


# Best practices for Ansys engineering innovation with next-gen CAE infrastructure



STRATEGIC PARTNER



# Executive summary: Meeting the demands of a new manufacturing age

**Manufacturers of all sizes are struggling to gain competitive advantage as they navigate ever-changing economic headwinds. As consumers become more demanding, and as products become smarter and more complex, advanced computer-aided engineering (CAE) simulation is essential to get quality products to market faster.**

**The dramatic increase in artificial intelligence (AI) workloads and the effort to design for sustainability are putting a strain on compute resources and engineering productivity. Meanwhile, manufacturers are juggling cost-efficiency pressures with changing product requirements and the need to support key initiatives such as electrification and personalization at the point of manufacture.**

**Innovative technology solutions are the key to overcoming these challenges.**

Whether it's measuring the velocity of air through a jet engine, analyzing autonomous vehicle safety, calculating heat transfer between circuit board components, testing interoperability of parts, or gauging a device's power efficiency, the ability of engineers to build better products depends on knowing how their designs will perform in the real world.

CAE is a critical tool in assessing and optimizing design fidelity, eliminating the need to build expensive physical prototypes by instead creating highly accurate computer-generated models, for operational cost savings and faster time to market.

However, ensuring a high level of fidelity requires ever-increasing design iterations and simulation complexity. CAE centers are under pressure to take on new workloads and techniques—such as multiphysics, digital twins, and the use of machine learning (ML)—to meet industry and government standards as well as market demand.

And modern product design is not limited to developing single components in silos. Design teams must collaborate across disciplines to model complex systems and optimize workflows, and engineers are being asked to support business decision-making by exploring the feasibility of approaches such as personalization and on-demand manufacturing.

To meet these challenges, leading manufacturers are adopting AMD-powered infrastructure solutions from Hewlett Packard Enterprise, which are proven to enhance the development of innovative products and services. HPE and AMD combine industry-leading compute and storage technology with top-performing CAE software from Ansys, speeding time to results and making it easy to support and scale a wide range of Ansys applications.

**This guide highlights essential best practices to help manufacturers accelerate product design and delivery by empowering next-level Ansys digital engineering with secure, reliable high performance compute infrastructure.**



# What's inside

Navigating CAE infrastructure requirements	4
Finding the optimal cloud strategy	8
Maximizing the potential of AI-augmented simulation	10
Achieving cost-effective sustainability	12
Conclusion	15



# Navigating CAE infrastructure requirements



## **Business challenge**

To get better-performing products to market faster while balancing sustainability and cost pressures, manufacturers need compute systems with exceptional capacity and performance to power their Ansys CAE simulation applications.



CAE and advanced simulation applications from Ansys are used throughout manufacturing, from automotive and aerospace to consumer products and industrial manufacturing, as an essential pathway to innovation. These applications are becoming more compute intensive as models increase in size and complexity, and as data sources expand, thanks in part to AI.

Multiple factors make it challenging to keep up with the computational demands of CAE, including data-heavy workloads in areas such as finite element analysis, computational fluid dynamics, and electromagnetics; increasing pressure to create lightweight or energy-efficient designs without compromising durability and quality; safety regulations and environmental standards calling for new product designs; and business mandates to leverage cloud while maintaining security, control, and cost transparency.

This is a perfect storm of new requirements—and it comes at a time when most CAE centers are already operating at capacity. While design and engineering teams struggle to keep pace with time pressures, outdated systems struggle to support evolving application needs. To meet the evolving requirements of Ansys CAE applications, manufacturers must adopt new infrastructure solutions that deliver a range of critical features and capabilities:

- High-frequency processors with high core counts help to complete more simulations faster
- High memory bandwidth and high cache per core to help maximize throughput
- High I/O performance for storage and network connections
- Low network latency and high network bandwidth to enable parallel simulations
- Industry-specific software tools that reduce design cycle times and optimize costs

In addition, CAE systems must offer high levels of reliability, availability, and serviceability to help maximize productivity and minimize downtime costs.



**Risk:** Relying on workstations or older commodity servers to meet evolving CAE requirements is not a viable long-term strategy. Performance liabilities can lead to bottlenecks in bringing designs to market—engineering work can't get done fast enough, hindering design innovation and impacting competitive edge. And ROI degrades as inefficient systems go underutilized.



