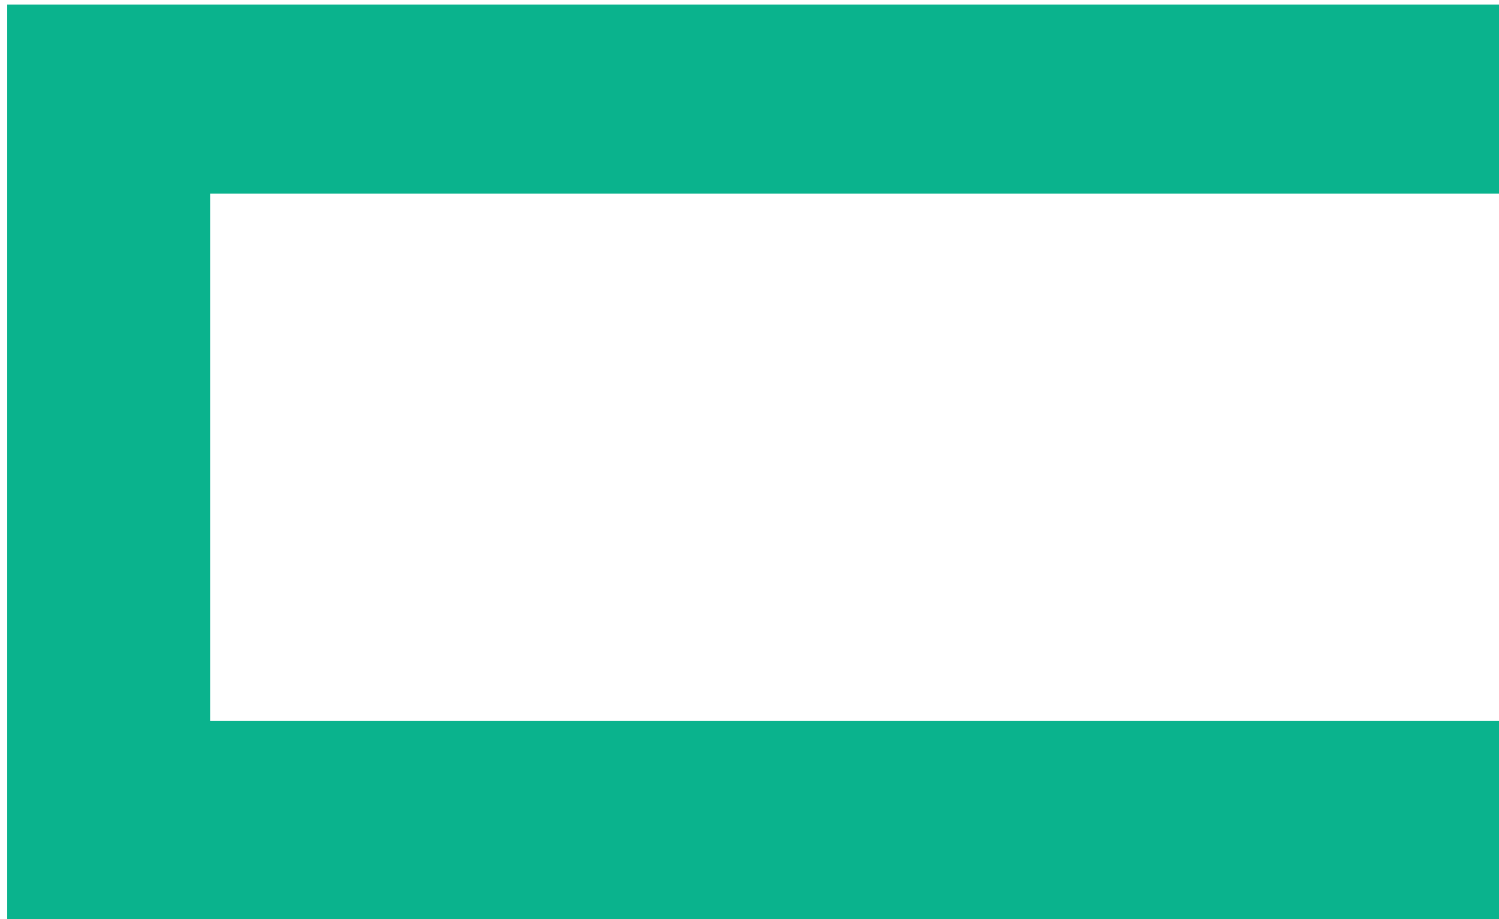
TABLE OF CONTENTS

INTRODUCTION

Few industries are more competitive than modern electronics manufacturing and chip design. Consumers expect devices to be faster, cheaper, and more reliable with each generation. Whether large or small, electronics manufacturers rely on electronic design automation (EDA) to enable these improvements.

High performance computers are used in all phases of the EDA cycle from system-level design, to logic, to analog design, to simulation and layout. Even for mid-sized projects, verifying proper device functionality is one of the largest challenges faced by chip designers. As engineers make changes to a design, they need to run extensive computer simulations to verify functionality. By most estimates, regression testing and verification account for roughly 80% of simulation workloads in modern electronic design environments.[1] Given the enormous cost of committing a design to silicon, projects must be error-free before tape-out. The performance and capacity of the EDA simulation environment directly affects product quality, time-to-market, downstream support costs, and IT costs—all impacting the bottom line.

