Technical white paper

Accelerating modern computer-aided engineering workloads

Boost design and engineering productivity with HPE ProLiant Compute XD230 Systems and Intel Xeon 6 Processors with MRDIMM technology

Contents

Executive summary 3
The evolving CAE landscape 3
New infrastructure requirements 4
HPE ProLiant Compute XD230 System 5
Powered by Intel Xeon 6 processors 6
    MRDIMM technology 7
    Intel Speed Select technology 7
The proof is in the performance 8
    Higher CAE throughput 9
    The impact of MRDIMMs 11
Purpose-built for high-performance data centers 13
    Intel oneAPI Toolkit 13
    HPE Performance Cluster Manager 14
    HPE can help find the optimal solution for CAE workloads 14
Conclusion: Why HPE and Intel for CAE? 14
Learn more at 15

Executive summary

Manufacturers rely on computer-aided engineering (CAE) tools to support a wide range of engineering simulations, including finite element analysis (FEA), dynamic analysis, and computational fluid dynamics (CFD). HPE tested different configurations of HPE ProLiant Compute XD230 Systems powered by Intel® Xeon® 6 processors with Performance-cores (P-cores) across a range of standard models and solvers to evaluate their performance against servers powered by 5th Gen Intel Xeon processors.

The results demonstrate substantial performance gains:

- 2.57x higher average throughput for Ansys® Fluent® across nine standard CFD models.[1]
- 2.42x average throughput improvement across four CFD solvers and 24 models.[1]
- 2.09x higher average throughput for compute-intensive FEA explicit (crash) simulations using Ansys® LS-DYNA®.[1]

These improvements are driven by a combination of architectural advances and new technologies in HPE ProLiant Compute XD230 Systems and Intel Xeon 6900P-series processors. These include higher core counts, increased cache capacity, additional memory channels, and significantly greater memory bandwidth enabled by Multiplexed Rank DIMM (MRDIMM) technology.

The evolving CAE landscape

Faced with intensifying global competition, manufacturers of all sizes are under increasing pressure to innovate faster and develop more capable products without sacrificing quality. For designers and engineers, this translates into larger models, higher-fidelity simulations, and a greater reliance on numerical simulations over physical prototyping.

Engineers are moving beyond single-domain analyses toward coupled multiphysics simulations that model structural, thermal, fluid, and electromagnetic behavior simultaneously. This shift helps teams identify design flaws earlier in the development cycle and avoid costly late-stage engineering changes or field failures. AI-augmented workflows and surrogate models are further reshaping the landscape, enabling organizations to explore early design concepts at a pace that was not feasible with traditional simulation methods.

[1] Please see details of the benchmark results conducted by HPE in the "Higher CAE throughput" section of this paper.

Computational Fluid Dynamics (CFD) and Finite Element Analysis (FEA) are among the most computationally demanding workloads in modern HPC environments. CFD is used for multiple applications, including modeling aerodynamic, thermal, and heat-transfer problems. Complex CFD simulations that model multiphase flows, combustion, and fluids in rotating machines can run for days, even in clustered environments. FEA is commonly used to model static structures and systems under various dynamic load conditions.

These escalating computational requirements place significant demands on the IT teams that support CAE environments. IT managers must balance multiple factors, including performance, software license costs, rack density, power draw, and cooling capacity. At the same time, infrastructure complexity is increasing, driven by heterogeneous compute resources, hybrid cloud integration, evolving storage architectures, and a growing portfolio of simulation applications.

Compounding these challenges is a shortage of experienced system administrators. As a result, organizations require high-performance computing (HPC) platforms that deliver both high performance and operational simplicity—enabling IT teams to deploy, manage, and scale environments efficiently—and to focus on their work rather than day-to-day cluster management.

New infrastructure requirements

The trends reshaping the CAE landscape are driving new requirements for HPC infrastructure. While raw performance is critical for improving engineering productivity and reducing simulation times, infrastructure selection is increasingly influenced by software licensing models and workload characteristics.

Many commercial CAE applications are licensed on a per-core or token basis. As a result, maximizing the productivity of each licensed core is essential. Faster cores can complete simulations more quickly, reducing license checkout time and improving overall throughput—effectively lowering the cost per simulation.

To address these requirements, organizations need infrastructure that delivers:

- High per-core performance to maximize solver throughput and license utilization
- High memory bandwidth to accelerate memory-bound CAE workloads
- Large memory capacity and cache to support increasingly large and complex models

Some CAE workloads are compute-bound, while others are constrained by memory bandwidth, latency, or I/O performance. As a result, versatile compute nodes, combined with flexible cluster architectures that match the right resources to each workload, are essential for maximizing overall system utilization and efficiency.

In addition to performance, operational considerations remain critical. Organizations require high-performance platforms that are easy to deploy, manage, and scale in clustered environments. HPE ProLiant Compute XD230 Systems, powered by Intel Xeon 6 processors, are purpose-built to meet these evolving requirements.

HPE ProLiant Compute XD230 System

The HPE ProLiant Compute XD230 System is an advanced high-performance computing platform designed to deliver exceptional performance and flexibility for HPC and CAE workloads. Delivered in a compact 1U form factor, the system supports the latest Intel Xeon 6900P-series processors and provides a scalable architecture ideal for the most demanding CAE environments.

