Raising the Bar for High-Performance Manufacturing Solutions

Performance-optimized solutions for Computer-Aided Engineering (CAE) and Electronic Design Automation (EDA)

Executive Summary

Manufacturers of all sizes must make smarter, more complex, and higher-quality products to gain a competitive advantage. However, their inherent challenges include reducing costs and risks, increasing productivity, driving sustainability, and shortening time-to-market.

Lenovo helps manufacturers overcome these obstacles by providing a high-performance infrastructure of servers, storage, workstations, and services from the edge to the data center. For example, the Lenovo ThinkSystem servers, powered by AMD EPYC™ processors, are designed to help reduce constraints on the number, size, and complexity of simulation models to:

- Produce faster time-to-results
- Improve energy efficiency and total cost of ownership (TCO)[1]
- Deliver better power and cooling solutions to enhance reliability, density, and performance[1]
- Integrate with several complementary immersive manufacturing solutions for artificial intelligence, design, and digital twins throughout the product lifecycle.

Engineering Simulation is Critical for the Manufacturing Industry

To compete effectively, manufacturers must make smarter, more complex, and higher-quality products while reducing costs and risks, increasing productivity, advancing sustainability goals, and shortening time-to-market. To address these demands, manufacturers and their suppliers use a wide range of Engineering Simulation solutions like CEA and EDA applications that allow users to design and verify components and products in a virtual, risk-free environment, minimizing the need for expensive physical prototypes and tests and driving more collaboration and innovation (Figure 1 – left).

Hardware manufacturers typically develop their solutions with Independent Software Vendors (ISVs). Ansys®, Siemens®, and Dassault Systèmes® are among the major CAE ISVs that provide manufacturers software applications for workloads like Structural, Thermal, Fluids, Crash, Electromagnetics, and Multiphysics. For EDA solutions, semiconductor manufacturers often use applications from Cadence®, Synopsis®, and other ISVs (Figure 1 – right).

Figure 1: Engineering Simulation: Benefits and Applications

Companies are investing in high-performance infrastructure with the best-performing processors and systems to maximize the value of these game-changing applications. However, as the adoption of engineering simulation grows and models get larger and more complex, deploying and operating this infrastructure for these diverse applications can be very challenging.

Infrastructure Challenges with Growing Engineering Simulation Workloads

Several factors drive growth (Figure 2 – left): growing awareness of the value of engineering simulation and the widespread availability of ISV applications for designing and developing more complex products. More recently, Artificial Intelligence (AI) and Machine Learning (ML) methods based on neural networks (known as Deep Learning, or DL) are accelerating the process and improving its accuracy by iteratively refining the insights gained from multiple simulation runs that are typical in design optimization loops.

It significantly automates the process and makes complex engineering tools easier for more users. Many engineering simulation workloads are also moving to the cloud.

However, with the continuing explosion of data, the growing use of multiphysics simulations, and the need to better integrate across the manufacturing value chain and product lifecycle and improve sustainability, there is a need to accelerate all CAE/EDA computational processes even more.

Many manufacturers realize that their CAE and EDA solutions are expensive, inflexible, and not optimized for these environments and future needs (Figure 2 – right). Their high-performance computing (HPC) systems are often underutilized or inadequate to meet the various demands of their business and the vast volumes of data. Moreover, ISV software licensing costs increase, as does the overhead of installing, configuring, and managing high-performance infrastructure. And that's assuming manufacturers can source and retain employees with the correct skillsets to maintain the system.

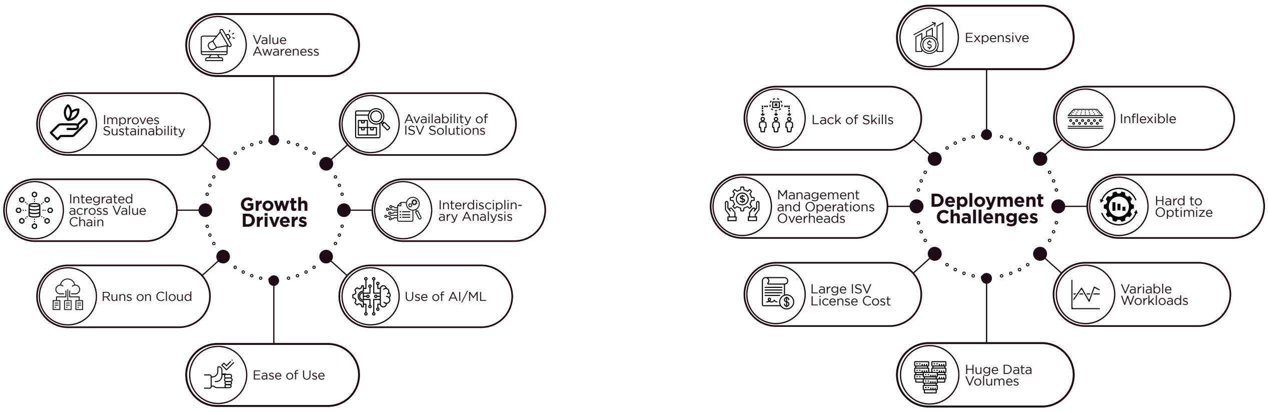

Figure 2: CAE/EDA Growth Drivers and Deployment Challenges

These deployment challenges are precisely the reasons to have the most modern, fully supported workload-optimized infrastructure that can fully utilize the licensing costs, is optimized for diverse CAE and EDA workloads, and helps manufacturers produce more innovative, high-quality products faster. However, there are several additional technical infrastructure challenges (Figure 3 – left):

- Inadequate processor frequency or core counts, requiring massive, often expensive scale-out solutions for many time-critical tasks
- Insufficient memory capacity and bandwidth and low ratios of cache per core hurt compute performance
- Difficult to optimize end-to-end including compute, I/O, networks, and visualization
- Lack of adequate data security procedures
- Poor energy efficiency necessitates expensive direct liquid cooling for some processor SKUs with the highest frequency and core counts, which typically deliver the best performance.