EDA firms compete based on the effectiveness of their design environments. In this brief, we explain how high-performance HPE Apollo 2000 Gen10 Plus Systems and HPE ProLiant servers powered by 2nd and 3rd generation AMD EPYC™ processors can provide a decisive advantage to electronics manufacturers. HPE systems can help customers increase simulation capacity, improve throughput and productivity, and reduce TCO in EDA server farms.

CUSTOMER CHALLENGES

Device simulation becomes more difficult as designs become larger. As the number of registers and memory in a device increases (call this "n"), the number of states to be modeled increases exponentially (2^n). System on a chip (SoC) designs are frequently in the range of hundreds of millions or even billions of gates, making verification more challenging with each product generation as designs become more complex.

In addition to size and complexity, reliability and security are important considerations. Products such as sensors for autonomous vehicles, embedded control systems, and medical devices need to work flawlessly. This demands higher levels of verification coverage and increased simulation to ensure quality and reliability.

Figure 1 illustrates the challenge faced by EDA design centers. Bringing innovative new products to market and improving reliability requires more simulation capacity. However, firms simultaneously face pressure to shorten design cycles to meet time-to-market objectives with limited budgets for hardware and software.

[1] Based on estimates from HPE's internal VLSI design environment, 2020.

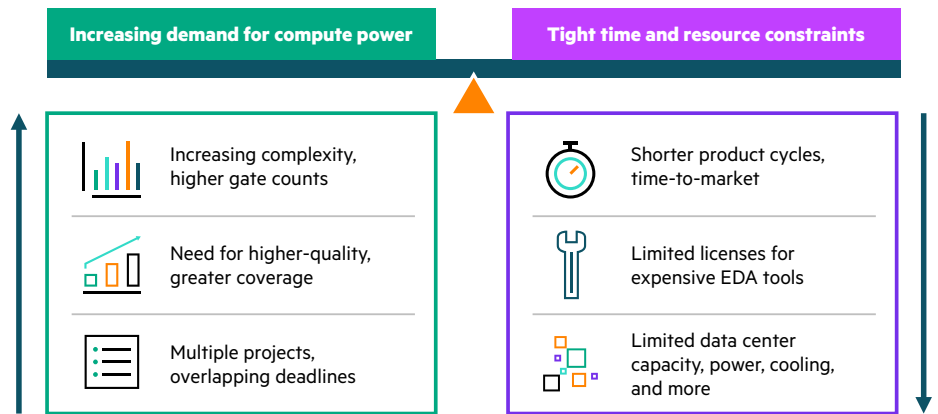| Increasing demand for compute power	Tight time and resource constraints
Increasing complexity, higher gate counts	Shorter product cycles, time-to-market
Need for higher-quality, greater coverage	Limited licenses for expensive EDA tools
Multiple projects, overlapping deadlines	Limited data center capacity, power, cooling, and more

FIGURE 1. EDA firms need more simulation capacity but face tight resource constraints

EDA software tools need to simulate multiple aspects of device functionality over different periods. Chip designers typically run tools from leading EDA vendors, including Cadence®, Synopsys®, and Mentor Graphics®. Workloads are diverse, with some simulations running for minutes or hours, while others run for days, or even weeks on large server farms.

Table 1 describes some typical verification workloads and their characteristics.

TABLE 1. Different types of EDA verification workloads

Category	Verification type	Description
Digital abstractions	Gate Level Simulations (GLS)	Models may consist of billions of gates. Simulation runtimes can range from hours to weeks, depending on the model and simulator.
	Register Transfer Level (RTL)	Models typically consist of millions of lines of C-like code executing 10K–100K simulated cycles per second (cps). Runtimes range from seconds to multiple days.
	Transaction Level Model (TLM)	Models consist of up to 1 million lines of C++-like code running 10K–1M simulated cps. Runtimes range from seconds to hours.
Analog abstractions	Transistor Level (SPICE)	Models consist of analog primitives: resistors, capacitors, transistors, and others.
	Verilog-AMS/ VHDL-AMS (CAMS)	Models consist of behavioral code and operate on voltage and current values in an analog simulator to solve a network.
	System-Level Verification	Models consist of C-like code executed on a digital simulator.

Many EDA tools are single-threaded. To optimize throughput and server utilization, customers tend to run many simulations on multicore servers. To achieve high throughput, customers need:

