Hewlett Packard Enterprise + intel

Technical white paper

Check if the document is available in the language of your choice.

HARDWARE-SOFTWARE COLLABORATION DRIVES THE FUTURE OF CAE

HPE and Intel help optimize performance for computer-aided engineering workloads, increasing efficiency for manufacturers

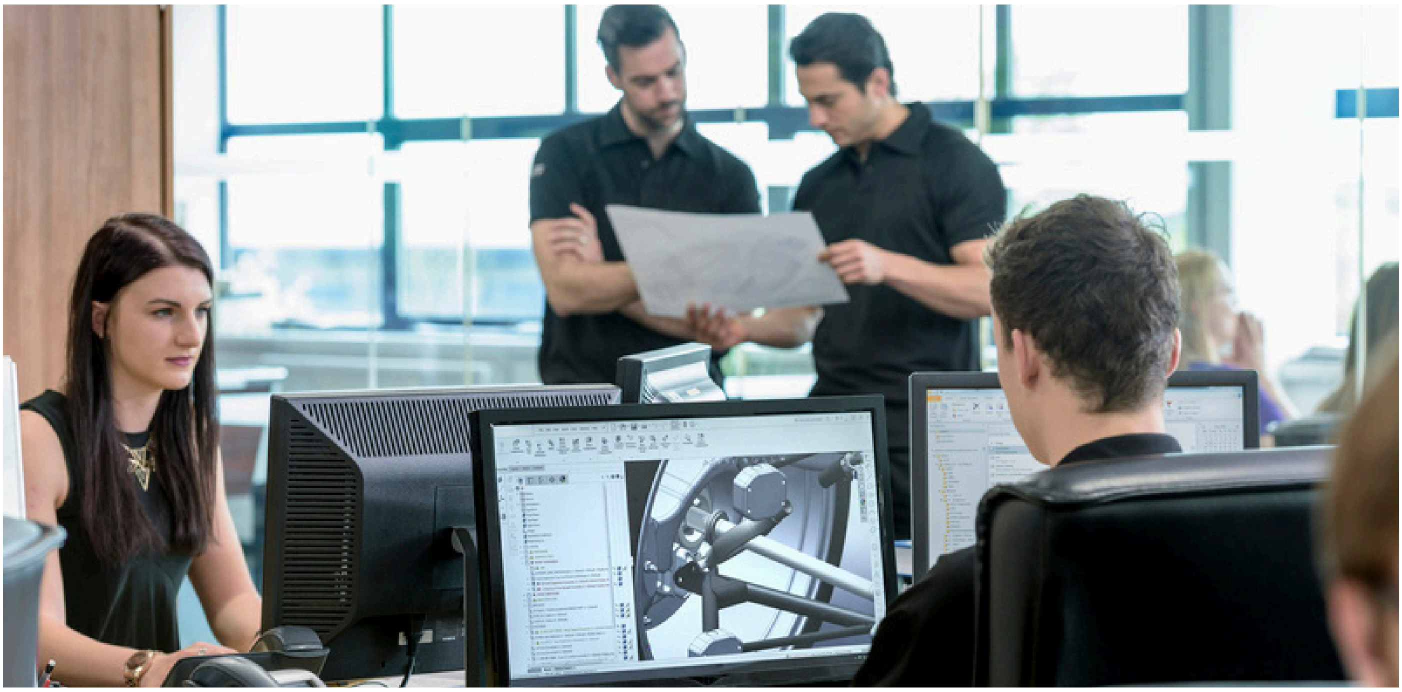

CONTENTS

EXECUTIVE SUMMARY

Today's product design engineers challenge the limits of computing with computer-aided engineering (CAE) applications. By accelerating the process of pinpointing design flaws and producing higher-quality, better-performing products, these applications enable manufacturers to improve designs and reduce time to market.

The key to improving CAE performance and reaching key insights faster is a combination of specific 3rd Generation Intel® Xeon® Scalable processors suited for CAE workloads, the HPE Apollo 2000 Gen10 Plus scale-out server, and close collaboration between Intel®, Hewlett Packard Enterprise (HPE), and CAE application vendors.

This white paper helps manufacturers decide which processor is most appropriate for their CAE needs, provides more information on the HPE Apollo 2000 Gen10 Plus System, and illustrates several examples of the performance gains manufacturers can expect by modernizing their data center infrastructure.

INTRODUCTION

As manufacturers continue their journey toward Industry 4.0, computer-aided engineering (CAE) is becoming an even more crucial element in the design and testing of a variety of products. Experts predict the CAE software and hardware market to reach a market value of USD 8.7 billion by 2026.[1]

Whether for automobiles, aircraft, ocean-going vessels, trains, or the machines and tools used to build them, computers can help engineers more efficiently design, analyze, and manufacture their products. By using high-performance computing (HPC) clusters to run simulations and solve complex equations, CAE helps reduce development time and cost while simultaneously helping improve product quality and, ultimately, customer satisfaction.

CAE is a broad term that encompasses several disparate disciplines. These include computational fluid dynamics (CFD), structural mechanics/finite element analysis (FEA), computational electromagnetics (CEM) analysis, and structural analysis for noise, vibration, and harshness (NVH). But all these CAE use cases require compute, memory, and input/output (I/O) to varying degrees (see Table 1). To reap optimal benefit from CAE applications, manufacturers need a modern HPC infrastructure that can balance performance needs with other considerations such as power consumption and software licensing expenses.

To achieve their CAE goals—well-designed products, reduced time to market, and reduced total cost of ownership (TCO)—manufacturers seek a combination of the following:

- Hardware (processors and servers) that is performant but cost, space, and energy-efficient
- Software that is optimized to take advantage of hardware features

Hardware or software alone cannot drive the necessary CAE performance. Rather, it is the combination of the right hardware and software that leads to optimal CAE. HPE and Intel bring a deep CAE application expertise to the market that can help create a winning HPC environment for CAE.

The strategic alliance between HPE and Intel signifies a long-term commitment to collaborating with independent software vendors (ISVs) to improve their CAE applications and help manufacturers achieve business success. HPE is one of the world's leading CAE platform vendors, while Intel processors are specifically designed for CAE's compute-, memory-, and I/O-intensive workloads. HPE and Intel are working with a variety of ISVs to develop an ecosystem of optimized CAE software. Major CAE applications supported and optimized include Altair RADIOSS and Altair FEKO; Ansys Fluent, Ansys Mechanical, Ansys LS-DYNA and Ansys HFSS; MSC Nastran; Siemens Star-CCM+; OpenFOAM® and SIMULIA Abaqus FEA®; and ESI PAM-CRASH—to name just a few. Decades of community involvement by HPE and Intel leads to a better out-of-the-box experience.

TABLE 1. Computer-aided Design (CAE) Use Cases and Typical Workload Profiles

Discipline	Typical Workload Profile
Computational Fluid Dynamics (CFD)	Memory- and compute-intensive
Structural Modeling/Finite Element Analysis (FEA)	Compute-intensive
Mechanical Engineering/Noise, Vibration, and Harshness (NVH)	Static: Compute- and memory-intensive Dynamic: Memory- and input/output (I/O)-intensive

[1] prnewswire.com/news-releases/global-computer-aided-engineering-cae-market-to-reach-8-7-billion-by-2026--301321641.html

WORKING TOGETHER TO ADDRESS CAE PERFORMANCE CHALLENGES

Through close collaboration and co-innovation, Intel and HPE are helping manufacturing engineers to:

- Reduce time to results by solving larger and more complex problems with greater accuracy.
- Provide more performance in a smaller footprint, which reduces data center floor space requirements and lowers energy costs.
- Minimize system downtime and unplanned outages through improved system reliability.
- Identify performance bottlenecks and reduce costs by performance testing key CAE applications on select Intel CPUs.

Choosing the right processor for the job

Intel's latest processors for HPC workloads are 3rd Generation Intel Xeon Scalable processors. Manufacturers can choose from a variety of SKUs so that the performance is optimized for the right resource constraints. Compared to previous-generation Intel Xeon Scalable processors, the 3rd Gen processors offer improvements across all the critical CAE vectors: compute, memory, and I/O.