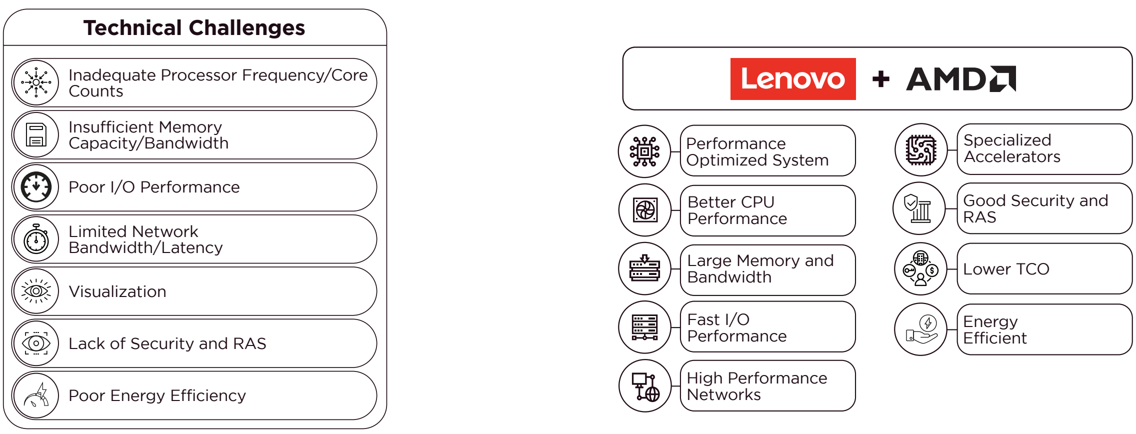

Figure 3: CAE/EDA Infrastructure Technical Challenges and Features of Lenovo + AMD Solutions

Lenovo ThinkSystem servers powered by AMD EPYC™ processors deliver the optimal architecture for CAE/EDA and help reduce constraints on the number, size, and complexity of simulation models while helping provide faster time to results. In addition, with these AMD CPU-based systems, engineers can improve design quality and prototype performance and significantly reduce total cost of ownership (TCO) by using fewer servers to do the same work, helping reduce power consumption and related emissions.

The right part of Figure 3 depicts the key features of these Lenovo ThinkSystem servers using AMD processors that significantly improve engineering simulation applications performance and help manufacturers perform more realistic, interdisciplinary simulations to drive more significant innovation and productivity.[2]

In addition to the best performance, manufacturers need real-time, highly reliable end-to-end infrastructure to process, store, and analyze massive amounts of data. Data movement costs can be substantial. So, most of the data (even for visualization) must be processed where it is created and not transferred to another remote location.

Delays in accessing the information or poor performance can hurt productivity when the engineering simulation footprint grows, such as when companies with multiple design centers worldwide collaborate and utilize the same data set to make decisions. Security concerns regarding proprietary algorithms or data often make public clouds unsuitable for sensitive engineering workloads.

To overcome these challenges, manufacturers need high-performance CAE/EDA systems with flexible purchasing options that offer cloud-like economics that can be deployed at the edge, in the cloud, or on-premises. Lenovo provides these integrated solutions for CAE and EDA with a broad portfolio of workstations, servers, storage, software, and services.

High-level Architecture of Lenovo CAE/EDA Solutions Powered by AMD

Lenovo CAE/EDA solutions use a foundation of expertly engineered on-premises high-performance infrastructure (Figure 4) ranging from workstations to the edge to the data center.

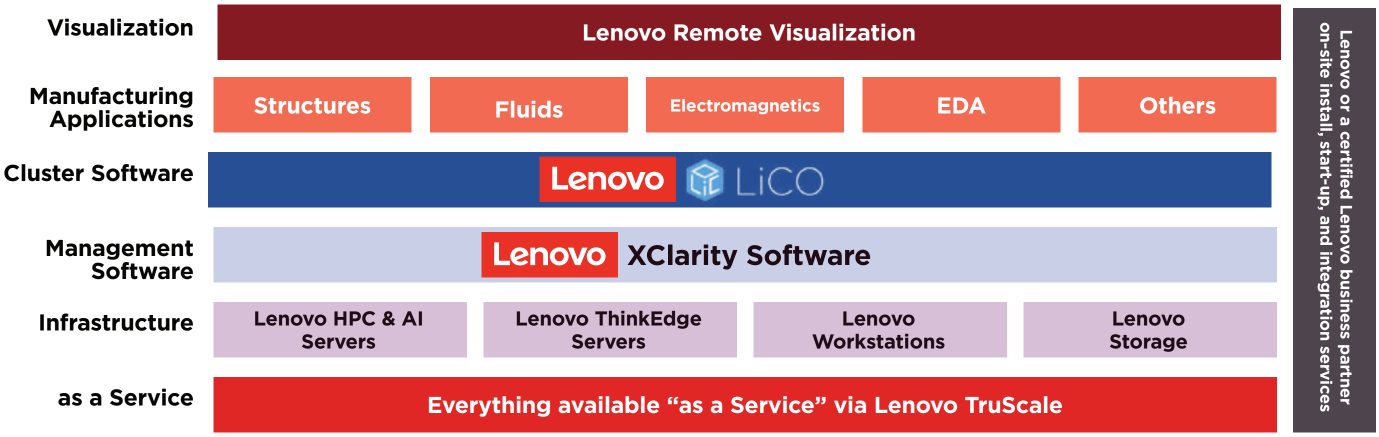

Figure 4: Lenovo CAE/EDA Solutions Portfolio

Key components of the Lenovo Solution Portfolio, starting at the infrastructure layer, include:

- **Lenovo Performance-optimized Servers:** Highly reliable, scalable, high-performance servers accelerate CAE/EDA significantly. This Lenovo portfolio of servers with AMD processors includes the GPU-rich Lenovo ThinkSystem SR675 V3, the ultra-dense Lenovo ThinkSystem SD665 V3 and Lenovo ThinkSystem SD665-N V3, and the slim-profile Lenovo ThinkSystem SR645 V3. Leveraging Lenovo's Neptune™ liquid cooling technologies, these systems range from direct water cooling for CPUs and GPUs to liquid-enhanced systems where liquid is used to augment standard air cooling.
- **Lenovo ThinkEdge Servers:** Deliver purpose-built and secure platforms suitable for compute-intensive and latency-sensitive applications like Lenovo's flagship edge-AI server, Lenovo ThinkEdge SE455 V3 deployed outside traditional data centers.
- **Lenovo Workstations:** ThinkStation P Series workstations deliver powerful performance. The ThinkStation P620 with its AMD Ryzen™ Threadripper™ PRO is ISV-certified, energy-efficient, and highly versatile.
- **Lenovo Storage:** Direct-Attached Storage JBODs and expansion units provide flexible, cost-effective, high-capacity storage and are ideal for space-constrained environments and cost-sensitive customers. Lenovo ThinkSystem DE Series All-Flash Arrays are designed for extreme performance with up to 2.0M IOPS and sub-100 microsecond latency and include industry-leading, enterprise-proven availability and security features

- **Lenovo XClarity:** This family of software simplifies and automates the deployment and management of Lenovo infrastructure so clients can focus on their high-value projects.
- **Lenovo Intelligent Computing Orchestration (LiCO):** Reduces the complexity of using a massive HPC cluster and simplifies application deployment, operation, and acceleration.
- **CAE/EDA Applications:** Lenovo has teamed with the industry's leading CAE software ISVs, including Ansys®, Dassault Systèmes®, Siemens®, Cadence®, Synopsis®, and others, to perform Structural Analysis, CFD, Electromagnetics, EDA, and more.
- **Lenovo Remote Visualization:** Provides reliable and secure access to graphics-intensive applications anytime, anywhere. Instead of issuing new expensive workstations to all design staff, IT can deploy less expensive enterprise or consumer-class personal computers. In addition, IT departments can maintain security and keep costs down by using remote virtualization hosted in an internal data center or from the cloud. Remote visualization performs intensive graphics operations on a high-end graphics server and generates a 2D pixel version that users can receive quickly.