• High clock frequencies

• Large amounts of physical memory

• Large amounts of L3 cache per simulation

• Low latency and high bandwidth to cache and memory

EDA SOFTWARE LICENSING

A specific challenge faced by electronic manufacturers is the high cost of software tools. Software license costs for EDA environments are typically much higher than hardware costs. Because of this cost disparity, IT administrators tend to be much more concerned about using software resources efficiently than maximizing server utilization. The cost of engineering talent is also an important consideration, and organizations need to maximize their productivity. Figure 2 provides a simplified view of a typical design environment.
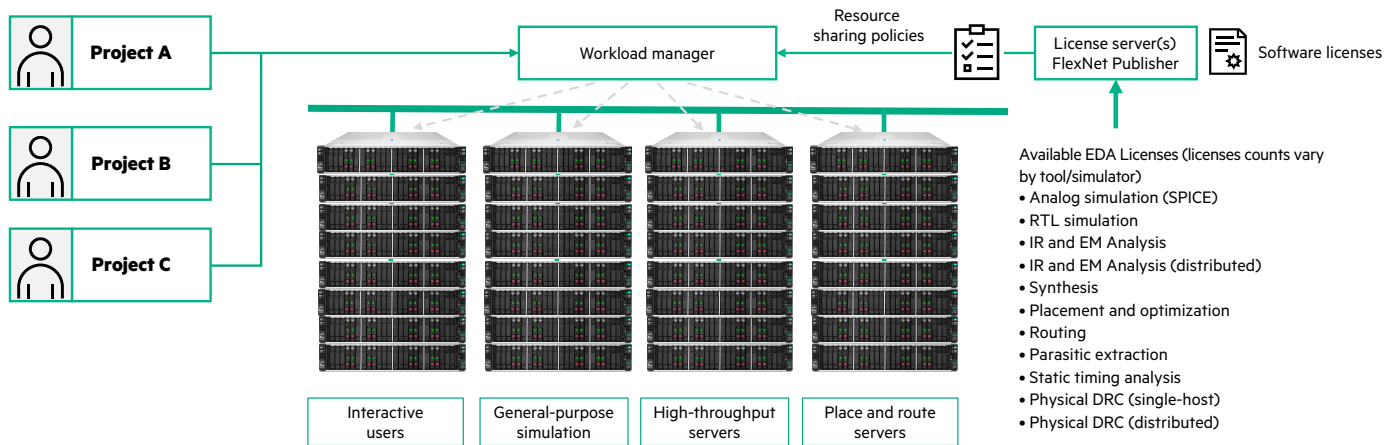

FIGURE 2. A typical EDA environment

Project teams typically work on multiple and sometimes overlapping designs and need access to EDA software tools and servers to run them on. Flexera® FlexNet Publisher (formerly FLEXlm) is used in EDA environments to meter and manage software licenses. As each tool runs, it contacts a license server and checks out a license. Tools return license features when executions are complete. In some cases, a simulation may consume multiple license features.

A single license for a verification tool can cost multiple thousands of dollars per year.[2] Regression tests can involve millions of discrete simulations and completing these quickly requires a large number of licenses. For high-demand tools, a design environment may have hundreds of license features. Overall license costs can easily exceed USD 1M annually for a single tool.

Because licenses are expensive, design firms have a strong incentive to keep these licenses fully utilized. Workload management software plays a critical role, coordinating with license servers, and scheduling various batch and interactive jobs. The scheduler seeks to ensure that project deadlines are met and that resources are shared according to policy, optimizing both licenses and infrastructure resources.

Not only is it important to minimize idle time for licenses, but it is also essential to use the licenses efficiently by running simulations as quickly as possible. A key metric for EDA firms is the number of simulations run per day per license. High-value tools need to execute on server nodes that deliver the highest possible throughput to maximize cost-efficiency.

AMD EPYC processors deliver exceptional performance and scalability for EDA workloads.

- World's first 7 nm x86 server CPU

- Highest available core count to maximize parallelism

- World's first PCIe® Gen4 capable x86 server CPU

- Eight memory channels per socket

- World's first x86 server processor with DDR4 3200 MT/s memory support

- Leadership L3 cache per core

THE "EPYC" ADVANTAGE

Built on 7 nm technology, AMD EPYC processors bring together high core counts, large memory capacity, extreme memory bandwidth, large cache sizes, and massive I/O with the right ratios to enable exceptional HPC workload performance. For EDA users, this can translate into higher-quality designs, reduced regression runtimes, and better license utilization.

While AMD EPYC processors are the choice of next-generation exascale supercomputers,[3] they are also highly affordable, often delivering superior performance to alternative processors while easily fitting within the budgets of design environments of all sizes.

AMD EPYC 7002 series processors were introduced in August 2019. These processors were a game changer, delivering industry leading clock frequencies, latency, memory bandwidth, and cache per core making AMD EPYC a preferred processor for EDA workloads.[4] The latest AMD EPYC 7003 series processors extend this leadership even further, offering exceptional single core performance with a consistent feature set across the stack. Both of these processor families deliver excellent performance and customers can select the processor that best meets their needs.

AMD EPYC 7003 series

AMD EPYC 7003 series processors, introduced in March 2021, offer several advantages over the 7002[5] series. Among these advantages are:

- A unified 8 core cache complex sharing a single 32 MB L3 cache per Core Complex Die (CCD) providing up to twice the amount of directly accessible L3 cache per core with low latency[6]

- Up to a 19% improvement in instructions per cycle (IPC)[7]

- A faster Infinity Fabric™, clocked at 1600 MHz enabling synchronous transfers with the 3200 MT/sec DDR4 memory

- Advanced chip-level security enhancements (SME, SEV-ES, SEV-SNP)

AMD EPYC 7003 series processors are a drop-in upgrade, fully compatible with AMD 7002 series systems.[8] Customers can deploy systems with either 7002 series processors or 7003 series processors depending on their needs.

An ideal architecture for memory-intensive EDA workloads

The unique architecture shown in Figure 3 is the key to the EPYC processor's throughput advantage. The 9-die system on a chip (SoC) features 8 CCDs providing up to 8 cores and 32 MB of cache per CCD. This design places large amounts of L3 cache close to compute cores delivering optimal throughput for clock and cache sensitive RTL and verification workloads. The advanced 7 nm process enables clock frequencies to scale to up to 4.10 GHz with max. boost enabled, helping minimize ISV application license checkout time, and enabling users to get more productivity from expensive ISV license features.[9]

While other processors share relatively small amounts of L3 cache across multiple cores, AMD EPYC processors offer up to 256 MB of L3 cache and provide a direct path between each core and associated L3 cache to speed throughput and help reduce latency.[10] This combination of more L3 cache per core and direct channels to cache combines to deliver exceptional throughput.