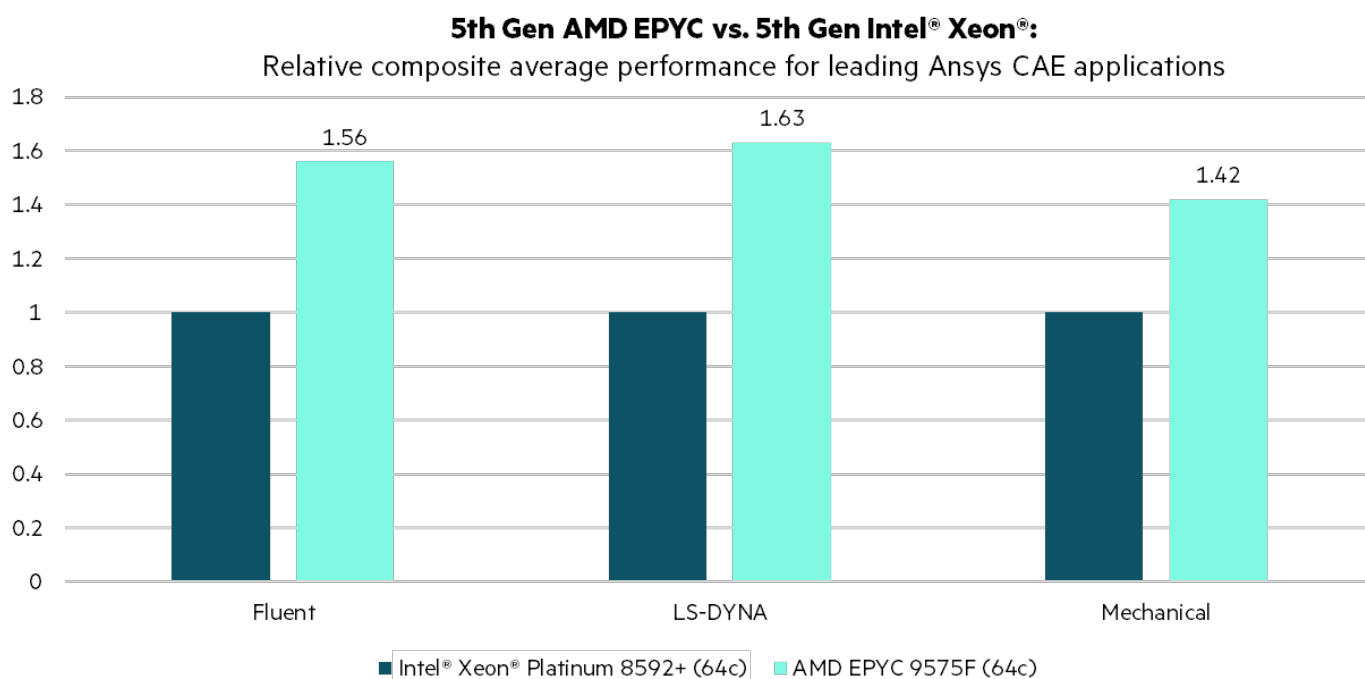
**Best practice:** Understand whether your CAE applications are performing at industry-standard levels, and work with hardware and software experts to determine whether legacy systems should be upgraded or replaced. Ensure your chosen infrastructure provides the reliability and scalability to accommodate new job types and requirements. Know which technologies and providers your competitors choose to work with, and why.

Leading manufacturers choose high performance solutions from HPE and AMD to meet the rising demands of CAE. HPE offers a broad portfolio of systems featuring the latest AMD EPYC™ CPUs and AMD Instinct™ GPUs to provide a decisive competitive advantage with the most powerful available technology.



HPE ProLiant servers featuring the latest AMD EPYC™ CPUs deliver up to 35% higher performance and 25% better energy efficiency compared to previous models, with record-setting benchmark results that surpass competitors in energy efficiency and performance across a range of workloads.<sup>1</sup> HPE ProLiant servers fit easily into existing data center environments and are performance-optimized, capable of powering CAE insights and innovation from edge to cloud. For manufacturers requiring the highest standards of performance, HPE high performance computing (HPC) and AI systems, powered by AMD, provide massive amounts of memory with flexible power and cooling options. These systems feature supercomputing-level capabilities at the right scale so manufacturers can take on the fastest-growing CAE workloads, empowering engineers to push product design boundaries like never before.