Compute:

- Up to 40 cores in a standard socket
- Enhanced per-core performance, including a 20 percent boost in instructions per clock (IPC)
- Wide range of frequency, feature, and power levels
- Intel Speed Select Technology (Intel SST), which provides fine-grain control over CPU performance that can help to optimize TCO
- Built-in HPC and artificial intelligence (AI) acceleration with Intel Advanced Vector Extensions 512 (Intel AVX-512) and Intel Deep Learning Boost (Intel DL Boost)

Memory:

- Increased memory capacity with up to eight channels—up to 6 TB of system memory per processor
- Enhanced memory performance with support for up to 3200 MT/s DIMMs (two DIMMs per channel)
- Increased L1 and L2 cache
- Faster internode connections with three Intel Ultra Path Interconnect links at 11.2 GT/s
- Support for Intel® Optane™ persistent memory 200 series

I/O:

- Support for PCI Express (PCIe) Gen4 and up to 64 lanes (per socket) at 16 GT/s
- Intel Optane solid state drives (SSDs), with consistently high performance and up to 100 drive writes per day (DWPD)
- Support for a wide range of network fabrics

With more than 30 3rd Gen SKUs available, choosing the most appropriate processor for a particular CAE discipline or workload may seem a bit daunting. Manufacturers can use Table 2, along with the subsequent detailed discussions, to guide their processor choices (actual choices will vary by customer need).

TABLE 2. General Processor Recommendations for Computer-aided Engineering (CAE) Disciplines and Workloads

Discipline	Workload Priorities	Typical Software Applications	Recommended SKUs
Computational Fluid Dynamics (CFD)	High memory capacity and bandwidth Moderate frequency needs Highly scalable to hundreds of thousands of cores	Altair AcuSolve Ansys CFX Ansys Fluent Converge OpenFOAM® Siemens Simcenter STAR-CCM+ SIMULIA PowerFLOW	6346 (16c, 3.1/3.6 GHz, 36 MB, 205W) 6336Ya (24c, 2.4 GHz, 36 MB, 185W or 12c, 2.9 GHz, 36 MB, 150W or 8c, 3.1 GHz, 36 MB, 140W) 8358b (32c, 2.6/3.4 GHz, 48 MB, 250W) 8360Y (36c, 2.4/3.1 GHz, 54 MB, 250W or 32c, 2.5/3.2 GHz, 54 MB, 250W or 24c, 2.6/3.3 GHz, 54 MB, 220W)
Structural Modeling/Finite Element Analysis (FEA)	High core counts and frequencies Low memory requirement	Ansys LS-DYNA Explicit Altair RADIOSS SIMULIA Abaqus/Explicit ESI PAM-CRASH	8352Y (32c, 2.2 GHz, 48 MB, 205W or 24c, 2.3 GHz, 48 MB, 185W or 16c, 2.7 GHz, 48 MB, 185W) 8358 (32c, 2.6/3.4 GHz, 48 MB, 250W) 8362 (32c, 2.8/3.6 GHz, 48 MB, 265W) 8360Y (36c, 2.4/3.1 GHz, 54 MB, 250W or 32c, 2.5/3.2 GHz, 54 MB, 250W or 24c, 2.6/3.3 GHz, 54 MB, 220W)
Mechanical Engineering/Noise, Vibration, and Harshness (NVH)	High memory capacity Medium to high frequencies and input/output (I/O) Negligible benefits from higher cores	Altair OptiStruct Ansys Mechanical CDH/AMLS SIMULIA Abaqus/Standard MSC Nastran Siemens Simcenter Nastran	6334 (8c, 3.6/3.7 GHz, 18 MB, 165W) 6346c (16c, 3.1/3.6 GHz, 36 MB, 205W) 6342c (24c, 2.8/3.5 GHz, 36 MB, 230W)

NOTES

a "Y" SKUs support Intel Speed Select Technology–Performance Profile 2.0 (Intel SST-PP).

b Especially suited for Fluent applications.

c Especially suited for AMLS and Abaqus/Standard applications.

Computational Fluid Dynamics (CFD)

TABLE 3. CFD Workload Characteristics

Performance Factor	Importance
Memory capacity	High
Memory bandwidth	High
Fabric latency	High
Core count	Medium
Frequency	Medium

CFD applications are used for simulating the flow of air, fluids, heat, and viscous material, and have many uses in a wide variety of industries—including aeronautic science, drag simulation in car shape design, and jet and thermal flow in engine design. Workloads are typically analyses of complex, often unstructured meshes with tens to hundreds of millions of cells. The data usually requires preprocessing, which can affect runtimes and the quality of results. CFD applications, either open source or per-core-licensed applications, generally run on medium-to-large clusters with 16-512 nodes.

To achieve optimum performance, CFD applications need plenty of memory capacity and memory bandwidth (see Table 3). In fact, some codes use close to the maximum available memory bandwidth more than 90 percent of the time. On the other hand, CFD workloads are not so dependent on high core frequency, although some codes may benefit from higher clock speeds and sustained Turbo performance, and others may benefit from a larger number of cores. For CFD, I/O performance is generally not a critical performance factor.

Four 3rd Gen Intel Xeon Scalable processor SKUs are well-suited for running CFD workloads (see Table 2 for number of cores, frequencies, and other details):

- **6346.** This processor is an **excellent entry-level CAE SKU** for performance per core, with strong core density and performance.
- **6336Y.** This processor also has strong core density and performance and is a good choice for a general-purpose CAE cluster. It offers **configuration flexibility** with three performance profiles, which enables manufacturers to choose a configuration with low thermal design power (TDP) for use in power-constrained environments.
- **8358.** This processor is a **good choice for codes that benefit from higher core counts**, such as OpenFOAM® and Ansys Fluent.
- **8360Y.** Like the 6336Y, this processor offers three performance profiles and is a good choice for a general-purpose CAE cluster. With more cores than the 6336Y, the 8360Y can be **used for codes that demand more compute**.

Finite Element Analysis (FEA) for Structural Mechanics and Crash Analysis

TABLE 4. FEA Workload Characteristics

Performance Factor	Importance
Frequency	High
Core count	High
Fabric latency	High
Memory bandwidth	Medium

Structural analysis concerns stress analysis on components and assemblies. Structural applications generally use FEA in two main numerical simulation approaches:

- **Implicit FEA**—Implicit analysis is used for longer duration, relatively static problems in which time dependency of the solution is not an important factor, such as analysis of forces on structures. Implicit analysis usually uses much longer time steps but can require more computational resources as the problem size increases.
- **Explicit FEA**—Explicit analysis is used for high-impact and short-duration simulations where each step takes into account forces like mass and inertia from the previous step. Examples include crash, impact, and blast simulations. Such “nonlinear” events are modelled to predict cascading damage to structural and component integrity. Major users of crash testing include automotive manufacturers, who use it to save money by reducing the need for real-world crash tests.

Workloads typically involve nonlinear FEA problems with large, high-fidelity models (1–15 million elements) and complex, dynamic load cases.

FEA applications are primarily per-core-licensed commercial codes, typically running on large clusters with datasets ranging from 10–20 GB. Depending on the industry, the number of cores varies from a few hundred cores per job (for example, in the auto industry) to thousands of cores per job (for example, in aerospace and large labs). Because repeatability and consistency of results are paramount, FEA workloads generally run on a fixed number of cores per job.

Unlike CFD workloads, FEA workloads don’t necessarily need high memory capacity, and memory bandwidth is of less importance as well. Instead, FEA workloads (see Table 4) require high clock speeds (3.4 GHz or more) and a core count (at least 32 cores on a two-socket system). These workloads are quite sensitive to network latency and fabric latency, but are not I/O-constrained.

Four 3rd Gen Intel Xeon Scalable processor SKUs are well-suited for running FEA workloads (see Table 2 for number of cores, frequencies, and other details):

- **8352Y and 8360Y.** With 32 cores and 36 cores respectively, these processors are tuned for a general-purpose CAE cluster but have the higher number of cores required by FEA workloads. They offer **configuration flexibility** with three performance profiles each, which enables manufacturers to choose a configuration with a TDP that suits their environment.
- **8358 and 8362.** Also offering 32 cores but different frequencies, these processors provide **good performance, especially for crash analysis**, such as LS-Dyna jobs.