- Industry leading 7 nm technology*
- Up to 64 cores/128 threads
- Up to 19% boost in instructions per cycle (IPC)
- Up to 4.10 GHz boost clock speed**
- Up to 256 MB L3 cache
- Up to 32 MB shared L3 cache per core (7003 series)
- Up to 16 MB shared L3 cache per core (7002 series)
- Faster Infinity Fabric
- 8 DDR4 memory channels per socket
- 3200 MT/s DDR
- 128 PCIe Gen4 lanes per socket
- Configurable NUMA operation
- Embedded AMD Secure Processor

* AMD EPYC is the world's first high-performance 7 nm x86 CPU powered by "Zen 2" or "Zen 3" processor cores. amd.com/en/press-releases/2018-11-06-amd-takes-high-performance-datacenter-computing-to-the-next-horizon

** Max. boost for AMD EPYC processors is the maximum frequency achievable by any single core on the processor under normal operating conditions for server systems
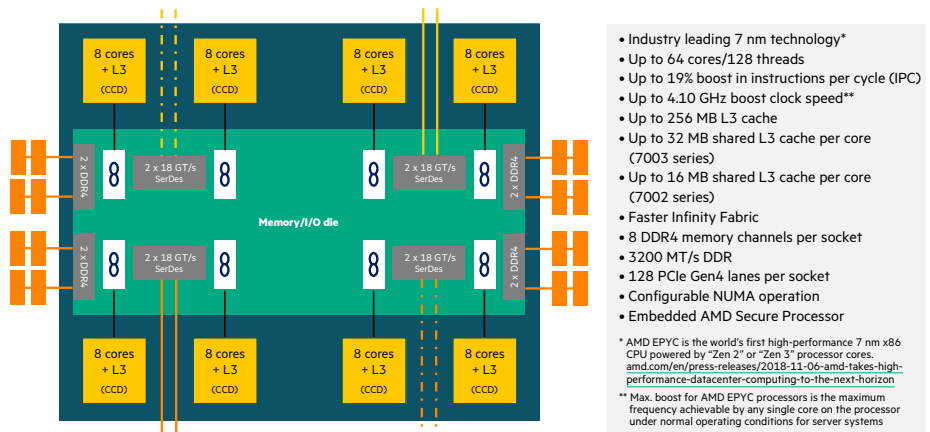
FIGURE 3. AMD EPYC high-level processor design

For EDA workloads, the high-frequency AMD EPYC 7Fx2 (7002 series) and 7xF3 (7003 series) processors shown in Table 2 will be of interest. These parts deliver leadership per-core performance while offering large amounts of L3 cache per core.

TABLE 2. AMD EPYC 7Fx2 and 7xF3 series high-frequency processors

EPYC model	Cores/threads	Base speed	Boost speed[11]	L3 cache	Power (Watts)	L3 cache per core
AMD EPYC 7002 Series						
7F72	24/48	3.20 GHz	Up to 3.70 GHz	192 MB	240	8 MB
7F52	16/32	3.50 GHz	Up to 3.90 GHz	256 MB	240	16 MB
7F32	8/16	3.70 GHz	Up to 3.90 GHz	128 MB	180	16 MB
AMD EPYC 7003 Series						
75F3	32/64	2.95 GHz	Up to 4.0 GHz	256 MB	280	8 MB
74F3	24/48	3.20 GHz	Up to 4.0 GHz	256 MB	240	10.7 MB
73F3	16/32	3.50 GHz	Up to 4.0 GHz	256 MB	240	16 MB
72F3	8/16	3.70 GHz	Up to 4.1 GHz	256 MB	180	32 MB

While performance for EDA applications will depend on the tool and design simulated, industry-standard benchmarks illustrate the advantage of AMD EPYC processors.
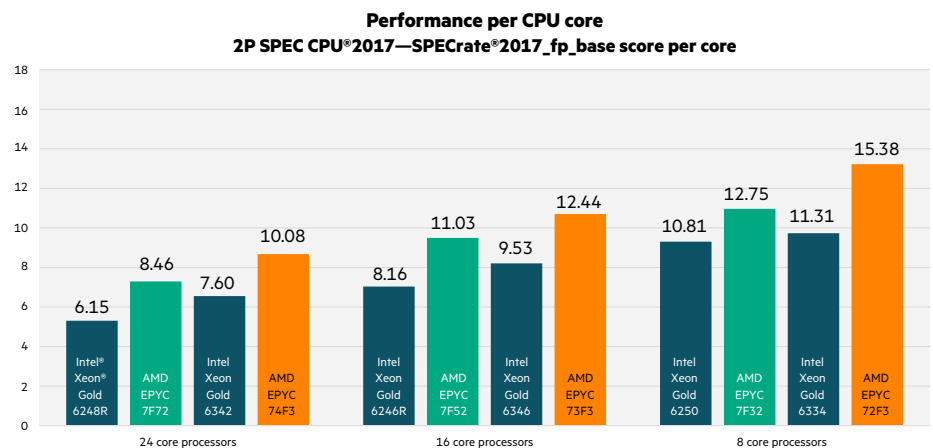
[11] Max. boost for AMD EPYC processors is the maximum frequency achievable by any single core on the processor under normal operating conditions for server systems. AMD EPYC-18

Performance per CPU core
2P SPEC CPU®2017—SPECrate®2017_fp_base score per core

FIGURE 4. EPYC 7Fx2 and 7xF3 series high-frequency parts vs. best-in-class competitors

Figure 4 shows relative SPECrate2017_fp_base scores per core comparing EPYC 7Fx2 and 7xF3 high-frequency parts to alternative processors with the similar core counts on dual-processor systems.[12] EPYC processors' superior performance is a result of high-frequency processors, fast DDR4 memory supporting 3200 MT/s, eight memory channels per processor, and large amounts of L3 cache per core. The green bar in Figure 4 represents high-frequency AMD EPYC 7002 series processors (7xF2) while the orange bar represents the latest EPYC 7003 series processors (7Fx3). In the case of AMD EPYC 7003 series processors, performance is further enhanced by the availability of a single 32 MB cache shared by all cores in a CCD and up to a 19% IPC increase due to enhancements in the "Zen 3" design.