High memory bandwidth—enabled by next-generation memory technologies—can significantly improve the performance of memory-bound workloads such as CFD. At the same time, HPE Integrated Lights-Out (iLO) simplifies deployment, monitoring, and remote management. The HPE ProLiant Compute XD230 also supports high-performance interconnects, including HPE Slingshot and InfiniBand, and integrates with HPE Performance Cluster Manager (HPCM) for cluster provisioning and management, effectively bringing supercomputer-class capabilities to commercial environments. Key features of the system are shown in Table 1.

Table 1. Key features of the HPE ProLiant Compute XD230 System.

HPE ProLiant Compute XD230 System	
Processor	Up to two Intel Xeon 6900P series CPUs Up to 128 cores / 256 threads per socket
OS support	Microsoft Windows, Red Hat® Enterprise Linux®, SLES™, Ubuntu®, VMware®[2]
Memory	Up to 24 DDR5 6400 MT/s DIMMs or 8800 MT/s MRDIMMs (1DPC) per server. Up to 3 TB per server[3]
Storage	M.2 NVMe SSD: 2 x 480 GB or 2 x 960 GB
Expansion slots	Up to 3 PCIe Gen5 x16 expansion slots, 2 x OCP 3.0 connectors, 2 x M.2 connectors with NS20 4i-u module option with hardware RAID
Interconnect options	HPE Slingshot (200 and 400 Gbps) InfiniBand NDR, Standard Ethernet
Power supplies	Up to 2 hot-plug power supplies HPE 2400W or 3200W options
Fans	7 fans per node
Cooling	Air-cooled or optional Direct Liquid Cooling (DLC)
Server management options	HPE iLO Advanced OpenBMC compatible HPE OneView Advanced HPE Performance Cluster Manager (HPCM)

HPE ProLiant Compute XD230 System front view

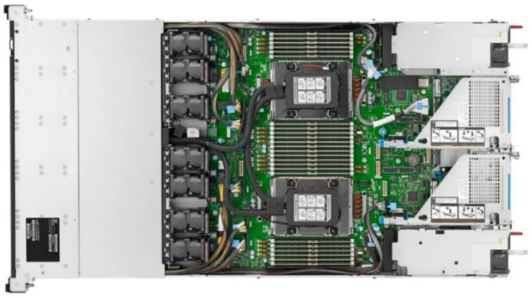

HPE ProLiant Compute XD230 System top view

---

[2] See the HPE Servers Support & Certification Matrices for details on supported operating environments.

[3] The XD230 supports 32GB, 64GB, 96GB, and 128GB RDIMMs, providing up to 3TB per server. With MRDIMMs, only 32GB, 64GB, and 96GB are supported due to mechanical height considerations.

The HPE ProLiant Compute XD230 System is engineered for sustained performance under demanding workloads. It supports multiple cooling options, including Closed-Loop Liquid Cooling (CLC) and Direct Liquid Cooling (DLC).[4]

Powered by Intel Xeon 6 processors

High-performance computing (HPC) workloads demand exceptional performance, flexibility, and the ability to handle diverse workloads—while balancing security, energy efficiency, and total cost of ownership (TCO). Intel Xeon 6 processors with P-cores deliver up to 2X the number of cores[5] compared to previous-generation processors, 2.36x higher peak memory bandwidth,[6] and built-in AI acceleration.[7]

Intel Xeon 6900P-series processors use a modular, tile-based design, as illustrated in Figure 1, that separates compute and I/O functionality. Each processor integrates up to three compute tiles containing CPU cores and shared cache, along with two I/O tiles that support memory controllers, PCIe connectivity, and system interconnects. The tiles are connected using the Intel Embedded Multi-Die Interconnect Bridge (EMIB), enabling high-bandwidth, low-latency communication within the processor package.

Intel Xeon 6980P processor:[8]

- Up to 128 P-Cores and 256 threads
- 12 memory channels - DDR5 (6400 MT/s) or MRDIMMs (8800 MT/s)
- 504 MB low-latency cache (LLC)
- Max Turbo Frequency: 3.9 GHz[9]
- All Core Turbo Frequency: 3.2 GHz[10]
- Intel UPI Speed: 24 GT/s
- 96 PCIe® Gen5 lanes
- TDP: 500 W
- Intel AMX, Intel SSE4.2, Intel AVX, Intel AVX2, Intel AVX-512, Intel AVX 10.1

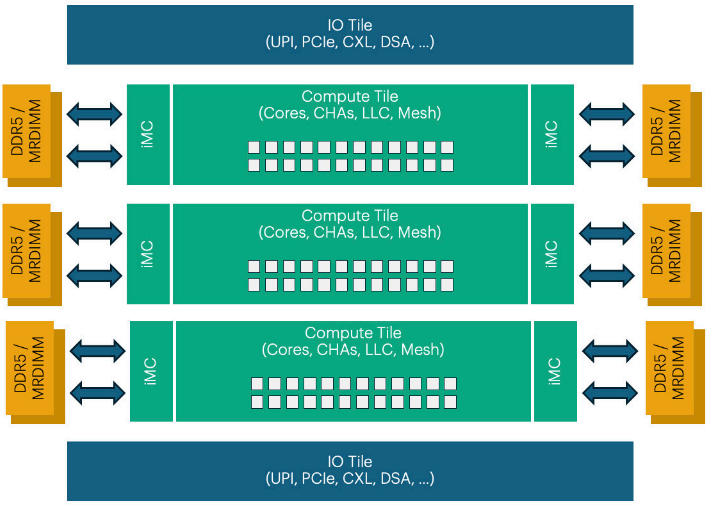

Figure 1. Intel Xeon 6980P logical block diagram.