Structural Analysis / Noise, Vibration, and Harshness (NVH)

TABLE 5. NVH Workload Characteristics

Performance Factor	Importance
Frequency	Medium, High
Core count	Medium
Memory capacity	High
Memory bandwidth	Medium, High
I/O performance	Medium, High

NVH applications simulate acoustics to pinpoint indicators associated with quality such as squeak and rattle, vibration issues and external and internal noise levels heard within a vehicle. NVH solvers can be static or dynamic; the latter models involve millions of finite elements (1–40 million degrees of freedom) with thousands of components and properties.

All NVH workloads can benefit from large memory capacity, but higher core counts provide negligible benefits (see Table 5). In fact, codes such as Ansys Mechanical should use lower core counts. The other performance factors are harder to generalize for NVH than for other CAE disciplines because they have varying performance requirements.

Typical systems range from two to 380+ nodes. In certain cases, jobs are run on a single node because traditional dynamic NVH solvers do not scale well over many cores. However, new scalable parallel solvers are being developed, such as the MUMPS (MUltifrontal Massively Parallel Sparse) solver.

The following 3rd Gen Intel Xeon Scalable processor SKUs are well-suited for running NVH workloads (see Table 2 for number of cores, frequencies, and other details):

- **6334.** This processor is an excellent entry-level CAE SKU for performance per core.
- **6346 and 6342.** Depending on a manufacturer's core density and core frequency requirements, either of these processors is a good choice when running AMLS or Abaqus Standard. The 6342 is best suited for Ansys Mechanical.

Deploying a scale-out platform for Computer-Aided Engineering

The previous discussion centered on choosing the appropriate processor for a specific type of CAE workload—which is important. But it is only part of the picture. A best-fit processor alone cannot solve all the challenges facing CAE data centers, such as keeping server footprint to a minimum, lowering TCO, and reducing the cost to the environment (TCE) by shrinking energy consumption.

The HPE Apollo 2000 Gen10 Plus System is a dense, multi-server platform that packs substantial performance and workload flexibility into a small data center space, while delivering the efficiencies of a shared infrastructure. It is designed to provide a path to scale-out architecture for traditional data centers, so enterprise customers can achieve the space-saving value of density-optimized infrastructure in a cost-effective and nondisruptive manner. This system offers twice the density of traditional 1U rack mount systems, helping data centers to maximize the use of valuable data center space.

The HPE Apollo 2000 Gen10 Plus System offers several unique benefits, making it a good choice for data center modernization as CAE workloads climb:

- **Plug-and-play.** Designed with standard form factor racks, cabling, and serviceability to plug and play in traditional data centers with little or no disruption to the infrastructure or operational practices.
- **Shared infrastructure.** Designed with shared power and cooling resources to drive higher levels of efficiency, compute density, and system scaling.
- **Highly secure.** HPE iLO5 provides end-to-end server security that starts with a silicon root of trust from HPE, includes daily firmware checks and alerts, the ability to recover the firmware to the last known good state if code is compromised and a high level of security and cryptography using commercial national security algorithms. The silicon root of trust from HPE starts protecting HPE servers early in the production process and all the way through the HPE Secure Compute Lifecycle that covers secure supply chain, secure runtime operations, data security in the server, regulatory compliance, all the way to safe disposal of data and hardware at end-of-life.
- **Flexible configuration options.** Trays are deployed in a cost-effective 2U chassis with configurations supporting up to four trays per 2U as well as offering a flexible backplane for multiple storage configurations. More servers can be easily added as CAE workloads expand. Direct Liquid Cooled options are available, enabling the use of dense memory configurations and the highest-bin Intel graphics processing units (GPUs) and Intel processors. Chassis, or groups of chassis, can be custom-configured to become affordable modular building blocks for customer-specific implementation at scale, and for future growth.

- **Storage and I/O flexibility.** Customers can optimize for performance or economy, to get the right compute for the right workload. Multiple storage chassis options are available, ranging from zero drives to 24 SFF drives with zoning. Multiple riser options provide the capability to have low-latency, high-bandwidth workloads.
- **Comprehensive manageability.** The platform comes with an extensive set of tools for node to rack management, such as the HPE Performance Cluster Manager (HPCM)—a fully integrated system management software offering administrators all the functionalities they need to manage their CAE clusters. The software provides system setup, hardware monitoring, and management (aggregating HPE iLO information), health management, image management, and software updates as well as power management for systems of any scale—all the way up to Exascale-level deployments. The HPCM reduces the time and resources spent administering HPC and CAE systems, lowering TCO, increasing productivity and providing a better return on hardware investments.

Manufacturers can choose to populate the HPE Apollo 2000 Gen10 Plus System with independently serviceable, 2U half-width, two-processor HPE ProLiant XL220n Gen10 Servers and/or HPE ProLiant XL290n Gen10 Servers:

- Latest generation Intel Xeon Scalable processors
- Eight channels per socket for a total of 16 memory channels and DIMM slots per server. With 128 GB DIMM support, each server can get up to 2 TB 3200 MT/s DDR4 SmartMemory)
- 2x PCIe Gen4 slots available for additional I/O and 1OCP 3.0 slot available
- Flexible storage options and support for 0–24 small form factor (SFF) SAS/SATA SSDs including zoning and up to six NVMe SSDs per node
- Two I/O slots for a choice of networking and clustering options including: 1/10/40 Gb Ethernet, 100 Gb/s EDR InfiniBand or Intel OPA, Fibre Channel and options for either one PCIe slot plus a FlexibleLOM or 2 PCIe slots
- Ability to accommodate up to eight GPUs per chassis or 4 per tray (HPE ProLiant XL290n Gen10 Server only)

Multiple chassis configuration options define the storage for the servers in the chassis:

- HPE Apollo n2600 Chassis—choose from several options: zero drives; 8 SFF hot-plug SAS/SATA/NVMe drives allocated equally across server nodes; 24 SFF hot-plug SAS/SATA drives only; or 24 SAS/SATA/NVMe drives allocated equally across server nodes.
- HPE Apollo n2800 Chassis—24 SFF hot-plug SAS Zoning—Supports flexible drive mapping, enabling custom drive allocations to match workloads and provide flexible storage density for various applications; an individual server can have 0–24 drives allocated.
- The chassis also have options for two 1600W or 3000W Platinum Power Supplies, with an N or N+N redundancy option.

HPE Services: Get the Most Out of Your Hardware Investment

HPE Pointnext Services help lower risks and costs by using proven best practices, automation, and methodologies that have been tested and refined by HPE experts through thousands of deployments globally. Advisory Services focus on business outcomes and goals, collaborating to design transformation and build a road map tuned to unique business challenges. Professional and Operational Services can be used to speed up time-to-production, boost performance, and accelerate business. Choose from the following services:

HPE Pointnext Tech Care goes beyond traditional support by providing faster access to product specific experts, an artificial intelligence (AI) driven digital experience, and general technical guidance to not only reduce risk but constantly search for ways to do things better.

HPE Pointnext Complete Care helps customers manage complex edge-to-cloud IT environments. This service enables customers to focus on data, critical workloads, and business outcomes. Customers can take advantage of these new capabilities standalone, or as part of the cloud experience we offer with the HPE GreenLake platform.

Other related services include Parts and Materials, Defective Media HPE Server Hardware Installation, HPE Installation and Startup, HPE Service Credits and HPE Education Services, HPE Factory Express and HPE Financial Services.

RESULTS: RUN MORE JOBS PER DAY TO IMPROVE DESIGN AND SPEED TIME TO MARKET[2]

Collaboration between HPE, Intel, and ISVs is producing impressive performance improvements across a wide variety of CAE workloads. These performance improvements enable product design engineers to run more jobs and iterations per day and use higher fidelity models. The results showcased here are industry-standard benchmarks that represent real-world workloads. The performance improvements illustrate the power of hardware and software vendors working together to provide the best possible outcome for customers. Intel, HPE, and ISV engineers pooled their expertise to optimize the various software applications to take advantage of Intel architecture. In particular, ISV developers used Intel oneAPI's open, unified programming model with Intel oneAPI Toolkit for HPC (built on familiar, proven CPU tools), to more easily optimize CAE codes for HPC environments.