AMD EPYC™ CPUs and AMD Instinct™ GPUs are the firepower behind leading HPE systems. AMD EPYC-based systems deliver market-leading performance for many of the most challenging CAE workloads (see Figure 1).<sup>2,3,4</sup> AMD EPYC processors include outstanding levels of core density, frequency, and memory bandwidth to support both compute- and memory-intensive CAE applications. AMD EPYC processors are a robust and reliable foundation for advanced simulations, delivering up to 37% better performance in AI and HPC compared to the previous generation.<sup>5</sup>



**Figure 1.** Performance comparison for AMD EPYC and Intel Xeon processors running Ansys applications

<sup>1</sup> [“HPE ProLiant delivers 48 world records in performance and energy efficiency for enterprise workloads,”](#) HPE, October 10, 2024.

<sup>2</sup> [“Ansys Fluent® on 5th gen AMD EPYC™ processors computational fluid dynamics,”](#) AMD, November 2024.

<sup>3</sup> [“Ansys LS-DYNA® on 5th gen AMD EPYC™ processors finite element analysis,”](#) AMD, November 2024.

<sup>4</sup> [“Ansys Mechanical® on 5th gen AMD EPYC™ processors finite element analysis,”](#) AMD, November 2024.

<sup>5</sup> 9xx5-001: Based on AMD internal testing as of 9/10/2024, geomean performance improvement (IPC) at fixed-frequency.



Investment in AMD-powered systems also reduces hardware footprint; in one test, the same amount of Ansys Fluent work was performed on 43% fewer servers, thanks to AMD EPYC versus an alternative CPU.<sup>6</sup>

AMD Instinct™ accelerators are designed to deliver leadership performance for both AI and HPC applications optimized to take advantage of GPUs. With industry-leading memory and TB/s bandwidth, they optimize performance and help reduce TCO. Instinct accelerates memory-intensive CAE codes—for example, Ansys recently integrated support for Instinct accelerators into Fluent, significantly enhancing simulation efficiency and power; with the Ansys Fluent GPU solver, simulations that once took weeks or months can now be completed in hours or days.<sup>7</sup>



<sup>6</sup> [amd.com/content/dam/amd/en/documents/epyc-technical-docs/white-papers/raising-the-bar-for-high-performance-solutions-in-the-public-sector.pdf](https://amd.com/content/dam/amd/en/documents/epyc-technical-docs/white-papers/raising-the-bar-for-high-performance-solutions-in-the-public-sector.pdf)

<sup>7</sup> "Ansys Fluent® adds AMD Instinct™ MI200 and MI300 Acceleration to Power CFD Simulations," AMD, 2024.



# Finding the optimal cloud strategy



## **Business challenge**

While cloud computing is a good fit for many applications, there are potential drawbacks for CAE workloads. The cost dynamics of high performance CAE systems are very different from traditional IT applications, given their high sustained utilization.



Digital transformation is fueling diverse and data-heavy workloads across a wide range of industries including manufacturing. The need for real-time analytics at scale has increased, along with the demand for supercomputing-level performance for engineering applications. But power and storage constraints are two major issues, which is why cloud solutions have become more appealing.

In CAE, using cloud resources can be helpful for proof of concepts and purchasing short-term or burst capacity. Cloud solutions can also help accelerate new application deployments. Cloud-based software-as-a-service (SaaS) solutions from independent software vendors (ISVs) such as Ansys are increasingly capable of supporting highly complex simulations.

Smaller businesses and startups often turn to public cloud, which can be advantageous since it enables faster time to productivity with fewer capital investments. However, the challenges and economics of IT environments differ for CAE workloads versus traditional enterprise applications. As companies grow in scale, this calculation can change rapidly.

A recent Forbes article reports that as much as 32% of organization's cloud spend is wasted<sup>8</sup>—reasons include underused and overprovisioned resources, lack of skills, and insufficient management controls. Ideally, companies need a solution that allows them to operate independently while taking advantage of the convenience and flexibility of the cloud.