Many EDA workloads are particularly sensitive to memory and L3 cache performance. Table 3A and 3B illustrate the unique advantages of AMD EPYC 7002 and 7003 series processors over comparable competitive offerings across multiple points of comparison.

TABLE 3A. AMD EPYC 7002 series processors provide superior memory capacity, bandwidth, and L3 cache per core

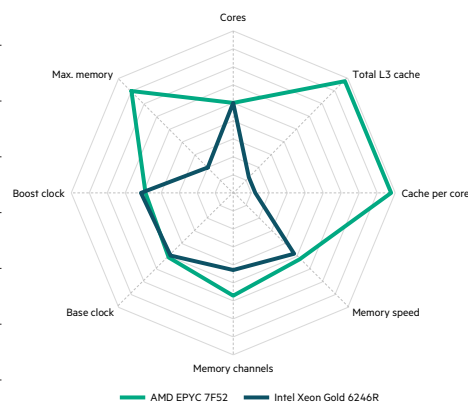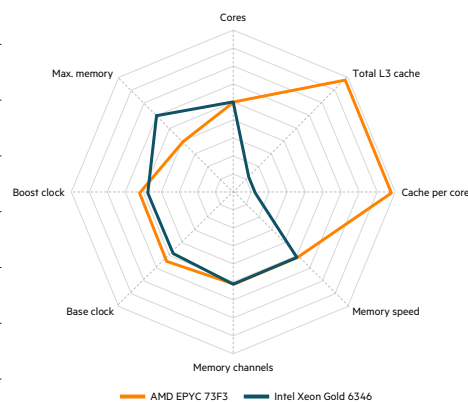|	AMD EPYC 7F52[13]	Intel Xeon Gold 6246R[14]
Number of cores	16	16
Total L3 cache	256 MB	35.75 MB
L3 cache/core	16 MB	2.23 MB
Memory speed	3200 MT/s	2933 MT/s
Memory channels	8	6
Base clock (GHz)	3.50 GHz	3.40 GHz
Boost clock (GHz)[15]	3.90 GHz	4.10 GHz
Max. memory	4 TB	1 TB

AMD EPYC 7F52   Intel Xeon Gold 6246R

TABLE 3B. AMD EPYC 7003 series processors provide superior clock speed, L3 cache, and cache per core

	AMD EPYC 73F3[16]	Intel Xeon Gold 6346[17]
Number of cores	16	16
Total L3 cache	256 MB	36 MB
L3 cache/core	16 MB	2.25 MB
Memory speed	3200 MT/s	3200 MT/s
Memory channels	8	8
Base clock (GHz)	3.50 GHz	3.10 GHz
Boost clock (GHz)[18]	4.00 GHz	3.60 GHz
Max. memory	4 TB	6 TB[19]

AMD EPYC 73F3   Intel Xeon Gold 6346

[12] SPEC, and SPECrate are trademarks of the Standard Performance Evaluation Corporation. All rights reserved. All stated results are as of June 2, 2021. See spec.org for more information. All benchmarks referenced were conducted on 2P systems, so the core counts referenced are across both processors. Configurations as follows: 2P Intel Xeon Gold 6248R (Total 48C) scoring 295 SPECrate2017_fp_base (295/48 = 6.15 score/core)—spec.org/cpu2017/results/res2020q3/cpu2017-20200915-23989.html 2P AMD EPYC 7F72 (Total 48C) scoring 406 SPECrate2017_fp_base (406/48 = 8.46 score/core)—spec.org/cpu2017/results/res2020q2/cpu2017-20200316-21224.html 2P Intel Xeon Gold 6342 (Total 48C) scoring 365 SPECrate2017_fp_base (365/48 = 7.60 score/core)—spec.org/cpu2017/results/res2021q2/cpu2017-20210510-26250.html 2P AMD EPYC 74F3 (Total 48C) scoring 484 SPECrate2017_fp_base (484/48 = 10.08 score/core)—spec.org/cpu2017/results/res2021q2/cpu2017-20210510-25992.html 2P Intel Xeon Gold 6246R (Total 32C) scoring 261 SPECrate2017_fp_base (261/32 = 8.16 score/core)—spec.org/cpu2017/results/res2020q4/cpu2017-20201023-24242.html 2P AMD EPYC 7F52 (Total 32C) scoring 353 SPECrate2017_fp_base (353/32 = 11.03 score/core)—spec.org/cpu2017/results/res2020q2/cpu2017-20200316-21248.html 2P Intel Xeon Gold 6346 (Total 32C) scoring 305 SPECrate2017_fp_base (305/32 = 9.53 score/core)—spec.org/cpu2017/results/res2021q2/cpu2017-20210510-26199.html 2P AMD EPYC 73F3 (Total 32C) scoring 398 SPECrate2017_fp_base (398/32 = 12.44 score/core)—spec.org/cpu2017/results/res2021q2/cpu2017-20210510-26213.html 2P Intel Xeon Gold 6250 (Total 16C) scoring 173 SPECrate2017_fp_base (173/16 = 10.81 score per core)—spec.org/cpu2017/results/res2020q3/cpu2017-20200915-23993.html 2P AMD EPYC 7F32 (Total 16C) scoring 204 SPECrate2017_fp_base (204/16 = 12.75 score per core)—spec.org/cpu2017/results/res2020q2/cpu2017-20200316-21244.html 2P Intel Xeon Gold 6334 (Total 16C) scoring 181 SPECrate2017_fp_base (181/16 = 11.31 score per core)—spec.org/cpu2017/results/res2021q2/cpu2017-20210510-26133.html 2P AMD EPYC 72F3 (Total 16C) scoring 246 SPECrate2017_fp_base (246/16 = 15.38 score per core)—spec.org/cpu2017/results/res2021q2/cpu2017-20210510-26277.html