Summary of testing and findings

We tested a variety of CAE benchmarks on a four-node cluster, comparing the performance of the 16-core, 3.4 GHz 2nd Gen Intel Xeon Scalable processor (35.75 MB L3 cache) and the 18-core, 3.0 GHz 3rd Gen Intel Xeon Scalable processor (39 MB L3 cache). We used the test results to determine how the new processor's improved memory, compute, and I/O features could benefit a range of CAE workloads. The benchmarks we used were Ansys standard benchmark models, not generic Message-Passing Interface (MPI) benchmarks. Working with Intel, the individual ISVs (such as Ansys and Dassault Systèmes SIMULIA Corp.) have tuned their applications to take optimal advantage of the Intel architecture. For each benchmark, we ran four iterations, varying the number of MPI tasks from 1x to 4x the number of cores.

Overall, the benchmark tests demonstrated that the majority of CAE workloads—CFD, FEA, or NVH—can support more jobs per day and reduced walltimes using the latest Intel Xeon Scalable processor. These results are directly related to the newer processor's additional cores, its IPC improvements, and its additional memory capacity and bandwidth. Running more jobs per day with lower walltime translates directly to accelerated design and testing, which in turn can lead to product improvements and faster time to market.

Ansys Fluent (computational fluid dynamics)

Fluent uses a cell-centered code capable of handling polyhedral mesh and cutcell meshes. It offers both pressure-based and density-based options. It is generally used for combustion, multiphase or chemically-reacting flows.

We ran over a dozen separate industry-standard benchmarks that ranged from a gasoline direct injection model to a landing gear analysis to external flows over an aircraft wing and various types of cars. Figures 1 and 2 show a sampling of the results, illustrating a 13–18 percent improvement in jobs per day, with consistently less walltime, using the latest Intel processor. The test also used Fluent's "`-platform=intel`" parameter, which provides optimizations specific to Intel architecture. With the 2021R2 version of Fluent, released in July 2021, Ansys and Intel worked together to further improve these optimizations.

[2] See appendix for configuration details.

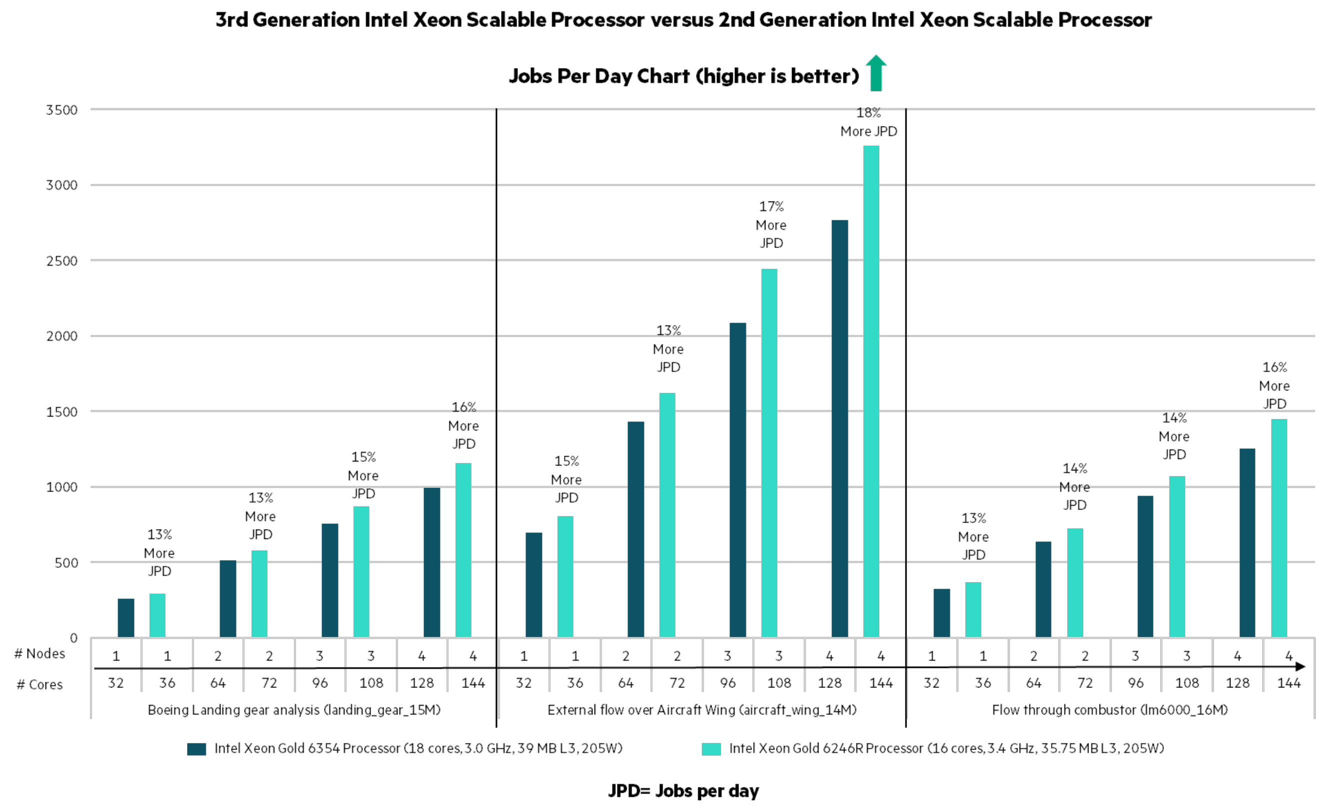

FIGURE 1. Ansys Fluent jobs-per-day results

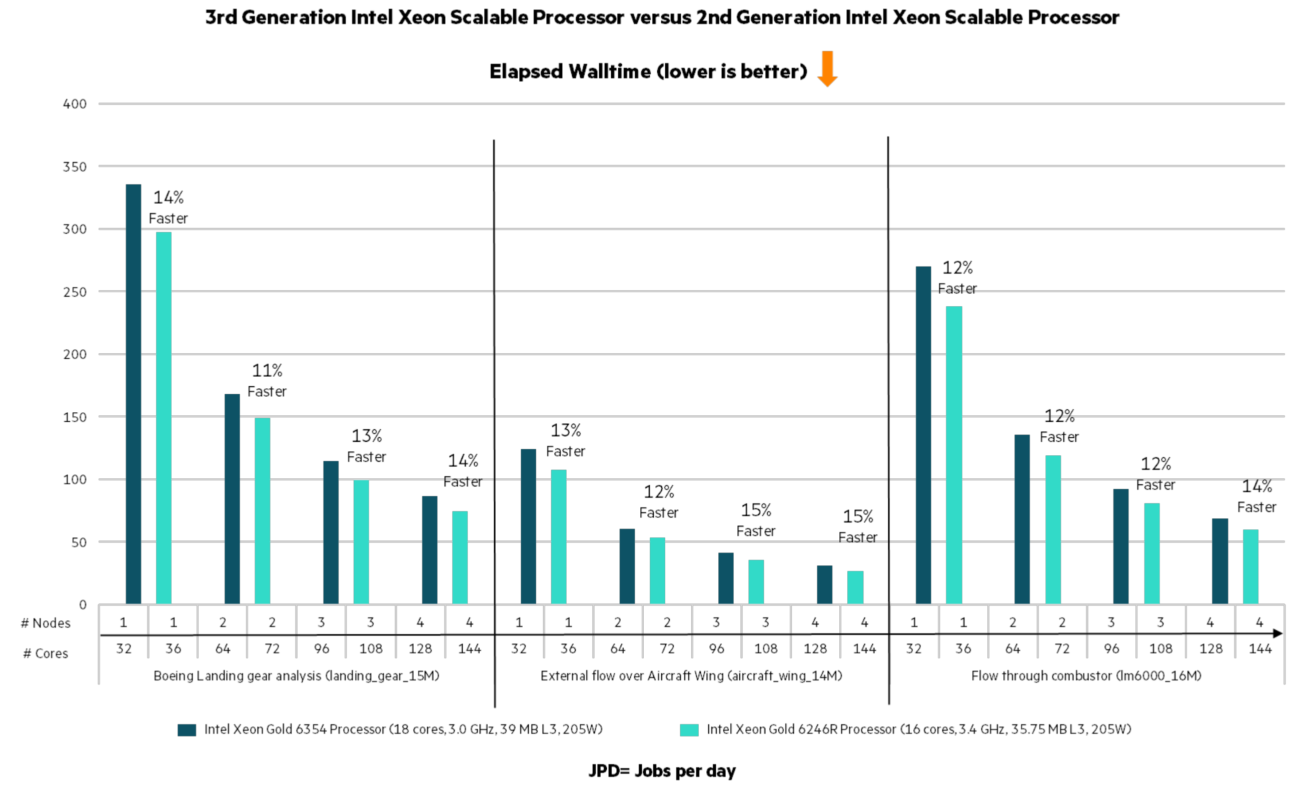

FIGURE 2. Ansys Fluent walltime benchmark results

Ansys CFX (computational fluid dynamics)

CFX is a fully implicit solver that requires quite a bit of storage. It uses a cell-vertex code, is pressure-based, and handles traditional tetra and hexa mesh topologies. CFX is often used in turbomachinery analyses.

