Hewlett Packard Enterprise | AMD

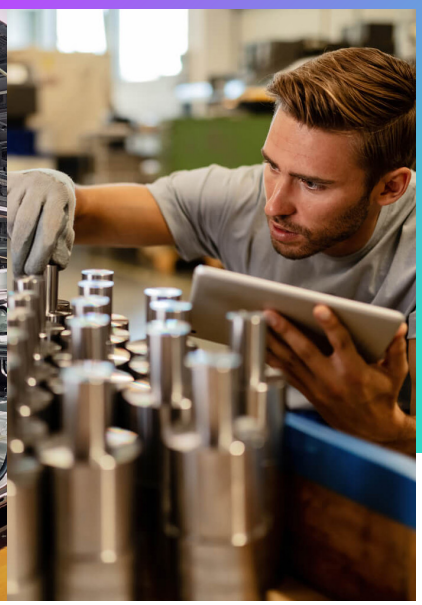

Driving manufacturing innovation with next-gen solutions for Ansys CAE workflows

Ansys

STRATEGIC PARTNER

The new age of manufacturing

Manufacturers are struggling to maintain or re-gain competitive advantage as they navigate post-pandemic economic and socio-political headwinds. As customers 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 need to design for sustainability, artificial intelligence (AI), and machine learning (ML) workloads, and a desire to improve the world's environmental and social footprint are putting a strain on engineering productivity and compute-intensive resources. These trends are fueling the need for innovation.

Manufacturers must find ways to tackle issues such as cost-efficiency, changing product requirements, and supporting personalization at the point of manufacture. Addressing these challenges demands more compute capability to optimize and validate engineering designs.

Whether it's measuring the velocity of air through a jet engine, calculating the transfer of heat between components on a circuit board, gauging the power efficiency of an electronic device, or analyzing any of millions of complex product design scenarios, the ability of engineers to build better products quickly and sustainably depends on knowing how product designs will perform in the real world.

CAE is a critical tool in bringing new products to market. CAE eliminates the need to build multiple physical prototypes of a product by replacing the resource-zapping physical models with computer-generated models that reduce operational costs and accelerate time to market. Advances in CAE have transformed manufacturing with a cutting-edge approach to modeling and simulation that boosts engineering productivity and profit margins and helps to manage lifecycle product costs.

However, optimizing a design to ensure a high level of fidelity now requires more iterations and even more complex simulations. Despite having finite data center capacity, CAE centers are under pressure to take on new workloads and employ new techniques to meet federal and industry standards — these include multiphysics simulation, ML, and the use of digital twins.

To meet these demands, manufacturers are adopting high-performance infrastructure with the processing power to meet all their CAE needs. Purpose-built solutions from Hewlett Packard Enterprise and AMD are proven to enhance the development of innovative products and services. HPE and AMD combine cutting-edge compute technology with CAE software from Ansys, making it simple to support and scale a wide range of applications. This brochure offers best practices to help manufacturers accelerate product design and delivery with CAE.

Table of contents

Best practice 1

Understanding CAE infrastructure challenges

Business challenge

Manufacturers need exceptional performance and processing capacity to enable simulation technology that supports collaboration, speeds up testing phases in development, and accelerates R&D cycles.

CAE and advanced simulation are used in multiple industries, ranging from automotive and aerospace to consumer products and industrial manufacturing. Engineers use CAE applications to test wind resistance (aerodynamics), assess durability of materials, ensure interoperability between parts and components, and optimize product quality.

As models increase in size and complexity, simulations are becoming more compute intensive. Several factors make it extremely challenging to keep up with the demands of CAE:

- Data-heavy workloads (such as computational fluid dynamics, finite element analysis, implicit and explicit simulations, and electromagnetic simulations) requiring more higher memory and more L3 cache
- Increasing pressure to create lightweight or energy-efficient designs without compromising durability and quality
- Complying with regulations regarding safety, electromagnetic interference, and environmental standards call for new product designs and engineering simulation models
- Moving to a hybrid cloud model without compromising IT security, control, or 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. Design and engineering teams struggle to keep pace with business demands, which increasingly include sustainability targets while containing costs. They need an infrastructure that is highly agile, reliable, and secure to achieve faster results for large simulations.