  In addition, server-side rendering considerably speeds up the process of using graphics in remote sessions.
- **Lenovo or Certified-partner Services:** Lenovo and an extensive ecosystem of highly specialized HPC services partners can deliver the integrated Lenovo stack depicted above. They also provide onsite installation and start-up services to integrate this into a client's work environment, including installing CAE/EDA applications.
- **"As a Service":** Subscribe to innovation that scales with you with Lenovo TruScale, which provides end-to-end delivery, management, and refresh services, meaning your IT teams don't have to lift a finger when they deploy new devices and scale their IT infrastructure.

At the core of this high-level architecture are Lenovo servers, such as the Lenovo ThinkSystem SR645 V3, using the AMD EPYC™ processors (Figure 5) with advanced chip-level security features that deliver excellent performance for sensitive mission-critical manufacturing applications.

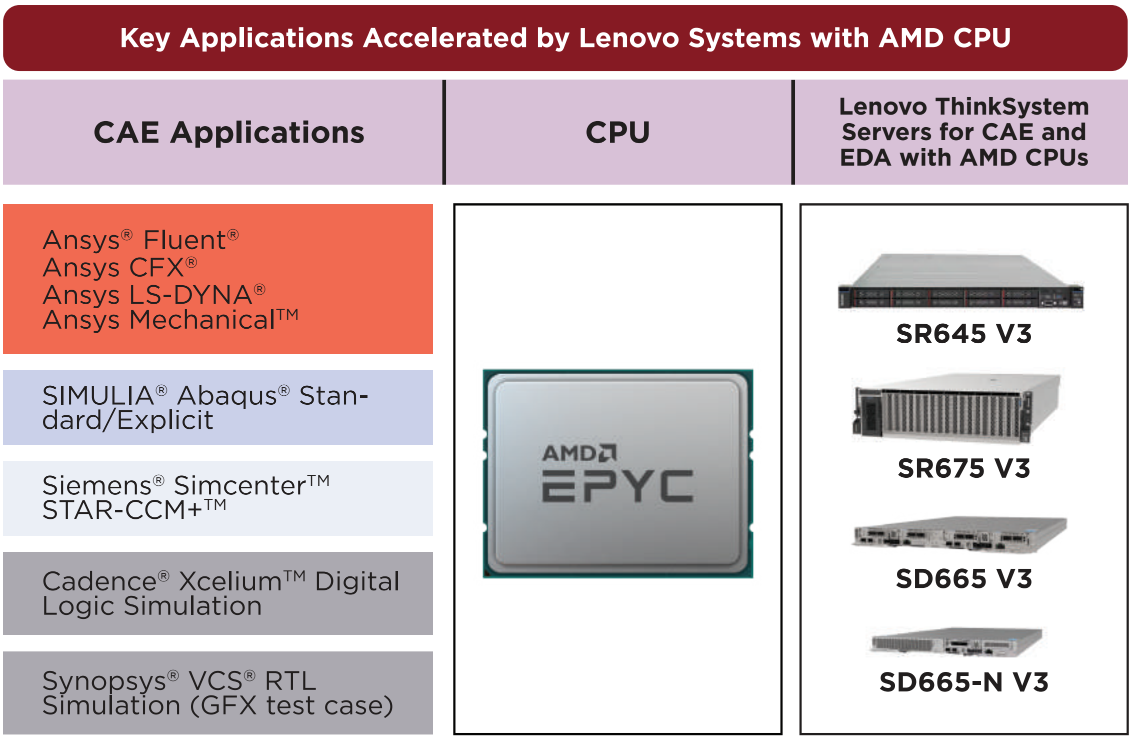

Figure 5: Key CAE/EDA Applications on Lenovo ThinkSystem Servers with AMD EPYC™ Processors

AMD Improves Performance of CAE/EDA and Recommended Lenovo Configurations

Compared to the prior generation, the latest AMD EPYC™ processors achieve better performance for CAE/EDA applications with up to 50% more cores, higher frequencies, support for AVX-512 instructions, more memory bandwidth, and faster PCIe® and Infinity Fabric™ data transfer rate.[3]

Many Computational Fluid Dynamics (CFD), Explicit Finite Element Analysis (FEA), and EDA applications also benefit from AMD EPYC™ processors with 3D V-Cache™ technology, providing triple the L3 cache compared to standard AMD EPYC™ processors. In addition, optimizing applications with AMD compilers and libraries can help enhance performance further.

The following sub-sections describe comparative performance improvement of critical CAE and EDA ISV applications with the latest AMD EPYC™ processors over the previous generation or alternative CPUs from the competition **based on AMD testing in conjunction with the respective ISV partner**. In addition, based on significant testing and optimization expertise on various client workloads, Lenovo provides recommendations for performance-optimized configurations based on the specific application and use case.

CFD Applications: Ansys CFX handles complex fluid flow physics like turbulence, heat transfer, and reaction modeling. It is also well known for its exceptional capabilities to simulate rotating machinery problems. Ansys Fluent models fluid flow, heat and mass transfer, and chemical reactions and offers a modern, user-friendly interface streamlining the CFD process from pre- to post-processing within a single workflow window.

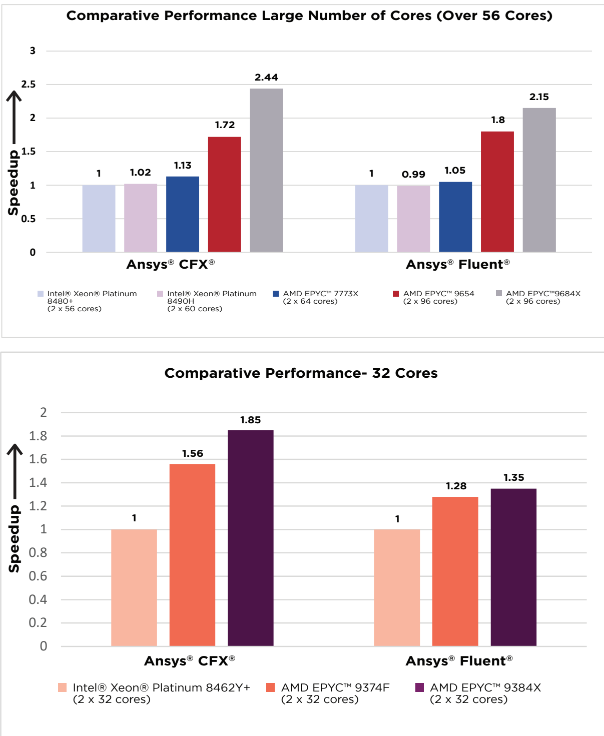

Figure 6: Excellent Performance for Ansys CFX and Ansys Fluent on the AMD EPYC™ 9004 Series Processors with 3D V-Cache[4, 5]

Figure 6 shows up to 144% performance gains for Ansys CFX and up to 115% for Ansys Fluent with the AMD EPYC™ processors with 3D V-Cache over a competitive alternative.