We ran benchmarks for a LeMans race car, a pump, and three air foils. All benchmarks ran better on the 3rd Gen Intel Xeon Scalable processor. Figures 3 and 4 show a sampling of the results, illustrating jobs-per-day improvements ranging from 19 percent to 29 percent and walltimes of up to 24 percent faster.

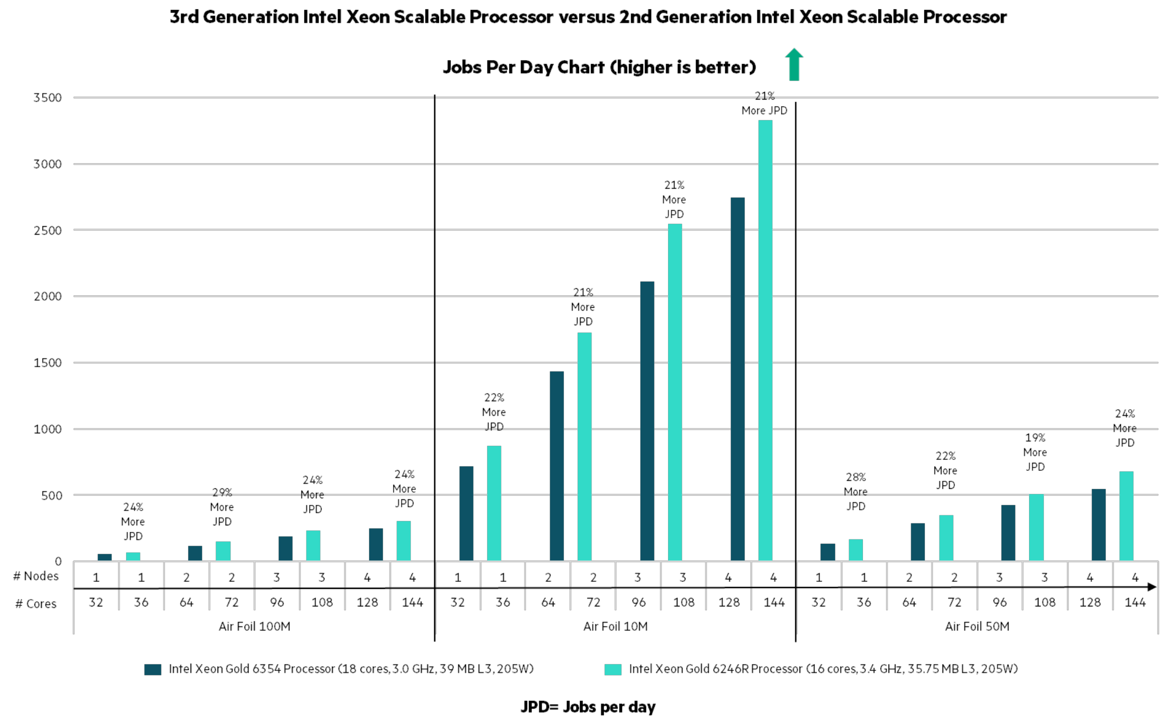

FIGURE 3. Ansys CFX jobs per day benchmark results

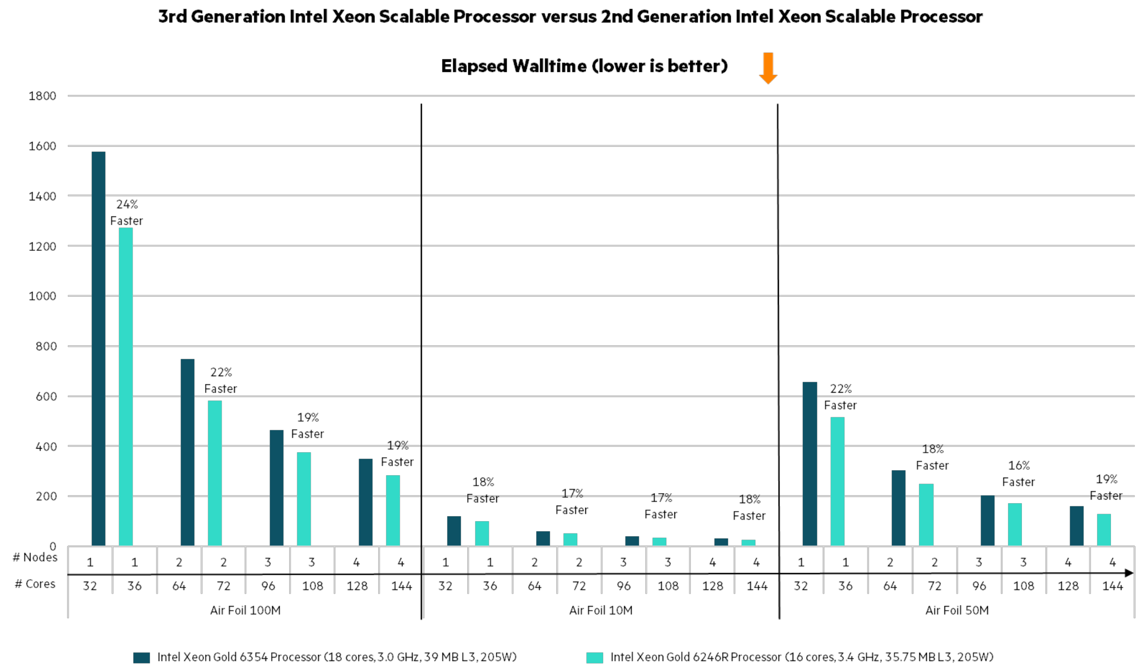

FIGURE 4. Ansys CFX walltime benchmark results

Ansys LS-DYNA (finite element analysis)

LS-DYNA is a popular multiphysics simulation application. This explicit simulation software is used for drop tests, impact and penetration, smashes and crashes, occupant safety, and more.

We ran three crash simulations. The benchmarks all showed marked improvement in jobs per day on the 3rd Gen Intel Xeon Scalable processor—up to 41 percent. Walltimes were also consistently lower. Figures 5 and 6 show a sampling of the results.