Outdated technologies struggle to keep up with the evolving needs of these applications. Legacy infrastructure often lacks the speed and scalability to use sophisticated algorithms that consume high levels of compute. **This performance liability can lead to bottlenecks in bringing designs to market, resulting in underutilized technology and losing out to the competition.**

Manufacturers are choosing high performance systems from HPE and AMD to meet the rising demands of CAE. HPE offers a broad portfolio of systems featuring the latest AMD EPYC™ CPUs to provide a decisive competitive advantage.

HPE ProLiant Gen11 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 levels of performance in liquid-cooled configurations, the HPE Cray XD2000 and HPE Cray EX2500 provide massive amounts of memory with flexible power and cooling options. These systems featuring supercomputing technologies offer next-generation compute capabilities at the right scale, so manufacturers can take on the fastest-growing CAE workloads.

AMD EPYC CPUs and Instinct™ GPUs are the firepower behind leading HPE systems. These technologies behind leading HPE systems provide record-setting performance for many of the most challenging CAE workloads. The latest generation of AMD EPYC delivers up to 1.45x greater throughput compared to other top-of-bin processors, which improves engineering productivity and accelerates simulation results — all at a lower TCO.[1]

[1] SP5-009D: SPECrate®2017_fp_base based on published scores from spec.org as of Jan 11, 2023. Configurations: 2P AMD EPYC 9654 (1480 SPECrate 2017_fp_base, 192 total cores, spec.org/cpu2017/results/ res2022q4/cpu2017-20221024-32605.html) is 1.45x the performance of published 2P Intel® Xeon® Platinum 8490H (1020 SPECrate 2017_fp_base, 120 total cores, spec.org/cpu2017/ results/res2023q1/ cpu2017-20221206-33040.html). SPEC®, SPEC CPU®, and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See spec.org for more information. Based on AMD testing. Matrices are divided by the rated power consumption of a representative accelerated compute node including the CPU host and memory and four GPU accelerators.