Intel Xeon 6900P-series processors support configurable non-uniform memory architecture (NUMA) modes, including Sub-NUMA Clustering (SNC), which divides a processor into multiple NUMA domains.[11] This can improve locality and reduce memory latency for parallel solvers common in CAE environments.

In addition to high core counts and memory bandwidth, Intel Xeon 6900P-series CPUs support Intel® Advanced Vector Extensions (AVX-512), enabling a single instruction to operate on 512-bit vectors containing multiple data elements, significantly improving performance for applications that rely on floating-point and vector operations.

[4] While not required, DLC is recommended for XD230 systems populated with high TDP processor SKUs.

[5] Based on a comparison of a 128-core Intel Xeon 6980P and a 64-core 5th Gen Intel Xeon 8592+.

[6] Intel Xeon Platinum 8592+: 5,600 MT/s × 8 bytes × 8 channels = 358,400 MB/s = 358.4 GB/s. Intel Xeon 6960P with MRDIMMs: 8,800 MT/s x 8 bytes x 12 channels = 844,800 MB/s = 844.8 GB/s. 844.8 GB/s / 358.4 GB/s = 2.36.

[7] On-chip acceleration based on Intel Advanced Matrix Extensions (AMX) is designed to accelerate AI training and inference workloads

[8] See Intel Xeon 6980P Processor Specifications.

[9] Max Turbo Frequency is the maximum single-core frequency at which the processor is capable of operating using Intel Turbo Boost Technology and, if present, Intel Turbo Boost Max Technology 3.0 and Intel Thermal Velocity Boost.

[10] All Core Turbo Frequency is the frequency at which the processor can operate all its available cores concurrently.

[11] See the Intel Xeon 6 with P-cores Configuration and Tuning Guide for HPC Applications.

With 12 DDR5 memory channels per socket, Intel Xeon 6900P-series CPUs provide the memory bandwidth required to support high core counts and data-intensive CAE workloads. Each memory channel can support one or two DIMMs. However, the HPE ProLiant Compute XD230 uses a one DIMM per channel (1DPC) configuration, allowing each DIMM to operate at the highest supported data rates. The HPE ProLiant Compute XD230 System supports the Intel Xeon 6900P-series CPUs listed in Table 2, allowing customers to choose the processor configuration best suited to their workloads. For CAE workloads, the Intel Xeon 6960P processor with 72 cores is often considered a sweet spot.[12]

Table 2. Supported Intel Xeon 6900P-series processors in the HPE ProLiant Compute XD230 System.

Intel Xeon 6 processor	Cores/ threads	Base clock	Max turbo clock	Max memory	L3 cache (LLC)	Default TDP (watts)	DDR5 memory	MRDIMM memory
Intel Xeon 6980P	128/256	2.00 GHz	3.90 GHz	3 TB	504 MB	500	6400 MT/s	8800 MT/s
Intel Xeon 6978P	120/240	2.10 GHz	3.90 GHz	3 TB	504 MB	500	6400 MT/s	8800 MT/s
Intel Xeon 6972P	96/192	2.40 GHz	3.90 GHz	3 TB	480 MB	500	6400 MT/s	8800 MT/s
Intel Xeon 6960P	72/144	2.70 GHz	3.90 GHz	3 TB	432 MB	500	6400 MT/s	8800 MT/s
Intel Xeon 6952P	96/192	2.10 GHz	3.90 GHz	3 TB	480 MB	400	6400 MT/s	8800 MT/s

MRDIMM technology

Multiplexed Rank DIMMs (MRDIMMs), also referred to as Multiplexed Combined Rank DIMMs (MCRDIMMs), are an advanced DDR5 memory technology designed to increase bandwidth while maintaining compatibility with standard DDR5 DIMM sockets. MRDIMM modules incorporate an onboard buffer and multiplex multiple DRAM ranks onto a single memory channel, increasing effective memory bandwidth.

MRDIMM support is built into Intel Xeon 6900P-series processors, and each supports memory speeds of up to 8800 MT/s across 12 channels, delivering up to 844 GB/s of peak memory bandwidth per socket—approximately 37% higher than the maximum theoretical bandwidth for DDR5-6400.[13] For some memory-bandwidth-intensive CAE workloads, this can reduce memory bottlenecks and shift the optimal core count higher, enabling solvers to run more efficiently at higher core counts and resulting in higher throughput.

Intel Speed Select technology

Intel Xeon 6900P-series processors also support Intel® Speed Select Technology—Performance Profile (SST-PP), which enables system administrators to configure processors with different core counts, operating frequencies, and power profiles.[14] In supported configurations, this allows a processor—for example, a 72-core Intel Xeon 6960P—to run with fewer active cores at higher per-core frequencies. This capability is particularly useful for commercial software applications licensed on a per-core basis. By increasing per-core performance, organizations can complete simulations more quickly while using fewer licensed cores, improving overall license usage and reducing cost per simulation.

[12] Based on HPE internal recommendations.

[13] For MRDIMM at 8800 MT/s: 8,800 MT/s × 8 bytes (bus width per channel) × 12 channels = 844.8 GB/s per socket. For DDR5-6400: 6,400 MT/s × 8 bytes × 12 channels = 614.4 GB/s per socket. The uplift is calculated as 844.8/614.4 = 1.37 or 37%.

[14] For details, see Intel® Speed Select Technology – Performance Profile (Intel® SST-PP) Overview and Enabling Guide.