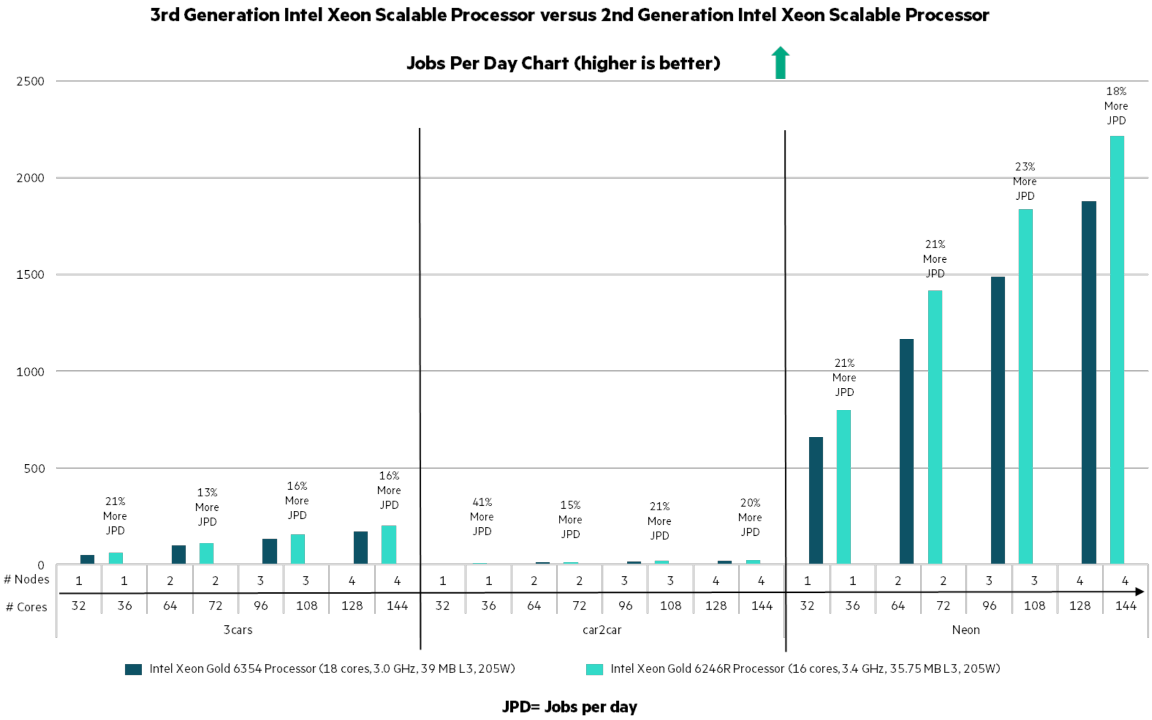

FIGURE 5. Ansys LS-DYNA jobs per day benchmark results

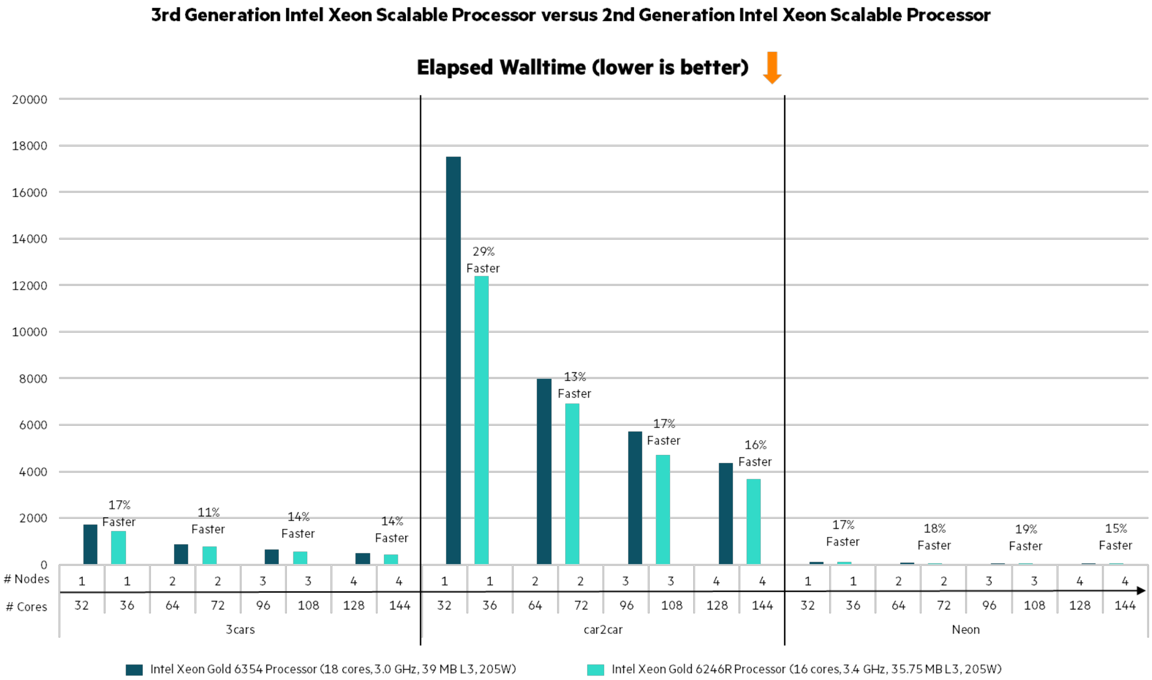

FIGURE 6. Ansys LS-DYNA walltime benchmark results

SIMULIA Abaqus/Explicit (finite element analysis)

Abaqus/Explicit is primarily used to analyze brief, high-speed, transient dynamic events. Such events include dropping a device, auto crashes and ballistic impact. It is particularly effective at analyzing severely nonlinear behavior.

We ran three benchmarks: airbag deployment, a thick plate under uniform pressure, and a car crash. Figures 7 and 8 show a sampling of the results, demonstrating significant performance improvement on the 3rd Gen Intel Xeon Scalable processor—up to 73 percent more jobs per day as well as much lower walltimes.