Best practice 2

Preparing for converged CAE and ML workloads

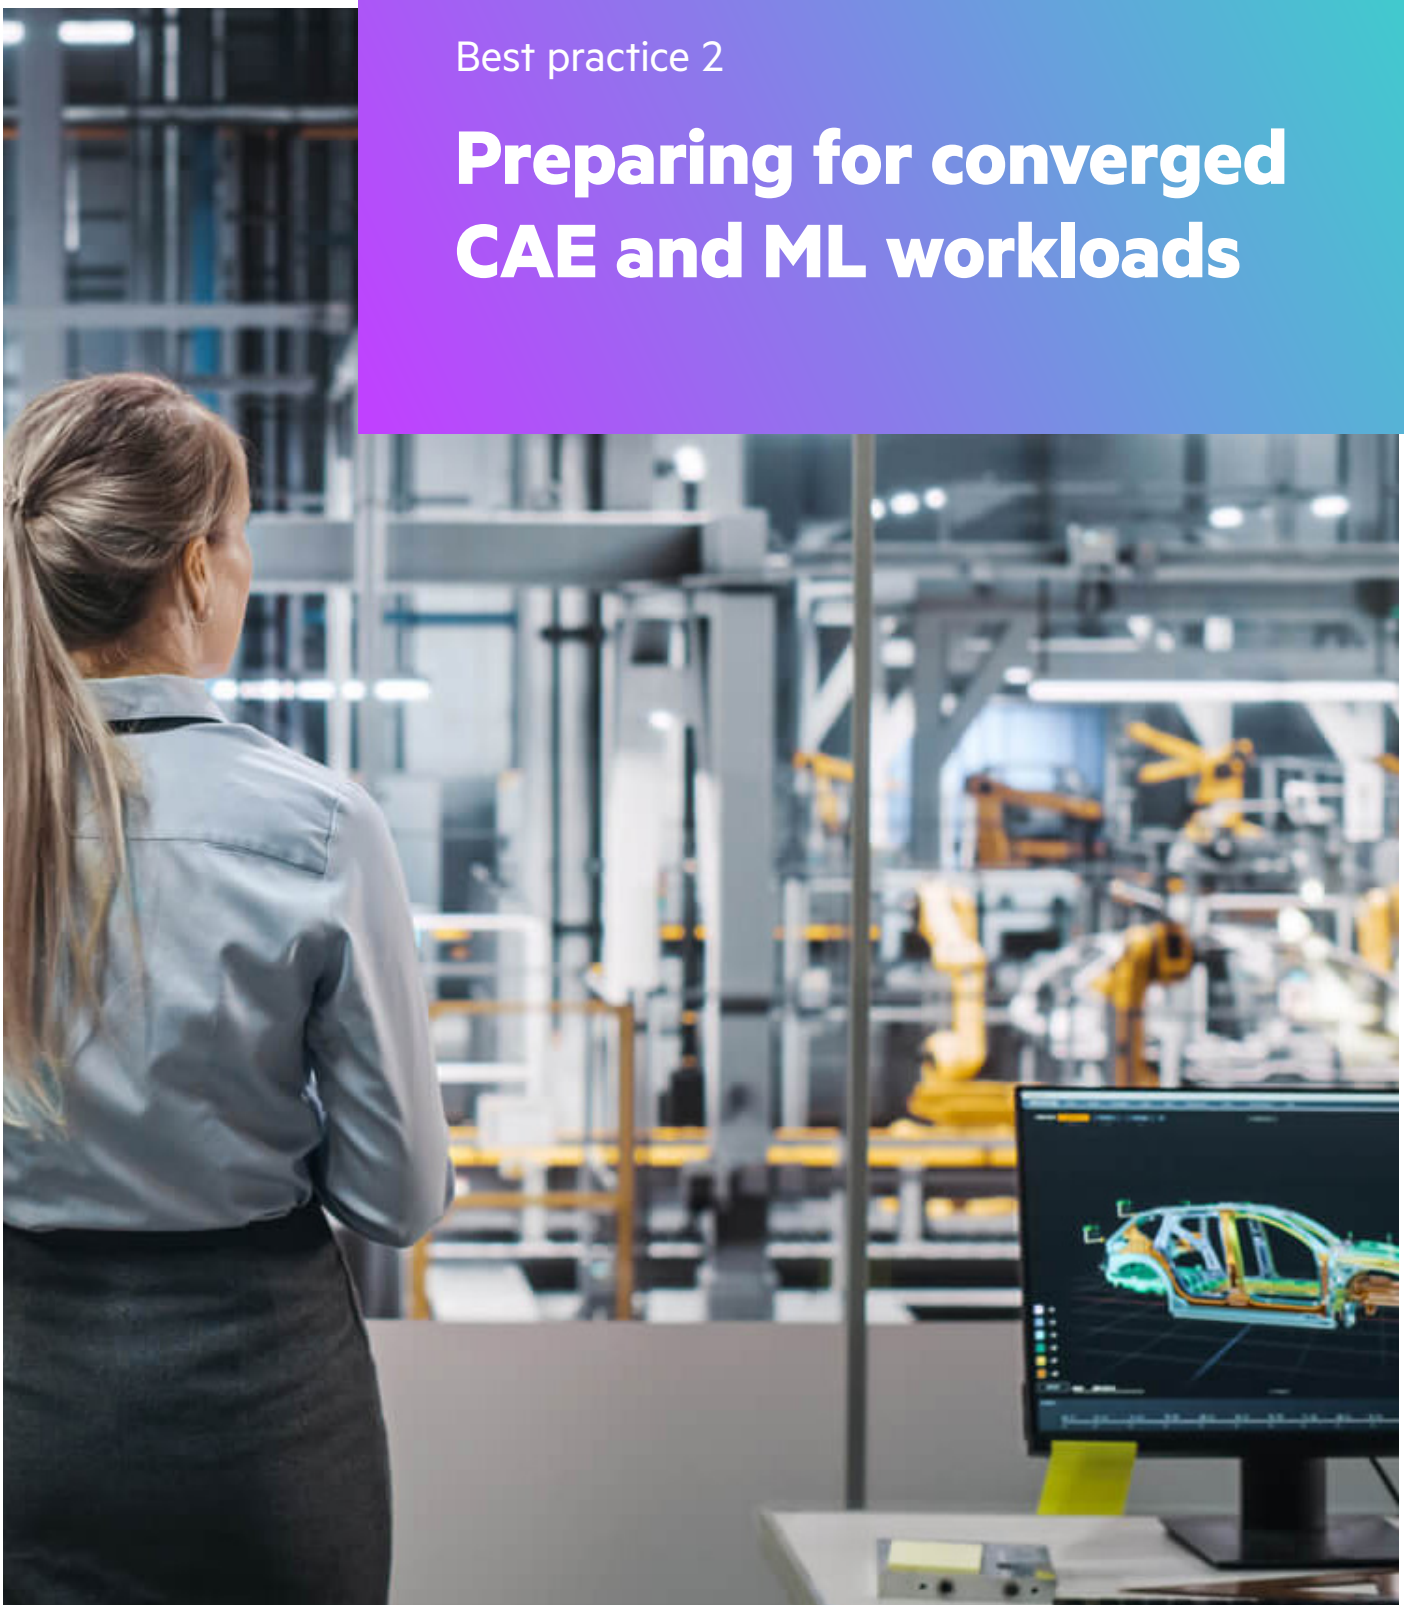

Business challenge

The use of compute-intensive workloads such as data analytics, AI, and ML is putting a strain on CAE resources. Manufacturers are reimagining compute to tackle new data challenges 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 for diverse simulations, immersive visualization, predictive analytics, and ML to pinpoint design flaws faster. Engineering IT managers are hard at work preparing CAE data centers to accommodate these new applications, which are exceptionally compute-intensive, in order to boost productivity.

Increasingly, product R&D teams are investigating the use of advanced analytics, data science, and AI and ML workloads to boost productivity. ML models are used for various purposes, from predictive maintenance to parameter optimization during simulations, to quality and defect management through the design and manufacturing process and product lifecycle. By replacing labor-intensive activities with intelligent automation, engineering teams can test with greater predictability and accuracy, often accelerating time to manufacture.

ML techniques in CAE are expected to proliferate in applications used throughout production. To meet these challenges, manufacturers must transform large-scale distributed computing environments to support increasingly converged workloads.