**Risk:** It's easy to overspend with the cloud, especially for workloads as demanding as CAE. Engineers often need training to come up to speed on using cloud interfaces efficiently. Public cloud users can experience performance and latency issues from having to run in a virtual vs bare-metal environment. In addition, security and compliance may be concerns for sectors where data privacy is paramount.



**Best practice:** A hybrid cloud solution for CAE to run Ansys applications can be a better fit and pose fewer risks than a cloud-only or cloud-first strategy. The hybrid approach gives manufacturers full control over their environments and reduces up-front costs while still providing the cloud resources they need.

Flexible deployment options from HPE GreenLake allow companies of all sizes to access high performance infrastructure and run Ansys simulation workloads in a cloud environment. Unlike public cloud, HPE GreenLake is a hybrid cloud by design—so users can deploy and manage resources across their private and public clouds while retaining control of data and flexibility over how they consume and manage their services.

HPE GreenLake combines scalability, security, visibility, ease of management, and pay-per-use\* predictability. And Ansys customers can take advantage of HPE's skilled experts to implement and operate the environment, helping reduce the cost and complexity of maintaining their own architecture.

Manufacturers choose HPE GreenLake to accelerate HPC-powered digitized development of next-generation products—in one case, a customer was able to boost simulation workloads with 3x faster performance.<sup>9</sup>



As the automotive industry accelerates its digital transformation and increasingly moves towards simulation instead of prototyping and performance evaluation testing, automakers are increasingly requesting mathematical models of tires which is a faster, more cost-efficient and sustainable process. This requires numerical models of the tire and our new HPC system (a 'seventh generation high performance computing [HPC] system, which is three times more powerful than its predecessor, as a service through HPE GreenLake') will advance this move towards digitized simulation.

Our design engineers are excited about the enhanced capabilities of the system, and we are looking forward to accelerating the training of our AI models.”

– **Tamotsu Mizutani**, Corporate Officer and Division General Manager, Technology Development Division, Toyo Tire<sup>10</sup>

\* May be subject to minimums or reserve capacity may apply

<sup>8</sup> [“What Companies Can Do About Cloud Spend Wastage,”](#) Forbes, April 22, 2024.

<sup>9, 10</sup> [“Toyo Tire selects HPE GreenLake cloud to accelerate HPC-powered digitized development of next-generation tires,”](#) HPE, August 22, 2024.

# Maximizing the potential of AI-augmented simulation



## **Business challenge**

The use of compute-intensive AI and ML workloads, while offering tremendous promise, is putting a strain on existing CAE hardware resources. Manufacturers must reimagine compute to tackle these new challenges and maintain competitive advantage as workloads converge.



Manufacturers are experiencing a steep incline in the use of data science, advanced analytics, and AI techniques across product development. Companies can now leverage massive amounts of data to augment CAE, pinpointing design flaws and reaching data-driven insights faster than ever, with the latest tools from software leaders such as Ansys.

Ansys AI-augmented simulation is revolutionizing engineering by delivering unmatched speed, innovation, and accessibility. With a powerful suite of ML capabilities spanning multiple physics domains, engineers can achieve greater accuracy and capture finer details in their simulations. For example, Ansys optiSLang AI+ enhances sensitivity analysis, optimization, and robust design with advanced ML methods while CFD AI+ enables turbulence modeling through AI-driven coefficient tuning. AI+ capabilities are also embedded within Ansys Granta, Ansys SynMatrix, Ansys Mechanical, Ansys Missions, and Ansys Electronics, providing engineers with intelligent, data-driven solutions across diverse applications.

These exceptionally compute-intensive applications require top-performing hardware, and engineering IT managers are hard at work preparing CAE data centers to accommodate new Ansys AI workflows along with traditional simulation jobs. HPC systems are a requirement to support both new AI and long-established CPU-based simulations.



**Risk:** Relying on commodity infrastructure to support AI initiatives often means leaving performance on the table, since workstations or older servers typically do not offer the high performance processors and I/O required for AI workloads.



**Best practice:** Invest in HPC systems with AI-tested processors and powerful interconnects to support fast data movement. Switch to high-efficiency CPUs to run traditional CAE jobs with fewer systems, allowing more data center space and energy to support your organization's growing AI workloads. Consult with your software and hardware partners about the right mix of CPU and GPU infrastructure to support your AI goals.