The proof is in the performance

In late 2025, HPE conducted an extensive set of benchmarks to characterize the performance of HPE ProLiant Compute XD230 Systems with Intel Xeon 6960P CPUs running popular CAE workloads. The tests involved 4 different server configurations, 8 solvers, and 40 models, evaluated across 3 clustered node configurations, yielding a total of 480 benchmark results. The following CAE applications were tested:

- CFD solvers: Ansys Fluent, Ansys CFX®, OpenFOAM®, and Simcenter™ STAR-CCM+®
- FEA/structural solvers: Ansys Mechanical and SIMULIA® Abaqus/Standard
- FEA/explicit (crash) solvers: Ansys LS-DYNA and DS SIMULIA Abaqus/Explicit

Each of the 40 models and 8 solvers was tested on the HPE system configurations shown in Table 3. Each server node had 960 GB of local NVMe storage and was connected to a shared NAS filer and a Lustre® file system for remote storage.

Intel Xeon Processor Advisor: To help customers assess the financial and business benefits of upgrading to Intel Xeon 6 processors, Intel offers the Intel Xeon Processor Advisor tool. Users can create a free account to access scenario advisors for on-premises and cloud environments, as well as an ROI calculator tailored to Ansys Fluent workloads. Users select a Fluent model representative of their workload and specify the 4th or 5th Gen Intel Xeon processors in their current cluster. The tool then estimates TCO impact, labor savings, and reductions in simulation time. For example, upgrading a 10-node cluster from 4th Gen Intel Xeon 8480+ CPUs to Intel Xeon 6980P CPUs with MRDIMMs is estimated to increase simulation throughput by 57% and save approximately USD 471K in labor costs over two years—not including software license savings or time-to-market benefits.[15]

---

[15] These results were obtained from the Xeon Process Advisor Ansys Fluent calculator based on the open_racecar_280M standard model, assuming a 10-node cluster with Intel Xeon 8480+ CPUs and DDR5 memory (1,120 fleet cores).

Table 3. HPE server configurations tested.

System configuration	System #1	System #2	System #3	System #4
HPE System/Server	HPE Cray XD2000 System (XD220v nodes)	HPE ProLiant Compute XD230 System	HPE ProLiant Compute XD230 System	HPE ProLiant Compute XD230 System
CPU	2 x Gen 5 Intel® Xeon® Platinum 8562Y+	2 x Gen 6 Intel Xeon 6960P	2 x Gen 6 Intel Xeon 6960P	2 x Gen 6 Intel Xeon 6960P
Clock (Base/Max Turbo)	2.8/4.1 GHz	3.1/3.9 GHz	2.7/3.9 GHz	2.7/3.9 GHz
Cores per socket/server	32c/64c	48c/96c (Intel SST-PP)	72c/144c	72c/144c
L3/Last Level Cache per socket	60 MB	432 MB	432 MB	432 MB
Memory	8 x 32 GB DDR5-5600 DIMMs per socket	12 x 32 GB DDR5-6400 DIMMs per socket	12 x 32 GB DDR5-6400 DIMMs per socket	12 x 32 GB 8800 MT/s MRDIMMs per socket
MPI Interconnect	InfiniBand HDR200	InfiniBand HDR200	InfiniBand HDR200	InfiniBand HDR200
MPI nodes tested (Intel MPI)	1,2, and 3 2P node(s)	1,2, and 3 2P node(s)	1,2, and 3 2P node(s)	1,2, and 3 2P node(s)

Higher CAE throughput

HPE ProLiant Compute XD230 Systems showed strong performance gains over previous-generation Intel-based systems across all models tested. The combination of Intel Xeon 6900P-series CPUs and MRDIMM technology was especially beneficial for memory-bandwidth-intensive CFD workloads. Average throughput improvements for the HPE ProLiant Compute XD230 Systems running Ansys Fluent, compared with the previous-generation HPE Cray XD2000 System, are shown in Figure 2.