Siemens Simcenter STAR-CCM+™ is a Multiphysics CFD application for the simulation of products operating under real-world conditions. It facilitates modeling and analysis of various engineering problems involving fluid flow, heat transfer, stress, particulate flow, electromagnetics, and related phenomena.

For Simcenter STAR-CCM+™, Figure 7 depicts the performance gains of up to 118% on the AMD EPYC™ processor with 3D V-Cache compared to the previous generation.[5]

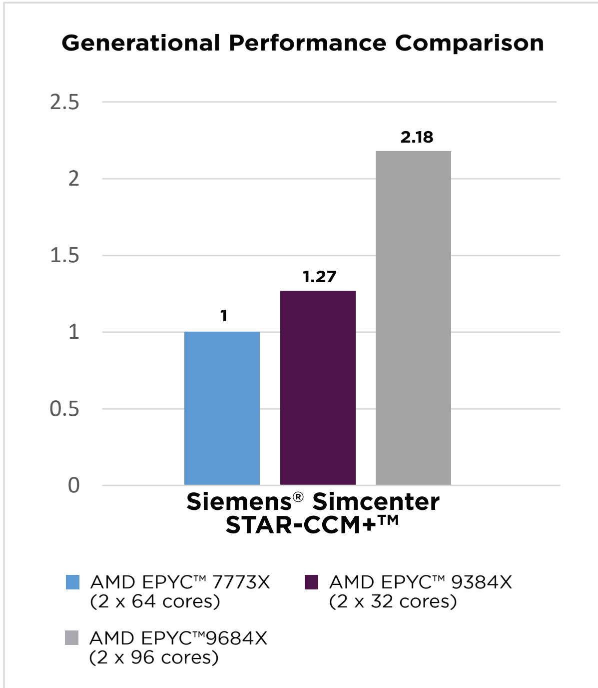

Figure 7: Excellent Performance for Simcenter STAR-CCM+ on the 4th Gen AMD EPYC™ Processors with 3D V-Cache[6]

Recommended Lenovo Configurations for CFD (Table 1): Lenovo ThinkSystem servers with AMD EPYC™ processors with AMD 3D V-Cache, 12 memory channels, and support for AVX-512 instructions can deliver high throughput per node for CFD applications since they benefit from multi-core parallelism and greater memory bandwidth.

Table 1: Sample Lenovo ThinkSystem Configurations for CFD

Suitable for	Server/Processor	Memory	Storage/Network
High Performance	• SR645 V3 (1U) • 2x AMD 9684X CPUs (96c, 2.2GHz, 400W TDP)	• 768GB Total RAM • 24x 32GB DDR5 4800MHz 2R DIMMs	• 1x480GB SATA Read Intensive • InfiniBand HDR100/Ethernet 100Gb 1-port adaptor
Balanced Performance	• SR645 V3 (1U) • 2x AMD 9384X CPUs (32c, 3.1GHz, 320W TDP)	• 384GB Total RAM • 24x 16GB DDR5 4800MHz 2R DIMMs	• 1x480GB SATA Read Intensive • 1 InfiniBand HDR100/Ethernet 100Gb 1-port adaptor
Performance per Watt per Dollar	• SR645 V3 (1U) • 2x AMD 9334 CPUs (32c, 2.7GHz, 210W TDP)	• 384GB Total RAM • 24x 16GB DDR5 4800MHz 2R DIMMs	• 1x480GB SATA Read Intensive • 1 InfiniBand HDR100/Ethernet 100Gb 1-port adaptor

Explicit FEA Applications: They are best suited for short-duration, highly nonlinear impact/deformation problems with complex contact conditions such as crash-worthiness/safety studies, metal forming, electronics drop tests, etc. Ansys LS-DYNA and SIMULIA Abaqus/Explicit are two leading applications widely used in the automotive and electronics industries.

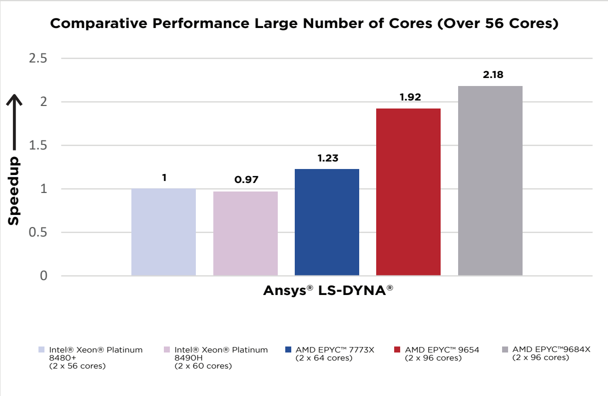

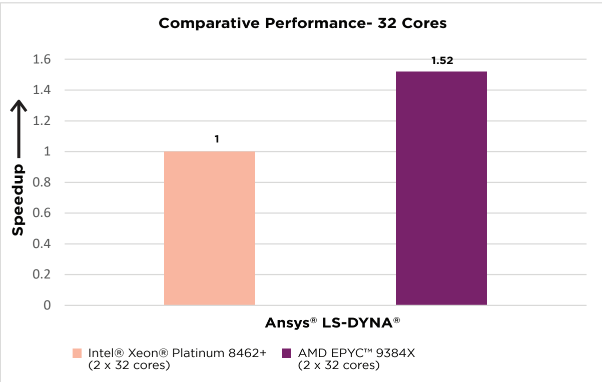

Figure 8: Excellent Ansys LS-DYNA Performance on the AMD EPYC™ 4th Gen Processor with 3D V-Cache[7]

For SIMULIA Abaqus/Explicit, Figure 9 shows up to 113% performance gains for the AMD EPYC™ 4th Gen processors with 3D V-Cache over the previous generation.

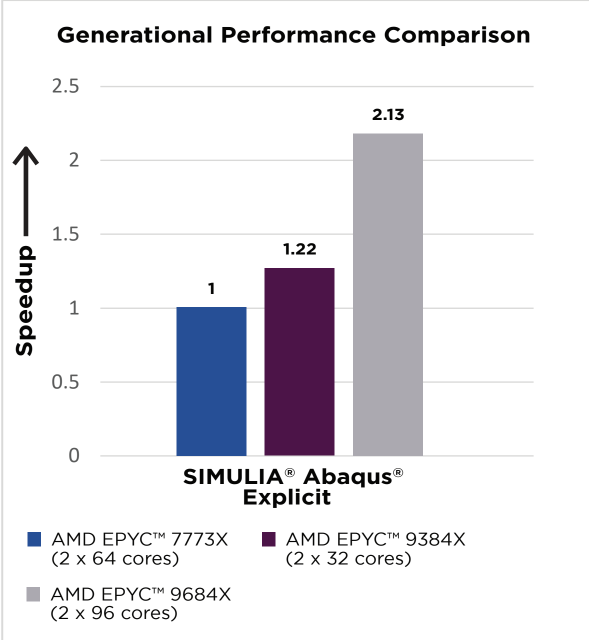

Figure 9: Excellent SIMULIA Abaqus/Explicit Performance on the AMD EPYCTM 4th Gen Processor with 3D V-Cache[8]

Recommended Lenovo Configurations for Explicit FEA (Table 2): Lenovo systems with medium- to large-core count 4th Gen EPYC™ processors with 3D V-Cache, high frequencies, and high cache-per-core and support for AVX-512 instructions offer very high performance per core to help efficiently utilize per-core software licenses.