[13] amd.com/en/products/cpu/amd-epyc-7f52

[14] ark.intel.com/content/www/us/en/ark/products/199353/intel-xeon-gold-6246r-processor-35-75m-cache-3-40-ghz.html

[15], [18] Max. boost for AMD EPYC processors is the maximum frequency achievable by any single core on the processor under normal operating conditions for server systems

[16] amd.com/en/products/cpu/amd-epyc-73f3

[17] ark.intel.com/content/www/us/en/ark/products/212457/intel-xeon-gold-6346-processor-36m-cache-3-10-ghz.html

[19] See Intel Xeon Gold 6346 Processor specs at ark.intel.com/content/www/us/en/ark/products/212457/intel-xeon-gold-6346-processor-36m-cache-3-10-ghz.html Note that 6 TB maximum memory assumes the use of Intel® Optane™ Persistent Memory. With DRAM, maximum memory capacity on the 6346 processor is 4 TB (same as the EPYC 73F3)

THE HPE APOLLO 2000 GEN10 PLUS SYSTEM

The HPE Apollo 2000 Gen10 Plus System is a dense, multiserver platform delivering tremendous performance, throughput, and workload flexibility in a small data center space footprint. Based on industry-leading AMD EPYC processors, HPE Apollo 2000 Gen10 Plus Systems deliver twice the density of traditional rack-mount servers. Each chassis supports up to four dual-processor HPE ProLiant XL225n Gen10 Plus hot-plug servers, each with 2 TB of high performance 3200 MT/s DDR4 memory in just two rack units (2U).

For EDA environments, HPE Apollo 2000 Gen10 Plus Systems provide the ideal blend of features. They offer exceptional simulation performance, expanded power capacity with 3000W power supplies, N+N redundant power, and increased thermal capacity and airflow to reliably support long-running, high-throughput EDA simulations.

HPE Apollo 2000 Gen10 Plus System

- Up to 4 x AMD EPYC-based servers per 2U chassis
- 2 x AMD EPYC 7002 or 7003 series processors per server
- Up to 64 cores and 128 threads per CPU
- 2 TB memory per server (8 TB in 2U)—16 x 128 GB
- Eight memory channels for superior throughput
- Up to 3200 MT/s DDR4 memory
- Up to 32 MB shared L3 cache per core (7003 series)
- Up to 16 MB shared L3 cache per core (7002 series)
- Hot Plug SFF SATA/SA and NVMe storage options
- Comprehensive management tools (APM/RCM)
- Security anchored in HPE iLO 5 and Silicon Root of Trust
- 2 x 3000W power supplies, N+N redundancy
- Enhanced thermal efficiency for HPC workloads
- Optional internal RAID controllers

HPE Apollo 2000 Gen10 Plus System front view with multiple storage options

HPE Apollo 2000 Gen10 Plus System real view with up to 4 hot-pluggable dual-processor servers per chassis for maximum density and flexibility

---

FIGURE 5. HPE Apollo 2000 Gen10 Plus features

With support for the full family of AMD EPYC 7002 and 7003 series processors, EDA IT administrators can configure systems to precisely meet workload demands. Customers can choose high-frequency EPYC 7F52 or 73F3 processors with fewer cores per processor to optimize per-core performance or select high-throughput parts such as the EPYC 7742 or 7763 processors with 64 cores.

Fast I/O is also critical for EDA server farms to ensure that file and network I/O do not emerge as bottlenecks. HPE Apollo 2000 Gen10 Plus Systems offer PCIe Gen4, providing twice the throughput of the previous generation.[20] HPE offers a variety of high-performance PCIe options, including 200 Gbps Mellanox® HPE HDR InfiniBand adapters,[21] multiport 100 Gb Ethernet adapters, and high-performance NVMe SSD drives. Multiple storage options are available inside the chassis ranging from 0 to 24 SFF SAS/SATA hard drives.

[20] PCIe 4.0 delivers 16.0 GT/s, twice the transfer speed of PCIe 3.0— en.wikipedia.org/wiki/PCI_Express

[21] HPE HDR InfiniBand adapters are based on standard Mellanox ConnectX-6 technology

[22] HPE ProLiant XL225n Gen10 Plus achieves 10 records on SPECpower_ssj2008— hpe.com/psnow/doc/a50001386enw

HPE PROLIANT SERVERS

For EDA customers that prefer 1U, single-processor systems, the HPE ProLiant DL325 Gen10 Plus v2 Server is an excellent solution. This server has modest power and cooling requirements and fits easily into most data center environments.