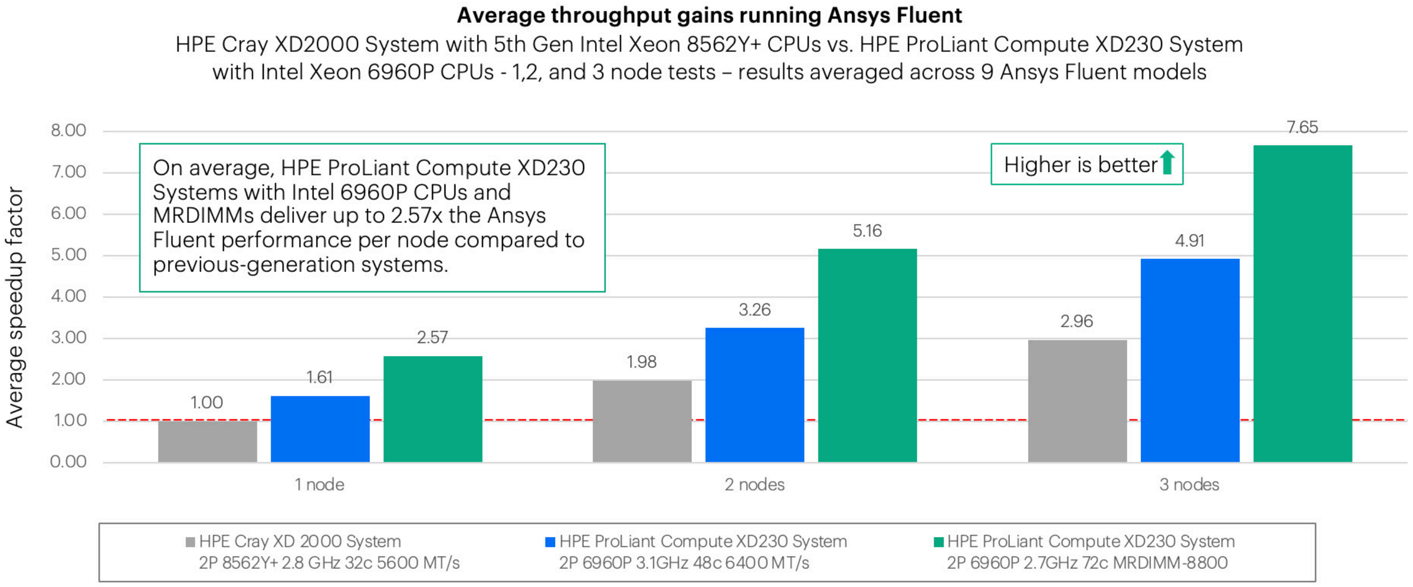

Average throughput gains running Ansys Fluent
HPE Cray XD2000 System with 5th Gen Intel Xeon 8562Y+ CPUs vs. HPE ProLiant Compute XD230 System
with Intel Xeon 6960P CPUs - 1,2, and 3 node tests – results averaged across 9 Ansys Fluent models

On average, HPE ProLiant Compute XD230 Systems with Intel 6960P CPUs and MRDIMMs deliver up to 2.57x the Ansys Fluent performance per node compared to previous-generation systems.

Figure 2. Average throughput gains running Ansys Fluent.

Across 9 standard Ansys Fluent models, the 2P HPE ProLiant Compute XD230 system with DDR5-6400 memory and 48 cores per socket was, on average, 1.61x faster than the baseline XD2000 system running 5th Gen Xeon CPUs in a single-node test.[16] The 2P XD230 system with 8800 MT/s MRDIMMs and 72 cores per socket delivered an average 2.57x throughput improvement over the baseline XD2000 in the single-node test.[17] Ansys Fluent throughput improvements on systems with MRDIMMs ranged from 2.40x faster running the Ansys standard f1_racecar_140m model to 2.81x faster running the standard aircraft_wing_14m model.[18]

Figure 3 shows the average performance gains by solver across the 40 models tested. All CFD solvers and models tested showed gains comparable to those observed with Ansys Fluent. Across all 24 CFD models and four solvers, the XD230 System with Intel Xeon 6960P CPUs (72c) and MRDIMM technology showed an average 2.42x throughput gain.

For structural simulations involving Ansys Mechanical and SIMULIA Abaqus/Standard, the results were mixed. The Xeon 6960P CPUs showed throughput improvements of 1.22x and 1.18x, respectively; however, in most of the Abaqus/Standard models, using MRDIMMs was counterproductive, delivering lower throughput than the system with DDR5-6400 memory.

Explicit FEA models used in applications such as crash and drop test simulations also benefited significantly from Intel 6960P CPUs. Here too, the benefits of MRDIMMs were solver- and model-dependent. The XD230 systems with 72-core 6960P CPUs and MRDIMMs delivered consistently higher throughput than the 48-core XD230 systems with DDR5-6400 memory for the explicit simulations.[19]