Table 2: Sample Lenovo ThinkSystem Configurations for Explicit FEA

Suitable for	Server/Processor	Memory	Storage/Network
High Performance	• SR645 V3 (1U) • 2x AMD 9684X CPUs (96c, 2.2GHz, 400W TDP)	• 768GB Total RAM • 24x 32GB DDR5 4800MHz 2R DIMMs	• 1x480GB SATA Read Intensive • 1 InfiniBand HDR100/Ethernet 100Gb 1-port adaptor
Balanced Performance	• SR645 V3 (1U) • 2x AMD 9384X CPUs (32c, 3.1GHz, 320W TDP	• 384GB Total RAM • 24x 16GB DDR5 4800MHz 2R DIMMs	• 1x480GB SATA Read Intensive • 1 InfiniBand HDR100/Ethernet 100Gb 1-port adaptor
Performance per Watt per Dollar	• SR645 V3 (1U) • 2x AMD 9334 CPUs (32c, 2.7GHz, 210W TDP)	• 384GB Total RAM • 24x 16GB DDR5 4800MHz 2R DIMMs	• 1x480GB SATA Read Intensive • 1 InfiniBand HDR100/Ethernet 100Gb 1-port adaptor

Implicit FEA Applications: Best suited for simulating static and structural dynamic linear and nonlinear events. Ansys Mechanical and SIMULIA Abaqus/Standard are two prominent applications widely used in the manufacturing industry for various applications from thermal-stress analysis, design optimization, structural dynamics, etc.

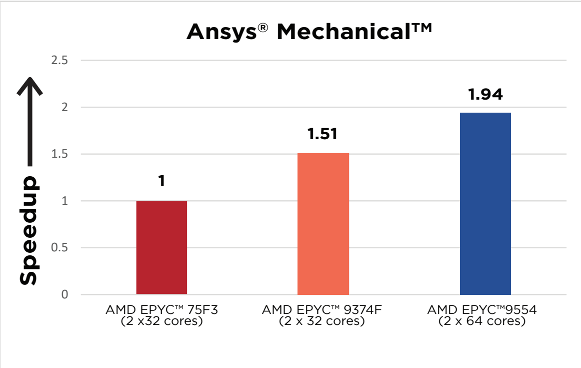

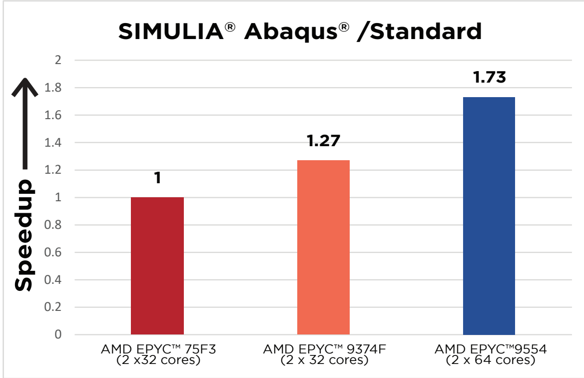

Figure 10: Excellent Performance on the AMD EPYC™ 4th Gen Processor for Ansys Mechanical[9] and SIMULIA Abaqus/Standard[10]

Figure 10 depicts the performance gains of up to 94% for Ansys Mechanical and up to 73% for SIMULIA Abaqus/-Standard on the AMD EPYC™ 4th Gen Processor compared to the previous generation.

Recommended Lenovo Configurations for Implicit FEA (Table 3): Lenovo ThinkSystem servers with lower-core count EPYC™ processors with high frequencies with support for AVX-512 instructions help efficiently utilize per-core software licenses and offer very high performance per core.

Table 3: Sample Lenovo ThinkSystem Configurations for Implicit FEA

Suitable for	Server/Processor	Memory	Storage/Network
High Performance	• SR645 V3 (1U) • 2x AMD 9384X CPUs (32c, 3.1GHz, 320W TDP)	• 768GB Total RAM • 24x 32GB DDR5 4800MHz 2R DIMMs	• 1x 2.5" 1.6TB SAS SSD (internal • 2x 2.5" 1.6TB NVMe PCIe (upgraded internal) • 1 InfiniBand HDR100/Ethernet 100Gb 1-port adaptor
Balanced Performance	• SR645 V3 (1U) • 2x AMD 9384X CPUs (32c, 3.1GHz, 320W TDP	• 768GB Total RAM • 24x 32GB DDR5 4800MHz 2R DIMMs	• 1x 2.5" 1.6TB SAS SSD (internal) • 3x 2.5" 1.6TB NVMe PCIe (upgraded internal) • 1 InfiniBand HDR100/Ethernet 100Gb 1-port adaptor

EDA Applications: Semiconductor manufacturers rely on EDA to compete and produce faster, cheaper, and more reliable devices with each generation. HPC and AI are used in most of the EDA workflow (Figure 11), with Verification and Regression typically consuming most of the compute capacity.

EDA workloads consist of two high-level design phases with a mix of sequential and threaded codes that are highly compute-intensive. The Front-End Design phase has mainly small files, which is more dominant. The Back-End phase typically has large files.

Front-End Design			Back-End Design			Manufacturing
Design Specification	**Verification/ Regression**	**Logic/Physical Synthesis**	**Physical Design**	**Physical Verification**	**Timing/SI/Power Signoff**	**Tape-out Release to Fab**
• Design entry (HDL or Transistor level)	Simulation • RTL/transistor • Gate Level • Formal/equivalency	Logic to Gates Mapping • Start converting code to actual circuits	Physical Implementation • Floor planning and block/cell placement • Routing	• DRC • LVS • ERC • Yeild checks and enhancement	• Parasitic Extraction • Timing • Noise/Signal Integrity • Power analysis (static/dynamic)	Released to foundry/fab • OPC • MDP • DRC/LVS • Yield analysis and enhancement

Figure 11: The EDA Workflow

Cadence Xcelium Digital Logic Simulation and Synopsys VCS RTL Simulation are two prominent applications widely used in the semiconductor device manufacturing industry for Verification. Both these EDA applications take full advantage of multicore processors to accelerate high-activity, long-cycle tests by allocating more cores at runtime.

Figure 12 depicts the performance gains of up to 40% for Cadence Xcelium Digital Logic Simulation and up to 73% for Synopsys VCS RTL Simulation on the AMD EPYC™ 4th Gen Processor with 3D V-Cache compared to a standard 4th Gen AMD EPYC™ Processor.