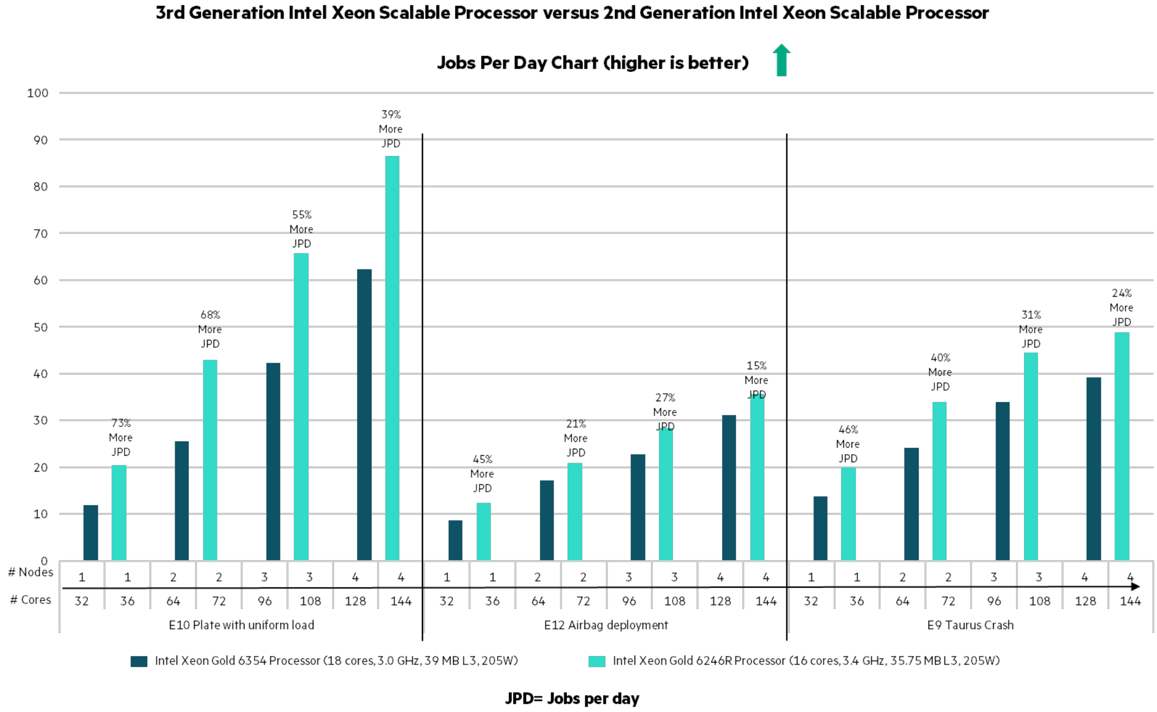

FIGURE 7. SIMULIA Abaqus/Explicit jobs per day benchmark results

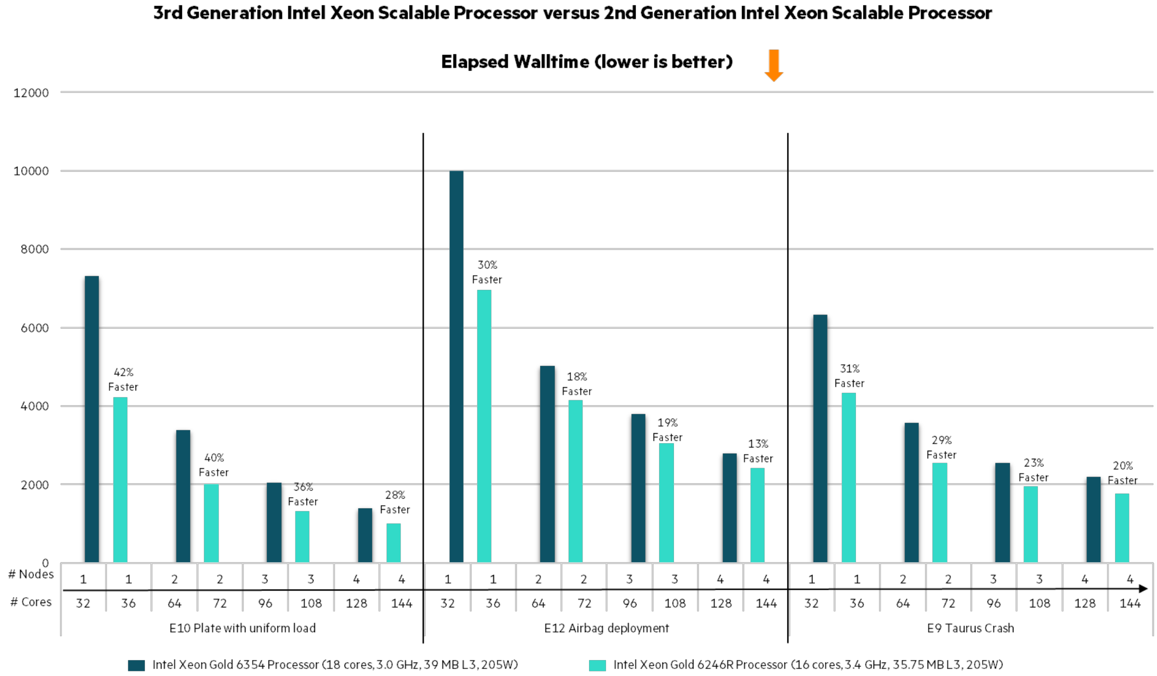

FIGURE 8. SIMULIA Abaqus/Explicit walltime benchmark results

SIMULIA Abaqus/Standard (mechanical engineering/NVH)

Abaqus is an implicit solver primarily used for static (steady-state) and low-speed dynamic events where highly accurate stress solutions are critically important. Examples include analyzing how well a gasket cover performs its job, or studying how well a tire rolls at a steady speed. Abaqus/Standard results can be fed into Abaqus/Explicit for further analysis. (The reverse is also true; Abaqus/Explicit results can be transferred to Abaqus/Standard.)

We ran seven benchmarks. 3rd Gen Intel Xeon Scalable processors enabled up to 53 percent more jobs per day with walltimes reduced by a similar amount. Figures 9 and 10 show a sampling of the results.

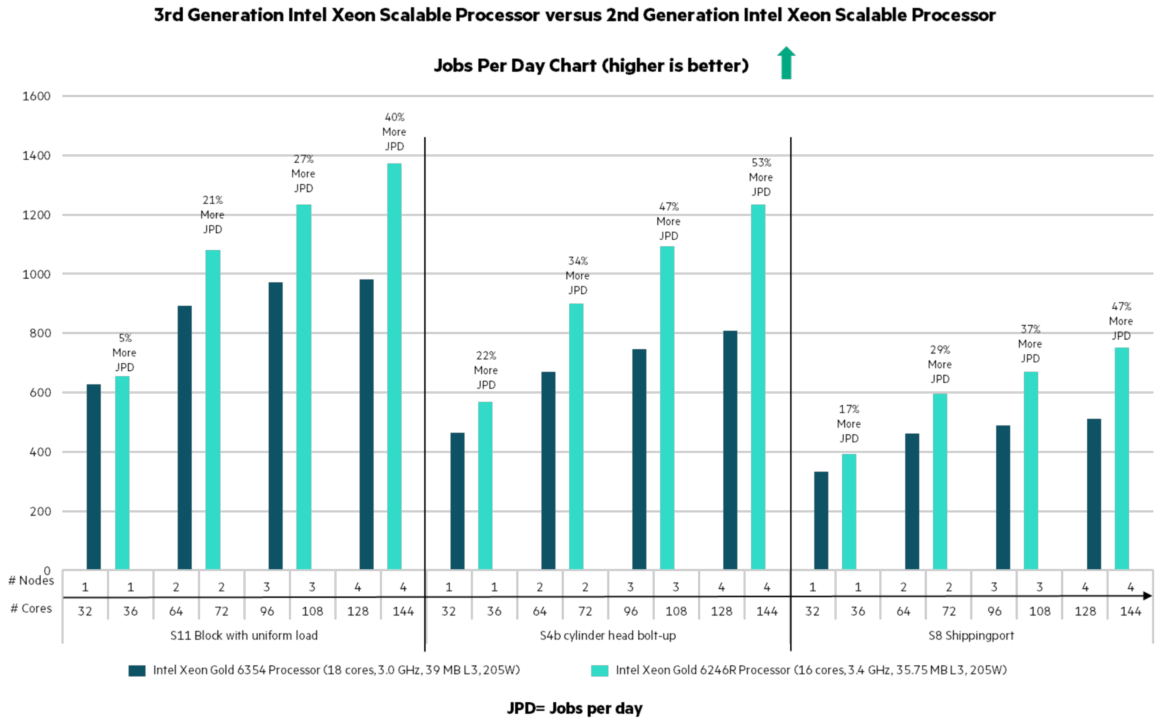

FIGURE 9. SIMULIA Abaqus/Standard jobs per day benchmark results

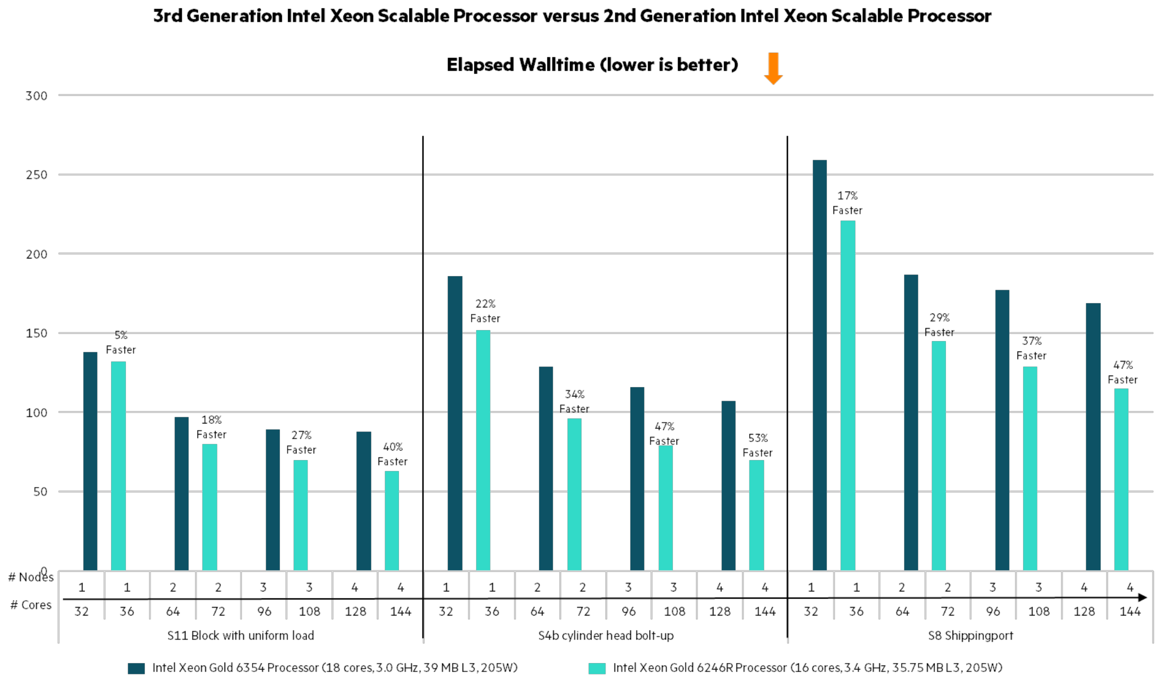

FIGURE 10. SIMULIA Abaqus/Standard walltime benchmark results

Intel AVX-512 Boosts Computer-aided Engineering (CAE) Performance

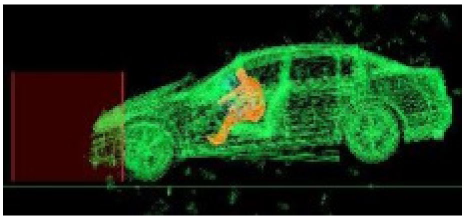

Vector mathematics is an integral aspect of computer-aided engineering (CAE). Millions of design variables in a particular model or solver are represented by mathematical vectors. These vectors are manipulated (added and multiplied) and related to each other (angles between vectors, for example) to create images such as the one here.