Fortunately, cost-effective high-performance solutions are available that can deliver greater compute in a smaller data center footprint and power envelope. HPE undertook a comprehensive series of internal Ansys standard benchmark tests, evaluating the latest 32-core 4th Gen AMD EPYC 9354 processor against a similar 3rd Gen AMD EPYC 7543 processor. The tests included Ansys Fluent, CFX, LS-DYNA, and Mechanical and were run on an HPE Cray XD2000 System comprised of four 2P clustered nodes (256 cores) and compared to a similar HPE Apollo 2000 System based on 3rd Gen EPYC technology with four nodes and 256 cores. The HPE Cray XD2000 delivered, on average, up to 1.52x more throughput than previous generation systems in the same data center footprint. As a result, manufacturers can complete over 50% more jobs in the same time frame as on the HPE Cray XD2000.[2]

Our joint solutions have the advantage of being easy to manage, maintain, and scale to save engineers from the burden of complex IT. Each platform is optimized for both traditional CAE workloads as well as new analytics, AI, and ML workloads. Engineers can also implement software tools that provide flexibility to run workloads on-premises or in the cloud, using hybrid cloud solutions with flexible cloud-bursting capabilities.[3]

[2] HPE internal testing — The Ansys tests were conducted during December 2022 through January 2023 comparing performance for an HPE Apollo 2000 Gen10 Plus platform with four 2P nodes to a similar HPE Cray XD2000 System. The HPE Apollo system had 8 x 3rd Gen AMD EPYC 9543 processors with a base clock of 2.8 GHz. The HPE Cray XD2000 system had 8 x 4th Gen AMD EPYC 9354 processors with a base clock of 3.25 GHz. Results may vary based on factors including silicon version, hardware and software configuration, and driver versions.