Fortunately, cost-effective high performance solutions are available that can deliver the greater compute required by AI in a smaller data center footprint and power envelope. HPE and AMD CPU-based solutions are capable of delivering outstanding performance for compute-intensive workloads, including AI-augmented tools such as the Ansys AI suite.

As established earlier in this document, running Ansys CAE jobs on AMD EPYC CPUs will also help maximize space and energy remains for AI jobs requiring GPU acceleration. For customers looking to shift CAE or AI workload to GPUs, HPE HPC and AI systems powered by AMD offer an excellent solution for GPU-intensive and GenAI workloads. Based on internal testing and looking at commonly used AI precisions, the AMD Instinct MI325X can provide up to a 1.3x advantage compared to competing accelerators for both compute and inference performance.<sup>11</sup> Built using the same technologies that power the world's most powerful and sustainable supercomputers, these systems also provide the lightning-fast I/O that is critical for CAE and AI workloads storage and networking.

<sup>11</sup> ["AMD Delivers Leadership AI Performance with AMD Instinct MI325X Accelerators,"](#) AMD, October 10, 2024.



# Achieving cost-effective sustainability



## Business challenge

Sustainability initiatives are running in parallel with existing product development, putting increased demands on computing capacity.



Top companies across various sectors have adopted sustainability practices, recognizing them as a force for positive change and a strategic advantage in business. Sustainability consistently remains a top 10 business priority, surpassing even productivity and efficiency in a recent Gartner survey.<sup>12</sup> No longer written off as one-and-done projects necessary for compliance, sustainability initiatives are invested in as a source of differentiation and a winning operating principle.

Nevertheless, the move to sustainability creates undeniable challenges as CAE data centers struggle to balance finite space, power, and cooling with ever-increasing demand. As energy consumption increases, current strategies become less effective. Sustainability initiatives regarding power consumption and associated carbon emissions are key considerations—not just for environmental goals but also to help reduce overall TCO.

Additionally, consumer demand for greener eco-friendly products, electric/hybrid vehicles, and higher levels of product and component recyclability and reuse are driving an increased need for advanced computer-based simulation.



**Risk:** To meet increasing CAE compute demands in order to innovate and stay competitive, manufacturers must maintain high performance data centers at increasing levels of capacity. The cost to maintain, let alone increase, this capacity poses a dilemma: find a cost-effective and sustainable solution, or fail to meet the needs of product design teams.



**Best practice:** Data center managers should seek servers that are proven to deliver maximum throughput per watt for the specific CAE applications they use. They also need dense, energy-efficient cluster systems that offer advanced cooling options to reduce the consumption of compute resources.

Manufacturers need ways to increase capacity without substantial data center upgrades or new capital spending—these goals go hand in hand with the mandate to operate more sustainably.

AMD EPYC processors power the industry's most energy-efficient servers, delivering up to 2.25x the energy efficiency compared to comparable servers.<sup>13</sup> AMD EPYC CPU-powered servers can help IT organizations meet requirements to reduce cost, energy consumption, and save physical space. For example, one customer achieved up to 20%–30% energy savings by upgrading to HPE and AMD-based servers.<sup>14</sup>

AMD is committed to ongoing gains in energy efficiency, achieving a 28x increase in energy efficiency for their CPUs and GPUs from 2021–2024.<sup>15</sup>

Even with efficient processors like EPYC, high performance systems consume massive amounts of energy—so sustainability plans must include a new look at hardware efficiency. While air cooling is adequate for some applications, compute-intensive CAE codes run best on highly dense, high-core count configurations. HPE offers systems with optional direct liquid cooling (DLC) enabling customers to increase power and data center efficiency while delivering greater simulation capacity and flexibility. HPE server racks connect directly to facility water supplies without the need for secondary plumbing. Tools such as HPE Performance Cluster Manager make cluster power usage easy to manage, helping eliminate common IT roadblocks.

<sup>12</sup> [“Gartner Survey Reveals 69% of CEOs View Sustainability as a Growth Opportunity,”](#) Gartner, June 6, 2024.