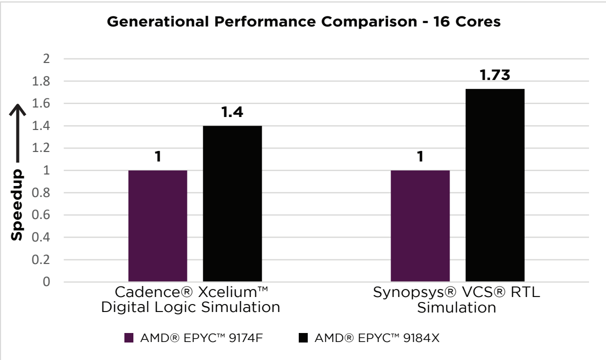

Figure 12: Higher Performance of Two Prominent EDA Applications with AMD 3D V-Cache Technology[11]

Recommended Lenovo Configurations for EDA (Table 4): Lenovo ThinkSystem servers with low core count, high base and boost frequency, large total L3 Cache, and dedicated L3 Cache/core help optimize per-core performance and lower expensive per-core software license costs. It is ideal for longer volume validation runs. Whereas for higher throughput per server, a higher core count is better.

Better Engineering Simulation Performance allows manufacturers to (Figure 13):

- Use fewer servers and less power, which helps lower carbon emissions and improve sustainability
- Maximize value from their application software license investment
- Lower the TCO of their solution
- Enhance the productivity of engineers and IT staff
- Improve the quality of simulations and solve larger, multi-physics problems
- Accelerate time-to-results to drive innovation, revenues, and profits.

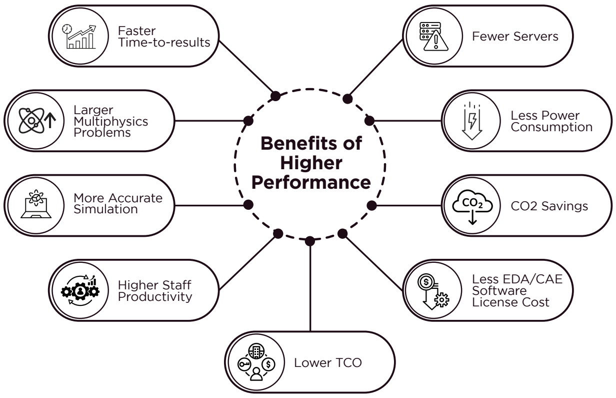

Figure 13: Benefits of Higher Performance for Engineering Simulation

In addition to getting performance-optimized systems, manufacturers can address their future requirements by partnering with Lenovo.

Table 4: Sample Lenovo ThinkSystem™ configurations for EDA

Suitable for	Server/Processor	Memory	Storage/Network
High Performance	• SR645 V3 (1U) • 2x AMD 9184X CPUs (16c, 3.55GHz, 320W TDP)	• 768GB Total RAM • 24x 32GB DDR5 4800MHz 2R DIMMs or 24x 64GB DDR5 4800MHz 2R DIMMs	• 1x 2.5" 1.6TB SAS SSD (internal) • 2x 2.5" 1.6TB NVMe PCIe (upgraded internal)
Balanced Performance	• SR645 V3 (1U) • 2x AMD 9384X CPUs (32c, 3.1GHz, 320W TDP)	• 768GB Total RAM • 24x 32GB DDR5 4800MHz 2R DIMMs or 24x 64GB DDR5 4800MHz 2R DIMMs	• 1x 2.5" 1.6TB SAS SSD (internal) • 3x 2.5" 1.6TB NVMe PCIe (upgraded internal)

Lenovo: Optimized for Today and Built for a Future of Energy Efficiency As technology rapidly evolves, manufacturers must integrate new technologies and workloads efficiently and seamlessly, often within resource, budget, capital, and power restrictions. Lenovo systems are optimized for today's CAE/EDA applications and designed for the future of manufacturing, where energy-efficiency issues are front and center.

Lenovo Innovative Energy Efficient Cooling: As processor frequencies and the number of cores increase to deliver the best performance, it is critical to cool these systems efficiently to avoid system overheating issues that cause shutdowns, slower performance, and potential data loss. EPYC™ 9004 Processor SKUs. If air-cooling is not feasible, clients can use other liquid-cooling technologies in the Lenovo Neptune portfolio.

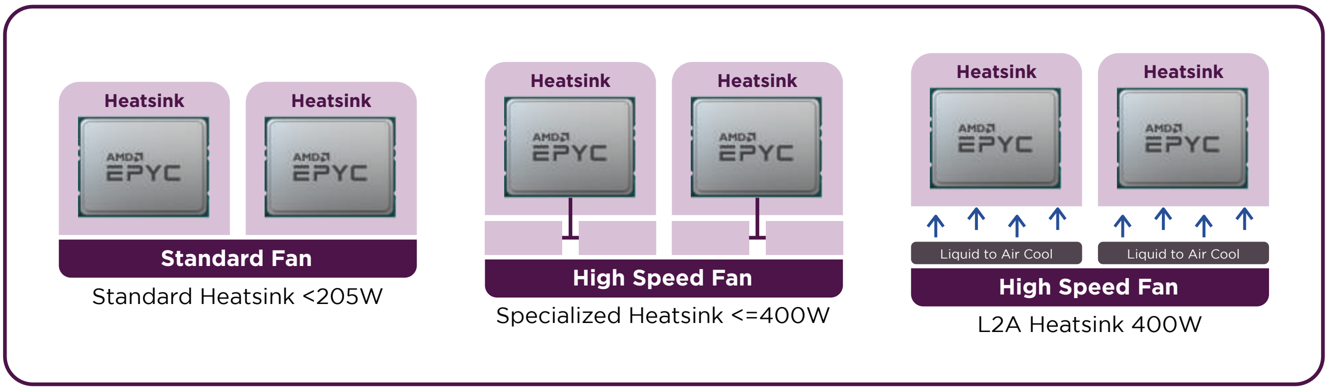

Figure 14: Schematic of Innovative Air-Cooling Options for the Lenovo ThinkSystem SR645 V3

To drive innovation and compete aggressively, manufacturers are integrating CAE/EDA (Figure 15) across the product lifecycle into a digital workflow from Product Development to Production to Operation. Other key technologies include Computer-Aided Design/Manufacturing (CAD/CAM), Augmented/Virtual Reality (AR/VR), the Internet of Things (IoT), AI/ML, and Digital Twins, which are digital replicas of business processes or actual components, systems, factories, cities, or even the entire planet.

Manufacturers can design and perform simulations in a digital twin to optimize business processes, improve results, drive sustainability, and gain a competitive advantage.

Lenovo Solutions: For a Digital Twin, as described earlier, manufacturers can use Lenovo Mobile or Desktop Workstations for CAD, Lenovo HPC and AI and Edge Servers for AI/ML, and the [Lenovo ThinkReality platform](), which provides a proven, scalable, and streamlined path from proof of concept to productivity for immersive AR/VR applications.