[3] Cloud bursting is a standard feature of most workload management solutions including Altair® PBS Professional®, IBM LSF, Slurm, and Univa Grid Engine®.

Best practice 3

Avoiding pitfalls of over-reliance on public cloud

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 systems for CAE are very different from traditional IT applications, given their high sustained utilization.

Digital transformation is fueling diverse and data-heavy workloads across manufacturing operations. The need for real-time analytics at scale has increased, along with the demand for supercomputing-type performance for CAE in variety of edge-to-cloud environments. But power and storage constraints are two major issues.

Using cloud resources is helpful for proof of concepts and purchasing short-term or burst capacity. It can also help accelerate new application deployments. Cloud computing has an important role in CAE simulation. Cloud-based software as a service from independent software vendors (ISVs) such as Ansys is increasingly capable of supporting highly complex simulations. Operating in the cloud has clear benefits:

- Deploying specialized infrastructure on demand
- Paying only for the resources they use
- Avoiding the cost and complexity of operating on-premises infrastructure

Public cloud can be advantageous for smaller companies and startups because it enables them to get started faster with fewer capital investments. However, the challenges and economics of IT environments for CAE workloads differ from those for traditional enterprise applications. As companies grow in scale, this calculation can change rapidly. **A recent HashiCorp-Forrester report found that 94% of enterprises are overspending in the cloud due to a combination of factors such as underused and overprovisioned resources, a 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.

A hybrid cloud solution for CAE 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 simulation workloads. HPE GreenLake for HPC is a private cloud — offered in both on-premises and colocation models — that makes it easier and faster to deploy converged CAE/AI workloads on high-performance clusters. The solution is turnkey with predictable, transparent costs, and continuous monitoring to enable capacity rightsizing plus the ability to scale up or down on demand.

By purchasing a packaged cluster environment that is fully integrated, manufacturers can avoid the integration risks associated with deploying solutions independently. Companies can choose to manage the infrastructure themselves or to own and control the infrastructure while having it managed for them by HPE.

HPE GreenLake for HPC enables more manufacturers to access CAE capabilities, so companies can focus on creating high-quality products faster and more cost-effectively than before.

Best practice 4

Achieving manufacturing sustainability goals

Business challenge

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

CAE data center managers face challenges such as finite space, power, and cooling. They need to find ways to increase capacity without substantial data center upgrades and new capital spending. Manufacturers are also concerned with operating more sustainably. Worldwide, companies must meet strict environmental, social, and governance (ESG) standards in order to remain compliant. Sustainability initiatives regarding power consumption and associated carbon emissions are key considerations — not just for environmental sustainability goals but also to help reduce overall TCO. Additionally, the consumer demand for greener ecofriendly products, electric/hybrid vehicles, and higher levels of product and component recyclability and reuse are driving an increased need for advanced computer-based simulation.

To lower their carbon footprint and create more sustainable processes, data center managers need servers that deliver maximum throughput per watt to minimize power and cooling requirements. They also need dense, energy-efficient systems that reduce the consumption of data center resources.

While air cooling is adequate for some applications, for highly dense configurations, HPE offers systems with optional direct liquid cooling (DLC). DLC allows companies to increase power density and data center efficiency while delivering more 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 liquid cooling easy to manage, helping eliminate common IT roadblocks.

AMD announced an ambitious goal to deliver a 30x increase in energy efficiency for AMD EPYC CPUs and AMD Instinct accelerators used for AI training by 2025.[4] Now, engineers can run more concurrent simulations per socket and still get results faster, meaning fewer nodes are required to deliver the same simulation throughput. AMD offers a Greenhouse Gas Emissions TCO Estimation Tool that can be used to estimate the potential savings and emission reductions with various AMD EPYC CPUs.