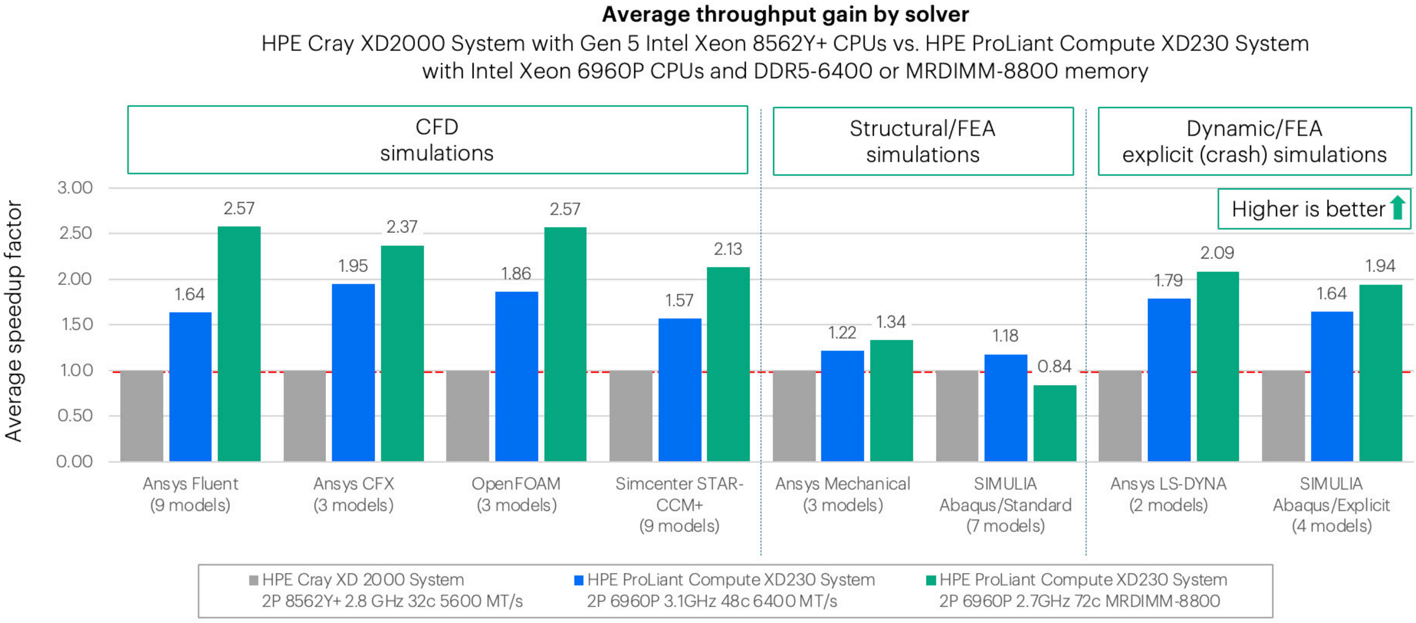

Figure 3. Average throughput gains by simulation type and solver across tested configurations.

---

[16] The Intel Xeon 6960P DDR5-6400 configuration was tested with 48 cores enabled via Intel SST because prior testing suggested that this is a sweet spot for many CAE codes. Using SST, the base clock increased from 2.7 GHz to 3.1 GHz. This would be a reasonable optimization in production environments.

[17] SST core reduction testing was not performed on the HPE ProLiant Compute XD230 systems with MRDIMMs. Those configurations, therefore, ran all 72 cores at a base frequency of 2.7 GHz. The rationale is that MRDIMMs can remove bottlenecks and make it practical to operate with higher core counts.

[18] Detailed results, including all model tests, are available from HPE.

[19] Tests to measure the impact of MRDIMMs showed that some of these gains were due to the higher core counts on the Intel Xeon 6960P operating with all 72 cores active rather than the presence of MRDIMMs.

The impact of MRDIMMs

To better understand how MRDIMMs affect simulation throughput compared to DDR5-6400 memory, all 40 models were run on identical server configurations with dual Intel Xeon 6960P CPUs, having all 72 cores active. Figure 4 shows the average improvement using MRDIMMs across the eight solvers tested. For each solver, the average, minimum, and maximum throughput gains are shown across the 40 models and 1-, 2-, and 3-node configurations tested.

In almost all CFD test cases, using MRDIMMs significantly improved simulation throughput. The only exception was a 3-node test involving the Ansys Fluent Standard open_racecar_280m model, where the XD230 systems equipped with MRDIMMs were ~9% slower than the equivalent test on systems with DDR5-6400 memory. In the 3-node test with the oil_rig_7m model, MRDIMMs boosted performance by a dramatic 2.63x.

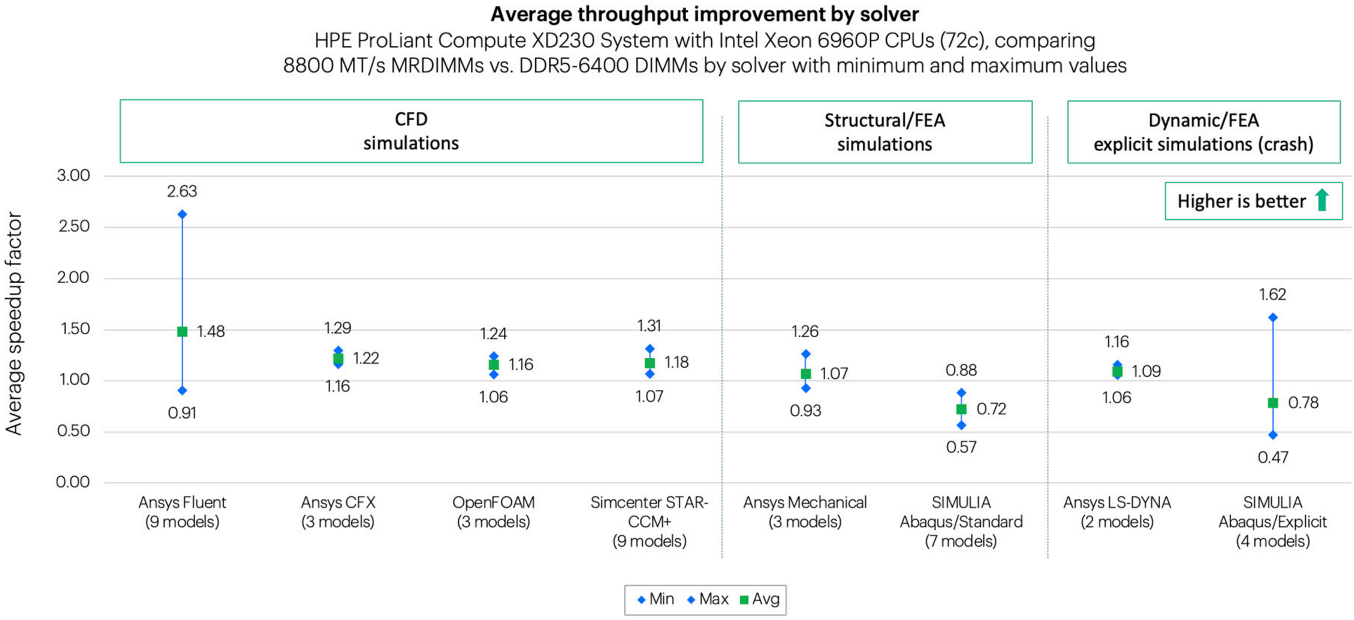

Figure 4. The impact of 8800 MT/s MRDIMMs vs DDR5-6400 DIMMs on simulation throughput.