Intel Advanced Vector Extensions 512 (Intel AVX-512) instructions represent a significant leap to 512-bit single instruction, multiple data (SIMD) support. Programs can pack eight double-precision or 16 single-precision floating-point numbers, or eight 64-bit integers, or sixteen 32-bit integers, within the 512-bit vectors—per clock cycle. With up to two 512-bit fused-multiply add (FMA) units, Intel AVX-512 enables processing of twice the number of data elements that Intel AVX/Intel AVX2 can process with a single instruction and four times that of Streaming SIMD Extensions (SSE).[3]

Intel AVX-512 instructions are important because they offer higher performance for the most demanding computational tasks, like computational fluid dynamics (CFD) and finite element analysis (FEA). Intel AVX-512 instructions offer a high level of richness in the design of the instructions. Intel AVX-512 features include

- 32 vector registers (often referred to as ZMM registers) each 512 bits wide
- Eight dedicated mask registers
- 512-bit operations on packed floating-point data or packed integer data
- Embedded rounding controls (override global settings)
- Embedded broadcast
- Embedded floating-point fault suppression
- Embedded memory fault suppression
- New operations
- Additional gather/scatter support
- High-speed math instructions
- Compact representation of large displacement value
- Ability to have optional capabilities beyond the foundational capabilities

It is interesting to note that the 32 ZMM registers represent **2KB of register space**.

With the need for greater computing performance continues to grow across all manufacturing segments, Intel is working with independent software vendors (ISVs) to optimize their code to take advantage of Intel AVX-512 features. These helps the ISVs the ability to help manufacturers accelerate performance for their CAE workloads.

[3] intel.com/content/www/us/en/developer/articles/technical/intel-avx-512-instructions.html

CONCLUSION

HPE, Intel, and CAE application ISVs are collaborating to provide an optimal platform for product design. Memory enhancements in the latest generation of Intel Xeon Scalable processors, along with acceleration technology like Intel AVX-512, provide the raw horsepower for resource-hungry CFD, FEA, and NVH workloads. ISV and Intel engineers work closely to optimize code so that it can take advantage of Intel architecture. And by deploying their CAE applications on HPE's robust, scale-out HPE Apollo 2000 Gen10 Plus system, manufacturers can minimize server footprint, lower overall costs, and reduce energy consumption.

Industry-standard benchmarks show that this workload-optimized combination of hardware and software can substantially improve CAE workload performance, as measured by number of jobs per day and elapsed walltime. These performance improvements can be a significant differentiator for manufacturers, because they can bring higher quality products to market more quickly.

LEARN MORE

- HPE Apollo 2000 Gen10 Plus System
- HPE High Performance Computing for Manufacturing brochure
- Manufacturers Tackle Challenging Modeling and Simulation Problems with HPC brochure
- Accelerate Results and Enhance Product Quality with High-Performance Computing eGuide
- 3rd Generation Intel Xeon Scalable processors
- Intel AVX-512
- Intel oneAPI

To discuss your next CAE cluster's needs and to review the full set of CAE benchmark results, contact your HPE representative.

APPENDIX: BENCHMARK SYSTEM INFORMATION

Table A1 provides the baseline (Intel® Xeon® Gold 6246R processor) and updated (Intel Xeon Gold 6354 processor) configurations for the test results included in this white paper. Table A2 provides the details about which benchmark versions were used in the testing.

TABLE A1. Hardware and Firmware Configurations

Component	Baseline Configuration (2nd Generation Intel Xeon Scalable processor)	Updated Configuration (3rd Generation Intel Xeon Scalable processor)
Date of testing	April, 2020	September, 2021
Who performed testing	HPE	HPE
Number of Nodes	4	4
CPU	2x Intel Xeon Gold 6246R processor (16 cores, 3.4 GHz/4.1 GHz	2x Intel Xeon Gold 6354 processor (18 cores, 3.0 GHz/3.6 GHz
Base Frequency	3.4 GHz	3.0 GHz
Boost Frequency	4.1 GHz	3.6 GHz
Cores per Node	16 x 2 = 32	18 x 2 = 36
L3/Last Level Cache	35.75 MB	39 MB
L3 Cache per Core	2.23 MB	2.16 MB
Memory Channels	6	8
Memory per Node	192 GB	512 GB
System memory total GB (slots/GB/speed)	192 GB (12x/16 GB/2933 MHz)	512 GB (16x/32 GB/3200 MHz)
Boot Drive	Diskless	Diskless
Storage	Lustre Parallel File System	Lustre Parallel File System
BIOS	Version: SAED1229 Release Date: 01/22/2019 BIOS Revision: 5:14	Version: U47 Release Date: 08/18/2021 BIOS Revision: 1.52
Thermal Design Power (TDP)	205W	205W
Interconnect	InfiniBand HDR	InfiniBand HDR
Intel Hyper-Threading Technology (Intel HT Technology) on/off	Off	Off
Intel Turbo Boost Technology on/off	On	On
Intel Volume Management Device (Intel VMD) enabled/disabled	Disabled	Disabled
OS and version / kernel	Red Hat Enterprise Linux® (RHEL) 7.6	RHEL 7.9
Other relevant information such as drivers, settings, etc.	HPC Workload Profile	HPC Workload Profile

TABLE A2. Software Configurations

Software Application and Ansys Benchmarks	Version Used with 2nd Generation Intel Xeon Scalable processor Testing	Version Used with 3rd Generation Intel Xeon Scalable processor Testing
Ansys Fluent • **Boeing Landing Gear Analysis (landing_gear_15M)** • **External Flow Over Aircraft Wing (aircraft_wing_14M)** • **Flow Through Combustor (lm6000_16M)**	2020 r1	2021 r2
Ansys CFX • **Air Foil 100M** • **Air Foil 10M** • **Air Foil 50M**	2020 r1	2021 r2
Ansys LS-DYNA • **3cars** • **car2car** • **Neon**	R9_3_1 avx2	R11_2_2_x64_avx13
Abaqus/Standard • **S11 Block with Uniform Load** • **S4b Cylinder Head Bolt-up** • **S8 Shipping Port**	2018 HF14	2021 HF6
Abaqus/Explicit • **E10 Plate with Uniform Load** • **E12 Airbag Deployment** • **E9 Taurus Crash**	2018 HF14	2021 HF6

Test setup

The HPE Application Engineer used a script to initialize each test run using Altair PBS Professional as a workload manager.

Output files were collected and stored.

Number of MPI Tasks:

3rd Generation Intel Xeon Scalable processor: 36, 72, 108, 144

2nd Generation Intel Xeon Scalable processor: 32, 64, 96, 128

Performance varies by use, configuration, and other factors. Learn more at intel.com/PerformanceIndex.

Performance results are based on testing by HPE as of April 2020 and September 2021 and may not reflect all publicly available security updates. See configuration disclosures for details. No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

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.

Your costs and results may vary.

Intel technologies may require enabled hardware, software, or service activation.

LEARN MORE AT

hpe.com/us/en/compute/hpc-manufacturing.html

a50005577ENW