Ansys emphasizes the role of engineering simulation in enabling sustainable design, clean technologies, and carbon footprint reduction. Combined with HPE and AMD solutions, Ansys promotes the use of simulation tools for sustainable engineering practices, also providing resources, training, and educational programs that bring engineers the knowledge and skills to integrate sustainability into their design and development processes. Manufacturers can benefit from various data center sustainability services:

- Energy and sustainability discovery workshops
- Energy and sustainability baselining and road map services
- Water usage analysis
- Energy efficiency analysis service
- Data center sustainability trade-off analysis

[4] Includes high-performance CPU and GPU accelerators used for AI training and HPC in a 4-accelerator, CPU-hosted configuration. Goal calculations are based on performance scores as measured by standard performance metrics (HPC: LINPACK DGEMM kernel FLOPS with 4K matrix size; AI training: lower precision training-focused floating-point math DGEMM kernels such as FP16 or BF16 FLOPS operating on 4K matrices) divided by the rated power consumption of a representative accelerated compute node including the CPU host + memory and four GPU accelerators.

Best practice 5

Deploying a purpose-built solution for CAE

Business challenge

Manufacturers need exceptional performance and processing capacity enabled by the power of HPE, AMD, and Ansys.

Manufacturers run a variety of CAE workloads depending on their industry, mix of tools, and business specifications. To meet the evolving requirements of CAE, manufacturers need a solution that delivers the following capabilities:

- High frequency processors with high core counts to complete more simulations faster
- Large memory capacity, high memory bandwidth, and high ratios of 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 development costs
- Reliability, availability, and serviceability to help maximize productivity and minimize downtime costs

Companies can address these challenges with integrated technologies from HPE and AMD. We help companies build the optimal solution for CAE workloads that provides supercomputing-class performance at the right scale.

AMD EPYC processors include up to 128 cores with 12 memory channels per core. This helps ensure optimal performance for both compute- and memory-intensive CAE applications. HPE systems built on the latest AMD EPYC and Instinct processors are a robust foundation for advanced simulations, delivering a 52% simulation throughput increase per server compared to other generations.[5] With flexible power and cooling options, our platforms can reduce energy bills by up to 87% with minimal facility disruption. This innovation helps to future-proof the infrastructure against new application workloads that will require increased use of high-end processors. These benefits flow to the bottom line in the form of enhanced productivity and savings.

On top of these powerful platforms, Ansys offers an extensive application portfolio to streamline each stage of product development. Ansys CAE software contains sophisticated numeric modeling tools and robust solvers to deliver fast, accurate insights for almost every engineering application. CAE simulations include mechanical, fluid, electronic, and embedded software components to test all possible physical conditions that exist in the real world to create high-value products.

Various Ansys HPC licensing options allow companies to scale to whatever level simulations require — from entry-level parallel processing for a single engineer to virtually unlimited parallel capacity. For larger user groups, Ansys supports multiple parallel processing simulations that are highly scalable for the most challenging projects.

Now, manufacturers can invest in purpose-built solutions for the challenges of CAE. HPE, AMD, and Ansys recognize the value of high-performance infrastructure that delivers the best throughput per dollar, and together, we are enabling design teams to work easier and innovate faster.

[5] HPE internal testing — The Ansys tests were conducted during December 2022 through January 2023 comparing performance for an HPE Apollo 2000 Gen10 Plus platform with four 2P nodes to a similar HPE Cray XD2000 system. The HPE Apollo system had 8 x 3rd Gen AMD EPYC 9543 processors with a base clock of 2.8 GHz. The HPE Cray XD2000 system had 8 x 4th Gen AMD EPYC 9354 processors with a base clock of 3.25 GHz. Results may vary based on factors including silicon version, hardware and software configuration, and driver versions. At the time of this writing, the latest instances are C6A and R6A compute- and memory-optimized instances available in the AWS Cloud based on 3rd Gen EPYC processors.

Conclusion

As the scale and scope of CAE continue to evolve, manufacturers need reliable partners with deep computing and manufacturing expertise. HPE and AMD with leading ISV Ansys provide a comprehensive portfolio of high-performance systems, world-class accelerators, tailored software, 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. As CAE becomes an even more integral part of the supply chain and product lifecycle, HPE systems powered by AMD EPYC and Instinct processors featuring Ansys software delivers excellent CAE application performance for today 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

HPE.com/us/en/supercomputing

Ansys.com/HPE

Ansys.com/AMD

Chat now (sales)

a00133670ENW

Hewlett Packard Enterprise