The results in Figure 4 show that in both static and dynamic simulations using Ansys Mechanical and Ansys LS-DYNA, MRDIMMs improve performance, but the gains are smaller than for CFD workloads.

In Ansys Fluent testing using the standard oil_rig_7m benchmark model, MRDIMMs delivered a 2.63x throughput gain—a surprisingly high result, given that peak theoretical bandwidth differs by only 37% between the two memory technologies (as explained earlier). The average throughput improvement of 1.48x across all Ansys Fluent models exceeded the theoretical 37%. The likely explanation is that memory latency, not bandwidth alone, is a significant contributor to solver performance.

To verify this, another test was run to measure memory latency under load on a server equipped with DDR5-6400 memory, compared with an identical server using 8800 MT/s MRDIMMs.[20] Figure 5 shows results obtained using the Intel® Memory Latency Checker (MLC). As the level of memory read activity increased, the MRDIMMs showed significantly lower latency. The MRDIMMs also delivered higher read bandwidth before latency became an issue. As the chart in Figure 6 illustrates, as memory becomes a bottleneck, requests can queue in the memory controller, causing CPU cores to idle while waiting for memory. Because MRDIMMs can sustain higher read bandwidth at low latency, processor cores are less likely to idle, helping to explain the non-linear performance gains achievable with MRDIMMs.

MRDIMMs should be used where they are most effective: While MRDIMMs improved throughput across nearly all CFD simulations and many structural workloads, the results for SIMULIA Abaqus/Standard and Abaqus/Explicit were unexpected—both showed MRDIMMs underperforming DDR5-6400. One possible explanation is thermal clock-frequency throttling. The higher per-DIMM power draw of MRDIMMs, combined with all 72 CPU cores active, may have increased system temperatures and reduced CPU core frequencies, offsetting the memory bandwidth advantage.

Customers will benefit from HPC clusters in which some server nodes have DDR5-6400 memory and others have MRDIMMs, with memory configured as a selectable resource type in the workload manager. This will allow jobs to be dispatched to nodes with the memory type that yields the highest simulation throughput and the shortest license feature checkout time, depending on the application and model.

---

[20] A 2P system with 128-core Intel Xeon 6980P CPUs was used for the loaded latency test to illustrate the relationship between read latency and bandwidth with MRDIMMs and DDR5-6400 memory. Intel Xeon 6960P CPUs were not available at the time this test was conducted.

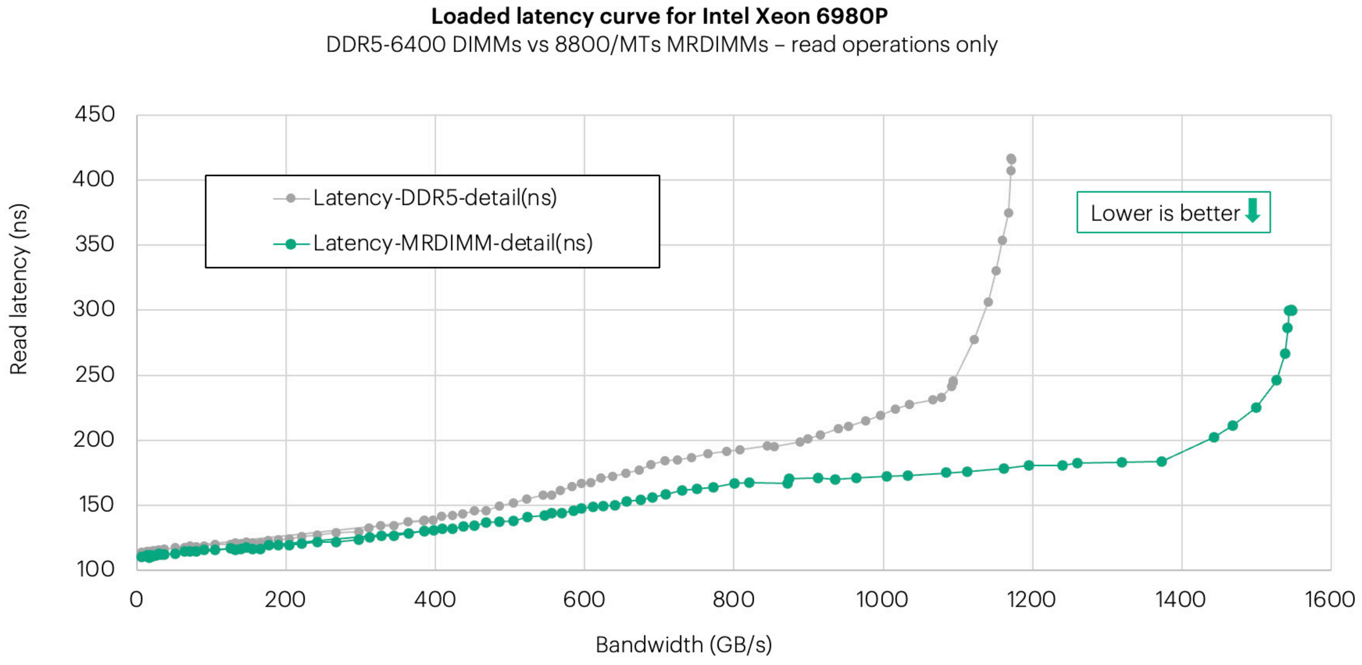

Figure 5. Loaded latency curve for Intel Xeon 6980P.