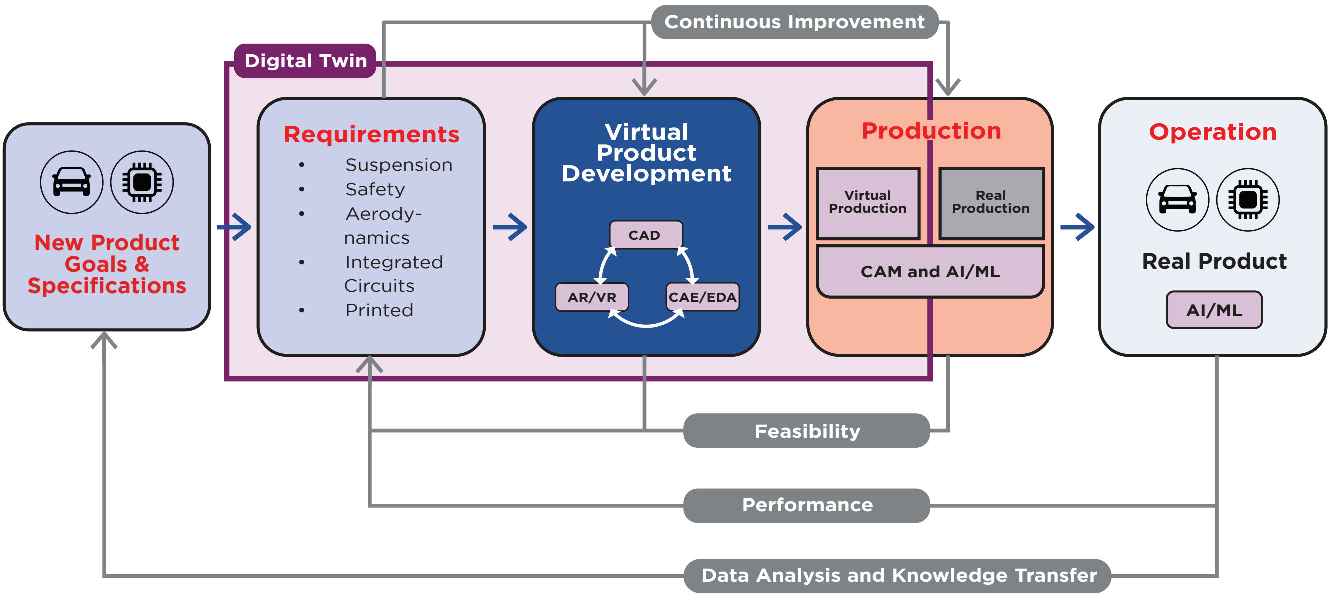

Figure 15: High-level End-to-end Development/Manufacturing Process and the Role of the Digital Twin

Get Started with Lenovo AMD HPC Manufacturing Solutions

Manufacturers rely on engineering simulation to design and refine their products. They cannot afford performance problems, delays, or downtime. Therefore, support must be proactive, carried out by technical specialists who work closely with the customer and deeply understand their environment.

As part of their Lenovo contract, manufacturers can receive a dedicated technical account manager or system admin to serve as their single point of contact. Whether onsite, working remotely, or a mixture of both, support professionals can quickly pinpoint and resolve any issues and ensure the engineering simulation environment runs optimally 24/7.

However, Lenovo goes way beyond specialized technical support. Lenovo's end-to-end service for EDA and CAE applications includes initial consultation, analysis, and configuring the right environment through ongoing cooling assessment and monitoring/maintenance services to billing and administration.

The Lenovo and AMD Advantage for Engineering Simulation

Lenovo and AMD help manufacturers maximize their return on investment (ROI) for engineering simulation, accelerate time to value to drive innovation and productivity, and drive sustainability by delivering:

- **Performance-Optimized Systems:** ThinkSystem Servers powered by AMD EPYC™ processors deliver the best performance([4, 5, 6, 7, 8, 9, 10, 11]) for many of the most demanding engineering simulation applications from a broad range of prominent ISVs widely used in the manufacturing industry.
- **Leadership in Energy Efficiency:** Lenovo's leadership in data center power and cooling technology and several innovative and unique air- and liquid-cooling solutions, including Neptune™ liquid cooling technologies.
- **Enterprise-level Support:** Systems are tested, validated, and optimized for performance, manageability, security, and scalability. Lenovo, or a certified business partner, provides onsite installation, start-up, integration, and proactive monitoring and remediation of any CAE/EDA operational issues.
- **A Complete Portfolio of Solutions:** Customers have an easy path from workstations to the most scalable ThinkSystem Servers. These systems come with a full range of storage, software, and comprehensive services that provide excellent performance, reliability, and security for a customer's IT environment from the edge to the data center to the cloud.
- **Solid Roadmap with Continuing Innovation:** The latest AMD EPYC™ processors with 3D V-Cache significantly boost performance for many engineering simulation problems over the previous generation. Likewise, Lenovo is delivering many leading-edge technologies in cloud computing (TruScale) and AR/VR (ThinkReality), which address manufacturers' future energy efficiency and immersive experience needs.

Finally, as every manufacturer has unique requirements, working with Lenovo and AMD will help them achieve a CAE/EDA solution tailored to their evolving workload, workflow, and workforce needs.

Optimize Your Engineering Environment

Please contact your Lenovo representative or visit
http://www.lenovo.com/amd-infrastructure

[1]All performance scores are estimates based on AMD internal testing in April and May 2023 and is shown in 'jobs/day/server'. TCO Analysis based on the AMD EPYC™ Bare Metal Greenhouse Gas TCO Tool v9.32 Pro. AMD processor pricing based on 1KU price as of May 2023. Intel® Xeon® Scalable CPU pricing from https://ark.intel.com as of April 2023. All pricing is in USD. This is a power only TCO with a time frame of 3-years. OpEx is power only, it does not include costs for real estate, admin and software. Power cost modeled @ $0.128/kWh with rack power of 10kW / rack having 10% reserved for non-server power use. The model uses a 1.7 PUE. Values are for USA. See endnote SP5TCO-045.