<sup>13</sup> SP5-011F: SPECpower\_ssj® 2008 comparison based on published 2P server results as of 1/12/2024. Configurations: 2P 128-core AMD EPYC 9754 (36,210 overall ssj\_ops/W, 2U, [spec.org/power\\_ssj2008/results/res2024q1/power\\_ssj2008-20231205-01347.html](https://spec.org/power_ssj2008/results/res2024q1/power_ssj2008-20231205-01347.html)) is 2.25x the performance of best published 2P 64-core Intel Xeon Platinum 8592+ (16,106 overall ssj\_ops/W, 2U, [spec.org/power\\_ssj2008/results/res2024q1/power\\_ssj2008-20231205-01349.html](https://spec.org/power_ssj2008/results/res2024q1/power_ssj2008-20231205-01349.html)).

<sup>14</sup> [“NovoServe Addresses Sustainability Goals with AMD,”](#) AMD.

<sup>15</sup> [“AMD nearly beats 30x power efficiency goal a year early — AMD’s new AI servers are 28.3 times more efficient than 2020 versions,”](#) Tom’s Hardware, December 11, 2024.





With our thoughtful approach to hardware and software co-design, we are excited about the possibilities ahead, where we see a path to massive energy efficiency improvements within the next couple of years.”

– **Sam Naffziger**, Senior Vice President, AMD Corporate Fellow and Product Technology Architect

With AMD-powered HPE systems, engineers can now run more concurrent simulations per socket and still get results faster, meaning fewer nodes are required to deliver the same simulation throughput. In the SPECpower® benchmark widely used to measure energy efficiency, AMD processors delivered approximately 2.75x the energy efficiency of a comparable competitive system.<sup>16</sup> This means manufacturing customers can achieve dramatic performance and efficiency gains by carefully selecting processor technologies for their server infrastructure.

HPE also provides resources, training, and educational programs that bring IT engineers the knowledge and skills to integrate sustainability into their data center design and development processes. A few of HPE's program offerings include:

- Energy and sustainability discovery workshops
- Energy efficiency and water usage analysis services
- Data center sustainability trade-off analysis
- HPE resource usage monitoring service
- Energy and carbon reporting service

<sup>16</sup> [“AMD Unveils Powerful EPYC CPUs, Ryzen AI MAX+ Laptop Processor, and Space-Qualified Versal AI Edge SoC,” StorageReview, March 18, 2025.](#)





# Conclusion

Manufacturers are not infrastructure experts. The decision-makers who review and approve system spend need reliable input from both application and hardware providers to ensure their capital expenditures will meet their own unique needs, now and in the future.

HPE and AMD, in collaboration with Ansys, provide a comprehensive portfolio of high performance systems, world-class accelerators, tailored software, pre-sales and deployment expertise, and outstanding ecosystem of performance-optimized CAE applications to help manufacturers reduce costs, improve quality, ramp up productivity, and speed time to market.

Companies worldwide are already using CAE with solutions from HPE and AMD to transform how they work, relying on HPE and AMD as trusted advisors. As CAE becomes an even more integral part of the supply chain and product lifecycle, and as AI-augmented CAE becomes a standard operating practice, HPE systems powered by AMD processors featuring Ansys software will continue to deliver outstanding application performance for today's and tomorrow's challenges.

This is your opportunity to work smarter, innovate faster, and solve your most complex CAE problems. Let us help you achieve new levels of profit and performance in the next era of manufacturing.

## Learn more at

[HPE.com/us/en/solutions/manufacturing](https://HPE.com/us/en/solutions/manufacturing)

[HPE.com/us/en/supercomputing](https://HPE.com/us/en/supercomputing)

[Ansys.com/HPE](https://Ansys.com/HPE)

[Ansys.com/AMD](https://Ansys.com/AMD)

[HPE.com/us/en/solutions/amd.html](https://HPE.com/us/en/solutions/amd.html)

Visit **HPE.com**

 **Chat now (sales)**

  
**Hewlett Packard  
Enterprise**

© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

AMD and the AMD Arrow logo are trademarks of Advanced Micro Devices, Inc. Intel Xeon and Intel Xeon Platinum are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. All third-party marks are property of their respective owners.

a00133670ENW, Rev. 1