Purpose-built for high-performance data centers

HPE and Intel have a long history of HPC systems engineering, joint development, tuning, and production-level validation. Aurora, currently ranked third on the November 2025 HPC Top500 list, is a collaboration between Intel, Argonne National Laboratory, and HPE.[21] As of early 2026, 60% of the world's top ten fastest supercomputers use HPE servers, and over half of the Top500 supercomputers are powered by Intel processors.[22]

Even in small- and medium-sized commercial environments, the experience gained in deploying frontier-class systems matters. Advances in hardware, compilers, interconnects, and management software all cascade into commercial offerings from HPE and Intel, resulting in high-performance systems that are easy to deploy and manage.

While high-performance processors and servers are critical, maximizing CAE productivity requires a complete, clustered environment that includes networking, storage subsystems, and HPC management software. Together, HPE and Intel provide a full-stack HPC solution that includes the products, software tools, and expertise required to maximize CAE productivity, on-premises and in hybrid cloud environments.

The following sections describe the key components of the HPE and Intel HPC solution.

Intel oneAPI Toolkit

The Intel oneAPI Toolkit is an open, standards-based HPC development platform used by both customers and independent software vendors (ISVs) to deliver fast applications that scale across clusters. The toolkit provides essential components, including the Intel® MPI Library and the Intel oneAPI Math Kernel Library (oneMKL), along with compilers, profiling, and performance analysis tools to optimize application performance. Applications linked with oneMKL 2025.0 or later automatically benefit from optimized vectorized and matrix instructions implemented in Intel Xeon 6900P-series processors, including AVX-512 and AMX. Similarly, Intel MPI Library includes built-in tuning profiles that are automatically loaded at runtime. These profiles configure communication algorithms to optimize performance without requiring manual configuration.[23]

---

[21] See details of the Aurora exascale supercomputer at Argonne National Laboratory delivered by HPE and Intel.

[22] See the November 2025 Top500 list.

[23] For details, see the Intel MPI Library Developer Reference for Linux OS.

When deploying parallel solvers on Intel Xeon 6900P-series processors with SNC enabled, Intel's process and thread pinning library for MPI rank-to-core mapping (IPL2) generates optimal MPI rank-to-core mappings that account for the asymmetric NUMA topology created by SNC mode, making it far easier to optimize performance. Finally, Intel provides a comprehensive Configuration and Tuning Guide for Intel Xeon 6900P-series processors, along with HPC application recipes that document best practices for popular CAE applications, including the solvers mentioned in this document.[24]

HPE Performance Cluster Manager

HPE Performance Cluster Manager (HPCM) is an integrated cluster management solution for Linux-based HPC environments, scaling from departmental clusters to supercomputers. HPCM simplifies software deployment through automated bare-metal provisioning, image management, and software updates, while providing comprehensive hardware monitoring, power management, and real-time visibility into cluster health, storage, and workload scheduler status. By integrating with leading workload managers, including PBS Professional and Slurm, and consolidating cluster administration into a single platform, HPCM can reduce management overhead and total cost of ownership—freeing teams to focus on their core priorities.

HPE can help find the optimal solution for CAE workloads

HPE works closely with Intel and ISV partners to run a variety of internal benchmarks that compare systems based on Intel Xeon processors with competitive CPUs using standard CAE applications. HPE can help customers find the optimal solution for their needs by sharing details of standard benchmarks. HPE can also assist customers who wish to benchmark their own application workloads. For additional details on available benchmarks or to inquire about running your own benchmarks, contact your HPE representative.[25]

Conclusion: Why HPE and Intel for CAE?

HPE ProLiant Compute XD230 Systems powered by Intel Xeon 6900P-series processors give manufacturers a measurable performance advantage at a time when simulation throughput directly drives competitive outcomes. Together, HPE and Intel provide an integrated HPC solution that includes the products, software tools, and expertise required to maximize CAE productivity and efficiency, on-premises and in hybrid cloud environments.

Compared to servers based on 5th Gen Intel Xeon processors, HPE ProLiant Compute XD230 Systems powered by Intel Xeon 6900P-series CPUs with MRDIMM technology deliver dramatic gains across the demanding CFD and FEA workloads that manufacturers rely on:

- 2.57x the average throughput for Ansys Fluent simulations
- 2.42x better average throughput across a broad set of CFD workloads and solvers
- 2.09x higher average throughput for FEA explicit (crash) simulations with Ansys LS-DYNA

Using HPE servers with the latest Intel Xeon 6900P-series processors, manufacturers can:

- Reduce simulation cycle times for complex CFD and FEA models, supporting more design iterations within the same development window
- Improve license efficiency by completing simulations faster with fewer core-hours consumed
- Simplify deployment and reduce administrative overhead with a complete, proven HPC solution backed by HPE and Intel global support

As simulation workloads grow more complex and time-to-market pressures intensify, HPE and Intel deliver the performance, scalability, and simplicity that manufacturers need to stay competitive.

---

[24] See the Intel Xeon 6 with P-cores Configuration and Tuning Guide for HPC Applications.

[25] HPE may request that customers sign a non-disclosure agreement NDA to obtain detailed benchmark results.

Learn more at

[HPE ProLiant Compute XD230 QuickSpecs]

[Product Brief: Intel Xeon 6 Processors]

[Intel Xeon 6 Processor Advisor]

Visit HPE.com

Chat now

a00159151enw

HEWLETT PACKARD ENTERPRISE

hpe.com