For physical design workloads that require large amounts of memory, either HPE ProLiant DL365 Gen10 Plus or HPE ProLiant DL385 Gen10 Plus v2 Servers are good choices. Both servers support up to 8 TB of memory, critical for memory-intensive EDA applications such as placement and routing.[23]

HPE ProLiant Servers	
HPE ProLiant DL325 Gen10 Plus v2 Server • 1U/1P Server • Single AMD EPYC 7003 series processor • 4 TB memory per server (16 x 256 GB LRDIMM) • Multiple chassis options with up to 10 SFF/4 LFF drives	HPE ProLiant DL325 Gen10 Plus v2 Server front view
HPE ProLiant DL365 Gen10 Plus Server • 1U/2P Server • Single or dual AMD EPYC 7003 series processors • 8 TB memory per server (32 x 256 GB LRDIMM) • 8 SFF drives and optional SFF or NVMe drive bay options	HPE ProLiant DL365 Gen10 Plus Server front view
HPE ProLiant DL385 Gen10 Plus v2 Server • 2U/2P Server • Single or dual AMD EPYC 7003 series processors • 8 TB memory per server (32 x 256 GB LRDIMM) • Multiple chassis options with up to 36 SFF drives	HPE ProLiant DL385 Gen10 Plus v2 Server front view

FIGURE 6. HPE ProLiant Servers for EDA workloads

The HPE ProLiant DL325 Gen10 Plus v2, HPE ProLiant DL365 Gen10 Plus, and HPE ProLiant DL385 Gen10 Plus v2 Servers run the latest 3rd generation AMD EPYC processors. For customers running EPYC 7003 series processors, minimum operating system requirements apply. Supported Linux operating environments include Red Hat Enterprise Linux (RHEL) 8.3, SUSE Linux Enterprise Server (SLES) 12 SP5, and SLES 15 SP2.[24]

Comprehensive Server Security and Management

For security conscience design environments, HPE Apollo and ProLiant Systems provide runtime firmware validation that authenticates critical firmware at start up. Only HPE offers industry-standard servers with firmware anchored into silicon with HPE iLO 5[25] and Silicon Root of Trust. Tied into the Silicon Root of Trust is the AMD Secure Processor, a dedicated security processor embedded in the AMD EPYC system on a chip (SoC).

Customers can also take advantage of optional HPE Apollo Platform Manager (APM), a rack-level power and system management solution for HPE Apollo servers providing an enhanced graphical interface for ease of system management.[26] An optional HPE Apollo 2000 Rack Consolidation Module kit allows HPE iLO aggregation at the chassis level that can be daisy-chained to connect to a top of rack (ToR) management switch.

HPE Performance Cluster Manager (HPCM) is a complete integrated cluster management solution for HPE Apollo systems. HPCM provides system setup, hardware monitoring, and management (aggregating system metrics + remote management from HPE iLO) and cluster health management, image management, and software updates as well as power management.

[23] With the 256 GB LRDIMMs, memory transfer speed is limited to 2933 MT/sec when two DIMMs are installed per memory channel (required to install 8 TB). 3200 MT/sec transfer speeds are supported with one LRDIMM per channel and with smaller DIMM types. hpe.com/psnow/doc/a50000674enw

[24] HPE ProLiant DL385 Gen10 Plus v2 server QuickSpecs—h20195.www2.hpe.com/v2/getdocument.aspx?docname=a50002560enw

[25] HPE iLO is a remote server management processor embedded in the system boards or HPE ProLiant servers providing "lights-out" operation

[28] h20195.www2.hpe.com/v2/Getdocument.aspx?docname=c04111481

HPE Apollo 2000 Gen10 Plus Systems deliver sustained high-performance across multiple cores.

EDA users can reduce regression runtime, maximize license utilization, and reduce TCO by delivering more simulation capacity with a smaller data center footprint.

PERFORMANCE WHERE IT MATTERS

Internal HPE benchmarks conducted using EPYC 7002 and 7003 series processors using an EDA industry register transfer level (RTL) software package shows the relative performance of various HPE dual-processor servers running the same block level model. Figure 7 illustrates the dramatic performance advantage that HPE Apollo System 2000 Gen10 Plus Systems and HPE ProLiant Gen10 Plus Servers bring to EDA environments.[27]

Typically, with RTL simulations and other EDA applications, multiple simulator instances are run simultaneously on the same server, each on a dedicated processor core. Because software licenses are a precious resource, EDA users need to maximize per-core performance to minimize license checkout time. Semiconductor firms look for processors that deliver the best per-core performance while simultaneously supporting the most concurrent simulations to maximize resource utilization.
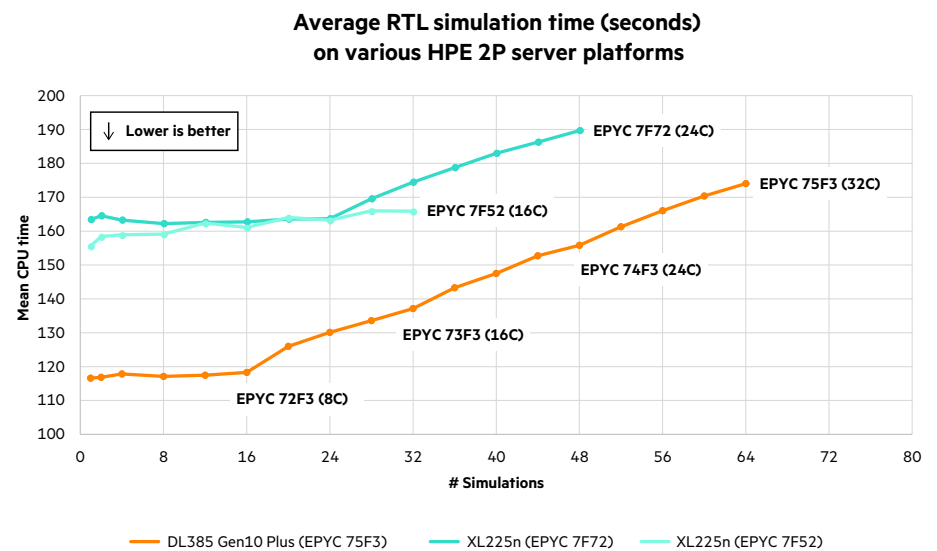
Average RTL simulation time (seconds) on various HPE 2P server platforms