SP5TCO-045:
This scenario contains many assumptions and estimates and, while based on AMD internal research and best approximations, should be considered an example for information purposes only, and not used as a basis for decision making over actual testing. The Bare Metal Server Greenhouse Gas Emissions TCO (total cost of ownership) Estimator Tool compares the selected AMD EPYC™ and Intel CPU based server solutions required to deliver a TOTAL_PERFORMANCE score of 165000 Ansys Fluent-pump2 using the performance metric based on the published scores for Intel and AMD EPYC CPU based servers. This estimation reflects a 3-year time frame. This TCO model does not include any networking or storage that is external to the servers. PUE is a function of data center hardware power; power for devices external to the server is not included in this analysis. The CPU & Chassis Power in this analysis is modeled at 100% of TDP for EPYC powered servers, and at 100% for the Intel based servers. This analysis compares a 2P AMD EPYC_9384X powered server with a Ansys Fluent-pump2 score of 14481.53 per server; link and score details - based on AMD internal testing May 2023 on a reference platform: OS version RHEL 8.7; Kernel Version 4.18.0-425.3.1.el8.x86_64; BIOS version 1007D | TYM1009B; SMT Off, Network card: (MT4119 - MCX512A-ACAT) CX512A -ConnectX-5 fw 16.35.2000; InfiniBand: (MT4123 - MCX653106A-HDAT) ConnectX-6 VPI fw 20.35.2000; Storage: SAMSUNG MZQL21T9HCJR-00A07 NVMe.; to a 2P Intel Platinum_8462Y+ based server with a Ansys Fluent-pump2 score of 8006.2 per server; link and score details - based on AMD internal testing June 2023 using a Lenovo server with: OS version RHEL 8.7; Kernel Version Linux 4.18.0-425.3.1.el8.x86_64; BIOS version ESE110Q-1.10; Network cards:(MT4119 - MCX512A-ACAT) CX512A -ConnectX-5 fw 16.35.2000; InfiniBand: (MT4123 - MCX653106A-HDAT) ConnectX-6 VPI fw 20.35.2000; Storage: SAMSUNG MZQL21T9HCJR-00A07 NVMe. With a Ansys Fluent-pump2 Requests / Sec (RPS) of 14481.53 for AMD EPYC powered server versus 8006.2 for the Intel based server, AMD does 1.81 times the Requests per Second per server. With a RPS for EPYC of 14481.53 vs 8006.2 for the Intel solution, this means that EPYC is 81% faster than the Intel based server. Due the large variation in these costs these components: Admin costs are not included as part of this analysis. Real estate costs are not included as part of this analysis. OPERATING COSTS - OpEx: Both AMD EPYC and Intel based servers use the same cost for the following elements of the analysis: internal storage $380; server rack size of 42; space allowance per rack of 27 sq feet; ; cost per kW for power $0.128; power drop per rack of 12kW; and a PUE (power usage effectiveness) of 1.7 - PUE is Power Usage Effectiveness. With the AMD solution over the 3-years of this analysis, the EPYC powered servers are estimated to use: 531338.4 kWh with a total cost of $68011, for a total 3-year Power Only OpEx cost of $68011. For the Intel solution over the 3-years of this analysis, the Intel based servers are estimated to use 854790.6 kWh with a total cost of $109413, for a total 3-year Power Only OpEx cost of $109413. The annual Power Only OpEx for the AMD solution is $13801 less that the Intel solution making the AMD solution 38% less expensive annually. The annual power cost for the AMD servers is $22670 and the annual cost for the Intel servers is $36471 making the AMD solution $41402 or 38% less than the Intel based solution. AMD uses 531338.4 kWh over the life of the analysis where Intel uses 854790.6 kWh. The AMD solution uses 323452.2 or 38% less power. HARDWARE COSTS - CapEx: The EPYC powered solution is estimated to take: 12 total 2P AMD EPYC_9384X powered servers at a hardware only acquisition cost of $20285 per server, which includes $5529 per CPU, total system memory of 1536GB, which is 24GB of memory / core and a total system memory cost using DDR5 DIMMs of $6672; internal storage cost of $380.

The EPYC servers all use a 2S, 2RU server chassis, at $2175 per chassis. The EPYC servers all have a total of 2 Optional IO cards using a total power for all cards of 24.6 watts per hour and a total cost for all cards of $0. This makes the total CapEx (hardware cost) for the AMD solution $243420. The Intel based solution is estimated to take: 21 total 2P AMD Platinum_8462Y+ powered servers at a hardware only acquisition cost of $18933 per server, which includes $5945 per CPU, total system memory of 1024GB, which is 16GB of memory / core and a total system memory cost using DDR4 DIMMs of $4448; internal storage cost of $380. The EPYC servers all use a 2S, 2RU server chassis, at $2175 per chassis. The Intel servers all have a total of 2 Optional IO cards using a total power for all cards of 24.6 watts per hour and a total cost for all cards of $0. This makes the total CapEx (hardware cost) for the Intel solution $397593. The total CapEx for the AMD solution is $154173 less that the Intel solution making the AMD solution 39% lower CapEx than the Intel based solution. TCO ANALYSIS: Delivering 165000 estimated score of Ansys Fluent-pump2 performance produces the following estimated results: the AMD EPYC solution requires 43% fewer servers [1-(AMD server count / Intel server count)]; 43% fewer sockets; 43% fewer cores; 50% less space [1-(AMD rack count / Intel rack count)]; 38% less power [1-(AMD power cost / Intel power cost)]. For the 3-years of this analysis, AMD's TCO (OpEx: $68011.3152 plus CapEx of $243420) is $311431. The Intel TCO (OpEx: $109413.1968 plus CapEx of $397593) is $507006. This makes the AMD EPYC solution TCO $195575 or 39% lower that the Intel TCO.GREENHOUSE GAS: The AMD EPYC_9384X powered servers save ~323452.2kWh of electricity for the 3-years of this analysis. Leveraging this data, using the Country / Region specific electricity factors from the '2020 Grid Electricity Emissions Factors v1.4 – September 2020', and the United States Environmental Protection Agency 'Greenhouse Gas Equivalencies Calculator', the AMD EPYC powered server saves ~120.69 Metric Tons of CO2 equivalents, which equals 133.04 US tons or 44.35 US tons of CO2 equivalents savings annually.

[2]https://www.lenovo.com/us/en/resources/streams/hpc-manufacturing/?storypopup=hpc-manufacturing-amd-ansys-lenovo-solution-brief

[3]Lenovo Ansys (amd.com)

[4]https://www.amd.com/system/files/documents/amd-epyc-9004x-pb-ansys-cfx.pdf

[5]https://www.amd.com/system/files/documents/amd-epyc-9004x-pb-ansys-fluent.pdf

[6]https://www.amd.com/content/dam/amd/en/documents/epyc-business-docs/performance-briefs/amd-epyc-9004x-pb-siemens-star-ccm.pdf

[7]https://www.amd.com/system/files/documents/amd-epyc-9004x-pb-ansys-ls-dyna.pdf

[8]https://www.amd.com/system/files/documents/amd-epyc-9004x-pb-dassault-abaqus-explicit.pdf

[9]Ansys® Applications Generational Performance Uplifts with 4th Gen AMD EPYC™ Processors

[10]SIMULIA® Applications (amd.com)

[10]SP5-050: EDA RTL Simulation comparison based on AMD internal testing completed on 4/13/2023 measuring the average time to complete a graphics card test case simulation. comparing: 1x 16C EPYC™ 9384X with AMD 3D V-Cache Technology versus 1x 16C AMD EPYC™ 9174F on the same AMD "Titanite" reference platform. Results may vary based on factors including silicon version, hardware and software configuration and driver versions.