FIGURE 7. HPE servers with AMD EPYC processors offer exceptional throughput and scalability for EDA workloads

For each server platform shown in Figure 7, the average time required to complete the RTL simulation was measured for various numbers of simultaneous simulations. The number of simulations was limited by available physical cores on the underlying processors. For example, the 3rd generation EPYC 75F3 processors has 32 cores, so a maximum of 64 concurrent simulations (2 x 32) could be run on a dual-processor server.

Dramatic performance gains for EDA verification
The results illustrated in Figure 7 are dramatic. Some key results are highlighted below:

• At 16 concurrent simulations, the 3rd generation EPYC 75F3-based server provided **36% better throughput** than the 2nd generation EPYC 7F52 processor.[28]

• With 3rd generation EPYC processors (32 core EPYC 75F3), customers can deploy up to **60% more simulations per server** than 2nd generation processors (16 core EPYC 7F52) without reducing throughput, dramatically reducing infrastructure and facilities costs.[29]

Usually, the more simultaneous simulations run on the same server, the lower the performance. This occurs for various reasons, including sharing of L3 cache, limited memory bandwidth, and NUMA effects as cores access memory on remote processors. The relatively flat line for the EPYC 7F53 in Figure 7 (the orange line) showcases a key advantage of the AMD EPYC architecture as simulations scale from one to sixteen simulations. Because cores in each "Zen 3" complex share a full 32 MB of cache, contention does not begin until eight simulations run on each processor. Even after this, scalability remains excellent, allowing EDA firms to run additional simulations on the same processor with minimal performance degradation.

[27] HPE has conducted internal testing using AMD EPYC 7002 and 7003 series processors as well as alternative processors using an EDA industry register transfer level (RTL) software package. Customers may be briefed on details of HPE's internal testing including processors, software, and detailed server configurations under a non-disclosure agreement.

[28] The 2nd generation EPYC 7F52 on average runs 16 simultaneous simulations in 161 seconds vs. 118 seconds for the 3rd generation EPYC 75F3—a 36% improvement.

[29] The 2nd generation EPYC 7F52 on average runs 32 simultaneous simulations in 166 seconds. The 3rd generation EPYC 75F3 runs 52 simultaneous simulations in just 161 seconds, representing 63% more workload (52 sims/32 sims) per server in approximately the same elapsed time.

Customers looking for the very best verification performance can choose where they want to operate along the curves in Figure 7, balancing absolute performance against density and resource use efficiency.[30] For example, customers running tools where license features are expensive and in short supply might deploy HPE servers with EPYC 72F3 processors and run eight simulations per processor for the best possible performance. Customers that need excellent performance while running more concurrent simulations per server may opt for processors with higher core counts, such as the EPYC 73F3, 74F3, or 75F3.

With HPE Apollo 2000 Gen10 Plus Systems and HPE ProLiant Servers based on AMD EPYC processors, EDA users can:

• Reduce regression runtimes to maximize productivity

• Enable high verification throughput to improve design quality

• Maximize EDA software licenses utilization to minimize cost

• Significantly reduce data center footprint by running more simulations per server

PURPOSE-BUILT FOR EDA WORKLOADS

Whether large or small, silicon design firms are dealing with multiple challenges, including increasing design complexity, time-to-market pressures, and the high cost of engineering talent and software tools. Electronic devices increasingly require more thorough verification as new applications demand higher levels of reliability and safety.

HPE servers powered by AMD EPYC processors provide an important new tool and added flexibility for organizations needing to improve the productivity and efficiency of their chip design environments.

By deploying HPE Apollo Gen10 Plus Systems or HPE ProLiant Servers, customers can:

• Accelerate the design process to meet time-to-market pressures

• Improve product quality and meet more stringent reliability requirements with the capacity to run more simulation and verification workloads within available timeframes

• Maximize value from limited IT budgets by deploying cost-effective, higher throughput systems that deliver improved server farm utilization, more efficient software license utilization, and better engineering productivity

LEARN MORE AT

hpe.com/servers/apollo2000

hpe.com/us/en/servers/proliant-servers.html

To learn more about AMD EPYC 7003 series processors, please visit amd.com/en/processors/epyc-7003-series.

[30] The benchmark results represent by the bottom (orange) line in Figure 7 were achieved using the 3rd generation EPYC 75F3 processor with 32 cores. Other processor SKUs in the same family are labeled in black text along this chart to illustrate that customers can choose different processors and achieve similar performance depending on the number of simulations they wish to run per server. For example, using an EPYC 73F3 (16 core) part would allow for a maximum of 32 simultaneous simulations on a 2P server.

**Make the right purchase decision.
Contact our presales specialists.**

Chat   **Email**   **Call**

Get updates

**Hewlett Packard
Enterprise**